The Cost of the Cold-Start Problem on Airbnb

Florian Dendorfer and Regina Seibel*

March 26, 2026

Abstract

On platforms with peer-to-peer reviews, new products encounter the so-called "cold-start" problem: consumers who book and review an unreviewed product generate information that benefits future consumers, but do not internalize this externality. Yet how large this externality is, and even whether it persists at all, depends on sellers' endogenous responses. This paper quantifies the informational externality underlying the cold-start problem and evaluates policies that target it. Based on data from Airbnb listings in Manhattan, we estimate a structural model with short-lived consumers and long-lived hosts who make dynamic entry, exit, and pricing decisions. We decompose the value of the first review and find that hosts already set low introductory prices to attract early bookings, partially internalizing the externality. Thus, the immediate consumer surplus gain per review is small. However, entry and exit dynamics amplify the externality by orders of magnitude. Faster learning strengthens selection, causing high-quality listings to survive longer and low-quality listings to exit sooner, which greatly improves market composition. Finally, we evaluate two targeted policies, a consumer transfer for booking unreviewed listings and a host bonus for obtaining the first review, each under three funding modes.

Keywords: Cold-start problem, digital platforms, experimentation, market dynamics, product reviews, social learning.

JEL: L11, L15, L83, L86, L88, D83

1 Introduction

On digital platforms, a consumer who books a product and leaves a review generates valuable information for future consumers but has no reason to account for this benefit. This externality is largest for new products with few reviews, where each additional piece of information has the

*Dendorfer: Competition Bureau Canada. Seibel: Rotman School of Management, University of Toronto. Email: florian.dendorfer@cb-bc.gc.ca, and regina.seibel@rotman.utoronto.ca. We thank Victor Aguirregabiria, Heski Bar-Isaac, Lanier Benkard, Stefan Bühler, El Hadi Caoui, Gregory Crawford, Florian Ederer, Liran Einav, Gautam Gowrisankaran, Philip Haile, Brett Hollenbeck, Mitsuru Igami, Armin Schmutzler, and conference participants at CRESSE 2023, MaCCI 2023, the Swiss IO Day 2023, DICE 2024, the 17th Digital Economics Conference 2025, the Economics of Platforms Seminar, IIOC 2025 and the Alpine IO Symposium 2025 for helpful comments. The views expressed are from a personal perspective and do not represent the views or the position of the Commissioner of Competition of the Government of Canada.

greatest impact on beliefs about quality. The resulting "cold-start" problem (Che and Hörner, 2018; Kremer, Mansour and Perry, 2014) means that new products are discovered too slowly: social learning is inefficiently slow. Platforms recognize this. Airbnb suggests hosts offer a 20% discount to their first guests, Amazon invites selected consumers to review new products for free through its Vine program, and on Taobao, sellers can reward buyers for truthful feedback (Li, Tadelis and Zhou, 2020).[1]

But how large is this externality, and can these interventions effectively address it? The answer depends on how sellers respond. New sellers may already set low introductory prices to attract reviews and build a competitive advantage, similar to a learning-by-doing strategy (Cabral and Riordan, 1994, 1997), partially internalizing the externality through their own pricing decisions.[2] At the same time, faster information provision changes which sellers survive in the market, affecting product variety and welfare through channels that go beyond social learning.

In this paper, we quantify the informational externality underlying the cold-start problem and evaluate targeted policies that address it. Specifically, we develop a structural model of the Airbnb market in Manhattan, New York, where hosts make dynamic pricing, market entry, and exit decisions. Since the quality of a listing is initially unknown, peer-to-peer reviews help reveal it over time. Using our model, we first measure the value of the first review, decomposing it across pricing regimes to isolate how the supply side shapes the externality. We then evaluate two policies that target the externality from different sides of the market: a consumer transfer that compensates guests for booking unreviewed listings, and a host bonus that rewards hosts for obtaining their first review. Each policy is analyzed under three funding modes (platform-funded, revenue-neutral with host financing, and revenue-neutral with consumer financing) to understand how the cost allocation mediates the welfare effects.

We begin by analyzing the informational externality in a stylized two-firm model with logit demand, an outside option, and endogenous pricing. We show that under uniform prices, a review is unambiguously valuable to consumers (pure information effect). Under static Nash pricing, rent extraction reduces the gain but does not eliminate it. Under dynamic pricing, the entrant engages in introductory pricing to encourage reviews, which benefits consumers in the pre-review state. This creates a tension: when a review arrives, consumers lose the introductory discount. For sufficiently patient firms, this introductory pricing effect can offset the information value, so that consumers prefer the pre-review state. We discuss how entry and exit dynamics, absent from the stylized model, amplify the welfare effects of reviews in the structural model.

Next, we present an empirical model of the Airbnb market that incorporates supply dynamics. On the demand side, we employ a discrete choice framework in which guests choose which of the available listings to book based on the expected quality of each listing. Guests update their prior beliefs according to Bayes' rule, using the observed number of reviews and ratings.

[1] Airbnb: https://www.airbnb.ca/resources/hosting-homes/a/how-to-set-a-pricing-strategy-15, Amazon: https://sell.amazon.com/tools/vine.

[2] Bergemann and Välimäki (1997, 2000) incorporate endogenous pricing and find that the speed of social learning is too high, not too low. Vellodi (2024) demonstrates that allowing for endogenous pricing fully alleviates the cold-start problem in his model with endogenous entry.

Importantly, guests do not consider the impact of their potential reviews on future guests when making booking decisions. We assume that if guests choose to review, they do so truthfully. On the supply side, there is a pool of hosts, some of whom are inactive. Active hosts face marginal costs and set rental prices taking into account their effect on the probability of receiving a new review, and therefore future demand. Hosts also make entry and exit decisions: active hosts decide whether to remain active or exit the market to avoid operating costs, while inactive hosts can enter at the beginning of each period, incurring an entry cost. Given the large number of Airbnb hosts, we approximate the symmetric Markov Perfect equilibrium using the oblivious equilibrium concept introduced by Weintraub, Benkard and Van Roy (2008).

We estimate demand and supply separately. For the demand estimation, we use the Generalized Method of Moments. Our estimates highlight the substantial impact of ratings both on guests and hosts. On average, a good review increases a listing's occupancy rate by 3.4%, whereas the estimated average price elasticity is -2.3. The impact of reviews is most pronounced for unreviewed listings, where it is about 50% larger than the average effect. For entrant listings without reviews, we estimate a prior quality expectation of 4.15 stars on a five-star scale. The observed mean rating of 4.69 stars reflects a selection effect: high-quality listings are more likely to persist in the market.

For the supply estimation, we use a nested maximum likelihood approach. In the inner loop, the marginal costs are estimated and the model is solved for a given set of entry and operating costs. The outer loop determines the cost parameters that maximize the likelihood of the observed distribution of listing types and reviews. Our estimates reveal substantial entry costs for hosts, averaging approximately half of their lifetime profit. Based on our estimates of marginal and operating costs, we calculate an average variable profit margin of around 53%. The model solution yields rental prices, occupancy rates, average ratings and exit rates that closely align with the observed data. Our findings demonstrate that reviews influence host and consumer behavior in a manner consistent with empirical observations.

In our counterfactual analysis, we first measure the informational externality. In the short run, we give every unreviewed listing its first review (in expectation) and decompose the resulting consumer surplus gain per review across three pricing regimes: uniform prices (pure information effect), static Nash prices (heterogeneous markups, no dynamic incentives), and dynamic Nash prices (investment pricing). The pure information effect is positive. However, under dynamic Nash pricing, hosts already set low introductory prices to attract bookings, partially substituting for the externality. As a result, the immediate consumer surplus gain per review is modest under dynamic pricing. But the short-run exercise abstracts from entry and exit entirely, and it is precisely through these margins that the externality becomes large. In the long run, when entry, exit, and the full state distribution adjust, we find that each additional first review generates consumer surplus gains of over $900, orders of magnitude larger than the pure information effect of just $2. The dominant channel behind the amplification is that faster learning strengthens selection: high-quality listings are reviewed sooner, survive longer, and accumulate more bookings, while low-quality listings are identified and exit more quickly. Faster learning also intensifies competition thereby reducing prices and leads to moderately higher

entry.

We then evaluate two targeted policies. The first is a consumer transfer: a lump-sum payment τ to consumers who book an unreviewed listing, which enters utility additively and increases demand for entrants. The second is a first-review bonus: a lump-sum B paid to hosts upon receiving their first good review, which enters the host's pricing first-order condition and incentivizes lower introductory prices. Each policy is analyzed under three funding modes: platform-funded, revenue-neutral with host financing (via a participation fee F on entrant hosts), and revenue-neutral with consumer financing (via a per-booking surcharge F_c). The host bonus is somewhat more effective than the consumer transfer, as it directly alters hosts' pricing incentives. However, the choice of funding mode matters at least as much as the choice of policy instrument. Under platform funding, consumer surplus and total surplus gains are largest, but the cost to Airbnb revenue makes implementation unlikely. Under host financing, the participation fee screens out low-value entrants, concentrating the market among hosts who generate more information per booking. This design is most effective at addressing the cold-start problem specifically, even though it reduces overall listing variety.

We conclude that the supply-side response, through pricing, entry, and exit, mediates the welfare effects of any policy targeting the cold-start problem. Hosts' investment pricing partially internalizes the externality, but does not eliminate it. Exit and entry play a major role in shaping the huge welfare gains of faster learning. Policies that target the externality can generate meaningful welfare gains, but their effectiveness depends critically on how they interact with hosts' dynamic incentives and on who bears the cost. These findings generalize to any digital platform that relies on peer-to-peer reviews and where sellers make dynamic decisions.

