Entry and Advertising on Digital Platforms*

Evidence from a Large E-Commerce Platform

Yizhou Jin and Zhengyun Sun

March 11, 2026

Abstract

Millions of businesses join digital platforms each year, yet little is known about their success rates or how platforms manage their entry. Using data from a leading Chinese e-commerce platform, we show that fewer than 1% of new, ID-verified sellers become successful small-to-medium-sized businesses. Analysis of search sessions reveals that consumers prefer entrant products over incumbent ones. Further, increasing entrant exposure on the margin increases overall purchase rate of a search session, indicating that entrants are under-explored compared to the social optimum. Finally, we revisit a prior intervention on the platform that increased entrant advertising. We find that it promoted entry with minimal impact on screening efficiency. A model of platform-mediated entry explains our findings and predicts how platform maturity, entrant quality, and advertising jointly determine optimal exploration.

Keywords: digital platforms, e-commerce, entry funnel, advertising, consideration set
JEL codes: L26, M53, D22

*Jin: University of Toronto (yizhou.jin@utoronto.ca); Sun: Uber (zhengyunsun@uber.com). We thank Dr. Long Chen and the team at our data providing firm. We thank Harvard Business School, Harvard Economics Department, and the University of Toronto for generous financial support. We thank Jie Bai, Heski Bar-Isaac, Emily Breza, Shawn Cole, Andrey Fradkin, Joshua Gans, Hong Luo, Asim Khwaja, Gabriel Kreindler, Michael Kremer, Song Ma, Ariel Pakes, Gautam Rao, Yulu Tang, Yuan Tian, Christopher Woodruff, Chuan Yu, Florian Zettelmeyer, participants at the Utah Winter Business Economics Conference, the Federal Trade Commission, the American Economic Association Meeting and various seminars. Ability to publish is not contingent on results (data usage agreement contact: carolina_harvey@harvard.edu).

1 Introduction

Every day, tens of thousands of new sellers appear on major e-commerce platforms.[1] Platform designers must decide how to assess entrant quality, often without reliable signals such as historical reviews (Tadelis, 2016; Farronato, 2025). When the volume of new offerings is small, platforms can *explore* quality by steering consumers to "sample" them or by soliciting bids from sellers on advertising keywords (Chen and He, 2011; Athey and Ellison, 2011). However, as the volume of new entrants grows and platform competition intensifies (Aridor et al., 2020), these mechanisms can become increasingly constrained, a challenge compounded by AI advances that widen the data advantage of incumbents over entrants (Milgrom and Tadelis, 2018).

Experiments on e-commerce and service platforms have revealed that entrants may be systematically under-explored: when exogenously sampled by researchers, new suppliers saw higher subsequent growth and survival rates (Pallais, 2014; Bai et al., 2020). Recent work suggests that many entrants rationally bid below the cost of acquiring reputation (Cui, 2025). Similarly, established businesses tend to advertise far more than entrants (Moshary, 2021), even though newer and lesser-known firms may benefit more from advertising (Dai, Kim and Luca, 2023; Sahni and Zhang, 2024). However, these findings draw on new businesses that have already earned revenue—likely a small fraction of all potential entrants.[2] More generally, no systematic account of entry outcomes on these platforms yet exists, nor is it clear whether the binding constraint is insufficient exploration or insufficient quality.

In this study, we use administrative data from a major e-commerce platform to document the systematic under-exploration of entrants and analyze a prior experiment to measure how advertising helps mitigate this initial disadvantage. We then develop a model of platform-mediated consumer search and seller advertising to formalize these findings, showing how organic screening, market maturity, and advertising frictions jointly affect the platform's exploration of entrant quality.

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

Second, entrants are subject to more intense screening than incumbents and are under-explored. To establish this, we construct a consideration-set panel from clickstream data associated with search sessions, each linking a consumer to a set of search keywords and the products clicked. Consumers are more likely to purchase from entrants within a consideration set; entrants also see fewer returns and more repeat purchases. Across nearly identical consideration sets, each additional entrant product raises the session's purchase

[1] Major platforms rarely disclose entry volumes or survival rates. Secondary sources estimate that in 2020, about 1.3 million new sellers joined Amazon Marketplace, 6.8 million joined Taobao and Tmall, and YouTube added 7 million new channels, per Marketplace Pulse and Socialblade.

[2] "The top 1% of sellers have been selling on Amazon for years and drive at least half of the sales volume. The thousands that join daily don't have the same experience, many never get a single order" (Marketplace Pulse, 2024). On TikTok Shop, "new entrants remain restricted from reaching creators… It's a classic catch-22: you need affiliates (creators) to grow, but you can't access them until you've already grown" (Perez-Tenessa, Trendio Blog, 2025).

probability by 7%: the platform would expand the market by steering more search toward entrants.

Third, to understand how advertising affects entry, we leverage an experiment on the platform that provided free online training to over two million new sellers (Jin and Sun, 2026). Training had precise null effects on all main seller strategies except advertising (+2.3%) and generated a 1.3% increase in traffic among treated sellers, entirely driven by ads during the period overlapping training. Per 10,000 potential entrants, this produced 28 additional sellers with meaningful traffic and 11 with positive revenue. Repeating the consideration-set analyses across treatment and control groups, we find that advertising promoted entry with minimal impact on screening efficiency: the marginal entrants reached through advertising were slightly lower quality but remained substantially superior to incumbents.

To discipline our interpretation of these findings, we extend the equilibrium model of Athey and Ellison (2011) to characterize how platform ranking screens entrants through organic search alone, and how this screening interacts with exploration and advertising. The model yields three main insights: (1) diminishing returns make wide-but-shallow exploration optimal—a small exploration rate applied broadly identifies superstars without deep per-entrant investment, and declines as the platform matures; (2) this exploration dilutes baseline screening, so high observed entrant conversion rates can indicate a lack of exploration; and (3) advertising can substitute for exploration as a screening device, but incompletely; information frictions and incumbent blocking preserve the need for platform-directed exploration.

Related Literature. This study contributes to the literature on entry, exploration, and advertising on digital platforms. We are, to our knowledge, the first to document the entry funnel on any major digital platform and to explicitly compare entrants and incumbents using search session data. Closely related is Bai et al. (2025), who show how congestion slows the organic discovery of high-quality sellers. We complement their findings with both a more macro (the entry funnel) and a more micro (search sessions) lens. We also explicitly model how platform ranking, exploration, and advertising jointly steer consumer search of new sellers. Moreover, we construct consideration-set data to control for unobserved demand heterogeneity, mapping directly to stylized search models such as Athey and Ellison (2011) and demonstrating a novel use of clickstream and pre-purchase data (Honka, Seiler and Ursu, 2024; Ursu, Seiler and Honka, 2025).

By identifying exploration as a key bottleneck driving entry difficulty, our work bridges to the literature on social learning and experimentation on platforms under quality uncertainty.[3] A central prediction is that platforms should "start small," gradually scaling exposure as credibility builds (Che and Hörner, 2018). Our entry-funnel and consideration-set data provide direct validation: exploration has indeed widened at the top of the funnel but shallowed over time, yet in practice is more constrained than the model anticipates. Our work connects to a large body of work on the design of reviews (Chevalier and Mayzlin, 2006; Luca, 2016; Tadelis, 2016).[4] However, our results suggest that most entrants on our platform never accumulate enough

[3] A distinct literature examines platforms' ranking, auction, and recommendation design under complete information of seller quality (theory: Hagiu and Wright (2024); Nocke and Rey (2024); Anderson and Renault (2025); Janssen et al. (2026) and empirics: Dinerstein et al. (2018); Lee and Musolff (2021); Lam (2021); Farronato, Fradkin and MacKay (2023); Waldfogel (2024)). In these settings, platforms control the visibility of product substitutes, distort organic rankings, or use position auctions not to acquire information, but to exploit known product characteristics to maximize rent extraction, coordinate seller positioning, or induce search. These strategies have implications for but are not driven by new-seller entry.

[4] Normatively, suppressing incumbent reviews (Vellodi, 2018), obfuscating early negative signals (Papanastasiou, Bimpikis and Savva, 2018), or offering price subsidies (Newberry, 2016; Dendorfer and Seibel, 2025) can further improve the efficiency of learning.

traffic to generate the reviews these mechanisms require.

Further, we provide large-scale evidence that advertising can promote entry and substitute, to a limited degree, for past consumer experience (Shapiro, 1982). Our focus on entry extends the literature on how advertising interacts with other platform design levers, including organic ranking (Berman and Katona, 2013) and commissions (Bar-Isaac and Shelegia, 2022; Yu, 2024). It also complements work on the signaling value of ads, which has yielded mixed results that "may not be true for small and new entities" (Blake, Nosko and Tadelis, 2015).[5] Finally, we complement research on cold-start subsidies in ad auctions (Ye et al., 2023) by demonstrating the broader relevance of such mechanisms for entry when most entrants never advertise.

After describing our setting and the entry funnel (Section 2), we present consideration-set evidence for the under-exploration of entrants (Section 3), leverage an experimental variation in entrant advertising to analyze its effect on entry (Section 4), develop a model of platform-mediated search and advertising that interprets the empirical findings (Section 5), and conclude (Section 6).

2 Background and the Entry Funnel

2.1 Platform Business Model and Data

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

Platform business model. Most sellers on the platform are third-party retailers. The platform does not directly sell to consumers in almost all sectors, with some notable exceptions such as fresh groceries and financial services. Our analysis excludes these sectors so that the platform only intermediates—and does not compete with—sellers. In this setting, platform revenue primarily comes from in-platform advertising (approximately 80%),[6] with the remaining share made up by sales commissions for large sellers and seller services such as return shipping insurance and advanced data analytics (Jin and Sun, 2019).

Almost all traffic to new sellers comes from search: consumers arrive at the platform and search for a desired product; the platform then parses the search query into keywords and ranks a set of related products and sellers.[7] Sellers can boost their rankings organically or through in-platform advertising. Strategies for boosting organic traffic mainly aim to increase conversion rates and positive reviews, for example by giving discounts or improving product listings (Lewis, 2011).

Sellers can advertise in-platform via sponsored product listings by submitting a daily budget, targeted keywords, and a per-click bid for each keyword and customer segment. In each consumer search session, the platform ranks all bids associated with the search keywords, determines the winning bids, and charges the corresponding sellers. The associated products are then displayed in sponsored positions regardless of

[5] Sahni and Nair (2020); Long, Jerath and Sarvary (2022) have also found evidence in favor of ads signaling, contrasting with Moshary (2021). Existing work tends to rely on experiments within established brands or buyer-level randomization that primarily affects those brands.

[6] This is not unique to our setting. For example, Google and Taobao both make over 80% of their revenues from ads in 2022. Since 2017, Amazon Marketplace has seen a 52% annualized increase in ads revenue.

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

their organic search ranking. This is similar to the mechanisms analyzed in Edelman, Ostrovsky and Schwarz (2007). Although advertising is available to all sellers, adoption among new sellers during their first year is only 7% (Jin and Sun, 2026), whereas the majority of established sellers advertise.[8]

Data. Our analyses use platform administrative data. The first dataset spans 2015 to 2019, including proportional aggregate statistics on all new sellers on the platform.[9] We describe each statistic in detail below as we introduce the entry funnel.

A second consideration-set dataset is constructed from an experiment on the platform (Jin and Sun, 2026) that covered all new sellers starting May 6, 2019. We describe the data and the experiment in Sections 3.1 and 4.1, respectively.

Although we observe the actions and outcomes of sellers and buyers in great detail, the platform's proprietary algorithms—such as product ranking within a search session or the bids in sponsored listings—are not available to us.

2.2 The Entry Funnel: Defined and Quantified

Prospective sellers undergo a rigorous verification process to register on the platform. The platform requires a national ID or business registry number and cross-references it against the national digital identity database. The platform also collects a small deposit to deter spam. Sellers that successfully registered are considered **potential entrants**. In the following year, those having earned non-zero traffic are considered **explored**; those with non-zero revenue are considered **entrants**. On the intensive margin, we define entrants that are ranked at the top 10% among all sellers by monthly revenue as small-and-medium-sized businesses (**SMBs**).[10] The entry funnel from potential entrants to established brands is illustrated by Figure 3 panel (a).

In 2019, the platform had tens of thousands of potential entrants per *day*. As shown in Table 1, per 10,000 potential entrants, just over half received traffic (were "explored"), and 15% became entrants, producing a long tail of zero-revenue sellers similar to that studied in Dubé, Hortaçsu and Joo (2021). In the first month after registration, only 62 became SMBs, and about two-thirds remained twelve months later. Figure 3 panel (b) redraws the illustrative funnel in panel (a) to reflect the realistic odds.

Over time, the entry funnel widened at the top: new sellers were more likely to receive at least some traffic. However, the depth of exploration diminished. By 2019, far fewer sellers had ever received over 100 unique visitors in their first year. Given that the average seller-level conversion rate is well below 1%, 100 visits provide far too little experimentation to produce a granular quality signal. Widening the surface of exploration spreads traffic more thinly, allowing only those with the strongest early signals to emerge. The lower parts of the funnel mirror this pattern: entry and SMB rates grew initially but sharply contracted by

[8] On Amazon, nearly three-quarters of brands and sellers use Sponsored Products advertising (Jungle Scout *Amazon Advertising Report*, 2024). Given that our platform derives an even larger share of revenue from advertising, incumbent participation is likely at least as high.

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

[10] This threshold is set by the platform; the corresponding revenue levels are dynamic and proprietary. However, reasonably, they correspond to monthly sales of around one thousand Yuan. More importantly, they highlight the comparison with established sellers: given the large stock of established sellers, we can assume that 10% of them are SMBs. The entry funnel statistics in this section are derived from aggregated administrative data; our data agreement does not permit us to disclose further detail.

Table 1: Entry Funnel from 2015 to 2019

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

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

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

These patterns are consistent with theories such as Che and Hörner (2018), which predict broad but shallow exploration: "starting small" by directing minimal traffic to untested products. This prediction mirrors what we find in Table 1: the platform's exploration has widened at the top of the funnel but "shallowed" over time. However, as entry volume has grown, the shallowing has become severe enough that most explored sellers do not receive enough traffic to generate meaningful quality signals through consumer search.[11]

3 Measuring Entrant Under-Exploration: Consideration-Set Analysis

The narrow entry funnel suggests that, despite the ease of setting up an online store, most potential entrants will find it hard to access consumers and grow their revenues. The narrowing of the entry funnel over time further disadvantages later potential entrants compared to their predecessors. Our focus henceforth is on the exploration of potential entrants; we refer to all potential entrants simply as "entrants" or "new sellers."

The exploration of entrants is ultimately achieved through consumer search steered by the platform. A

[11] Che and Hörner further show that sustained exploration requires the platform to "pool" recommendations, concealing whether a product is shown for proven quality or merely for exploration. This pooling strategy is particularly constrained when signals such as cumulative sales volumes are publicly visible, as on our platform.

key empirical challenge is demand heterogeneity: consumers who search entrants may differ in unobservable ways from those who search incumbents. We address this heterogeneity using consideration-set data that allow us to compare new and incumbent sellers searched by the same consumer in the same session. We first compare purchase likelihood within consideration sets containing both new and incumbent sellers; we then compare across nearly identical sets—matched on consumer, search keywords, session size, and date—that differ in whether they contain new sellers. Since we do not observe search rankings or advertising prices, the goal of our analysis is to identify the state of entrant under-exploration rather than structurally separating the roles of platform steering and consumer search behavior.