The paper is organized as follows. Section 2 provides a comprehensive review of the relevant literature. Section 3 develops the theoretical intuition for the informational externality in a stylized two-firm model. In Section 4, we introduce the empirical model. In Section 5, we present the data. Section 6 and Section 7 contain the estimation procedure and results for the demand and supply side of the model respectively. In Section 8, we describe the model fit. In Section 9, we describe our counterfactuals and present our counterfactual results. Section 10 concludes.

2 Literature

There is an extensive theoretical literature on social learning. Bergemann and Välimäki (1996, 1997, 2000) embed strategic experimentation into a market setting with endogenous prices. Che and Hörner (2018) and Kremer et al. (2014) employ a mechanism-design perspective and find that a recommender system which occasionally recommends "ex-ante unappealing" products can be "socially valuable because some of them are ultimately worthy of consumption" (p. 872, Che and Hörner, 2018). Hagiu and Wright (2020) analyze the platform-optimal level of experimentation which differs from the seller-optimal level of learning if sellers have market power and the platform takes into account buyer surplus. Kovbasyuk and Spagnolo (2023) and Vellodi (2024) incorporate market entry and find that the cold-start problem acts as an entry barrier.

Both papers suggest that limiting information on established sellers may alleviate these barriers and increase surplus.[3] To our knowledge, we are the first to provide an empirical model of social learning through reviews and estimate welfare effects of inducing consumers to explore new products. Apart from our paper, to our knowledge, only Pallais (2014) assesses the cold-start problem empirically. In contrast to our paper, Pallais uses an experimental setup in a labor market context.

Multiple studies estimate the impact of reviews on sales (Anderson and Magruder, 2012; Chevalier and Mayzlin, 2006), revenues (Luca, 2016), and exit rates (Cabral and Hortacsu, 2010). Additionally, several studies attribute substantial consumer welfare gains to online rating systems (Fang, 2022; Lewis and Zervas, 2016; Reimers and Waldfogel, 2021; Wu, Che, Chan and Lu, 2015). Bao, Fang and Osborne (2024) analyzes how quality disclosure via reviews affects entry and exit dynamics. We contribute to this literature by confirming the significance of ratings on Airbnb, demonstrating that unreviewed Airbnb listings are booked less frequently, are cheaper, and are more likely to be discontinued than reviewed ones.

A large body of empirical literature on Airbnb reviews examines the incentives for guests to write reviews and the sources of review bias (Fradkin and Holtz, 2022; Fradkin, Grewal and Holtz, 2021; Proserpio, Xu and Zervas, 2018; Zervas, Proserpio and Byers, 2021). Fradkin and Holtz (2022) highlight that, since consumers bear the cost of writing a review but do not receive all the benefits, online reviews are likely underprovided. Airbnb has also been the subject of extensive structural modeling. Most of these studies focus on the housing market (Almagro and Domínguez-Iino, 2025; Calder-Wang, 2021) or hospitality market (Farronato and Fradkin, 2022). Huang (2022) examines pricing behavior on Airbnb and finds that pricing frictions lead to a substantial welfare loss. Rossi (2024) explores the role reviews have in motivating sellers to exert effort, and examines how competition mediates this relationship.

We also contribute to the literature on estimating dynamic oligopoly models. Aguirregabiria, Collard-Wexler and Ryan (2021) provide an excellent recent overview. Most existing studies focus on oligopolistic games involving entry, exit, and sometimes innovation (Barwick and Pathak, 2015; Collard-Wexler, 2013; Igami and Uetake, 2020; Kellogg, 2014; Takahashi, 2015; Wollmann, 2018). Another strand in this literature incorporates dynamic considerations into pricing decisions. For instance, in Huang (2022), Williams (2022) and Hortaçsu, Natan, Parsley, Schwieg and Williams (2023), dynamic pricing arises from intertemporal price discrimination involving perishable products. In other settings, dynamic pricing reflects the notion that price does not only influence current demand but also acts as an investment. This phenomenon can be driven by network effects (Dubé, Hitsch and Chintagunta, 2010), learning-by-doing (Besanko, Doraszelski and Kryukov, 2014, 2019; Besanko, Doraszelski, Kryukov and Satterthwaite, 2010; Sweeting, Jia, Hui and Yao, 2022), switching costs (Chen, 2016), entry deterrence (Sweeting, Roberts and Gedge, 2020) or learning about product quality (Ching, 2010).

Similar to Ching (2010), we consider a setting where the prices moderate the speed at which consumers learn about products. However, Ching (2010) does not explicitly consider the inefficiency due to social learning and analyzes the market for prescription drugs. Unlike the above

[3] Butler, Carbone, Conzo and Spagnolo (2020) support this intuition experimentally in a public procurement context.

papers, we solve our model using the oblivious equilibrium concept developed in Weintraub et al. (2008) and Weintraub, Benkard and Van Roy (2010) which is particularly suitable in our setting of many small firms. The oblivious equilibrium has been applied in various studies (Bao et al., 2024; Brancaccio, Kalouptsidi and Papageorgiou, 2020; Chen and Xu, 2023; Frechette, Lizzeri and Salz, 2019).

3 The Informational Externality

Before presenting the structural model, we analyze the informational externality in a stylized setting. Consider an entrant (E) and an incumbent (I) supplying products of quality $\omega_j \in [0,1]$. The incumbent's quality ω_I is publicly known. The entrant's quality ω_E is unknown, with prior expected quality $\omega_E^0 \in (0,1)$. If a consumer purchases E, she learns E's quality and may share it in a review. The first review fully resolves uncertainty, updating beliefs to ω_E^g (good review, probability ω_E^0) or ω_E^b (bad review), where $\omega_E^g > \omega_E^0 > \omega_E^b$ and $\omega_E^0 \cdot \omega_E^g + (1-\omega_E^0) \cdot \omega_E^b = \omega_E^0$.

Each period, a consumer chooses among E, I, and an outside option. The expected indirect utility of product j is $u_{ij} = \omega_j - p_j + \epsilon_{ij}$, where p_j is the price and ϵ_{ij} is an i.i.d. Gumbel taste shock; the outside option yields $u_{i0} = \epsilon_{i0}$. We analyze the consumer surplus value of a review, $\Delta CS \equiv \omega_E^0 \cdot CS(\omega_E^g) + (1-\omega_E^0) \cdot CS(\omega_E^b) - CS(\omega_E^0)$, under three pricing regimes. All proofs are in Appendix A.1.

Uniform prices (pure information effect). When both products are priced identically ($p_E = p_I = \bar{p}$), the consumer surplus value of a review isolates the pure information effect.

Lemma 1. *Under uniform prices, $\Delta CS^U > 0$: expected consumer surplus conditional on a review is strictly higher than consumer surplus without a review.*

The result follows from Jensen's inequality: consumer surplus is convex in the entrant's expected quality, and a review is a mean-preserving spread in beliefs. Under uniform prices, a review allows demand to reallocate toward higher-quality products at no price premium.

Static Nash pricing ($\delta = 0$). When firms set profit-maximizing prices, each firm j charges $p_j^* = 1/(1-q_j)$. A better-reviewed entrant charges a higher markup, extracting part of the informational rent. However, the outside option limits this extraction.[4]

Lemma 2. *Under static Nash prices, $0 < \Delta CS^* < \Delta CS^U$: a review is still valuable to consumers, but the gain is smaller than under uniform prices.*

Dynamic pricing ($\delta > 0$). We extend the model to an infinite horizon. Once the entrant receives its first review, prices revert to static Nash permanently. Until then, firms are forward-looking and may deviate from static Nash to influence the review probability. The entrant always engages in introductory pricing ($p_E^0 < p_E^*$), lowering its price to encourage purchases

[4] The outside option is essential: without it, the entrant's rent extraction makes consumer surplus concave in quality, so that reviews would reduce expected consumer surplus.

and reviews. The incumbent's profit is concave in ω_E^e, so for moderate discount factors the incumbent is hurt by reviews in expectation and also lowers its price. Both adjustments benefit consumers in the pre-review state. When a review arrives and prices revert to static Nash, this introductory discount is lost.

Proposition 1. *The consumer surplus value of a review under dynamic pricing satisfies* $\Delta CS^0 \leq \Delta CS^*$, *with equality when* $\delta = 0$*:*

$$\Delta CS^0 = \underbrace{\Delta CS^*}_{\text{information value}} - \underbrace{\left(CS^0 - CS^*(\omega_E^0)\right)}_{\text{introductory pricing effect } (\geq 0)}.$$

The introductory pricing effect is increasing in δ*. For* δ *sufficiently small,* $\Delta CS^0 > 0$.

The introductory pricing effect captures how much consumers benefit from dynamic pricing relative to static Nash. For patient firms, this benefit can exceed the information value, making $\Delta CS^0 < 0$: consumers prefer the pre-review state with its lower prices to the post-review state with its resolved uncertainty but higher markups (Figure A.1 in Appendix A.1 illustrates this using numerical simulations of the dynamic equilibrium). This highlights a fundamental tension: the same forward-looking behavior that partially internalizes the externality (through introductory pricing) also reduces the consumer surplus value of resolving uncertainty.

Implications for the structural model. The theoretical analysis motivates three empirical quantities that we decompose in Section 9.1.1: the pure information value of a review (uniform prices), the rent-extraction reduction (static Nash), and the introductory pricing offset (dynamic Nash). However, the two-firm model cannot capture the entry and exit margin, which turns out to be the dominant channel through which reviews affect consumer welfare in practice.

To see why, consider what happens when a review accelerates learning about the entrant's quality. If the entrant turns out to be good, it attracts more demand, earns higher profits, and is more likely to remain in the market, adding a high-quality option to consumers' choice sets. If the entrant turns out to be bad, it loses demand and exits sooner, freeing capacity for new entrants who may be of higher quality. In both cases, faster learning improves the composition of active products: good products survive longer and bad products are replaced more quickly. This selection effect is absent from the stylized model, where the entrant and incumbent are permanently present regardless of the review outcome.

Moreover, the prospect of faster learning changes entry incentives. If reviews arrive quickly, an entrant's initial disadvantage is short-lived, making entry more attractive. This increases the flow of new products into the market, which benefits consumers through greater variety. Conversely, if reviews are slow, the cold-start period is prolonged, deterring entry and reducing variety. The welfare consequences of these dynamic responses can be large, since consumers value having a diverse set of options. In the structural model of Section 9, we find that the long-run welfare gains from an additional review, operating through entry, exit, and market composition, are orders of magnitude larger than the short-run per-period effects analyzed here.

Beyond entry and exit, the structural model features additional equilibrium effects that the

two-firm setting cannot capture. With many competing listings, a review for one listing affects demand for all others through competitive spillovers: a well-reviewed entrant steals market share not just from one incumbent but from the entire set of active listings. The resulting price and quantity adjustments propagate through the market, creating general equilibrium effects on profits, entry incentives, and the steady-state distribution. Furthermore, the structural model allows for multiple reviews, so that learning does not end with the first review but continues as listings accumulate a track record. The option value of future reviews shapes pricing at every state, not just at the entry state. These interactions between competition, multi-period learning, and dynamic entry and exit generate a rich set of equilibrium feedbacks that motivate the full structural analysis.

4 Model

4.1 Setup

In this section, we present the empirical model and apply it to the context of Airbnb. We assume that the true quality of each listing is unknown, but reviews serve as publicly observable signals of quality. As in the stylized model of Section 3, information is symmetric: Airbnb guests and hosts have the same information and, therefore, have the same beliefs about listing quality at all times. The time horizon is infinite, and hosts maximize the discounted sum of future profits through their pricing, entry, and exit decisions.

Indirect utility and demand. – There exists a set $\mathcal{J}$ of listings, indexed by j. Not all listings are active, so we define the subset of active listings in period t as $\mathcal{A}_t \subset \mathcal{J}$. Each period, there is a new set $\mathcal{I}_t$ of guests, indexed by i. Let N_{jt} denote the number of reviews and K_{jt} the number of good reviews of listing j accumulated up to period t. Each listing's true quality ω_j (the probability of a successful stay) is drawn independently from a $Beta(a, b)$ prior distribution. This specializes the general framework of Section 3 to the beta-binomial conjugate setting: after observing K good reviews out of N total, the expected quality of listing j is $(a + K_{jt})/(a + b + N_{jt})$.

Guests have unit demand and are risk-neutral. Let u_{ijt} denote a guest i's indirect expected utility of renting Airbnb listing $j \in \mathcal{A}_t$ in time period $t \in \{1, .., +\infty\}$.

$$u_{ijt} = \beta + \gamma \frac{a + K_{jt}}{a + b + N_{jt}} + (1 + f)\alpha p_{jt} + \xi_{jt} + \zeta_{igt} + (1 - \sigma)\epsilon_{ijt} \tag{1}$$

β is the intercept coefficient which only applies to Airbnb listings. γ reflects the utility value of having a successful stay, and α is the rental price coefficient, while p_{jt} is the rental price and f is the fee Airbnb levies on consumers. ξ_{jt} captures unobserved (to the econometrician) listing characteristics. We allow for different levels of substitutability between Airbnb listings among each other, and between Airbnb listings and the outside good by placing them in different nests. ζ_{igt} is the idiosyncratic group preference, and σ is the nesting parameter. ϵ_{ijt} is the idiosyncratic taste shock which is assumed to be distributed i.i.d according to a Type-I extreme value distribution, and ζ_{igt} is distributed such that $\zeta_{igt} + (1 - \sigma)\epsilon_{ijt}$ is also extreme value distributed.

We assume that the unobserved demand shock is additively separable in a time-constant listing characteristic ξ_j and an i.i.d. time-varying demand shock $\tilde{\xi}_{jt}$, i.e. $\xi_{jt} = \xi_j + \tilde{\xi}_{jt}$. Note that the transitory demand shock $\tilde{\xi}_{jt}$ affects prices and quantities in period t and thereby indirectly the transition probability. However, in any period, the transitory demand shocks from previous periods contain no relevant information about demand conditional on the current state. Hence, the state x_{jt} of listing j in t is $(K_{jt}, N_{jt}, \xi_j) \in X$ where $X = \{(K, N, \xi) \in \mathbb{N}_0^2 \times \mathbb{R} : K \leq N, N \leq \bar{N} \in \mathbb{N}_0\}$. We can write $u_{ijt} = \nu(p_{jt}, x_{jt}, \tilde{\xi}_{jt}) + \zeta_{igt} + (1-\sigma)\epsilon_{ijt}$ To keep the supply estimation tractable, we will later winsorize the number of reviews at 50, hence $\bar{N} = 50$, and discretize the unobserved listing characteristic parameter ξ_j, see Section 6.2.

We normalize the mean utility of taking the outside good, which we interpret as booking a hotel rather than an Airbnb, to zero. We abstract from guests arriving sequentially in the run-up to t to make bookings, as this would immensely complicate characterizing the demand system.[5] Rather, in each t, a discrete number of guests arrive simultaneously in the market to book accommodation for the duration of the period. Guest arrival $|\mathcal{I}_t|$ follows a Poisson process with mean μ_t, which may vary over time. Listings are capacity-constrained; at most, one consumer can rent listing j in t. Equation (3) characterizes the demand q for j in t. It equals the probability that at least one of the guests arriving in t wants to book j. If more than one guest wants to book j in t, we assume that one guest successfully books the listing while all remaining ones consume the outside good. Hence, demand for Airbnb listing j at time t is given as follows:

$$ccp_j(p_t, x_t, \tilde{\xi}_t) = \frac{\exp\left(\frac{\nu(p_{jt}, x_{jt}, \tilde{\xi}_{jt})}{1-\sigma}\right)}{D_{gt}^{\sigma}(1 + D_{gt}^{1-\sigma})}, \quad \text{where} \quad D_{gt} = \sum_{j' \in \mathcal{A}_t} \exp\left(\frac{\nu(p_{j't}, x_{j't}, \tilde{\xi}_{jt})}{1-\sigma}\right). \tag{2}$$

$$q_j(p_t, x_t, \tilde{\xi}_t) = 1 - \exp\left(-\mu_t ccp_j(p_t, x_t, \tilde{\xi}_t)\right) \tag{3}$$

where p_t, x_t and $\tilde{\xi}_t$ are vectors containing all prices, states, and transitory demand shocks of firms active in period t.

State transitions. – If a listing j is booked in t, with probability $v_r(x_{jt}) \in [0, 1]$, the guest accurately reports its experience (success, failure) in a review. We refer to $v_r(x_{jt})$ as the conditional review probability (as it is conditional on a booking). [6] The probability ρ_{j0} that a listing's state does not change, captures both the case where the listing is not booked as well as the case where the guest fails to leave a review. The listing receives a good review with probability ρ_{jg} or a bad review with probability ρ_{jb}, depending on the listing's quality prior. Hence $\rho_{jg} + \rho_{jb} = 1 - \rho_{j0}$ is the unconditional review probability.

If the review is good, both N and K increase by one in $t+1$. If the review is bad, N increases by one, but K does not. The possible transitions are illustrated in Figure 1.

[5] If consumers arrive sequentially, expected demand does not have a simple, closed-form solution because the current set of available listings depends on past booking decisions. While it is possible to integrate the different booking sequences numerically, it is computationally expensive, see Conlon and Mortimer (2013).

[6] Note that we assume that reviews are truthful and, while we allow the probability of leaving a review to depend on the current state, it is exogenous. The incentive of consumers to leave a product review may be manifold and is the subject of analysis in many other scholarly contributions, such as Fradkin et al. (2021). Modeling this is beyond the scope of this paper.

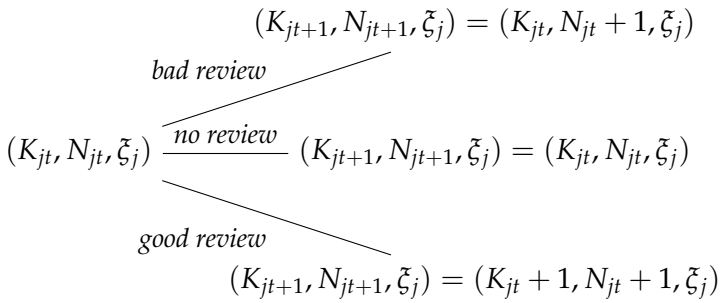


Figure 1: State transitions.