3.1 Consideration-Set Data

A Platform-Wide Experiment. The consideration-set data comes from an experiment on the platform (Jin and Sun, 2026), in which a random subset of all newly registered sellers were given access to a digital training program. The research window is roughly a three-month period starting from May 6, 2019, when the experiment started, until August 15, 2019. This dataset is appropriate in the present study of seller entry at large because we can source a representative subset of all new and incumbent sellers. The experiment also caused an exogenous shock to illustrate whether and how advertising can promote entry, which we study in Section 4. Nonetheless, the scale of the intervention and the magnitude of the treatment effect (detailed in Section 4.1) was small enough that it did not qualitatively alter the entry process.

The Dataset. We construct consideration-set data based on records from search sessions during the research window. In each session, we observe an anonymized consumer identifier and the keywords parsed from the search query. The *consideration set* consists of all products the consumer had clicked on, each corresponding to one seller. We also observe the product the consumer bought, if any. For a given consumer, we aggregate search sessions with the same keyword combination that occurred on the same calendar day. Our approach differs from prior papers such as Lam (2021), in which the consideration set probabilities are structurally derived from an optimal sequential consumer search process in the absence of clickstream data.

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

On the seller level, we observe some time-invariant characteristics, including the registration date (which categorizes the seller as new or incumbent), treatment assignment, location based on mailing address, and some basic characteristics of the owner. We also observe a set of time-varying variables describing sellers' actions or states, including the number of product listings, and a composite seller rating across products. To account for heterogeneity across product categories, we develop "sector" dummies based on sellers' product portfolios.[12]

Sampling Procedures. Two features of the raw data require sub-sampling: the sheer volume (over 10 million visits per day to new sellers alone) and extreme skewness (new sellers appeared in less than 0.1%

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

of consideration sets). Starting from a random subsample of consumers who visited new sellers during the research window, we pull their search sessions—each consisting of a consumer, parsed search keywords, products clicked, and any resulting transaction—filtered to popular keyword combinations. We then match in additional sessions by non-sampled consumers with identical keyword-date-size combinations, enabling the across-consideration-set comparisons in Section 3.3 via interacted fixed effects. We restrict to sessions with at least three products to focus on genuine consumer searches.[13] Full details and a schematic illustration of the sampling procedure are in Appendix B.1 and Figure 9.

Summary Statistics. Our dataset contains 6.0 million consideration sets (Table 2), each consisting of 4.6 sellers on average. A total of 65,246 distinct keywords were searched across all major sectors. Despite over-sampling, 82% consideration sets still had no new sellers. These are omitted when analyzing differential seller purchase likelihoods within consideration sets in Section 3.2 ($N = 1,084,228$); we introduce them back in when analyzing differential outcomes across consideration sets in Section 3.3. There, when analyzing refunds and repeat purchases, we focus on consideration sets that result in a purchase ($N = 289,258$).

Consumers who visited new sellers are observably different. In particular, they seem to search more sellers regardless of treatment status (Figure 8). This illustrates how demand heterogeneity (such as consumer search cost or purchase intent) can bias quality inference using raw purchase or conversion rates.

Table 2: Summary Statistics: Consideration Set Sample

Statistic	Mean	St. Dev.	Min	Pctl(25)	Pctl(99)	Max
Number of Sellers Browsed	4.6	2.7	3	3	12	55
% of New Sellers	0.57	8.7	0	0	33	100
Purchased (on the same day)	0.180	0.390	0	0	0	1
Purchased from New Sellers	0.001	0.033	0	0	0	1
Pay Amount (USD)	4.8	21.2	0	0	753	2,050
Avg. Seller Price Level (USD)	34.4	86.3	0.0	8.8	582	9,674
Avg. Seller Rating	13.0	2.9	0	11.5	20	20

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

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

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

$$Y_{js} = \tau \mathbf{New}_j + \mathbb{X}_j \gamma + \alpha_s + \varepsilon_{js} \tag{1}$$

The outcome variable is an indicator for a consideration set s, or a search session, resulting in a purchase from seller j. The coefficient of interest, τ, measures the differences in purchase likelihood between new

[13] Sessions with one or two sellers make up over 42% of all sessions yet account for only 1.9% of purchases, likely reflecting "passive searches" (Ursu, Simonov and An, 2022) and test queries.

sellers, defined as those that registered during our research window, and incumbents. α_s, the consideration-set fixed effects, soak up any unobserved differences across consumers and across each consumer's search sessions.

Table 3, Panel A, reports the estimates: within a consideration set, new sellers had significantly higher purchase probability than incumbents, by 1.3 percentage points (7.6% of the incumbent mean). This is inconsistent with substantial bandit exploration, under which an entrant would be explored only when the expected quality of incumbents is even lower.

3.3 Across Consideration-Set Analysis: with vs. without New Sellers

The analysis above compares seller outcomes within the same consideration sets, controlling for unobserved demand heterogeneity. However, requiring within-set variation filters the sample to consideration sets containing both new and incumbent sellers. In this section, we remove this filter and compare across *similar* consideration sets that may or may not contain new sellers.

Market expansion effects. Lastly, we relax all data filters and include a majority of consideration sets that neither contains new sellers nor resulted in a purchase. Doing so capitalizes on a novel property of our data: the ability to observe the *potential* market, defined as all consumers who searched at least one seller regardless of whether they made a purchase. This lets us measure potential "market expansion effects"—the degree to which marginally increasing the search likelihood of entrants expands the number of matches.

Refunds and repeat purchases. We next examine whether the higher purchase probability of new sellers is accompanied by post-purchase quality differences. Focusing on consideration sets that resulted in purchases, we test whether buying from new sellers is associated with different rates of refunds or repeat purchases.

Unlike Equation 1, where the unit of analysis is a seller within a consideration set, the unit here is the consideration set itself: each observation corresponds to a search session that resulted in a purchase. The main independent variable is an indicator of whether the consideration set contains new sellers, $\mathbf{New}_s$, and we replace the consideration-set fixed effects with consumer fixed effects $\alpha_{i(s)}$ and search *keyword-size-date* interacted fixed effects IFE_s, as introduced in Section 3.1. We also average over observable characteristics of all sellers in the consideration set to obtain $\mathbb{X}_s$—number of products offered, average price, consumer star rating, and average order size—and further control for differences across search sessions. The dependent variable is an indicator for return or refund request, or an indicator for the consumer making repeat purchases from the seller from which the original purchase was made.

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

Table 3, Panel B, reports the estimates. Across nearly identical search sessions, buying from new sellers is associated with lower likelihood of returns and higher likelihood of repeat purchases. These results strengthen the case for entrant under-exploration.

Table 3: Regression Results: Consumer-Side Analyses

	(1)	(2)	(3)	(4)
Panel A: Purchase Conditional on Search				
	Has Purchased		Log Spending	
Is New Seller	0.013^{***}	0.003^{***}	0.056^{***}	0.010^{***}
	(0.001)	(0.001)	(0.004)	(0.004)
Incumbent Mean	0.17		0.71	
R^2	0.053	0.073	0.054	0.078
Panel B: Post-Purchase Experience				
	Is Returned		Has Repeat Purchased	
Has New Seller	-0.006^{***}	-0.004^{***}	0.006^{***}	0.004^{**}
	(0.002)	(0.001)	(0.002)	(0.002)
Incumbent Mean	0.078		0.064	
R^2	0.398	0.463	0.301	0.378
Panel C: Market Expansion Effects				
	Has Purchased		Log Spending	
Number of New Sellers	0.003^{***}	0.002^{***}	0.011^{***}	0.008^{***}
	(0.001)	(0.001)	(0.003)	(0.003)
Incumbent Mean	0.041		0.17	
R^2	0.053	0.073	0.054	0.078
Seller Controls	N	Y	N	Y

Note: Standard errors in parentheses. Panel A reports estimates from Equation 1. Panel B reports estimates from Equation 2. Panel C reports estimates from Equation 3. Odd (even) columns exclude (include) seller controls: number of products offered, average price, consumer star rating, and average order size.
Panel A: fixed effects: consideration-set. Unit: seller. Sample: consideration sets with ≥ 1 new seller. $N = 1{,}084{,}228$. Panel B: fixed effects: consumer, intent-effort-date. Unit: consideration set. Sample: consideration sets resulting in purchases. $N = 289{,}259$. Panel C: fixed effects: consumer, intent-effort-date. Unit: consideration set. Sample: all consideration sets. $N = 6{,}007{,}329$. * p<.1, ** p<.05, *** p<.01.

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

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

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

Table 3, Panel C, shows that across similar consideration sets, searching new sellers instead of incumbents increased purchase probability by 0.3 percentage points (7.3% of the incumbent-only baseline) and dollar spent. These results imply that marginally increasing entrant search can expand the market, confirming an informational entry barrier in the search process.

Discussion: Entrant Under-Exploration. Despite receiving very little traffic, entrants face more intense screening than incumbents: they convert at higher rates within the same consideration set. This is inconsistent with bandit exploration, which would randomly sample entrants and thus yield *lower* conversion than incumbents with better known quality. Moreover, because our fixed effects control for the number of searches in addition to demand heterogeneity, the difference in consideration-set-level conversion rates directly maps to differences in total surplus, a point that we revisit in Section 5.4 combined with our model. Steering consumers to more entrants would thus expand the market on the margin, indicating that the platform under-explores entrants relative to the social optimum.

When consumer-search-based exploration is this limited, theory prescribes that platforms rely on costly internally generated signals (Che and Hörner, 2018); in ad-supported platforms such as ours, paying for advertising, through which entrants effectively "purchase sampling," is a natural such mechanism (Shapiro, 1982). In Section 5, we extend the equilibrium model of Athey and Ellison (2011) to characterize how platform ranking leads to intense baseline screening of entrants, and how this screening interacts with platform exploration and advertising.

4 Can Advertising Promote Entry?

How advertising affects entry is ultimately an empirical question. While advertising can help entrants "purchase sampling" and overcome cold-start frictions, it can also exclude liquidity-constrained entrants, and platforms may foster an advertising arms race that compromises signaling value (Crémer, De Montjoye and Schweitzer, 2019). In practice, entrants rarely advertise, and at far lower rates than incumbents, making it important to understand whether increased advertising can promote entry.

4.1 A Training Experiment that Increased New-Seller Advertising

Jin and Sun (2026) studied a digital training program targeting new sellers on the platform. From May 2019 to June 2020, a subset of all new sellers received access to the program. Among potential entrants, trained ones saw 1.7% higher revenues, almost entirely driven by increased advertising: 2.3% on the extensive margin and 3.0% overall. Traffic from ads increased by 7.6% during the first three months after registration, when the vast majority of training occurred, and 4% over the full 9-month period.

We refer readers to the original paper for more details but highlight two points. First, in terms of seller strategies, a key result is a null one: despite a broad training coverage, training had precise zero effect on sellers' product variety, pricing, and ratings.[14] We thus treat this experiment as an exogenous shift in new-seller advertising.

Second, recent studies have questioned the validity of seller-level randomization on digital platforms due to potential SUTVA violations. However, the treatment affected entrants only, who, despite their large

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

number, had little impact on product or advertising equilibrium. Treated new sellers did not directly compete with control-group ones.

4.2 Are Entrants Still Under-Explored?

Although treated new entrants received more traffic and earned higher revenues, the aggregate impact on the entry funnel was diluted by low adoption and a low baseline advertising rate—only 7% of new sellers ever advertised, with an average of 1.57 sponsored products (Jin and Sun, 2026). We quantify this by regressing funnel-stage indicators on the treatment assignment, as shown in Table 4. Per 10,000 potential entrants, the experimental increase in advertising produced 28 additional sellers with meaningful traffic (from a baseline of 2,280) and 11 additional entrants with positive revenue (from a baseline of 1,525). The point estimates suggest effects at deeper funnel stages as well: per 10,000 potential entrants, roughly 1–2 additional sellers reached SMB status, with the month-9 survival effect statistically significant (from baselines of 50 and 40, respectively). Entrants that did enter also received more traffic and earned higher revenues, widening the intensive margin of the funnel.

Table 5, Panel C, which maps to Table 3, Panel C, tells a similar story: the market expansion effect (of marginally increasing entrant exposure) is still prominent in the treatment group, albeit slightly—and statistically insignificantly—weakened. To assess whether the additional advertising changed the composition or screening of entrants, we replicate the regression analyses in Section 3 but focus on testing the differences across the treatment and control groups. Table 5, Panels A and B, map to Table 3, Panels A and B, respectively. Directionally, they suggest that the quality of entrants searched by consumers deteriorated slightly with more advertising in the treatment group. The magnitude of the differential, however, is negligible and statistically insignificant (except a higher return frequency), meaning that the entrants in the treated group remain substantially superior (in terms of purchase probability conditional on search) compared to incumbents.

Table 4: Intent-to-Treat Effects on Entry Funnel Stages (All Potential Entrants)

	Dependent variable:				
	Has Traffic (1)	Has Traffic (100+ UV) (2)	Has Revenue (3)	SMB (Month 1) (4)	SMB (Month 9) (5)
Treatment	−0.0003	0.0028***	0.0011**	0.00013	0.00012**
	(0.0020)	(0.0010)	(0.0004)	(0.00009)	(0.00005)
Dep Var Mean	0.498	0.230	0.161	0.005	0.004
Adjusted R^2	0.172	0.176	0.181	0.264	0.268

Notes: Dependent variables correspond to funnel stages in Section 2.2. Each column reports $\hat{\beta}$ from $Y_{ics} = \beta\, Treatment_i + \alpha_c + \alpha_s + \varepsilon_{ics}$, with cohort and sector fixed effects (as in Jin and Sun (2026)). Standard errors clustered by seller. "Has Traffic (100+ UV)" requires cumulative UV > 900 over the nine-month window. "SMB" indicates monthly revenue exceeding the sector's 90th percentile. * p<.1, ** p<.05, *** p<.01.

Table 5: Regression Results: Consumer-Side Analyses (Treatment - Control T test)

	(1)	(2)	(3)	(4)
Panel A: Purchase Conditional on Search				
	Has Purchased		Log Spending	
Is New Seller	−0.0016	−0.00081	−0.0074	−0.0039
(Treatment - Control)	(0.0013)	(0.0013)	(0.0060)	(0.0059)
Panel B: Post-Purchase Experience				
	Is Returned		Has Repeat Purchased	
Has New Seller	0.0054**	0.0026*	0.0024	0.0019
(Treatment - Control)	(0.0026)	(0.0016)	(0.0024)	(0.0015)
Panel C: Market Expansion Effects				
	Has Purchased		Log Spending	
Number of New Sellers	−0.00046	−0.00023	−0.0025	−0.0018
(Treatment - Control)	(0.00044)	(0.00044)	(0.0020)	(0.0020)
Seller Controls	N	Y	N	Y

Note: Standard errors in parentheses. Panels A, B, and C correspond to Table 3, Panels A, B, and C, respectively. Odd (even) columns exclude (include) seller controls.
Panel A: fixed effects: consideration-set. Unit: seller. Sample: consideration sets with ≥ 1 new seller. $N =$ 1,084,228. Panel B: fixed effects: consumer, keyword-size-date. Unit: consideration set. Sample: consideration sets resulting in purchases. $N = 289{,}259$. Panel C: fixed effects: consumer, keyword-size-date. Unit: consideration set. Sample: all consideration sets. $N = 6{,}007{,}329$. * p<.1, ** p<.05, *** p<.01.