Equation (4) summarizes the transition probabilities.

$$\begin{aligned} \rho_{j0}(p_t, x_t, \tilde{\xi}_t) &= 1 - v_r(x_{jt}) q_j(p_t, x_t, \tilde{\xi}_t) \\ \text{and } \rho_{jg}(p_t, x_t, \tilde{\xi}_t) &= v_r(x_{jt}) q_j(p_t, x_t, \tilde{\xi}_t) \frac{a + K_{jt}}{a + b + N_{jt}} \\ \text{and } \rho_{jb}(p_t, x_t, \tilde{\xi}_t) &= v_r(x_{jt}) q_j(p_t, x_t, \tilde{\xi}_t) \left(1 - \frac{a + K_{jt}}{a + b + N_{jt}}\right) \end{aligned} \tag{4}$$

The host's problem. – For the remainder of the paper, it is convenient to define the industry s_t as a vector that contains, for each state $x \in X$, the number of firms in state x at time t. Hence, we can write $q_j(p_t, x_t, \tilde{\xi}_t) = q_j(p_t, x_{jt}, s_t, \tilde{\xi}_t)$. The state space is then $\mathcal{S} = s \in \mathbb{R}_0 : \sum_x s(x) < \infty$. Each host may only operate a single listing.[7] Consequently, the full set of listings is the full set of (both active and inactive) hosts $\mathcal{J}$. An active host earns expected variable per-period profit $\pi(p_{jt}, x_{jt}, s_t, \tilde{\xi}_{jt}, c_{jt}) = \mathbb{E}_{\tilde{\xi}_{-j}, p_{-j}}\left[(\text{period length in days}) q_j(p_t, x_{jt}, s_t, \tilde{\xi}_t)(p_{jt} - c_{jt})\right]$, where c_{jt} is the marginal cost of renting out a listing. In addition, a host incurs operating cost ϕ_{jt}, which is a fixed cost. ϕ_{jt} reflects that a host cannot use their apartment themselves for the time it is listed on Airbnb. This holds irrespective of whether the listing is booked or not. There is also a fixed cost κ_{jt} from entering the market and becoming active as the host must convert their apartment to an Airbnb rental. All costs are random variables that may vary over time. Each period, c_{jt} is i.i.d. drawn from a distribution G, $c_{jt} \overset{\text{i.i.d.}}{\sim} G$, such that $E[c_{jt}] = \bar{c}_j$. We allow for the fact that entry and exit decisions may happen less frequently than decisions on prices (e.g., because of contractual restrictions such as rental contracts). Specifically, we assume that, in each period, an active listing can make an exit decision and draw an operating cost only with probability r_e, and an inactive listing can make the entry decision and draw an entry cost with that same probability. Conditional on drawing any operating or entry cost, ϕ_{jt} and κ_{jt} follow exponential distributions with mean $\bar{\phi}_j$ and $\bar{\kappa}_j$ respectively. That is, $\kappa_{jt} \overset{\text{i.i.d.}}{\sim} Exponential(\frac{1}{\bar{\kappa}_j})$ and $\phi_{jt} \overset{\text{i.i.d.}}{\sim} Exponential(\frac{1}{\bar{\phi}_j})$. For ease of exposition, we will drop the j subscript from the mean cost parameters in the remainder of this section.

Timing. – In each t the timing of events is as follows.

1. With probability r_e, inactive hosts observe an i.i.d. entry cost draw κ_{jt} and decide whether to enter at the beginning of $t+1$. They pay κ_{jt} if they enter.

[7] 90% of hosts in our dataset operate a single listing. The average number of listings per host is 1.10.

2. Review outcomes are determined and active hosts observe their own state x_{jt}. With probability r_e they observe an i.i.d. cost draw of operating the listing, ϕ_{jt}. Conditional on the draw, they decide whether to stay active or exit.

3. If they stay active, they learn the industry state s_t and pay ϕ_{jt}.

4. Incumbent hosts observe an i.i.d. marginal cost draw c_{jt} as well as a transitory demand shock $\tilde{\xi}_{jt}$. They set their price taking into account the implications for future profits. They collect variable per-period profit $\pi(p_{jt}, x_{jt}, s_t, \tilde{\xi}_{jt}, c_{jt})$.

Note that our timing assumption implies that firms make exit decisions knowing their own state realization and that they may exit before the pricing decision.[8] They learn the industry state after all other active firms have also made their exit decision, if any, which is the payoff-relevant state.

4.2 Equilibrium Concept

We will first provide a full description of the strategy space and the value functions which form the baseline of any Markov perfect equilibrium (MPE). We then apply the oblivious equilibrium (OE) introduced by Weintraub et al. (2008) as an approximation of the MPE.

Similar to Weintraub et al. (2008), we assume that the equilibrium is symmetric and features real-valued functions that translate states into actions. We also make use of the fact that the exit strategy follows a cut-off rule: a firm j at time t exits if and only if $\phi_{jt} \leq e(x_{jt}, s_{t-1})$. A minor complication is that our price function has to take into account the marginal cost draw and thus our model allows for price heterogeneity across firms in the same state. The price function implies that each firm j at time t will set a price $\mathcal{P}(x_{jt}, s_t, \tilde{\xi}_{jt}, c_{jt})$. Denote the exit strategy as $e : X \times \mathcal{S} \to \mathbb{R}_+$ and the price strategy as $\mathcal{P} : X \times \mathcal{S} \times \mathbb{R}^2 \to \mathbb{R}_+$. Denote that strategy as $(\mathcal{P}, e) \in (\mathcal{M}, \mathcal{E})$, where $\mathcal{M}$ is the set of price strategies and $\mathcal{E}$ is the set of exit strategies. Similarly, denote an entry rate function $\lambda \in \Lambda$, where Λ is the set of entry rate functions and $\lambda : X_0 \times \mathcal{S} \to \mathbb{R}_+$ ($X_0 \subset X$ such that X_0 only contains states with $N = 0$). Finally, denote the expected variable per-period profit of firm j, conditional on its competitors following price strategy $\mathcal{P}$ and the firm itself follows strategy $\mathcal{P}'$ as $\mathbb{E}_{c,\tilde{\xi}}\left[\pi(p_{jt}, x_{jt}, s_t, \tilde{\xi}_{jt}, c_{jt}) | p_{jt} = \mathcal{P}', p_{-jt} = \mathcal{P}\right] = \pi_{\mathcal{P}',\mathcal{P}}(x_{jt}, s_t)$. Hence, the value of an active firm just before the realization of the marginal cost demand shock (stage 4 in the timeline above), given that it follows the strategy $(\mathcal{P}', e')$, while competing firms follow the strategy $(\mathcal{P}, e)$ and the entry rate is λ, is as follows:

$$V(x_{jt}, s_t |\, \mathcal{P}', \mathcal{P}, e', e, \lambda) = \pi_{\mathcal{P}',\mathcal{P}}(x_{jt}, s_t) + \delta\, \mathbb{E}_{x_{jt+1}, s_{t+1}}\left[V(x_{jt+1}, s_{t+1} | \mathcal{P}', \mathcal{P}, e', e, \lambda) - r_e Pr(\phi_{jt+1} < e'(x_{jt+1}, s_t)) \mathbb{E}\left[\phi_{jt+1} | \phi_{jt+1} < e'(x_{jt+1}, s_t)\right]\right] \tag{5}$$

$\delta \in (0, 1)$ denotes the discount factor. Equation (5) incorporates the host's decision to exit the market when the current operating cost exceeds the expected value of remaining in the market.

[8] We make this assumption to allow for a CCP approach in the marginal cost estimation, see Section 7.2.

A host who exits the market cannot re-enter and is replaced with an inactive host.

Following Weintraub et al. (2008), we assume that firms make decisions only based on their own state and the long-run industry state. While they make decisions without knowing the current industry state, their decisions collectively give rise to a long-run industry state consistent with their expectation. Hence, with some abuse of notation, we can drop the industry state as argument of the strategy functions $(\mathcal{P}, e)$ and the entry rate λ. If all firms follow a common strategy $(\mathcal{P}, e)$, the state evolves as an independent transient Markov chain. The long-run expected industry state is then $\tilde{s}_{\mathcal{P},e,\lambda}(x) = \lim_{t\to\infty} \mathbb{E}_{\mathcal{P},e,\lambda}[s_t(x)] \forall x$. The corresponding oblivious value function is given as follows:

$$\begin{aligned} V(x_{jt}|\, \mathcal{P}', \mathcal{P}, e', e, \lambda) = \pi_{\mathcal{P}',\mathcal{P}}(x_{jt}, \tilde{s}_{\mathcal{P},e,\lambda}) + \\ \delta\, \mathbb{E}_{x_{jt+1}} \left[V(x_{jt+1}|\mathcal{P}', \mathcal{P}, e', e, \lambda) - r_e Pr(\phi_{jt+1} < e'(x_{jt+1})) \mathbb{E}\left[\phi_{jt+1} | \phi_{jt+1} < e'(x_{jt+1})\right] \right] \end{aligned} \tag{6}$$

For simplicity, we denote $V(x|\mathcal{P}^*, e^*, \lambda^*) = V(x|\mathcal{P}^*, \mathcal{P}^*, e^*, e^*, \lambda^*)$. Definition 1 lists the oblivious equilibrium conditions in our context.

Definition 1. *The oblivious equilibrium is given by the following conditions:*

1. $V(x|\mathcal{P}^*, e^*, \lambda^*) \geq V(x|\mathcal{P}', \mathcal{P}^*, e^*, e^*, \lambda^*) \qquad \forall x \in X, \forall \mathcal{P}' \in \mathcal{M}$
2. $e^*(x) = V(x|\mathcal{P}^*, e^*, \lambda^*) \qquad \forall x \in X$
3. $\lambda_j^* = 1 - \exp\left(-\delta W_j \bar{\kappa}^{-1}\right) \qquad \forall j \in \mathcal{J},$
   where $W_j = V((0,0,\xi_j)|\mathcal{P}^*, e^*, \lambda^*) - r_e Pr(\phi_{jt} < e^*((0,0,\xi_j)))\mathbb{E}\left[\phi_{jt} | \phi_{jt} < e^*((0,0,\xi_j))\right]$

We refer the reader to Weintraub et al. (2008) for proofs on the existence of the cut-off rule as the optimal exit strategy and the existence of a stationary distribution over states in the environment described above. The oblivious equilibrium is a close approximation of the Markov Perfect Equilibrium in large, unconcentrated markets, where equilibrium industry states exhibit little variation and "observing the industry state and designing strategies [...] do not lead to significant increases in profit" for individual firms (p. 1248). We believe this is an adequate description of the market in our application, with there being hundreds of Airbnb listings in the market at any given time.

In the remainder of this paper, we will focus exclusively on the equilibrium value and strategy functions, and therefore omit $\mathcal{P}^*$, e^* and λ^* for simplicity. Hence, we denote $V(x) = V(x|\mathcal{P}^*, e^*, \lambda^*)$. Note that, in the oblivious equilibrium, the cut-off value for the firm's exit decision is $e^*(x) = V(x)$. Moreover, the exponential distribution of the operating cost allows us to provide a closed form solution of the exit rate $\chi(x_{jt}) = Pr(\phi_{jt} > e^*(x_{jt}))$, where $\chi(x) : X \to \mathbb{R}_+$:

$$\chi(x_{jt}) = \exp(-V(x_{jt})\bar{\phi}^{-1}) \tag{7}$$

It follows from Equation (7) that there is a one-to-one mapping between exit rates and expected values after the operating cost realization but before the marginal cost shock realization. We will exploit this fact when setting up our estimation routine, as we will see in Section 7.2.

Our distributional assumptions allow us to also simplify the firm value at the beginning of each period before the exit decision:

$$V(x_{jt}) - r_e Pr(\phi < e^*(x_{jt}))\mathbb{E}\left[\phi|\phi < e^*(x_{jt})\right] = V(x_{jt}) - r_e(1-\chi(x_{jt}))\bar{\phi}. \tag{8}$$

Hence, we can rewrite the oblivious value function in the oblivious equilibrium as follows:

$$V(x_{jt}) = \pi_{\mathcal{P}^*,\mathcal{P}^*}(x_{jt}, \tilde{s}_{\mathcal{P}^*,e^*,\lambda^*}(x_{jt})) + \delta\, \mathbb{E}_{x_{jt+1}}\left[V(x_{jt+1}) - r_e(1-\exp(-V(x_{jt+1})\bar{\phi}^{-1}))\bar{\phi}\right] \tag{9}$$

5 Data

Airbnb is the market-leading peer-to-peer platform for short-term accommodation. It enables hosts to rent their apartment or a room to guests, usually tourists. Since its founding in 2008, Airbnb has grown to feature more than four million Airbnb hosts worldwide, housing about 33 million guests annually on average.[9] Airbnb's annual revenue of roughly 6 billion US dollars in 2021 rivals the revenue of large hotel chains. With thousands of active Airbnb listings between 2016 and 2019, the time period we analyze, New York City was by far Airbnb's largest market among US cities. Due to the large number of *a priori* unknown listings, the crowd-based rating system of Airbnb is one of its key services to potential guests. Guests are encouraged to leave a review of the listing they stayed at on a one-to-five stars scale within two weeks after concluding their stay, which is then published for other potential guests to see. The average star rating and the cumulative number of reviews are prominently displayed on Airbnb's website.

We use data collected by AirDNA (2019), a data analytics company, on all Airbnb listings in Manhattan, New York, between January 2016 and December 2019.[10] Figure 2 (left) illustrates the spatial distribution of listings across Manhattan. During this time, the number of listings is relatively stable, supporting the idea that the market is in a stationary equilibrium (see Figure 2, right). We choose the time frame for two reasons: First, we wanted to exclude the period affected by the COVID-19 pandemic since the tourism industry was one of the most affected industries. Second, on September 5th, 2023, the City of New York mandated that short-term rentals must register with the Mayor's Office of Special Enforcement.[11] As a result, Airbnbs have all but disappeared as a form of accommodation for tourists.[12] Focusing on a time frame before early 2020 ensures that our analysis is not distorted by these major global or industry events.

We have rental prices in USD for each listing and day during the observation period and information on whether a listing was available, reserved, or blocked for private use by the host on a given day. In addition, we observe the number of reviews and the rating on a one-to-five stars scale for each listing in a roughly bi-weekly frequency.[13] Lastly, the dataset includes

[9] See https://news.airbnb.com/about-us/.

[10] See https://www.airdna.co/.

[11] See https://www.nyc.gov/site/specialenforcement/registration-law/registration.page.

[12] This measure aims to curb short-term home-sharing and support the long-term rental market. Calder-Wang (2021) finds that although Airbnb's presence generally contributes positively to welfare, it has the adverse consequence of increasing rents in New York City, thereby making most households worse off. Our model does not account for housing market effects and they are not covered by our analysis.

[13] The frequency of the ratings and reviews depends on the frequency of the AirDNA scraping. All other vari-

various listing attributes and amenities, which we use to define the relevant market. We focus on listings that offer the entire apartment for rent (as opposed to a single room), include at least one picture of the place, permit at most two guests and no pets, and feature one bedroom and one bathroom.

We suspect that, due to how AirDNA web scrapes the AirBnB website, the occupancy rate around the time a listing first becomes active is underreported. This is because in the month a listing becomes active it is treated as available, even though it might not be active yet. To address this issue, we drop observations that occur prior to the first booking of a listing, ensuring that we only consider data after the listing has definitively become active. Note that this approach excludes listings that are never booked during their lifespan.

We also drop listing-days with observed prices below the first and above the 99th percentile, corresponding to \$65 and \$550, respectively. We account for the cleaning fee hosts typically charge per stay by adding it to the daily price, divided by the average reservation length of 5.7 days. If we do not observe the cleaning fee for a listing, we assume it is equal to the average cleaning fee in the data sample.[14]

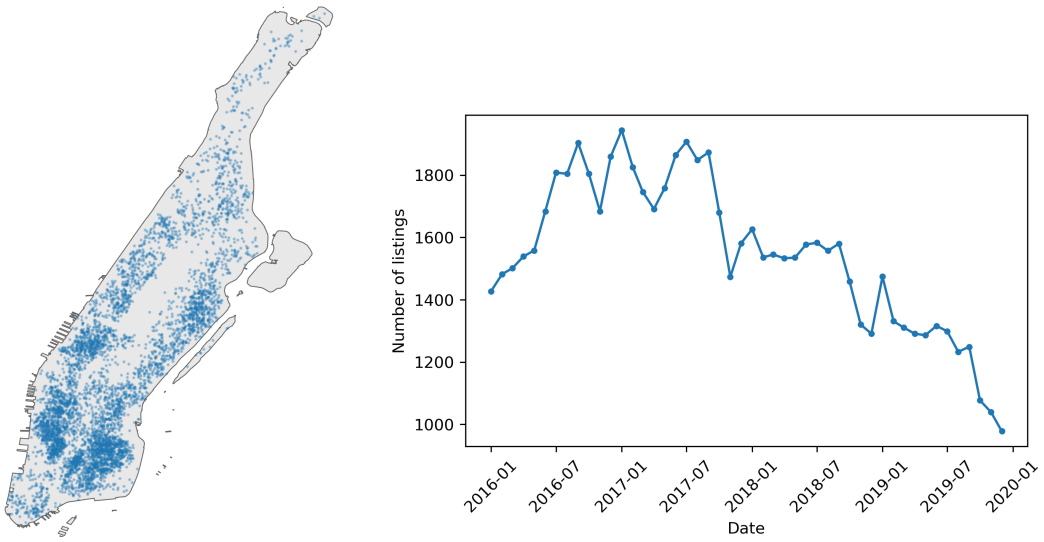


Figure 2: Number of listings over time and their location.

To estimate the supply-side of the model, we aggregate the data on a weekly level. A week is a sufficiently short time span to ensure a low probability of multiple bookings in one period; the average reservation length of a listing which is around 6 days. Moreover, there are very few weeks with more than one new review coming in for any listing (this only happens in around 5% of listing weeks in our data). Hence, we believe that aggregating at the weekly level is in line with the model assumption of receiving at most one review per period. We chose a higher level of aggregation for the demand estimation to limit measurement error in the market shares we measure.[15] After cleaning the data, we have 149,533 observations at the listing-week level.

Table 1 summarizes key variables in the data on the listing-week level. The average rental price is \$195.94 in our sample. Rental prices exhibit moderate variation. The majority of listings

ables are available at the daily level. There are less than two weeks between 40% of review observations. Only 1%

	N	mean	std	min	25%	50%	75%	max
Rental price (in $)	149,533	196.51	62.21	68.63	153.30	188.10	228.76	584.03
Occupancy rate (in %)	149,533	59.57	38.87	0.00	20.00	71.43	100.00	100.00
Number of reviews	149,533	18.43	18.26	0.00	3.00	11.00	32.00	50.00
Rating	131,374	4.69	0.50	1.00	4.60	4.80	4.95	5.00

Table 1: Summary statistics (listing-week level).

are priced between $150 and $280. On average, Airbnb listings are occupied 59.36% of the time and have been reviewed about 25 times. On average 2.99% of hosts in any given week exit and are no longer active in the subsequent week, while 2.59% have entered.[16] The average lifespan within our sample is 1.3 years, but most listings are active for less than a year.[17]

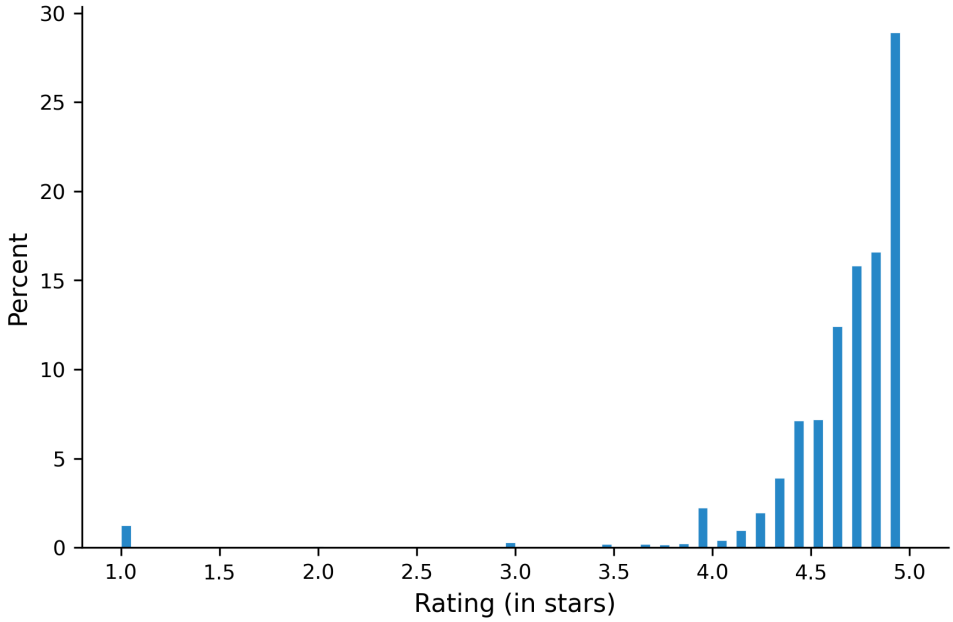


Figure 3: Rating distribution by number of reviews.