Overall, the consideration-set analyses suggest that the primary channel through which advertising promoted entry was expanding the *quantity* of exploration—more consumers searched entrant products—rather than altering its *quality*. The additional entrants reached by advertising were slightly lower quality than those who would have been searched without advertising, but remained superior to the average incumbents.

Despite the statistically significant per-entrant increases, the aggregate impact remains modest. The bottleneck is adoption: only 7% of new sellers ever advertised during their first year of operation. The 2.3% experimental increase on the extensive margin, while economically important, represents a small shift in adoption. Beyond information frictions, entrant advertising may be systematically suppressed by factors such as incumbent blocking in advertising, which we explore theoretically in Section 5 with a formal model.

5 A Theory of Platform-Mediated Entry

The preceding sections establish three empirical patterns. First, the entry funnel has narrowed as the platform has matured: while more potential entrants receive nonzero traffic, exploration depth has declined and fewer achieve sustained success—a pattern the model below characterizes as wide-but-shallow exploration driven by rising displacement costs. Second, conditional on appearing in a consumer's consideration set, entrants convert at higher rates than incumbents, and marginal increases in entrant exposure raise session-level purchase rates. Third, nudging new sellers to advertise can promote entry, yet only 7% of new sellers ever

purchased advertising, and even a large-scale experimental nudge raised the advertising rate by just 2.3%.

These patterns are difficult to fully reconcile ex ante. Entrants receive little traffic yet are subject to intense screening: how is such screening achieved when the average entrant cannot generate enough visits to produce a granular quality signal? Search categories vary enormously in maturity and competitiveness: how does this heterogeneity shape the efficiency of exploration and the platform's ability to screen entrants through consumer search? And if advertising can promote entry, why do so few entrants advertise? Can structural features of the advertising market other than information frictions, such as incumbents' bidding strategies and liquidity constraints facing entrants, provide a more complete explanation?

To discipline our interpretation of these patterns, we develop a search equilibrium framework under platform ranking by extending Athey and Ellison (2011) to incorporate entrant quality uncertainty and cold-start. The model rationalizes each pattern through a single exploration-exploitation trade-off, and clarifies when advertising can, and when it cannot, substitute for platform exploration.

5.1 Setup

Platform-Mediated Search. Consider a search session in which $J \geq 1$ incumbent products are relevant to the consumer's query, each corresponding to a single seller. Sellers differ by a one-dimensional price-inclusive quality $\delta_j \in [0,1]$ that captures the likelihood of meeting a consumer's specific need. This formulation follows Athey and Ellison (2011): given platform-mediated search, a seller's pricing is deterministic given its cost structure, so the one-dimensional representation is without loss of generality.[15] The platform knows each incumbent's quality δ_j perfectly and ranks all J products in descending order of quality.

The baseline timing is as follows. Nature draws J incumbent qualities from $U[0,1]$; the platform ranks them by quality; consumers then search the marketplace sequentially. Following Athey and Ellison (2011), each consumer draws a per-inspection search cost $s_i \sim U[0,1]$ and inspects positions sequentially, stopping upon matching with a seller (probability δ_j, yielding a payoff of 1 for both parties) or when the next position's quality falls below s_i. The only source of uncertainty is whether a product of quality δ will meet a given consumer's need. Under this setup, equilibrium quality and traffic at rank r are[16]

$$\theta^B(r,J) = \frac{J-r+1}{J+r} \tag{4}$$

$$\varphi^B(r,J) = \bar{s}^B \cdot \prod_{r'=1}^{r-1} \left(1-\theta^B(r',J)\right) \tag{5}$$

where the entry threshold is $\bar{s}^B = \theta^B(1,J) = J/(J+1)$. Larger sessions attract higher-quality top sellers ($\theta^B(1,5) \approx 83\%$; $\theta^B(1,20) \approx 95\%$). Each subsequent rank receives only the fraction of consumers who failed to match at all higher ranks (Figure 4). Consumer welfare is $W^B(J) = \frac{1}{2\bar{s}^B}\sum_r \varphi_r^B\,(\theta_r^B)^2$, a special case

[15] The quality-sorting equilibrium we analyze is one of multiple that can arise under platform-mediated search; see Armstrong and Zhou (2011) and Nocke and Whinston (2022) for analyses of optimal platform ranking rules. Footnote 19 provides a more formal derivation showing that all qualitative results extend when sellers optimally set prices.

[16] The unconditional expected quality at rank r among J draws from $U[0,1]$ is the order statistic $(J-r+1)/(J+1)$. Athey and Ellison (2011) show that in the cascade equilibrium, sellers that reach deeper positions face a selected consumer pool, yielding the conditional quality $(J-r+1)/(J+r)$. The two formulas coincide at $r=1$ ($\theta_1^B = J/(J+1)$) and diverge at deeper ranks—with the conditional formula decaying faster—but all qualitative results are robust to either specification.

of the Monotone Identity (Lemma A.1 in Appendix A).

Entrant and Platform Exploration. A single potential entrant has true quality δ, drawn from Beta$(2q, 2-2q)$, with mean $q \in (0,1)$ and support $[0,1]$. The platform observes the distribution but not δ; consumers learn quality only upon inspection. The platform publicly commits to an exploration rate $\kappa \in [0,1]$: with probability κ the entrant is placed at rank 1, otherwise it receives its organic rank based on prior q.[17] The marketplace takes one of three configurations:

(i) Exploration regime. The entrant at rank 1 displaces each incumbent one rank deeper. A consumer who enters inspects the entrant first, then continues through incumbent positions as long as $\theta^B(r,J) \geq s$. Setting the marginal consumer's surplus to zero, the entry threshold solves

$$(q - \bar{s}^E) + \frac{1-q}{\bar{s}^B} \sum_{r=1}^{k^E} \varphi^B(r,J)\left(\theta^B(r,J) - \bar{s}^E\right) = 0 \tag{6}$$

where $k^E = \max\{r : \theta^B(r,J) \geq \bar{s}^E\}$. The first term is the net value of inspecting the entrant; the sum is the option value of reaching deeper incumbent positions.

Standard bounding arguments yield $q < \bar{s}^E < \bar{s}^B$: evaluating (6) at $s = \bar{s}^B$ gives strictly negative surplus (since $q < \bar{s}^B$ and $\theta^B_r \leq \bar{s}^B$), while at $s = q$ the sum is strictly positive ($\theta^B_1 > q$). The marginal consumer therefore inspects strictly more positions than under baseline ($k^E \geq 1$): exploration reduces the entry threshold, but the option value of reaching deeper incumbents partially compensates. The entrant's higher pass-through ($1-q > 1-\bar{s}^B$) further offsets the entry reduction by channeling more traffic to deeper ranks.

Consumer surplus is linear in δ for each search cost s—the match probability δ and pass-through $(1-\delta)$ both enter linearly—and the entry threshold (6) does not depend on δ, so consumer welfare $W^E(\delta,J)$ is linear in δ and $\mathbb{E}_\delta[W^E(\delta,J)] = W^E(q,J)$.[18]

(ii) Organic regime. The entrant receives its organic rank $r_{\text{org}} = \lceil (J(1-q)+1)/(1+q) \rceil$. When $q < \theta^B(1,J) = J/(J+1)$—the typical case for new sellers—the entrant is placed below rank 1, with deeper placement as q decreases or J increases. A low-q entrant in a large session is buried at a position with negligible traffic. The same linearity argument applies: $\mathbb{E}_\delta[W^O(\delta,J)] = W^O(q,J)$, computed via the Monotone Identity (Appendix A, footnote 18).

[17] We model exploration as a κ probability of jumping to rank 1 primarily for tractability, so as not to break the assortative cascade. In practice, digital platforms frequently implement exploration through "soft" rank boosting rather than discrete top-ranking. However, because organic traffic decays precipitously at lower ranks, a moderate rank boost for a deeply buried entrant yields a negligible increase in traffic. To generate sufficient data to resolve quality uncertainty, a rank boost must be substantial enough to push the entrant near the top of the search cascade. Consequently, a smaller κ closely approximates the expected traffic volume of a meaningful rank boost. The primary conceptual divergence is that algorithmic boosting often leverages the entrant's prior q, whereas our baseline κ models uniform, blind exploration. In terms of allocative efficiency, this q-based exploration occupies an intermediate space: it is more efficient than blind exploration but remains strictly dominated by the frictionless advertising auction modeled in Section 5.3, which perfectly screens based on private, realized quality. We can still reflect the system's fundamental trade-offs without modeling it; see Appendix A.4 for a formal characterization. In a related context, Ye et al. (2023) implement exploration through shadow bids that subsidize new advertisers within ad auctions; our mechanism differs by directly altering the consumer search sequence when only a small fraction of new sellers participate in ad bidding.

[18] The per-consumer surplus under exploration is $\text{CS}^E(\delta,s) = (\delta - s) + \frac{1-\delta}{\bar{s}^B} \sum_{r=1}^{k(s)} \varphi^B(r,J)\,(\theta^B(r,J) - s)$, where $k(s) = \max\{r : \theta^B(r,J) \geq s\}$. Under organic ranking, $W^O(q,J) = \frac{1}{2\bar{s}^B} \sum_{r=1}^{J+1} \varphi^O_r(q)\,(\theta^O_r(q))^2$ via the Monotone Identity, with $\theta^O_{r_{\text{org}}}(q) = q$ and pass-through $(1-q)$ at the entrant's rank. Under full information, $W^*(\delta,J) = \frac{1}{2\bar{s}^{O*}(\delta)} \sum_{r=1}^{J+1} \varphi^{O*}_r(\delta)\,(\theta^{O*}_r(\delta))^2$, where $\bar{s}^{O*}(\delta) = \max(\delta, \bar{s}^B)$ and the entrant occupies its true rank $r^*(\delta)$. See Appendix A for derivations.

(iii) Full-information regime. If the platform could observe δ, it would place the entrant at its true rank $r^*(\delta) = \min\{r : \theta^B(r,J) < \delta\}$. Unlike the prior-based regimes, both $r^*(\delta)$ and the entry threshold $\bar{s}^{O*}(\delta) = \max(\delta, \bar{s}^B)$ depend on δ, so $W^*(\delta,J)$ is *not* linear—it has kinks at each threshold $\theta^B(r,J)$ where the entrant's true rank improves (footnote 18). Moreover, when $\delta > \theta^B_1$, the entry threshold rises to δ, drawing in additional consumers. Intuitively, the heavy skew of traffic toward top positions means that correctly promoting a high-δ entrant generates disproportionately more welfare gain than is lost by keeping a low-δ entrant at a deep rank.

Since $W^E(\delta,J)$ and $W^O(\delta,J)$ are both linear in δ (Appendix A), define

$$\text{MLC}(J,q) = W^O(q,J) - W^E(q,J) \geq 0 \tag{7}$$

$$L(J,q) = \mathbb{E}_\delta\big[W^*(\delta,J)\big] - W^O(q,J) \geq 0 \tag{8}$$

MLC is the *marginal learning cost*: the welfare loss from displacing the top incumbent with an unknown entrant, increasing in J and decreasing in q. L is the *learning gap*: the welfare gain from placing the entrant at its true rank, strictly positive whenever $q < \bar{s}^B$ due to the convexity of $W^*(\delta,J)$. Let $\varphi^O_{\text{org}} \equiv \varphi^O(r_{\text{org}},J)$ denote the traffic reaching the entrant's organic rank.

Proposition 5.1 (Diminishing Returns to Exploration)**.** *The per-session welfare gain from exploration is*

$$W(\kappa; J,q) - W^O = \kappa\big(MLB(\kappa) - MLC\big) \tag{9}$$

where $MLB(\kappa) \equiv \beta \cdot L \cdot \varphi^O_{org}/T(\kappa)$ *is the* marginal learning benefit*,* $T(\kappa) \equiv \varphi^O_{org} + \kappa(\bar{s}^E - \varphi^O_{org})$ *is total entrant traffic per consumer arrival, and* β *is the product of the per-interaction learning probability and expected future demand.* $MLB(\kappa)$ *is decreasing in* κ*. A unique* $\kappa^* \in (0,1]$ *exists whenever* $MLB(0) = \beta \cdot L > MLC$*.*

Proof in Appendix A. The learning term $\text{MLB}(\kappa) \cdot \kappa$ captures diminishing returns to exploration: since Bayesian precision scales linearly with total traffic, the organic baseline φ^O_{org} governs the rate of decay. An obscure entrant with negligible organic traffic gains a large proportional increase in data from initial algorithmic exposures, so uncertainty resolves rapidly and marginal value drops sharply. A more visible entrant already accumulates a steady information flow; achieving an equivalent proportional reduction in uncertainty requires a larger absolute volume of algorithmic traffic, yielding a more gradual decay in marginal benefit. The formal derivation from Bayesian variance reduction under a $\text{Beta}(2q, 2-2q)$ prior appears in Appendix A.

Diminishing returns imply that the platform's first units of exploration are its most valuable. Even with negligible κ, a superstar entrant—one whose realized quality δ far exceeds its prior q—reveals itself upon a single rank-1 placement: it matches consumers at a rate far above the typical entrant, producing an unambiguous quality signal regardless of how pessimistic the prior may be. The convexity of $W^*(\delta,J)$ ensures that the learning value L is strictly positive for any $q < \bar{s}^B$, so even very shallow exploration captures meaningful information from the right tail of the quality distribution. The platform can therefore screen effectively through broad, shallow coverage rather than concentrated deep exploration.

5.2 Aggregate Welfare and Optimal Exploration

To aggregate across sessions, the number of incumbents J follows a truncated Pareto distribution with shape α ($J_{\min} = 2$, $J_{\max} = 100$), following the long-tail structure documented in Goolsbee and Chevalier (2003). Cross-session entrant quality q follows a Beta$(2\mu, 2-2\mu)$ distribution with mean μ and support $[0,1]$. Heterogeneity in (J,q) across sessions provides an additional force toward interior aggregate exploration, since not all session types satisfy MLB$(0) >$ MLC. Three primitives govern the exploration decision. Within each session, β (learning amplifier) scales the value of information about δ, capturing both per-inspection learning efficiency and expected future demand. Across sessions, α (category maturity) controls incumbent competition through the distribution of J, while μ (average entrant quality) governs the cross-session distribution of q. We assume the platform chooses a uniform exploration rate κ across all search sessions. Since each session's entrant conversion rate equals q exactly (Lemma B.2 in Appendix A), the aggregate entrant conversion rate is a traffic-weighted average of q across sessions:

Proposition 5.2 (Exploration Dilutes Baseline Screening)**.** *The aggregate entrant conversion rate $\bar{C}_E(\kappa)$ is decreasing in κ.*

Proof in Appendix A. Even at $\kappa = 0$, assortative organic ranking screens entrants by concentrating traffic on high-q sellers and burying low-q entrants at deep ranks. This screening operates on the prior q, not the realized δ: an entrant with $q = 0.3$ receives the same deep organic rank regardless of whether its true quality turns out to be $\delta = 0.9$ or $\delta = 0.01$. Exploration disrupts this assortative structure by placing untested entrants at rank 1 regardless of their prior. As κ rises, traffic shifts toward low-q sessions that benefit disproportionately from rank-1 placement, diluting the aggregate conversion rate (Appendix A). The bottom rows of Figures 1 and 2 illustrate this numerically: at low κ, organic screening concentrates entrant traffic on high-q sessions, producing entrant conversion rates that can exceed the incumbent average; as κ rises, this gap narrows.