At 4.69 stars, the average rating is high. In comparison, at 3.8 stars, the average rating on TripAdvisor is substantially lower (Zervas et al., 2021). It is a priori unclear if this is mainly because highly rated listings remain in the market longer or because the listings' quality is generally high. Figure 3 shows the rating distribution in our sample. The distribution is left-skewed because most listings are rated four stars or higher. Listings with a lower rating tend to have fewer reviews. Almost all listings with a 1-star, three-star, or 3.66-star rating have less

observations lie more than three months apart.

[14] This leaves us with 7,803 unique listings and 2,263,124 listing-days, excluding blocked ones.

[15] Months refer to four-week periods and do not necessarily correspond to calendar months.

[16] We use the whole sample with dates before January 2016 and after December 2019 to determine entry and exit. Rates are relative to the active hosts in the market.

[17] Here, we measure lifespan as the number of weeks for which we observe a listing in our sample. Hence, it has a natural maximum of 4 years, or 208 weeks.

than four reviews.

To construct the state space, we back out the number of "good" reviews K as the number of five-star reviews required to achieve the observed rating if the remaining $N - K$ reviews were "bad" one-star ones. For example, for a listing with ten reviews and a rating of 3.8 stars, the implied number of good reviews is 7. We deal with missing information about the number of (good) reviews by filling in the most recent observed value. To keep a reasonably sized state space, we will censor the number of total reviews in the model at 50, i.e., $\bar{N} = 50$. For N exceeding 50, we find that changes in the posterior mean in response to additional reviews become negligible. Figure 4 illustrates this using our estimated prior (mean 4.15 stars, variance 0.04; see Section 6): as a listing receives a sequence of good or bad reviews, the posterior mean converges and the marginal impact of each additional review diminishes.

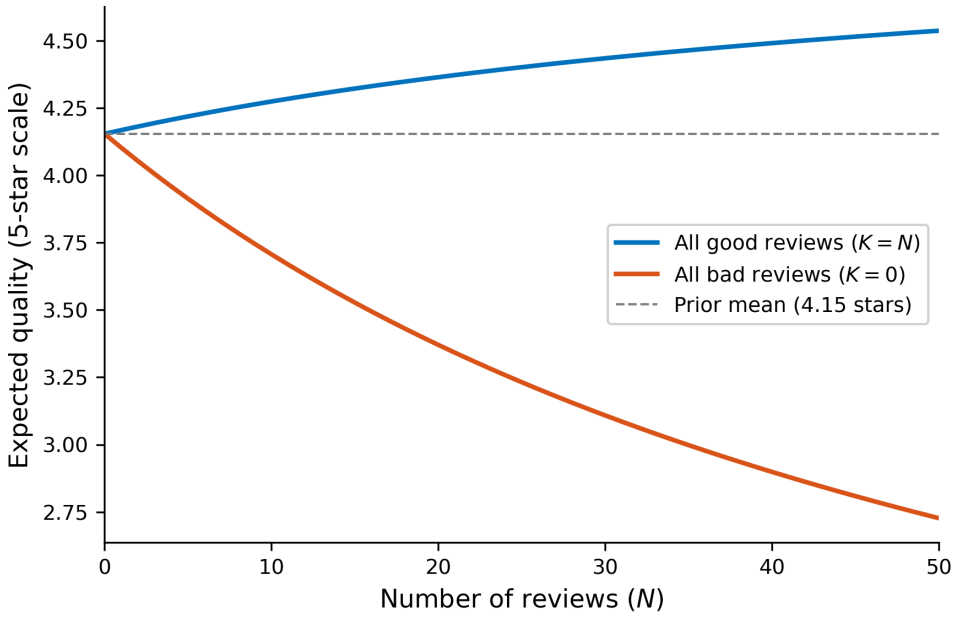


Figure 4: Posterior mean updating with good and bad reviews, using the estimated prior distribution ($a/(a+b) = 4.15$ stars on a 5-star scale).

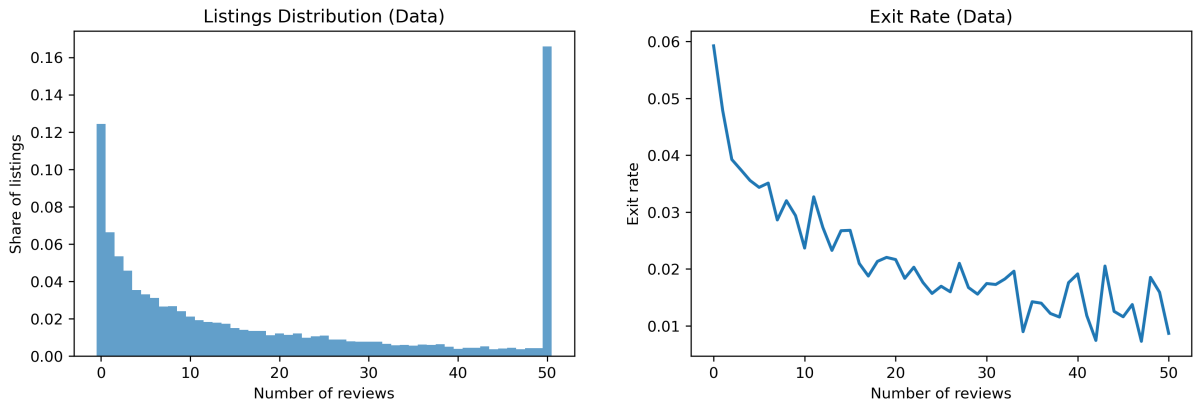


Figure 5: Number of listings and exit rates by review state.

Hosts of listings with many reviews and a good rating make more revenue and remain in the market, whereas they leave the market otherwise. Hence, our dataset contains very few or no observations for many states associated with small K, while most observations are concentrated around high-K states (see Figure 5, left). The right-skewed listing distribution of firms over (good) reviews is driven by a selection effect reminiscent of Jovanovic (1982). High-quality listings survive, whereas low-quality ones fail and exit. The listing distribution features a "pitch fork" shape, as most listings either have few or no reviews or – due to right-tail compression – have the maximum amount. Notably, around 14% of listings have no reviews and about 18% of reviewed listings have the maximal review count. We also observe the selection effect in hosts' exit behavior (see Figure 5, right). Compared to a reviewed listing, an unreviewed listing is around three times as likely to exit in a given week.

Figure 6 shows how occupancy rates and rental prices evolve over a listing's lifetime, after removing listing-level heterogeneity. We regress each outcome on listing fixed effects and age (week-of-life) dummies, and plot the estimated age coefficients. New listings start with lower occupancy rates and lower prices. Both increase over the first year as listings accumulate reviews and hosts learn about demand. The upward trajectory is consistent with the selection mechanism described above: listings that survive tend to be of higher quality, and as they gather reviews, consumers' willingness to pay increases.

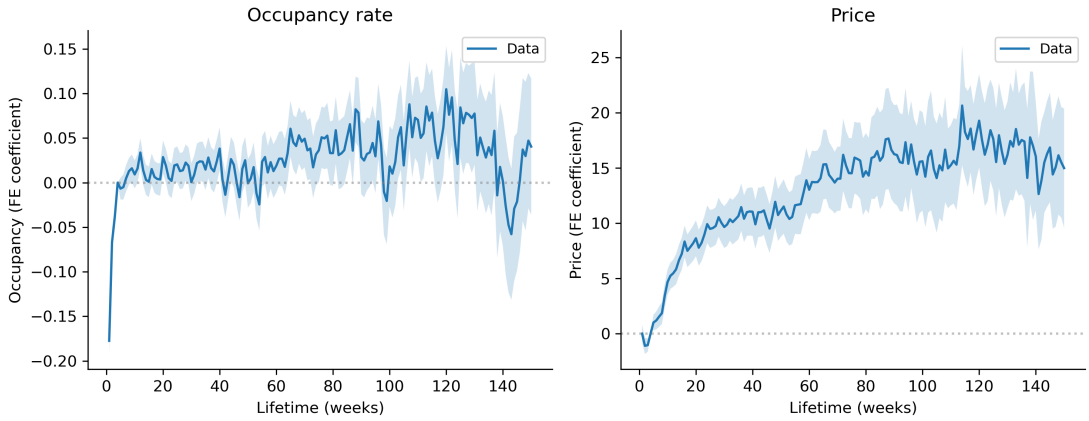


Figure 6: Within-listing lifecycle trajectories of occupancy rates and rental prices. Coefficients are from a regression of each outcome on listing fixed effects and week-of-life dummies (week 4 = reference).

To support the graphical intuition of Figure 5 and Figure 6, we regress the rental price, occupancy rate, and exit decision on the number of good reviews and the total number of reviews. Table B.1 in Appendix B.4 indicates that, within a listing, a higher number of good reviews is associated with a higher price, while bad reviews do not appear to have a significant impact on price. Similarly, more good reviews correlate with higher occupancy rates, whereas a larger number of bad reviews is associated with lower occupancy. As expected, the exit decision is negatively related to good reviews and positively to bad ones. Naturally, these results only reflect correlations, as all involved variables are equilibrium outcomes.