A benevolent platform maximizes aggregate welfare $\mathscr{W}(\kappa) \equiv \sum_{J,q} w(J,q) \cdot W(\kappa; J, q)$, where $w(J,q)$ is the joint probability of session type (J,q) under the Pareto–Beta distributions, by choosing κ^*. $\mathscr{W}$ is concave under standard regularity conditions (Appendix A), yielding a unique optimum.

Proposition 5.3 (Optimal Exploration)**.**

$$\frac{\partial \kappa^*}{\partial \beta} \geq 0, \qquad \frac{\partial \kappa^*}{\partial \alpha} \geq 0, \qquad \frac{\partial \kappa^*}{\partial \mu} \geq 0 \tag{10}$$

Proof in Appendix A. Optimal exploration increases with expected entrant quality μ, the learning amplifier β (which captures both the efficiency of learning and the expected volume of future consumers), and concentration on small sessions (higher α). Conversely, as a platform matures and incumbent competition intensifies (α falls, J increases), optimal exploration declines—consistent with the narrowing entry funnel documented empirically. The proof applies Milgrom and Shannon (1994) to aggregate per-session comparative statics $\partial \kappa^*/\partial J < 0$ and $\partial \kappa^*/\partial q > 0$.

A uniform κ distributes exploration thinly across all sessions rather than concentrating it on a few. Diminishing returns (Proposition 5.1) make this *wide-but-shallow* strategy efficient: a small κ applied to many sessions identifies superstars in the population of entrants without requiring deep per-session

exploration. As the platform matures and displacement costs rise, κ^* falls but remains positive, producing the pattern observed empirically: many entrants receive some traffic, yet few are explored deeply enough to emerge. This parallels the "starting small" result of Che and Hörner (2018), where the recommender explores many products at low per-product intensity to maintain credibility. In our setting, the constraint is informational—diminishing Bayesian precision gains—rather than incentive-based, but both mechanisms predict broad, shallow exploration.

Simulation. We illustrate Propositions 5.2 and 5.3 numerically. Figure 1 (top row) plots $\mathscr{W}(\kappa) - \mathscr{W}(0)$ for $\beta \in \{5, 10, 20\}$ under three specifications. The left panel uses the baseline calibration: quality parity ($\mu = 0.5$) with the Goolsbee and Chevalier (2003) session-size distribution ($\alpha = 1.5$). Figure 7 (Appendix A) plots the model-predicted distribution of consideration set sizes from the baseline equilibrium across three Pareto concentration levels, paralleling the empirical distribution in Figure 8. The center panel shows lower-quality entrants ($\mu = 0.35$): displacement costs increase and exploration requires higher β. The right panel shows a more mature platform ($\alpha = 1.1$, $\bar{J} = 5.9$): more incumbent competition raises MLC, reducing κ^* at every β. The bottom row illustrates Proposition 5.2: the entrant conversion rate declines monotonically as κ increases, reflecting the dilution of organic screening. Notably, under baseline organic screening alone ($\kappa = 0$), the entrant conversion rate can exceed the incumbent's (consistent with Panel A of Table 3) though this is not guaranteed in very mature markets or when entrant quality is low.

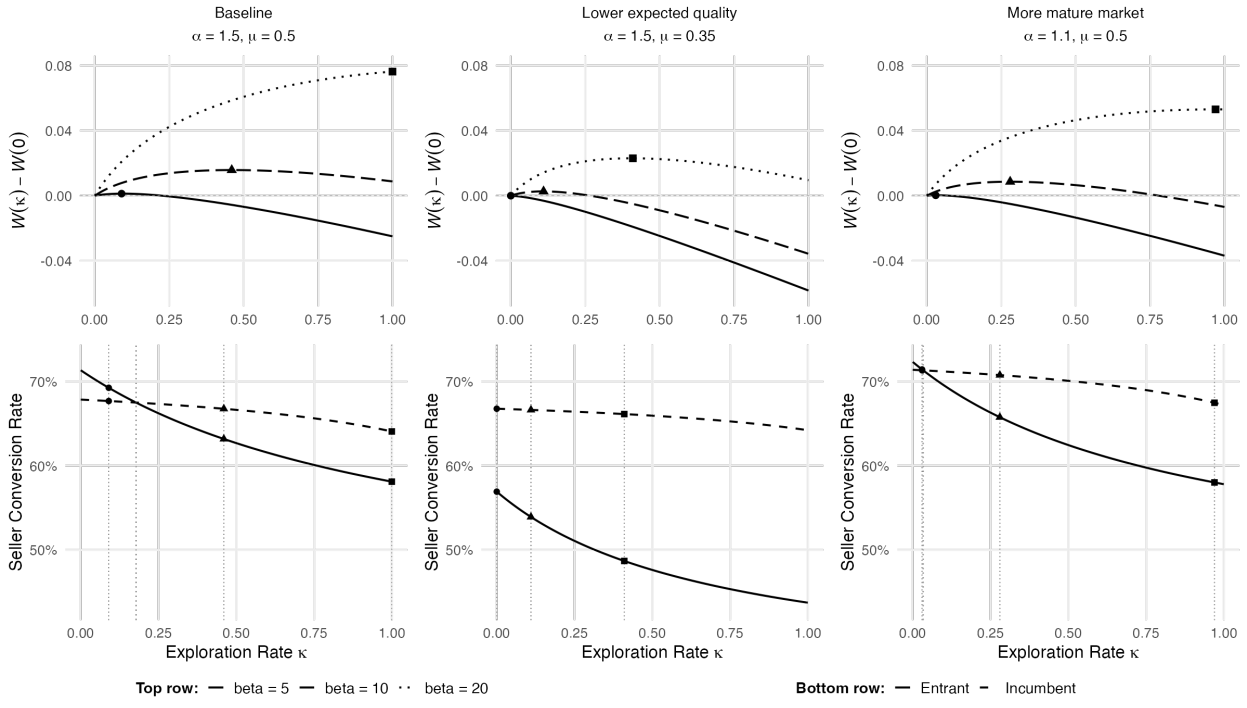


Figure 1: Welfare (top) and Seller Conversion (bottom)

Notes: Analytical computation over $J \in \{2, 3, \ldots, 100\}$ and $q \in \{0, 0.01, \ldots, 1\}$. κ evaluated on $[0, 1]$ in steps of 0.01. Dots mark κ^*. Bottom row: solid = entrant, dashed = incumbent; vertical dotted lines mark the crossing point and κ^* values.

5.3 Advertising

Assume the platform offers a^* as a posted price for rank-1 placement only to the entrant. We impose a *liquidity constraint*: the entrant can bid at most its current-period payoff from rank 1, consistent with empirical evidence that new sellers are financially constrained and rarely bid above short-run returns (Gao and Yu, 2023; Cui, 2025). The entrant's willingness to pay is

$$v(\delta) \;=\; (1-\kappa)\,(\bar{s}^E - \varphi^O_{\text{org}}) \cdot \delta \tag{11}$$

where $(1-\kappa)$ reflects the outside option under exploration and $(\bar{s}^E - \varphi^O_{\text{org}})$ is the traffic gain from moving to rank 1.[19] Since $v(\delta)$ is increasing in δ, the reserve price a^* induces a cutoff: entrants with $\delta \geq \delta^*$ purchase rank 1. The entry threshold when the entrant occupies rank 1 remains $\bar{s}^E$.[20]

Proposition 5.4 (Unraveling of Exploration)**.** *Frictionless advertising weakly dominates exploration: the aggregate social optimum sets $\kappa^* = 0$, with screening threshold $\delta^*(J,q) > q$.*

Proof in Appendix A. Since $v(\delta)$ is monotone in δ, the reserve price screens entrants by realized quality: those with $\delta \geq \delta^*$ self-select into the paid channel, while the remainder—for whom promotion is welfare-reducing—fall back to organic ranking. This renders blind exploration redundant ($\kappa^* = 0$). Moreover, $\delta^* > q$: at the prior mean, displacement cost is strictly positive while learning value is zero, so only above-average entrants are worth promoting. Formally, let

$$\Delta W(\delta) \;\equiv\; \beta\left[W^*(\delta,J) - W^O(q,J)\right] - \left[W^O(\delta,J) - W^E(\delta,J)\right] \tag{12}$$

denote the net welfare gain from promoting an entrant of realized quality δ, equal to learning value minus displacement cost. The welfare-maximizing reserve satisfies $\Delta W(\delta^*) = 0$.

Information frictions. Proposition 5.4 assumes every entrant can participate in the auction. In practice, participation requires awareness of the advertising channel, technical capability, and a budget—conditions that most new sellers fail to meet. Let $\lambda \in [0,1]$ denote the fraction of entrants that participates. Among participants, only those with $\delta \geq \delta^*$ are promoted; all others fall back to the κ mechanism. Whether $\kappa^* > 0$ depends on the severity of frictions.

[19] Since δ is price-inclusive quality, one might worry that higher δ from lower prices reduces bid capacity. With optimal pricing ($p^* = \delta'/2$ where δ' is inherent quality), per-consumer profit is δ^2 and $v(\delta) \propto \delta^2$—still monotone. All qualitative results are unchanged under this refinement. More generally, entrants may strategically underprice to build reputation (Aleksenko and Kohlhepp, 2023; Dendorfer and Seibel, 2025), which further reinforces the monotonicity of v in δ.

[20] Consumers observe the ranking but not the promotion mechanism. Under disclosure, two opposing forces arise. On the Bayesian side, consumers observing an ad-placed entrant infer $\delta \geq \delta^*$ and update to $\mathbb{E}[\delta \mid \delta \geq \delta^*] > q$, raising $\bar{s}^E$ and strengthening screening. On the behavioral side, Yu (2024) estimates that consumers penalize sponsored positions—equivalent to roughly 8% skipping the ad slot—which pushes $\bar{s}^E$ toward $\bar{s}^B$ and weakens screening.

Incumbent blocking. Now suppose the platform runs a second-price position auction for the top slot.[21] The rank-1 incumbent's total value of retaining its position is

$$v_{\text{inc}}(q,J) \;=\; \underbrace{\theta_1^B\left[\bar{s}^B - \bar{s}^E(1-q)\right]}_{\text{static displacement}} \;+\; \underbrace{\beta_I \int_{\theta_1^B}^{1} \theta_1^B\left[\bar{s}^B - \theta_1^B(1-\delta)\right] f(\delta \mid q)\, d\delta}_{\text{information blocking}} \tag{14}$$

The static term is the revenue loss from dropping to rank 2, where the incumbent receives traffic $\bar{s}^E(1-q)$ instead of $\bar{s}^B$. The dynamic term captures information blocking: if the platform observes $\delta \geq \theta_1^B$ during rank-1 placement, it permanently promotes the entrant; $f(\delta \mid q)$ is the Beta$(2q, 2-2q)$ density, and β_I, analogous to β, is the product of the per-interaction learning probability and the incumbent's expected future demand.

Under the liquidity constraint (11), whenever $v_{\text{inc}} > a^*$ the incumbent blocks all entrants with $\delta \in [\delta^*, \bar{\delta}_{\text{inc}}]$—those good enough to pass the welfare screen but not good enough to outbid the incumbent—where

$$\bar{\delta}_{\text{inc}} \;\equiv\; \frac{v_{\text{inc}}}{(1-\kappa)\,(\bar{s}^E - \varphi_{\text{org}}^O)} \tag{15}$$

is the quality at which the entrant's bid just equals the incumbent's.[22] For large J, $\bar{\delta}_{\text{inc}}$ can exceed 1, blocking all ad-based entry.[23]

Proposition 5.5 (Advertising Frictions Restore Exploration)**.** *Suppose $\kappa^*_{\text{no-ads}} > 0$ (the optimum from Proposition 5.1 absent advertising). Then $\kappa^* > 0$ with advertising whenever:*

(i) *Information frictions (no incumbent bidding): $\kappa^* > 0$ whenever $\lambda < \bar{\lambda}$, where*

$$\bar{\lambda} \;=\; \frac{\mathbb{E}_\delta[\Delta W(\delta)]}{\mathbb{E}_\delta[\Delta W(\delta)\cdot \mathbf{1}(\delta \geq \delta^*)]} \;\in\; (0,1) \tag{16}$$

(ii) *Incumbent blocking ($v_{\text{inc}} > a^*$): $\kappa^* > 0$ whenever $\lambda < \bar{\lambda}'$, where*

$$\bar{\lambda}' \;=\; \frac{\mathbb{E}_\delta[\Delta W(\delta)]}{\mathbb{E}_\delta[\Delta W(\delta)\cdot \mathbf{1}(\delta > \bar{\delta}_{inc})]} \;>\; \bar{\lambda} \tag{17}$$

If $\bar{\delta}_{\text{inc}} \geq 1$ (complete blocking), $\bar{\lambda}' = \infty$ and $\kappa^* > 0$ for any λ.

Proof in Appendix A. The threshold $\bar{\lambda}$ is close to 1 when sub-threshold displacement losses (from promoting entrants with $\delta < \delta^*$) are small relative to above-threshold learning gains, and approaches 0 when sub-threshold losses dominate. Incumbent blocking shrinks the set of entrants served by ads from $\{\delta \geq \delta^*\}$ to

[21] Each bidder computes its expected surplus integrating over the opponent's type; in the resulting two-bidder second-price auction with these expected values, truthful bidding is a dominant strategy. The incumbent bids based on expected pass-through $(1-q)$ rather than realized $(1-\delta)$.

[22] Relaxing the liquidity constraint introduces dynamic bidding: an entrant who knows $\delta > \theta_1^B$ anticipates permanent promotion and is willing to pay $v^{\text{dyn}}(\delta) = (1-\kappa)(\bar{s}^E - \varphi_{\text{org}}^O)\delta + \beta^E\,\delta(\delta - \varphi_{\text{org}}^O)$, which can exceed v_{inc} for high δ. Three forces interact: the entrant's worse outside option and private knowledge of δ favor outbidding, while the traffic structure ($\bar{s}^E < \bar{s}^B$) favors the incumbent. Figure 6 in the Appendix reports simulation results under dynamic entrant bidding.

[23] In practice, bid-level stochasticity—arising from keyword-specific valuations, budget shocks, or auction noise (as estimated by Yu, 2024)—would smooth the deterministic blocking boundary into a probabilistic transition. The qualitative prediction that incumbents block a range of welfare-improving entrants is unchanged.

$\{\delta > \bar{\delta}_{\text{inc}}\}$, raising the participation threshold: $\bar{\lambda}' > \bar{\lambda}$. Under complete blocking ($\bar{\delta}_{\text{inc}} \geq 1$), no participant can outbid the incumbent, so $\bar{\lambda}' = \infty$ and $\kappa^* > 0$ for any λ. The frictionless benchmark predicts $\kappa^* = 0$: advertising perfectly screens entrants, rendering blind exploration redundant. In practice, information frictions leave a mass of entrants unserved, and incumbent blocking shuts out socially valuable entrants who do participate. The platform's exploration rate κ is therefore the primary—and often the only—channel through which new sellers enter the marketplace.

Simulation. Figure 2 illustrates Propositions 5.4 and 5.5 numerically. The left panel sets $\lambda = 0.07$—matching the observed advertising participation rate on our platform—without incumbent bidding. Even at this low participation rate, advertising reduces optimal exploration: $\kappa^* = 0/0.25/1.00$ for $\beta \in \{5, 10, 20\}$, compared to $0.09/0.46/1.00$ without advertising (Figure 1), indicating that even 7% participation can provide enough screening to substitute or even supplant blind exploration. The center panel raises $\lambda = 0.20$. Advertising now completely substitutes for exploration ($\kappa^* = 0$ for all β), confirming Proposition 5.4.

The right panel introduces incumbent blocking at $\lambda = 0.20$, setting $\beta_I = \beta$. Under the liquidity constraint, the incumbent outbids all entrants whenever $\bar{\delta}_{\text{inc}} > \delta^*$, shutting down the advertising channel entirely. Exploration is restored at $\kappa^* = 0.09/0.45/1.00$ for $\beta = 5/10/20$, confirming Proposition 5.5.

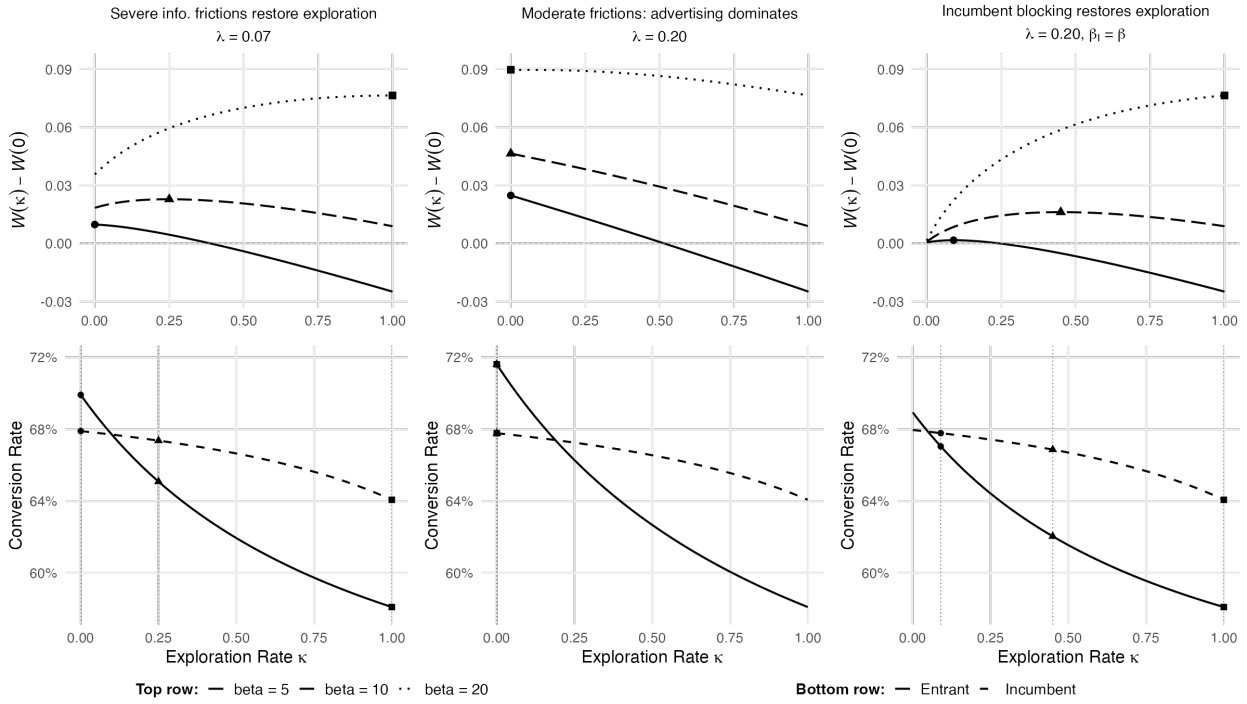


Figure 2: Advertising Frictions and Optimal Exploration

Notes: Left panel: $\lambda = 0.07$, no incumbent bidding. Center: $\lambda = 0.20$, no incumbent bidding. Right: $\lambda = 0.20$ with incumbent blocking ($\beta_I = \beta$). Top row: $\mathscr{W}(\kappa) - \mathscr{W}(0)$; dots mark κ^*. Bottom row: entrant and incumbent conversion rates at $\beta = 10$, computed as traffic-weighted averages across both the advertising and organic channels. Baseline calibration: $\alpha = 1.5$, $\mu = 0.5$.

5.4 Connection to Empirical Findings

The model highlights three forces behind the narrow and narrowing entry funnel (Table 1): the social benefit of exploration diminishes rapidly as the platform matures; advertising substitutes for exploration even when most entrants face information frictions (Figure 2 middle panel); and organic ranking already screens entry intensely (the conversion advantage of entrants over incumbents in Table 3A reflects this baseline screening; Proposition 5.2 shows that exploration dilutes it, reducing the average entrant conversion rate).

Despite these forces, exploration is often still optimal. The market-expansion test (Table 3, Panel C) provides direct evidence, and can be mapped to the model as a bounding exercise. The regression conditions on consumer and keyword-date-number-of-searches fixed effects, holding session type approximately fixed; within-cell variation in the number of entrants isolates the effect of entrant exposure on purchase probability. A positive coefficient implies that replacing one incumbent slot with an entrant raises match volume: a sufficient condition for the welfare effect ($\partial \mathscr{W} / \partial \kappa$) to be positive. This is because the regression captures only contemporaneous outcomes, omitting potentially long-lasting dynamic learning benefit. In addition, the experiment validates advertising's screening value, but simultaneously suggests that information frictions alone is unlikely to explain the low take-up rate of advertising.

Our approach complements Che and Hörner (2018), whose "starting small" result—exploring many products at low per-product intensity to preserve credibility—parallels the wide-but-shallow equilibrium that emerges here from diminishing Bayesian returns; Dendorfer and Seibel (2025), who structurally estimate Airbnb's cold-start problem; and Bai et al. (2025), who model congestion-driven misallocation on AliExpress. We simplify learning dynamics to focus on how exploration coexists with organic ranking and advertising—a choice guided by our consideration-set data, which allow us to recover the sequential-search environment of Athey and Ellison (2011).

6 Conclusions

A key promise of digital platforms is their ability to lower entry barriers and decentralize market access. This study documents that, on a dominant e-commerce platform, traditional barriers to entry have indeed been largely eliminated: tens of thousands of sellers register daily. Yet fewer than 1% become viable businesses. Our evidence suggests that the binding constraint has shifted to an informational one: entrants lack the transaction history needed to demonstrate quality, and the platform's capacity to experimentally resolve this uncertainty is insufficient at scale. The consideration-set evidence makes this concrete: conditional on being explored, entrants convert at rates that significantly exceed incumbents', and marginally replacing incumbent search with entrant search expands the market. The problem is not that there is a lack of high-quality entrants, but that most are never explored deeply enough to reveal it.

Our model provides a theoretical foundation for these patterns. As platforms mature, optimal exploration declines and the entry funnel narrows toward its top. While frictionless advertising can replace exploration, low participation limits its ability to address cold-start. Our model identifies information frictions and incumbent blocking as likely channels, though the data do not allow us to fully separate their contributions. As a result, many entrants are subject to intense screening based on organic ranking alone, which can be

nearly impossible to overcome in mature markets.

More generally, platform policies that prioritize steering consumers toward incumbents may be socially efficient yet perpetuate entrant under-exploration, even when multiple instruments are available (Franck and Peitz, 2024). Entrepreneurs, in turn, should not over-interpret past successful entry: as platforms mature, the entry funnel's dynamism declines.

We do not observe the platform's ranking algorithm or advertising prices, and thus cannot structurally separate platform steering from consumer search behavior as drivers of the under-exploration. Our model provides likely mechanisms driving both the narrowing entry funnel and the persistently low advertising rates, but we cannot fully test the model with the data at hand. Several extensions could advance both theory and empirics. First, exploiting exogenous variation in ranking algorithms, including soft rank boosting rather than the discrete top-placement we model, would allow estimation of the displacement cost of exploration along a continuous margin. Second, bid-level auction data could directly measure incumbent blocking and test whether defensive bidding accounts for low entrant ad adoption. Third, leveraging cross-platform differences in advertising design, such as the ease for entrants to begin advertising or their degree of shielding from incumbent competition, would illuminate when advertising substitutes for exploration. Finally, extending the model to incorporate learning-by-doing, where seller quality improves with experience, would address a channel our static framework omits: early exploration may generate not only information but also human capital, amplifying the long-run returns to platform intervention.

References

Aleksenko, Stepan, and Jacob Kohlhepp. 2023. "Reputational Underpricing." Available at SSRN 4380756.

Anderson, Simon P, and Régis Renault. 2025. "Search direction: Position externalities and position auction bias." *Review of Economic Studies*, rdaf051.

Aridor, Guy, Yishay Mansour, Aleksandrs Slivkins, and Zhiwei Steven Wu. 2020. "Competing bandits: The perils of exploration under competition." *arXiv preprint arXiv:2007.10144*.

Armstrong, Mark, and Jidong Zhou. 2011. "Paying for Prominence." *Economic Journal*, 121(556): F368–F395.

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

Bai, Jie, Maggie Chen, Jin Liu, and Daniel Yi Xu. 2020. "Search and Information Frictions on Global E-Commerce Platforms: Evidence from AliExpress."

Bai, Jie, Maggie X. Chen, Jin Liu, Xiaosheng Mu, and Daniel Yi Xu. 2025. "Stand Out from the Millions: Market Congestion and Information Friction on Global E-Commerce Platforms." Working Paper.

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

Bergemann, Dirk, and Juuso Välimäki. 2006. "Dynamic pricing of new experience goods." *Journal of Political Economy*, 114(4): 713–743.

Berman, Ron, and Zsolt Katona. 2013. "The Role of Search Engine Optimization in Search Marketing." *Marketing Science*, 32(4): 644–651.

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

Bolton, Patrick, and Christopher Harris. 1999. "Strategic experimentation." *Econometrica*, 67(2): 349–374.

Chen, Yongmin, and Chuan He. 2011. "Paid placement: Advertising and search on the internet." *The Economic Journal*, 121(556): F309–F328.

Chevalier, Judith A, and Dina Mayzlin. 2006. "The effect of word of mouth on sales: Online book reviews." *Journal of Marketing Research*, 43(3): 345–354.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Optimal design for social learning." *The Quarterly Journal of Economics*, 133(4): 2021–2067.

Crémer, J, YA De Montjoye, and H Schweitzer. 2019. "Competition Policy for the Digital Era, European Commission." *Publications Office of the European Union*.

Cui, Jingyi. 2025. "Bidding for Reputation." Job Market Paper, Yale University.

Dai, Weijia, Hyunjin Kim, and Michael Luca. 2023. "Frontiers: Which Firms Gain from Digital Advertising? Evidence from a Field Experiment." *Marketing Science*, 42(3): 429–439.

Dendorfer, Florian, and Regina Seibel. 2025. "The Cost of the Cold-Start Problem on Airbnb." Working Paper, University of Toronto.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan. 2018. "Consumer price search and platform design in internet commerce." *American Economic Review*, 108(7): 1820–59.

Dubé, Jean-Pierre, Ali Hortaçsu, and Joonhwi Joo. 2021. "Random-Coefficients Logit Demand Estimation with Zero-Valued Market Shares." *Marketing Science*, 40(4): 637–660.

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

Farronato, Chiara. 2025. "Digital platforms as information aggregators: Implications for their design and regulation." *Cambridge Handbook on Digital Platforms*.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay. 2023. "Self-preferencing at Amazon: evidence from search rankings." Vol. 113, 239–243, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Franck, Jens-Uwe, and Martin Peitz. 2024. "The Digital Markets Act and the Whack-A-Mole Challenge." *Common Market Law Review*, 61: 631–666.

Gao, Kailin, and Miaojie Yu. 2023. "Signaling and quality upgrading: Evidence from e-commerce certification in China." *Managerial and Decision Economics*.

Goolsbee, Austan, and Judith Chevalier. 2003. "Measuring Prices and Price Competition Online: Amazon and Barnes and Noble." *Quantitative Marketing and Economics*, 1(2): 203–222.

Hagiu, Andrei, and Julian Wright. 2024. "Optimal Discoverability on Platforms." *Management Science*, 70(11): 7770–7790.

Honka, Elisabeth, Stephan Seiler, and Raluca Ursu. 2024. "Consumer Search: What Can We Learn from Pre-Purchase Data?" *Journal of Retailing*, 100(1): 114–129.

Janssen, Maarten CW, Thomas Jungbauer, Marcel Preuss, and Cole Williams. 2026. "Search platforms: Big data and sponsored positions." *The Economic Journal*, ueag002.

Jin, Yizhou, and Zhengyun Sun. 2019. "Information Acquisition and the Return to Data." *Harvard University Working Paper*.

Jin, Yizhou, and Zhengyun Sun. 2026. "AI Training for Online Entrepreneurs: Evidence from a Large E-Commerce Experiment." *Available at SSRN*. `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4334216`.

Lam, H Tai. 2021. "Platform search design and market power." *Job Market Paper, Northwestern University*.

Lee, Kwok Hao, and Leon Musolff. 2021. "Entry into two-sided markets shaped by platform-guided search." *Job Market Paper, Princeton University*.

Lewis, Gregory. 2011. "Asymmetric information, adverse selection and online disclosure: The case of eBay motors." *American Economic Review*, 101(4): 1535–1546.

Long, Fei, Kinshuk Jerath, and Miklos Sarvary. 2022. "Designing an online retail marketplace: Leveraging information from sponsored advertising." *Marketing Science*, 41(1): 115–138.

Luca, Michael. 2016. "Reviews, reputation, and revenue: The case of Yelp.com." *American Economic Journal: Applied Economics*, 8(1): 1–24.

Milgrom, Paul, and Chris Shannon. 1994. "Monotone Comparative Statics." *Econometrica*, 62(1): 157–180.

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

Moshary, Sarah. 2021. "Advertising Effects in Equilibrium."

Newberry, Peter W. 2016. "An empirical study of observational learning." *The RAND Journal of Economics*, 47(2): 394–432.

Nocke, Volker, and Michael D. Whinston. 2022. "Concentration Thresholds for Horizontal Mergers." *American Economic Review*, 112(6): 1915–1948.

Nocke, Volker, and Patrick Rey. 2024. "Consumer search, steering, and choice overload." *Journal of Political Economy*, 132(5): 1684–1739.

Pallais, Amanda. 2014. "Inefficient hiring in entry-level labor markets." *American Economic Review*, 104(11): 3565–3599.

Papanastasiou, Yiangos, Kostas Bimpikis, and Nicos Savva. 2018. "Crowdsourcing exploration." *Management Science*, 64(4): 1727–1746.

Sahni, Navdeep S, and Charles Zhang. 2024. "Are consumers averse to sponsored messages? The role of search advertising in information discovery." *Quantitative Marketing and Economics*, 22(1): 63–114.

Sahni, Navdeep S, and Harikesh S Nair. 2020. "Does Advertising Serve as a Signal? Evidence from a Field Experiment in Mobile Search." *The Review of Economic Studies*, 87(3): 1529–1564.

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

Sun, Tianshu, Zhe Yuan, Chunxiao Li, Kaifu Zhang, and Jun Xu. 2021. "The value of personal data in internet commerce: A high-stake field experiment on data regulation policy." *Available at SSRN 3962157*.

Tadelis, Steven. 2016. "Reputation and feedback systems in online platform markets." *Annual Review of Economics*, 8: 321–340.

Ursu, Raluca, Andrey Simonov, and Eunkyung An. 2022. "Online advertising as passive search." *NYU Stern School of Business Forthcoming, Columbia Business School Research Paper*.

Ursu, Raluca, Stephan Seiler, and Elisabeth Honka. 2025. "The Sequential Search Model: A Framework for Empirical Research." *Quantitative Marketing and Economics*, 23(1): 165–213.