We combine the main AirDNA dataset with two additional data sources to estimate the mar-

ket size and the number of potential Airbnb listings in Manhattan. We allow the average daily number of guests arriving in the market to vary by month, based on occupancy data for New York City from NYC & Company.[18] Between 2016 and 2019, New York City sold between 2 and 3.5 million hotel nights per month, with approximately 84% of these bookings occurring in Manhattan. We assume that one in three travelers seeking accommodation consider Airbnb, which roughly aligns with the reported 4% average market share across 50 U.S. cities in Farronato and Fradkin (2022). The total number of potential listings (J) is based on the 2017 New York City Housing and Vacancy Survey.[19] By adding the survey estimate of relevant vacant units with the average number of available Airbnbs in the data, we arrive at an estimate of 24,522 apartments that could potentially be converted into Airbnb listings.[20]

6 Demand Estimation

6.1 Estimation Approach

We estimate the parameters governing demand by the Generalized Method of Moments (GMM) on data aggregated to the monthly level. Some listings have zero bookings in a given month (or, more rarely, are fully booked out), in which case we have to drop those observations for the demand estimation. We invert demand and back out ξ_{jt} to compute the moment conditions:

$$\nu(p_{jt}, x_{jt}, \tilde{\xi}_{jt}) = \ln\left(\frac{-\ln(1-B_{jt})}{\mu_t}\right) - \ln\left(1 + \frac{\sum_{j\in\mathcal{A}_t}\ln(1-B_{jt})}{\mu_t}\right). \tag{10}$$

We account for seasonality by allowing the daily arrival rate to vary by month as described in Section 5.

Without instrumental variables, our estimates are biased because the rental price and the number of (good) reviews are endogenous. To address price endogeneity, we use a cost shifter as an instrument for the rental price, and past transitory demand shifters as instrument for the number of reviews. Specifically, we use the average reservation length of a listing to capture the marginal cost of renting it out. The rationale is that longer reservations reduce the host's workload per night, as checking guests in and out and cleaning typically occur only once at the beginning or end of a stay. The instrument would be invalid if reservation length were systematically related to a listing's quality – for instance, if lower-quality listings tended to have longer reservations on average. We argue that this is unlikely. Presumably, guests adjust their demand on the extensive margin (e.g., not picking a low-quality listing) rather than on the intensive margin (e.g., deciding to stay three instead of four days because a listing is of low quality).

For the (good) number of reviews, we use the de-meaned (at the listing level) five-, six- and seven-month lagged occupancy rates, and their interaction with the rating as instruments for

[18] See https://assets.simpleviewinc.com/simpleview/image/upload/v1/clients/newyorkcity/FYI_Hotel_reports_March_2020_9afc1fca-79c1-47a9-8a21-fdefa3c4ffcb.pdf.

[19] See https://web.archive.org/web/20241204060848/https://www.census.gov/data/datasets/2017/demo/nychvs/microdata.html.

[20] Specifically, we consider vacant one-bedroom units with fewer than four rooms available for rent in Manhattan.

the (good) number of reviews. Since we adjust past occupancy rates by the listing's average, these rates reflect transitory demand shocks rather than persistent listing quality. For example, neighborhood festival, holidays, or university graduation ceremonies may temporarily boost occupancy rates in specific locations. The lagged occupancy rate would not be a valid instrument if such demand shocks persisted into the current period. This could occur if a listing or its surroundings experienced extended or permanent changes, observable to guests at the time of booking—for instance, if the apartment was remodeled or the surrounding area was upgraded. However, we believe this issue is minimal in our data, as such changes would likely result in sustained and significant price adjustments. In our sample, 18.8% of listings show no rental price changes, and the standard deviation of rental prices within a listing is only $5.83.

As is standard, we use the common instrument of the total number of listings active in a given period as an instrument for the within-group market share.

Following Ferrari and Cribari-Neto (2004) and Dickstein (2021), we do not estimate a and b directly, but estimate ψ and ι instead. They are defined in Equation (11). ψ determines the prior mean, whereas ι is closely related to the variance of the prior distribution.

$$\frac{1}{1+\exp(-\psi)} = \frac{a}{a+b} \quad \text{and} \quad \exp(\iota) = a + b. \tag{11}$$

This alternative formulation facilitates the estimation. Naturally, the coefficient of the rental price is identified by variation in the rental price. The identification of γ, ψ, and ι comes from within- and between-listing variation in the number of (good) reviews. ι pins down to what extent an additional review moves the posterior mean away from the prior mean. If ι is small, the prior belief is precise, and guests make only marginal adjustments to their beliefs after observing the rating. ψ depends on variation in the rating to be identified. The posterior mean responds more strongly to reviews if the rating differs greatly from the prior mean. In particular, the posterior mean increases by more as the rating improves if the prior mean is low. In this way, the relationship between the rating and the posterior mean allows us to estimate the prior mean. γ captures the impact of the posterior mean on guests' booking decisions. If the reviews have little effect on the occupancy rate of the listings, regardless of their current rating and review count, γ must be low.

6.2 Results

We present our demand estimates in Table 2. Column (1) shows our results for a plain logit model without nests, while, in column (2), we place Airbnb listings in one nest, with the outside good in a different nest. We find that the nesting parameter in column (2) is small and not significant, suggesting that the substitutability between Airbnbs is similar to the one between Airbnbs and the outside good. Nonetheless, we will focus on the point estimates of the nested logit specification in the remainder of this paper and use them in the supply side estimation. Based on the nested logit estimation results, ψ and ι imply a prior mean and variance of 4.15 stars and 0.04, respectively, and the coefficient of the posterior γ is large and significant at the 10%

level.[21] Figure 7 relates the estimated prior distribution to the rating distribution in the data. Recall that the rating is an imprecise signal of the true quality and features a larger variance. Also, the rating is subject to listing selection; highly rated listings are observed in the data, whereas poorly rated listings are not.

		Nested logit	
prior	ψ	1.316	(0.155)
	ι	4.103	(0.681)
rental price	α	-0.0045	(0.0008)
expected quality	γ	6.209	(3.706)
Nesting parameter	σ	0.630	(0.041)
Observations		24,012	

Note: Robust standard errors in parentheses.

Table 2: Demand estimation results.

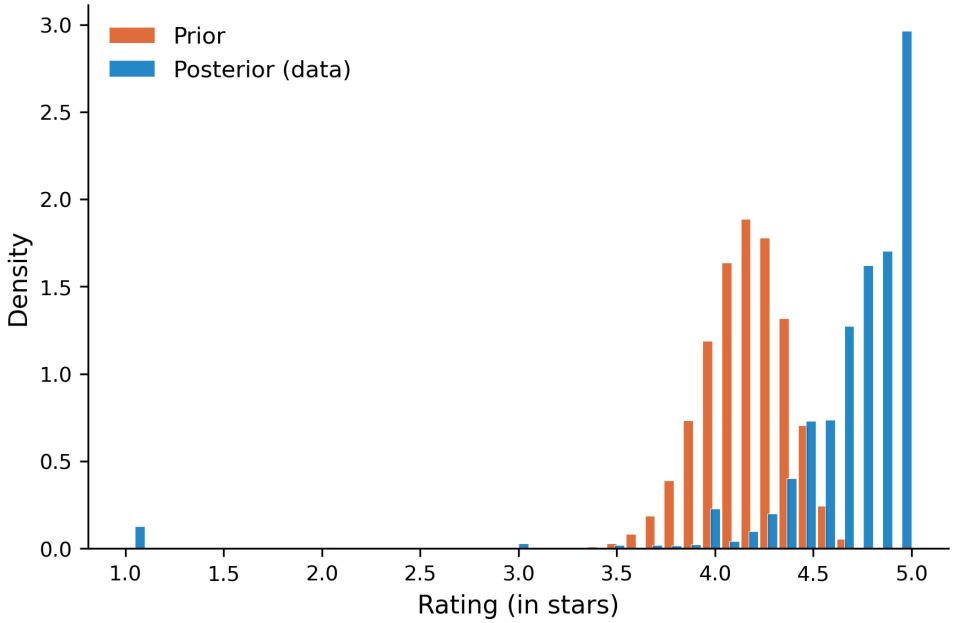


Figure 7: Estimated prior distribution and observed rating distribution.

The estimated rental price coefficient corresponds to an average listing-month price elasticity of -2.3. To compare, Farronato and Fradkin (2022) report an average price elasticity of the demand for Airbnb and hotel accommodations of -4.27. In Huang (2022), the average price elasticity of demand for Airbnb accommodation in San Francisco is -2.51. While our estimated elasticity is somewhat lower in absolute terms compared to the literature, we believe it is reasonable. It is possible that our estimate of the price coefficient reflects search frictions; if guests

[21] See Table B.2 in Appendix B.4 for the first-stage regression results of our demand estimation.

are not aware of all available listings, they would substitute to other options at a comparatively lower rate in response to a price increase. The constant β is estimated to be negative and large in magnitude, reflecting that most guests take the outside good (i.e., book a hotel) rather than book an Airbnb. This is in line with Airbnb's small market share of 4%, as reported in Farronato and Fradkin (2022).

	mean	std	min	25%	50%	75%	max
Own-price elasticity	-2.2715	0.9068	-8.0523	-2.7213	-2.1156	-1.6465	-0.0151
Good review semi-elasticity	3.4374	1.4922	0.0053	2.2745	3.5975	4.6771	9.0782

Table 3: Summary of elasticities.