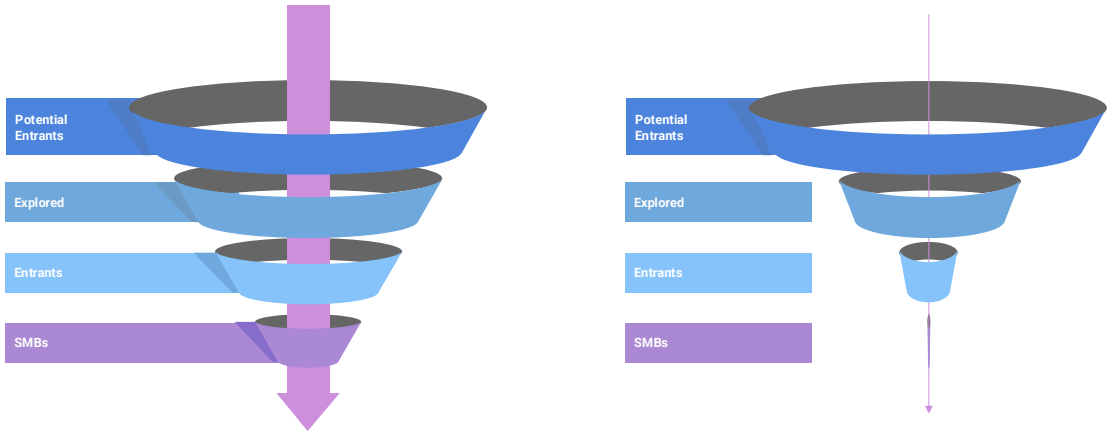
van der Vaart, A. W. 1998. *Asymptotic Statistics. Cambridge Series in Statistical and Probabilistic Mathematics*, Cambridge University Press.

Vellodi, Nikhil. 2018. "Ratings design and barriers to entry." *Available at SSRN 3267061*.

Waldfogel, Joel. 2024. "Amazon self-preferencing in the shadow of the digital markets act." National Bureau of Economic Research.

Ye, Zikun, Heng Zhang, Renyu Philip Zhang, Xin Chen, Zhiwei Xu, and Dennis J Zhang. 2023. "Cold Start to Improve Market Thickness on Online Advertising Platforms: Data-Driven Algorithms and Field Experiments." *Management Science*, 69(7): 3838–3860.

Yu, Chuan. 2024. "The Welfare Effects of Sponsored Product Advertising." *Available at SSRN 4817542*.



(a) Illustrative Funnel

(b) Proportional Funnel

Figure 3: Entry Funnel

A Model Proofs [For Online Publication]

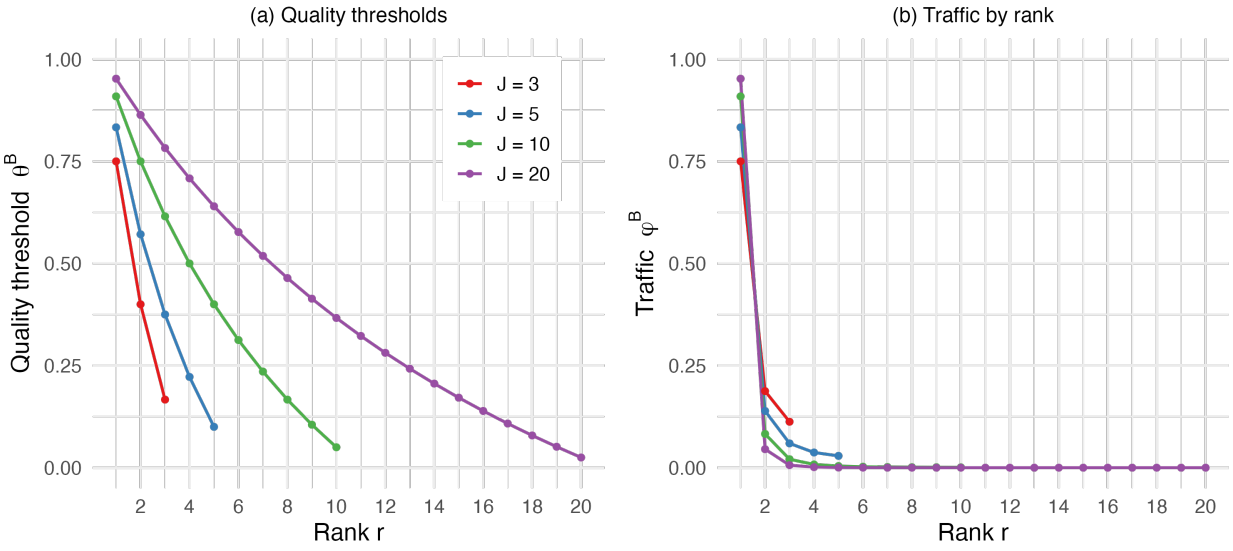


Figure 4: Baseline Equilibrium: Quality Thresholds and Traffic by Rank

Notes: Panel (a) plots the quality threshold $\theta^B(r,J) = (J-r+1)/(J+r)$ from equation (4). Panel (b) plots the traffic $\varphi^B(r,J)$ from equation (5). In larger sessions, top-ranked sellers have higher quality thresholds and traffic is more concentrated at rank 1, with rapid decay across lower ranks.

A.1 Setup

Lemma A.1 (Monotone Identity)**.** *For any monotone ranking ($\theta_1 \geq \theta_2 \geq \cdots$) with entry threshold $\bar{s}$, consumer welfare satisfies:*

$$W = \frac{1}{2\,\bar{s}} \sum_r \varphi_r\, \theta_r^2 \tag{18}$$

Proof. A consumer with cost s inspects rank r whenever $\theta_r \geq s$ (the expected payoff from one more inspection exceeds its cost). Under a monotone ranking, the entry threshold is $\bar{s} = \theta_1$ and consumer $s < \bar{s}$ inspects all ranks up to $k(s) = \max\{r : \theta_r \geq s\}$. Let $p_r = \prod_{r'=1}^{r-1}(1-\theta_{r'})$ denote the probability of reaching rank r without a prior match. Exchanging the sum and integral:

$$W = \int_0^{\bar{s}} \sum_{\{r:\,\theta_r \geq s\}} p_r\,(\theta_r - s)\,ds = \sum_r p_r \int_0^{\theta_r} (\theta_r - s)\,ds = \sum_r p_r \cdot \frac{\theta_r^2}{2}$$

Defining $\varphi_r \equiv \bar{s} \cdot p_r$, we obtain $W = \frac{1}{2\bar{s}} \sum_r \varphi_r\, \theta_r^2$, yielding (18). □ □

Derivation of Linearity of $W^E(\delta, J)$ and $W^O(\delta, J)$. Under the exploration ranking, the entrant is at rank 1 with realized quality δ and the J incumbents occupy ranks $2, \ldots, J{+}1$. Consumers enter based on the prior q (not the realized δ), so the entry threshold $\bar{s}^E$ from (6) is independent of δ. A consumer with cost $s < \bar{s}^E$ visits the entrant (match probability δ), then continues to incumbent ranks as long as $\theta^B(r, J) \geq s$. The per-consumer surplus is linear in δ for each s, since δ appears only through $(\delta - s)$ and $(1-\delta)$. Since the integration limits $[0, \bar{s}^E]$ are independent of δ:

$$W^E(\delta, J) = \int_0^{\bar{s}^E} \mathrm{CS}^E(\delta, s)\,ds = A + B\,\delta$$

for constants A, B depending only on q and J. Therefore $\mathbb{E}_\delta[W^E(\delta, J)] = A + Bq = W^E(q, J)$. The same argument applies to $W^O(\delta, J)$: the organic ranking uses expected quality q for stopping and entry, and δ enters linearly through the match probability and pass-through at rank r_{org}. Linearity of $W^E(\delta, J)$ carries over to the advertising setting of Section 5.3: under no-disclosure, consumers use prior q for the rank-1 entrant regardless of the promotion mechanism, so $\bar{s}^E$ remains independent of δ and the FOC (13) can be evaluated pointwise at $\delta = \delta^*$. □ □

Proof of Proposition 5.1 (Diminishing Returns). The learning term in (9) follows from Bayesian variance reduction under continuous information flow, a framework common to experimentation models (Bolton and Harris, 1999; Bergemann and Välimäki, 2006; Papanastasiou, Bimpikis and Savva, 2018).

The platform holds a Beta$(2q, 2-2q)$ prior over the entrant's true quality δ, with mean q and precision parameter 2. Each consumer who inspects the entrant generates a Bernoulli signal about δ. Under exploration, the entrant occupies rank 1 and receives traffic $\bar{s}^E$ per consumer arrival; under organic ranking, it receives traffic φ^O_{org}. Total entrant traffic is therefore $T(\kappa) = \varphi^O_{\text{org}} + \kappa(\bar{s}^E - \varphi^O_{\text{org}})$.

By the Bernstein–von Mises theorem (van der Vaart, 1998), as the number of consumer arrivals N grows large, the posterior distribution of δ converges to a Normal distribution whose variance equals the inverse of the total Fisher information. Since each inspection contributes independently to the Fisher information, the

total Fisher information scales linearly with $N \cdot T(\kappa)$, and hence the asymptotic posterior variance satisfies $\mathrm{Var}(\delta \mid \text{data}) \propto 1/T(\kappa)$, holding N fixed. This justifies treating $T(\kappa)$ as the *precision* of the platform's continuous information flow about δ, without requiring $T(\kappa)$ to be an integer count of signals.

Under a quadratic loss approximation—standard in Bayesian decision theory—the expected value of information is proportional to the reduction in posterior variance. The incremental learning value from exploration (relative to the organic baseline) is therefore

$$\Delta\mathscr{V}(\kappa) = C\left(\frac{1}{\varphi^O_{\text{org}}} - \frac{1}{T(\kappa)}\right) = \frac{C\,\kappa(\bar{s}^E - \varphi^O_{\text{org}})}{\varphi^O_{\text{org}}\,T(\kappa)}$$

for a scaling constant $C > 0$. To connect this to the model's economic primitives, normalize so that the marginal benefit of the first explored consumer equals $\beta \cdot L$. Since $T(\kappa)$ is linear in κ with $T(0) = \varphi^O_{\text{org}}$, we have $\frac{d}{d\kappa}\left[\kappa/T(\kappa)\right] = \varphi^O_{\text{org}}/T(\kappa)^2$, which gives

$$\left.\frac{d\,\Delta\mathscr{V}}{d\kappa}\right|_{\kappa=0} = \frac{C(\bar{s}^E - \varphi^O_{\text{org}})}{(\varphi^O_{\text{org}})^2} = \beta \cdot L \quad\Longrightarrow\quad C = \frac{\beta \cdot L \cdot (\varphi^O_{\text{org}})^2}{\bar{s}^E - \varphi^O_{\text{org}}}$$

Substituting, the $(\bar{s}^E - \varphi^O_{\text{org}})$ terms cancel:

$$\Delta\mathscr{V}(\kappa) = \frac{\beta \cdot L \cdot (\varphi^O_{\text{org}})^2}{\bar{s}^E - \varphi^O_{\text{org}}} \cdot \frac{\kappa(\bar{s}^E - \varphi^O_{\text{org}})}{\varphi^O_{\text{org}}\,T(\kappa)} = \frac{\beta \cdot L \cdot \varphi^O_{\text{org}} \cdot \kappa}{T(\kappa)}$$

The $(\bar{s}^E - \varphi^O_{\text{org}})$ cancellation means $\bar{s}^E$ enters the learning value only through $T(\kappa) = \varphi^O_{\text{org}} + \kappa(\bar{s}^E - \varphi^O_{\text{org}})$, as in (9). That C depends on φ^O_{org} reflects the normalization convention: β is calibrated so that the first explored consumer delivers $\beta \cdot L$ regardless of organic position.

Existence and uniqueness.. The marginal learning benefit $\widetilde{\text{MLB}}(\kappa) \equiv \beta \cdot L \cdot {\varphi^O_{\text{org}}}^2/T(\kappa)^2$ satisfies: (i) $\widetilde{\text{MLB}}(0) = \beta \cdot L > 0$; (ii) $\widetilde{\text{MLB}}(\kappa)$ is strictly decreasing in κ; and (iii) $\lim_{\kappa\to 1}\widetilde{\text{MLB}}(\kappa) < \text{MLC}(q)$ for $q < \bar{s}^B$. Define $g(\kappa) = \widetilde{\text{MLB}}(\kappa) - \text{MLC}(q)$. By (i) and (iii), $g(0) > 0$ and $g(1) < 0$. By (ii), g is strictly decreasing. The Intermediate Value Theorem guarantees existence of $\kappa^* \in (0,1)$ with $g(\kappa^*) = 0$; strict monotonicity implies uniqueness. When (iii) fails ($\widetilde{\text{MLB}}(1) \geq \text{MLC}$), $\kappa^* = 1$. □

A.2 Aggregate Welfare and Optimal Exploration

Lemma B.2 (Per-Session Entrant Conversion Rate)**.** *For any session (q, J) and any exploration rate $\kappa \in [0,1]$, the entrant conversion rate equals q:*

$$\frac{\mathbb{E}_\delta[\textit{entrant needs met}]}{\mathbb{E}_\delta[\textit{entrant traffic}]} = q$$

Proof. Under exploration, the entrant occupies rank 1 with traffic $\bar{s}^E(q,J)$; under organic ranking, it occupies rank $r_{\text{org}}(q)$ with traffic $\varphi^O_{\text{org}}(q,J)$. Both are independent of δ. The combined traffic is $T = \kappa\,\bar{s}^E + (1-\kappa)\,\varphi^O_{\text{org}}$ and needs met are $\delta \cdot T$. Since T does not depend on δ: $\mathbb{E}_\delta[\delta \cdot T]/T = \mathbb{E}[\delta] = q$. □

Proof of Proposition 5.2 (Exploration Dilutes Baseline Screening). Fix session size J. By Lemma B.2, the aggregate entrant conversion rate is a traffic-weighted average of q:

$$\bar{C}_E^J(\kappa) = \frac{\sum_q q \cdot T(q,J,\kappa)}{\sum_q T(q,J,\kappa)}$$

where $T(q,J,\kappa) = \kappa \bar{s}^E(q,J) + (1-\kappa)\,\varphi_{\text{org}}^O(q,J)$.

For $\kappa' > \kappa$, the traffic ratio $T(q,\kappa')/T(q,\kappa)$ is decreasing in φ_{org}^O:

$$\frac{\partial}{\partial \varphi_{\text{org}}^O} \frac{T(q,\kappa')}{T(q,\kappa)} = \frac{\bar{s}^E \cdot (\kappa - \kappa')}{[\kappa \bar{s}^E + (1-\kappa)\,\varphi_{\text{org}}^O]^2} < 0$$

Since $\varphi_{\text{org}}^O(q,J)$ is weakly increasing in q, the likelihood ratio $T(q,\kappa')/T(q,\kappa)$ is decreasing in q. By the monotone likelihood ratio property, the traffic distribution under κ' places relatively more weight on low-q sessions. Since q is the function being averaged, $\bar{C}_E^J(\kappa') < \bar{C}_E^J(\kappa)$. □

Lemma C.3 (Per-Session Comparative Statics)**.** *For each session (J,q) with $q < \bar{s}^B$, the per-session optimal exploration satisfies:*

$$\frac{\partial \kappa^*}{\partial J} < 0 \qquad \textit{and} \qquad \frac{\partial \kappa^*}{\partial q} > 0 \tag{19}$$

Proof. Both results follow by implicit differentiation of $\widetilde{\text{MLB}}(\kappa^*) = \text{MLC}(q)$.

$\partial \kappa^*/\partial J < 0$**.** The denominator $\partial \widetilde{\text{MLB}}/\partial \kappa$ is negative since $\widetilde{\text{MLB}}$ is strictly decreasing. For the numerator: $\partial \text{MLC}/\partial J > 0$ because additional incumbents raise welfare more under organic ranking (which is welfare-maximizing given the prior) than under exploration. Meanwhile $\partial \widetilde{\text{MLB}}/\partial J \leq 0$ since the information gain from correct ranking shrinks as incumbent quality rises. The numerator is positive; combined with the negative denominator: $\partial \kappa^*/\partial J < 0$.

$\partial \kappa^*/\partial q > 0$**.** The MLC channel is unambiguously negative: $\partial \text{MLC}/\partial q < 0$, since higher q improves the entrant at rank 1 and narrows the exploration–organic welfare gap. The MLB channel is ambiguous in general: higher q raises φ_{org}^O, which increases the intensive-margin learning rate, but may reduce the learning value $L(J,q) = \mathbb{E}[W^*(\delta)] - W^O(q)$ as the prior concentrates. However, the extensive margin of exploration (the slope at $\kappa = 0$, equal to $\beta \cdot L$) does not depend on φ_{org}^O and remains positive as long as the prior has positive variance. The result $\partial \kappa^*/\partial q > 0$ holds whenever the MLC reduction dominates any decline in learning value—a sufficient condition that is satisfied throughout the parameter space of Table 6. □

Concavity of $\mathscr{W}(\kappa)$.. From (9), the per-session welfare $W(\kappa;J,q)$ has second derivative

$$W''(\kappa) = -\frac{2\beta \cdot L \cdot \varphi_{\text{org}}^O\,(\bar{s}^E - \varphi_{\text{org}}^O)^2}{T(\kappa)^3} < 0$$

since $\bar{s}^E > \varphi_{\text{org}}^O$ and $L > 0$ for $q < \bar{s}^B$. Each per-session welfare function is strictly concave in κ. Since $\mathscr{W}(\kappa) = \sum_{J,q} w(J,q) W(\kappa;J,q)$ is a positively weighted sum of strictly concave functions, $\mathscr{W}$ is strictly concave, yielding a unique optimum.

Proof of Proposition 5.3 (Optimal Exploration). We establish each comparative static for the aggregate optimum κ^* that maximizes $\mathscr{W}(\kappa) = \sum_{J,q} w(J,q) \cdot W(\kappa; J, q)$.

$\partial \kappa^*/\partial \beta \geq 0$**.** The learning benefit $\beta \cdot L \cdot \varphi^O_{\text{org}} \cdot \kappa / T(\kappa)$ is increasing in β for every $\kappa > 0$, so the marginal benefit of exploration increases with β. By the monotone comparative statics theorem (Milgrom and Shannon, 1994), κ^* is weakly increasing in β.

$\partial \kappa^*/\partial \alpha \geq 0$**.** Higher α shifts Pareto weight toward smaller J. By Lemma C.3, per-session κ^* is decreasing in J, so smaller sessions have higher optimal exploration. The reweighting increases the aggregate marginal benefit at every $\kappa > 0$.

$\partial \kappa^*/\partial \mu \geq 0$**.** Higher μ shifts the quality distribution upward. By Lemma C.3, per-session κ^* is increasing in q. Higher q reduces MLC in every session, increasing the aggregate marginal benefit. □

Learning rate thresholds.. Table 6 reports the learning rate thresholds across the (α, μ) grid: interior exploration requires $\beta \in (\beta_{\text{low}}, \beta_{\text{high}})$.[24] Among the three primitives, average entrant quality μ is quantitatively dominant. At $\alpha = 1.5$, a platform with $\mu = 0.10$ requires $\beta > 34$ merely to begin exploring, while one with $\mu = 0.75$ begins at $\beta > 0.7$. Moreover, β_{high} is typically one to two orders of magnitude larger than β_{low}, implying that full exploration is almost never optimal—the interior regime dominates.

Table 6: Learning Rate Thresholds: Interior Exploration Requires $\beta \in (\beta_{\text{low}}, \beta_{\text{high}})$

μ	$\alpha = 2.5$ ($\bar{J} = 2.7$)	$\alpha = 2.0$ ($\bar{J} = 3.1$)	$\alpha = 1.5$ ($\bar{J} = 4.1$)	$\alpha = 1.2$ ($\bar{J} = 5.3$)	$\alpha = 1.1$ ($\bar{J} = 5.9$)
0.10	(25, 1016)	(28, 1189)	(34, 1502)	(41, 1831)	(44, 1982)
0.25	(7, 181)	(8, 213)	(10, 272)	(12, 335)	(12, 363)
0.50	(2.4, 12)	(2.7, 14)	(3.3, 16)	(3.8, 19)	(4.1, 21)
0.75	(0.4, 1.5)	(0.5, 1.9)	(0.7, 2.7)	(0.9, 3.5)	(1.0, 3.9)

Notes: Each cell reports $(\beta_{\text{low}}, \beta_{\text{high}})$ where $\kappa^* = 0$ for $\beta < \beta_{\text{low}}$ and $\kappa^* = 1$ for $\beta > \beta_{\text{high}}$. Cross-session entrant quality $q \sim \text{Beta}(2\mu, 2-2\mu)$. Columns ordered by decreasing α (increasing $\bar{J}$). See footnote 24 for derivation.

A.3 Advertising

Proof of Proposition 5.4 (Unraveling of Exploration). Fix a session (J,q) with $q < \theta^B_1$. Since $W^*(\delta, J)$ is piecewise linear and convex (slope increases at each rank threshold $\theta^B(r,J)$) and $W^O(\delta,J) - W^E(\delta,J)$ is linear, ΔW (12) is convex in δ. The FOC (13) gives $\Delta W(\delta^*) = 0$.

$\delta > \delta^*$**:** The derivative $\Delta W'(\delta^*) = \beta W^{*\prime}(\delta^*) - [W^{O\prime}(\delta^*) - W^{E\prime}(\delta^*)] > 0$ since W^* is increasing for $\delta > q$ and $W^O - W^E$ is strictly decreasing (higher δ reduces displacement cost). Convexity implies $\Delta W'$ is non-decreasing, so ΔW is strictly increasing on $[\delta^*, 1]$ and $\Delta W(\delta) > 0$ for all $\delta > \delta^*$.

[24] Evaluating the first-order condition of (9) aggregated across sessions at the boundaries $\kappa = 0$ and $\kappa = 1$ gives the thresholds analytically: $\beta_{\text{low}} = \sum_{J,q} w(J,q)\,\text{MLC}(J,q) \big/ \sum_{J,q} w(J,q) L(J,q)$ and $\beta_{\text{high}} = \sum_{J,q} w(J,q)\,\text{MLC}(J,q) \big/ \sum_{J,q} w(J,q) L(J,q) \left(\varphi^O_{\text{org}}(J,q)/\bar{s}^E(J,q)\right)^2$. At $\kappa = 0$, each unit of exploration is fully informative, so exploration begins when $\beta > \beta_{\text{low}}$. At $\kappa = 1$, organic traffic already provides substantial information, so full exploration requires $\beta > \beta_{\text{high}}$.

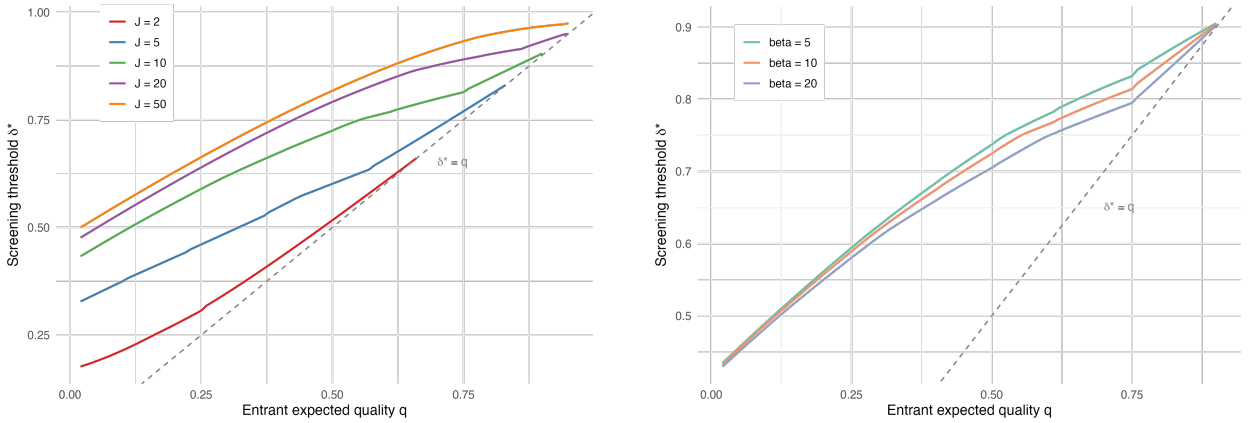
$\delta \in [q, \delta^*)$**:** At $\delta = q$, the learning value vanishes ($W^*(q,J) = W^O(q,J)$) while displacement cost is strictly positive, so $\Delta W(q) = -\text{MLC} < 0$. Since $q < \delta^*$, convexity gives $\Delta W(\delta) \leq \frac{\delta^* - \delta}{\delta^* - q} \Delta W(q) + \frac{\delta - q}{\delta^* - q} \Delta W(\delta^*) = \frac{\delta^* - \delta}{\delta^* - q} (-\text{MLC}) < 0$.

$\delta \in [0, q)$**:** At $\delta = 0$, $W^*(0,J) = W^B(J) \leq W^O(q,J)$ (adding the entrant at its organic rank weakly raises welfare) and $W^O(0,J) \geq W^E(0,J)$ (a zero-quality entrant wastes less search effort at a deep rank than at rank 1), so $\Delta W(0) \leq 0$. Convexity on $[0,q]$ gives $\Delta W(\delta) \leq \frac{q - \delta}{q} \Delta W(0) + \frac{\delta}{q} \Delta W(q) \leq \frac{\delta}{q} (-\text{MLC}) < 0$ for $\delta > 0$, and $\Delta W(0) \leq 0$ at the boundary. □ □

Analytical characterization of δ^*.. Since $W^E(\delta,J)$, $W^O(\delta,J)$ are linear and $W^*(\delta,J)$ is piecewise linear with $W^*(q,J) = W^O(q,J)$, the condition $\Delta W(\delta^*) = 0$ yields

$$\delta^* = q + \frac{\text{MLC}(q)}{\beta \cdot W^{*\prime} + |\text{MLC}'|} \tag{20}$$

where $W^{*\prime}$ is the slope of $W^*(\delta)$ on the piece containing δ^* (determined by the entrant's full-information rank at quality δ^*) and $\text{MLC}' = \frac{d}{d\delta}[W^O(\delta) - W^E(\delta)] < 0$. The numerator is the displacement cost at the prior; the denominator is the rate at which net benefit grows with δ—the sum of the marginal learning benefit and the marginal displacement cost reduction. Figure 5 plots δ^* as a function of q for several values of J.



(a) Varying J ($\beta = 10$)

(b) Varying β ($J = 10$)

Figure 5: Welfare-Maximizing Screening Threshold δ^*

Notes: δ^* solves $\Delta W(\delta^*) = 0$ for each (q, J, β). The dashed line is $\delta^* = q$; all curves lie strictly above it (Proposition 5.4). Kinks correspond to incumbent quality thresholds $\theta^B(r,J)$ where the entrant's full-information rank improves and $W^{*\prime}$ increases discretely. Panel (a): larger J raises the screening threshold because displacement cost increases with incumbent competition. Panel (b): higher β lowers δ^*, as stronger learning incentives justify promoting lower-quality entrants.

Proof of Proposition 5.5 (Advertising Frictions Restore Exploration). Part (i). At $\kappa = 0$, the marginal wel-

fare from raising κ is

$$\left.\frac{dW}{d\kappa}\right|_{\kappa=0} = \underbrace{(1-\lambda)\,\mathbb{E}_\delta[\Delta W(\delta)]}_{\text{non-participants}} + \underbrace{\lambda\,\mathbb{E}_\delta[\Delta W(\delta)\cdot\mathbf{1}(\delta<\delta^*)]}_{\text{sub-threshold participants}}$$

The first term is positive (since $\mathbb{E}[\Delta W(\delta)]>0$ by $\kappa^*_{\text{no-ads}}>0$). The second is negative: exploration is blind, so raising κ also promotes sub-threshold participants for whom $\Delta W(\delta)<0$. Writing $A=\mathbb{E}[\Delta W(\delta)\cdot\mathbf{1}(\delta<\delta^*)]<0$ and $B=\mathbb{E}[\Delta W(\delta)\cdot\mathbf{1}(\delta\geq\delta^*)]>0$, we have $dW/d\kappa|_{\kappa=0}=A+(1-\lambda)B>0$ iff $\lambda<(A+B)/B=\bar{\lambda}$. Since $A<0$ and $A+B>0$, $\bar{\lambda}\in(0,1)$.

Part (ii). When $v_{\text{inc}}>a^*$, the incumbent blocks entrants with $\delta\in[\delta^*,\bar{\delta}_{\text{inc}}]$. Among participants, only $\delta>\bar{\delta}_{\text{inc}}$ types are served by ads. Writing $B'=\mathbb{E}[\Delta W(\delta)\cdot\mathbf{1}(\delta>\bar{\delta}_{\text{inc}})]\leq B$, the marginal welfare from raising κ at $\kappa=0$ is $A+(1-\lambda)B'+\lambda\,\mathbb{E}[\Delta W(\delta)\cdot\mathbf{1}(\delta\in[\delta^*,\bar{\delta}_{\text{inc}}])]$. Since blocked entrants have $\Delta W>0$, this simplifies to $dW/d\kappa|_{\kappa=0}>0$ iff $\lambda<(A+B)/B'=\bar{\lambda}'$. Since $B'\leq B$, $\bar{\lambda}'\geq\bar{\lambda}$. When $\bar{\delta}_{\text{inc}}\geq 1$, $B'=0$ and $\bar{\lambda}'=\infty$, so $\kappa^*>0$ for any λ. □

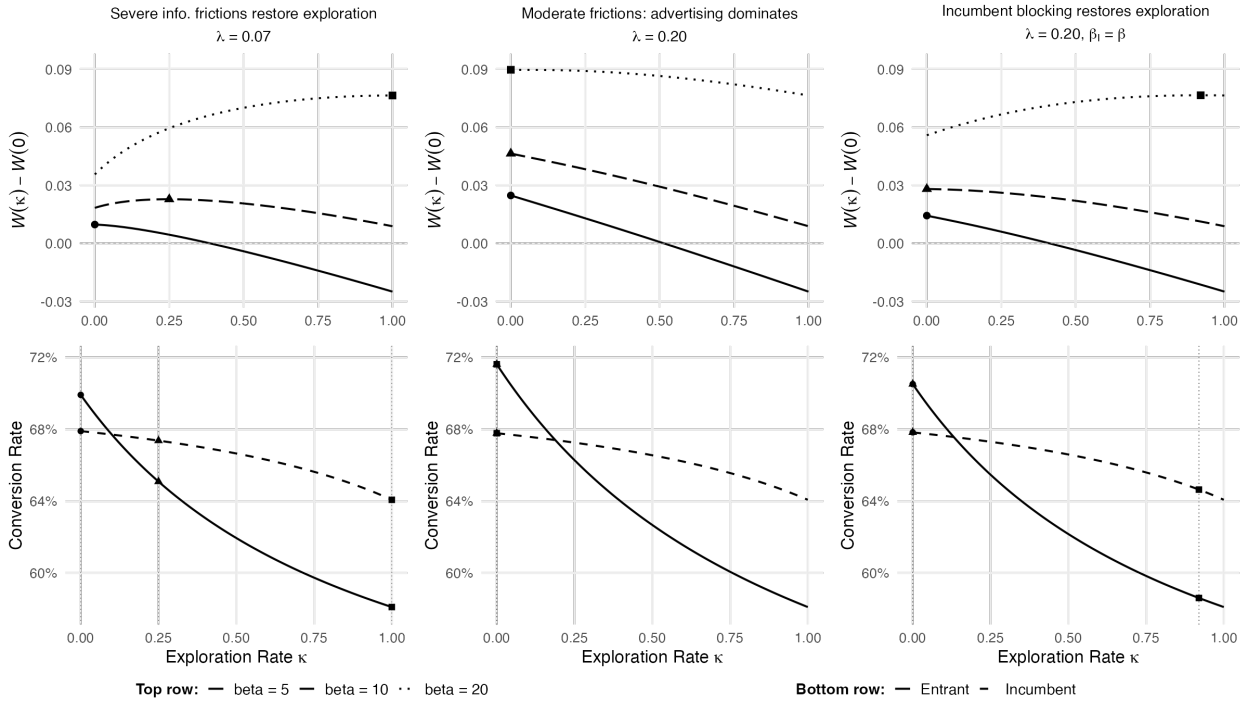


Figure 6: Advertising Frictions with Dynamic Entrant Bidding (Robustness)

Notes: Same specification as Figure 2, but relaxing the entrant liquidity constraint. Entrants with $\delta>\theta_1^B$ anticipate permanent promotion and bid $v^{\text{dyn}}(\delta)=(1-\kappa)(\bar{s}^E-\varphi^O_{\text{org}})\delta+\beta^E\,\delta(\delta-\varphi^O_{\text{org}})$, which can exceed v_{inc} for high δ. Blocking is partial: only entrants with $\delta<\theta_1^B$ are blocked. Exploration is restored only at high β ($\kappa^*=0.92$ at $\beta=20$), compared to the liquidity-constrained case (Figure 2) where blocking is more severe and exploration is restored at lower β.

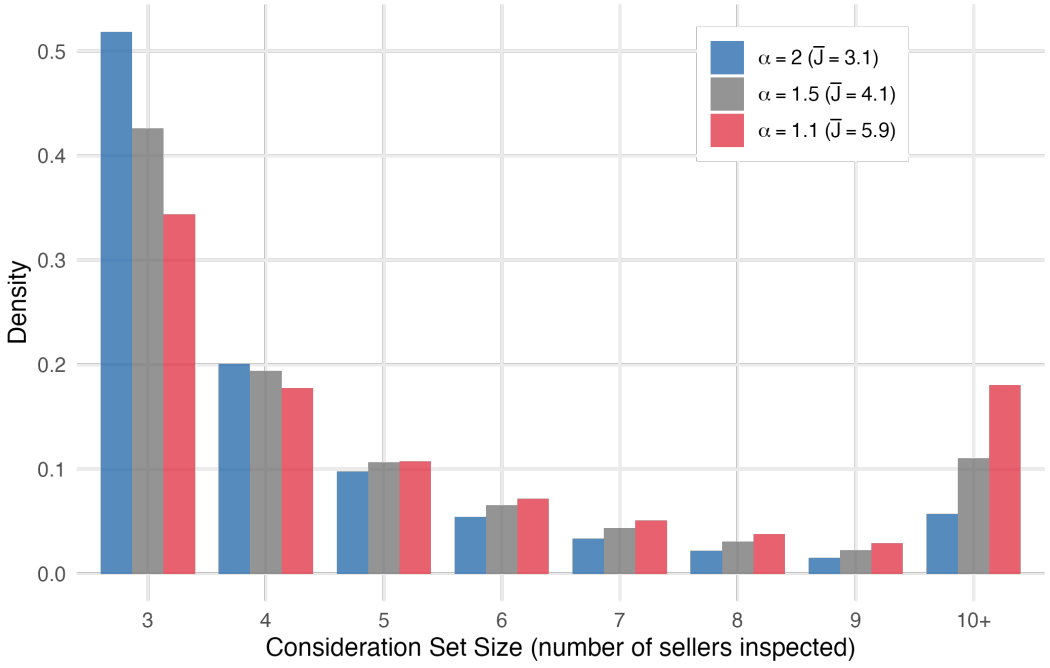


Figure 7: Model-Predicted Distribution of Consideration Set Sizes

Notes: Model-predicted density of consideration set sizes from the baseline equilibrium (equations (4)–(5)), aggregated over $J \in \{2, \ldots, 100\}$ with truncated Pareto weights at three concentration levels α. Filtered to sets of size ≥ 3 for comparability with the empirical distribution in Figure 8.

A.4 Extensions: Alternative Exploration Mechanisms

Soft-Boost Exploration.. The baseline model uses a binary mechanism: with probability κ, the entrant jumps to rank 1. In practice, platforms often implement exploration through deterministic rank boosting. Consider a "soft boost" in which the platform shifts the entrant from organic rank $r_{\text{org}}(q)$ to a boosted rank $r_b \leq r_{\text{org}}$, displacing incumbents at ranks $r_b, \ldots, r_{\text{org}}-1$ one position deeper.

Under soft boost, the entrant receives traffic $\varphi^O(r_b, J)$ rather than $\bar{s}^E$ (under rank-1 exploration) or φ^O_{org} (under organic ranking). Because traffic decays precipitously with rank—$\varphi^B(r,J) = \bar{s}^B \prod_{r'=1}^{r-1}(1-\theta^B(r',J))$—the traffic ratio $\varphi^O(r_b)/\varphi^O(r_{\text{org}})$ remains close to 1 for moderate boosts (r_b only slightly below r_{org}) and grows rapidly only as $r_b \to 1$. For a deeply buried entrant ($r_{\text{org}} \gg 1$), generating enough traffic to meaningfully reduce Bayesian posterior variance requires r_b close to 1, at which point the soft boost approximates our κ model.

When the platform conditions the boost on the prior quality q—setting $r_b(q)$ to boost high-q entrants more aggressively—the mechanism achieves higher allocative efficiency than blind κ-exploration. Formally, the welfare under q-dependent boosting satisfies

$$W^{\text{blind}}(\kappa) \;\leq\; W^{\text{boost}}(r_b(q)) \;\leq\; W^{\text{ads}}(\delta^*)$$

where the first inequality reflects the gain from targeting exploration by prior quality, and the second reflects the residual inefficiency from screening on q (observable) rather than realized δ (private). The q-based boost occupies an intermediate position in the efficiency hierarchy: more efficient than blind exploration (which ignores the prior entirely) but strictly dominated by advertising (which screens on private, realized quality). Because prior-directed boosting introduces a midpoint in this hierarchy without altering the fundamental trade-off between exploration and advertising, the baseline κ model captures the essential forces with fewer mechanical complications.

Quality-Weighted Reserve Prices.. The welfare-optimal reserve $a^*(q,J)$ from Proposition 5.4 already varies across session types through $\delta^*(q,J)$: sessions with higher q have lower screening thresholds (since displacement cost MLC is smaller), yielding lower reserve prices. In practice, however, platforms may not solve the full welfare optimization for each session. A simpler policy sets the reserve proportional to predicted conversion:

$$\tilde{a}^*(q) \;=\; c \cdot q \cdot (\bar{s}^E - \varphi^O_{\text{org}})$$

for a constant $c > 0$, inducing screening threshold $\tilde{\delta}^*(q) = c \cdot q/(1-\kappa)$. This q-weighted rule has two properties:

(i) Constant screening ratio. Higher-q entrants face higher absolute reserves but a constant reserve relative to their quality: $\tilde{\delta}^*/q = c/(1-\kappa)$ is independent of q. In contrast, the welfare-optimal $\delta^*(q)/q$ is decreasing in q, since high-q sessions have lower displacement costs and the platform screens less aggressively.

(ii) Welfare loss. The q-proportional rule cannot replicate the welfare-optimal reserve because $\delta^*(q)$ is not proportional to q (Proposition 5.4: $\delta^* > q$ with a gap that depends nonlinearly on MLC and β). For low-q sessions, the constant-ratio rule under-screens (admitting entrants with $\Delta W < 0$), while for high-q sessions, it over-screens (excluding welfare-improving entrants). The welfare loss from the q-proportional rule is bounded by the cross-session variance of $\delta^*(q)/q$, which is small when the learning value $\beta \cdot L$ varies little with q.

Shadow Bidding.. Ye et al. (2023) propose a "shadow bidding" mechanism: the platform adds a subsidy σ_j to new advertiser j's bid, where the optimal σ_j is determined by the dual variable of an exploration constraint. In our framework, this translates to the platform adding $\sigma(q) \geq 0$ to the entrant's bid $v(\delta)$, so the entrant obtains rank 1 whenever $v(\delta) + \sigma(q)$ exceeds the reserve or the incumbent's bid.

Equivalence without incumbent blocking. In the absence of incumbent bidding, adding subsidy σ to the entrant's bid is equivalent to lowering the reserve from a^* to $a^* - \sigma$. The welfare-optimal subsidy therefore implements the same screening threshold δ^* as the optimal reserve, and both achieve $\kappa^* = 0$ when all entrants participate ($\lambda = 1$).

Advantage under incumbent blocking. The mechanisms diverge when the incumbent bids. Under the reserve-price mechanism alone, the incumbent blocks entrants with $\delta \in [\delta^*, \bar{\delta}_{\text{inc}}]$ (Proposition 5.5). A shadow bid can partially undo this blocking: an entrant of quality δ now outbids the incumbent whenever

$$v(\delta) + \sigma(q) \;\geq\; v_{\text{inc}}(q,J) \quad\Longleftrightarrow\quad \delta \;\geq\; \frac{v_{\text{inc}} - \sigma(q)}{(1-\kappa)(\bar{s}^E - \varphi^O_{\text{org}})}$$

The shadow bid shrinks the blocking interval from $[\delta^*, \bar{\delta}_{\text{inc}}]$ to $[\delta^*, \bar{\delta}_{\text{inc}} - \sigma/\{(1-\kappa)(\bar{s}^E - \varphi^O_{\text{org}})\}]$, unblocking welfare-improving entrants with $\Delta W(\delta) > 0$ who would otherwise be deterred. The welfare-optimal subsidy equates the platform's marginal cost with the marginal gain from unblocking:

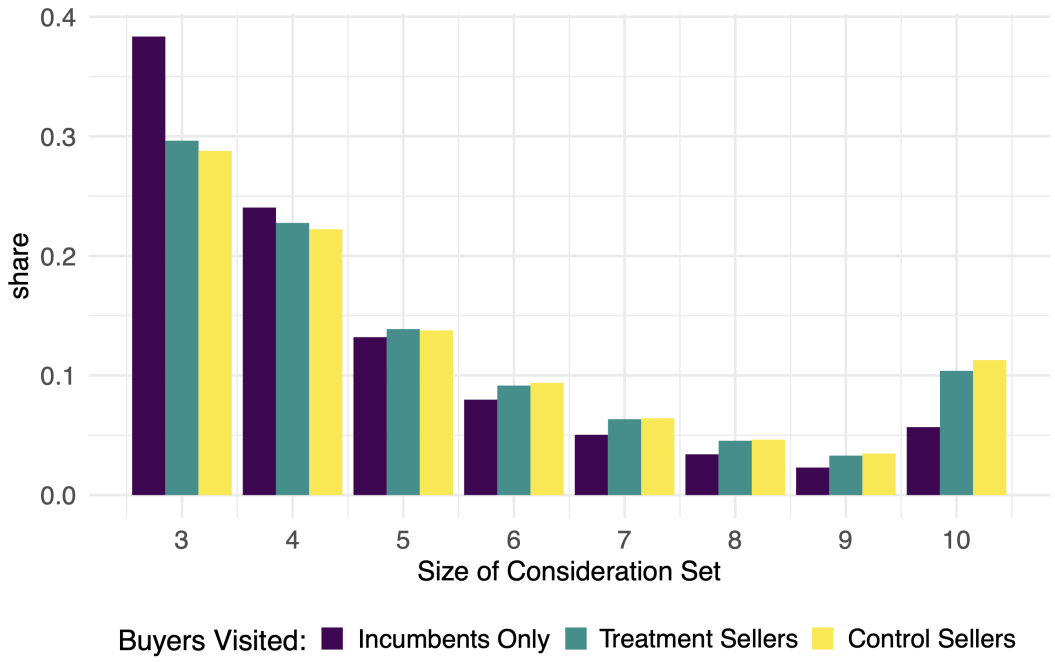
$$\Delta W\left(\bar{\delta}_{\text{inc}} - \frac{\sigma^*}{(1-\kappa)(\bar{s}^E - \varphi^O_{\text{org}})}\right) = \sigma^*$$

Unlike a simple reserve adjustment, the shadow bid is funded by the platform (as a transfer to the entrant's bidding account), creating a direct cost that must be weighed against the welfare gain from reduced blocking.

Exploration remains necessary. Shadow bidding, like reserve pricing, only serves participating entrants. For the $(1-\lambda)$ fraction of non-participants, the platform's exploration rate κ remains the only channel for entry. This contrasts with Ye et al. (2023), where all advertisers participate in the auction by construction, making shadow bidding sufficient. In our setting, the interaction between advertising (with or without shadow bids) and exploration is governed by Proposition 5.5: the lower the effective participation rate—whether due to information frictions ($\lambda < 1$) or incumbent blocking ($\bar{\delta}_{\text{inc}} > \delta^*$)—the larger the role of κ in mediating entry.

B Additional Tables and Figures [For Online Publication]

Figure 8: Distribution of the Size of Consideration Sets by Seller Type



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

B.1 Consideration-Set Data Construction

We construct consideration-set data from search sessions during our research window (May 6–August 15, 2019), sourced from the platform-wide experiment described in Section 4.1 and Jin and Sun (2026). The procedure, illustrated in Figure 9, involves several steps to address computational constraints and data skewness.

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

Keyword filtering. To address the skewness of keywords (new sellers appeared in less than 0.1% of consideration sets), we filter down to search sessions that contain popular keyword combinations, defined as those searched at least three times by consumers in the CCS throughout the research window.

Matching non-CCS consumers. Focusing on the CCS means leaving out a large group of consumers who did not visit new sellers at all during the research window. For each consideration set in the CCS, we pull additional sets conducted by non-CCS consumers that had the same keyword combinations, the same size (number of products visited), and occurred on the same date. We include up to 50 consideration sets per keyword-size-date combination for computational ease, subsampling randomly if the number of matching sets exceeds 50. This enables across-consideration-set comparisons via keyword-size-date interacted fixed effects.

Session size restriction. Many search sessions may not be genuine, especially those containing only one or two sellers. These make up over 42% of all sessions during the research window yet are responsible for only 1.9% of purchases. This likely reflects "passive searches" as studied in Ursu, Simonov and An (2022), where consumers are directed to product pages via external links that can be disguised as search sessions, as well as test searches by new consumers. We thus restrict all analyses to sessions with at least three products.

Figure 9: Procedures of Constructing the Consumer Sample