We also compute the semi-elasticity of demand if a listing receives a good review. On average, the occupancy rate increases by 3.4% in response to a good review. Figure 8 shows that the demand for listings with fewer reviews is more elastic. Guests who see an additional review will only slightly adjust their expectations of a listing's quality if the listing already has many reviews and the rating is precise. Conversely, observing a review has a much greater impact on guests' beliefs if the listing has few or no reviews and there was previously little information available to assess its quality. Accordingly, the average good review semi-elasticity in the data is 5.2 for unreviewed listings but only 1.5 for listings with 49 reviews. Table 3 provides summary statistics for the two elasticities we estimate.

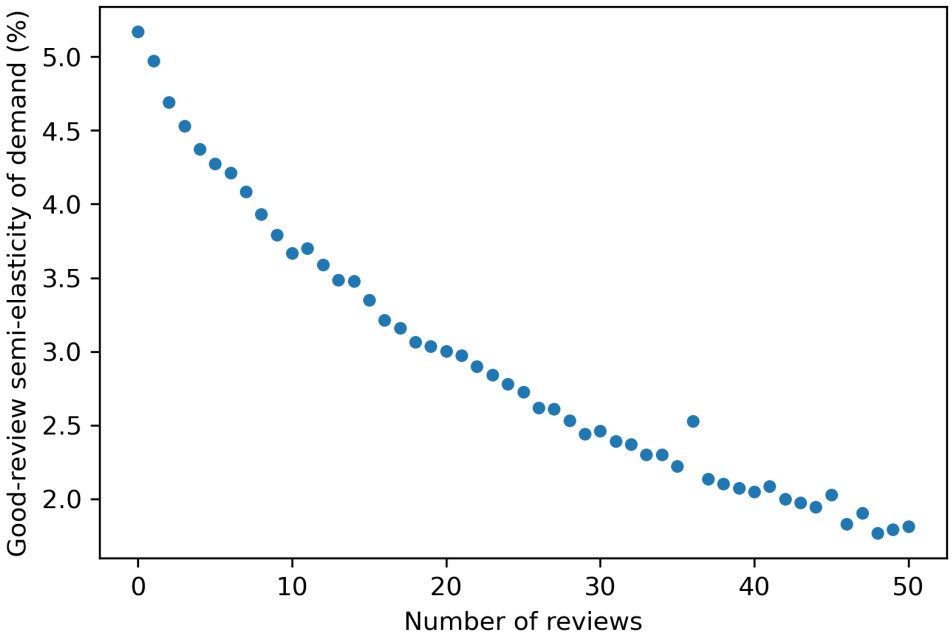


Figure 8: Semi-elasticity of booking probability with respect to one additional good review.

We are not aware of any studies that provide review elasticities for Airbnb to which we can compare our estimates. Fradkin and Holtz (2022) find no effect of reviews on Airbnb reservations, but this result is likely driven by the specifics of their field experiment. For comparison,

Reimers and Waldfogel (2021) estimate that book sales increase by approximately 2.6% if the book is reviewed in the *New York Times*.

The prior distribution parameter estimates, together with the estimates of the expected quality and the rental price, imply that guests' willingness to pay is *ceteris paribus* \$4.18 higher per day for a listing with one review and a 5-star rating compared to an unreviewed listing. An additional good review of a listing with 49 reviews and a four-star rating is worth \$2.49 per day to guests.

Using our demand estimates, we back out the demand shock ξ_{jt} for each observation. We compute $\bar{\xi}_j$, the average demand shock of each listing. To ensure a manageable state space, we define four listing types instead of allowing state transitions to be listing-specific. We compute the quartiles of the empirical distribution of the estimated $\bar{\xi}_j$s and assign type 1 through 4 to a listing depending on the quartile $\bar{\xi}_j$ falls into. Note that by distinguishing only four types, we choose to be parsimonious and limit the size of the state space in order to keep the model tractable. We report summary statistics for the different types in Table 4.

On average, higher types feature a higher occupancy rate. While a type 4 listing is, on average, booked 63% of the time, the average occupancy rate of a type 1 listing is 49%. The correlation between the type and the rating is very weak (Pearson correlation coefficient: -0.002). The number of reviews tend to increase in the type (0.09). At the same time, the rental price is highly correlated with the listing type (0.47). This correlation provides support for the notion that types are largely differentiated by characteristics that are observable to consumers, rather than the unobserved quality.

	$\bar{\xi}$	Price	Occupancy	Reviews	Rating
type 1	-0.36	\$157.91	49.84%	38.06	4.67 stars
type 2	-0.07	\$177.76	61.42%	32.03	4.68 stars
type 3	0.13	\$198.85	63.55%	21.66	4.71 stars
type 4	0.47	\$251.52	63.49%	13.47	4.70 stars

Table 4: Summary: listing types.

7 Supply Estimation

7.1 Calibrated Supply-Side Parameters

We do not estimate all the parameters of our model, but we calibrate some of them using data moments. Recall that a period is one week. For the discount factor δ, we choose 0.9988, implying an annualized interest rate of 6.7%. We allow the conditional review probability v_r to vary by the number of current reviews. Hence we estimate v_r as the inverse of the number of booked weeks required until an additional review is observed in the data for listings with less than 50 reviews. In this way, we determine that v_r varies between 38% and 60% depending on the number of reviews a listing has. It is important to note that we do not allow v_r to vary based on

the rating, the rental price, or whether the guest experiences a success or a failure.[22]

Given that rental contracts in the market typically operate on a monthly cycle, it is reasonable to assume that hosts make their entry and exit decisions on a monthly basis. Consequently, we set the probability of making an entry or exit decisions in any given week to be 25%, i.e., $r_e = 0.25$.

By the construction of listing types, the distribution of types across active listings is uniform. However, since we do not observe inactive listings, we cannot infer the unconditional distribution of listings. We assume that higher-quality types are less likely to occur; specifically, types three, two and one are twice, thrice, and four times as likely to occur than type four. The average number of listings of a certain type outside the market is J_l minus the total number of listings of that type in the market, where $l \in \{1, 2, 3, 4\}$. While this assumption affects the estimated mean cost parameters, we find that the average *incurred* entry and operating costs we estimate are fairly robust to the assumption of the listings-type distribution among potential listings. Specifically, if the share of a type among the total number of listings is larger (smaller), our average entry cost estimate is larger (smaller) as to offset the larger (smaller) number of potential entrants. Nevertheless, we caution the reader to interpret our estimation results with the assumption we made regarding the listing-type distribution in mind.

We find that the majority of the variance of the overall demand shock ξ_{jt} is explained by the time-constant type $\bar{\xi}_j$. Moreover, the variance of $\tilde{\xi}_{jt}$ is dwarfed by the marginal cost variation based on the estimation routine we describe below. Hence, for computational reasons, we only integrate over the marginal cost shocks in the model solution below and assume $\tilde{\xi}_{jt} = 0$ for all j and t.

Lastly, Airbnb charges – both in the model and in reality – a commission fee of 14.2%.[23]

7.2 Estimation Approach

We estimate the type-specific cost parameters of the model, denoted by $k = (\bar{\kappa}, \bar{\phi}, \bar{c})$, with $\bar{\kappa} = (\bar{\kappa}_1, \bar{\kappa}_2, \bar{\kappa}_3, \bar{\kappa}_4)$, $\bar{\phi} = (\bar{\phi}_1, \bar{\phi}_2, \bar{\phi}_3, \bar{\phi}_4)$ and $\bar{c} = (\bar{c}_1, \bar{c}_2, \bar{c}_3, \bar{c}_4)$, conditional on the demand parameter estimates $\theta = \{\beta, \gamma, \psi, \iota, \alpha, \bar{\xi}_1, \bar{\xi}_2, \bar{\xi}_3, \bar{\xi}_4\}$, with a nested Maximum Likelihood routine. In the inner nest, we estimate the marginal costs parameters $\bar{c}$ for a guess of the entry costs $\bar{\kappa}$ and operating costs $\bar{\phi}$ using GMM, and solve the model to arrive at the state distribution implied by the cost parameters. In the outer nest, we search for the entry and operating cost parameters which maximize the likelihood over the state distribution $s^*(x|k, \theta)$.

7.2.1 Inner nest

Marginal cost estimation We will back out marginal cost from the optimality condition for prices.[24] Assuming that competing firms follow the optimal pricing and exit strategies, and the entry rate is individually optimal as well, with some abuse of notation, we denote demand

[22] See Figure B.2 in Appendix B.4 for the empirical review probabilities in the data.

[23] See https://www.airbnb.ca/resources/hosting-homes/a/how-much-does-airbnb-charge-hosts-288.

[24] While this is standard for static models, we are not aware of previous work implementing the estimation approach we use for dynamic models of pricing.

for listing j with $q(p_j, x_j)$. We drop the j subscripts and use Equation (8) to express the firm value at the beginning of each period. Hence, the dynamic first-order condition of firm j can be rearranged as follows:

$$c = p - m(p,x) + \left(\frac{\partial q(p,x)}{\partial p}\right)^{-1} \delta \sum_{x'} \frac{\partial Pr(x'|p,x)}{\partial p}(V(x') - r_e(1-\chi(x'))\bar{\phi}) \tag{12}$$

where $m(p,x)$ is listing j's mark-up, i.e., $-q(p,x)\left(\frac{\partial q(p,x)}{\partial p}\right)^{-1}$. Note that all the components of the above cost equation are jointly identified by the demand parameters and the data, with the exception of $V(x')$. Nonetheless, we are able to estimate the value function conditional on a guess of the operating cost parameter, as we describe below.

We use the fact that, conditional on a candidate of the average operating costs, there is a one-to-one mapping between the exit rate and the expected value of a listing after the operating cost realization but before the realization of the marginal cost shock and the demand shocks: $\chi(x) = \exp(-V(x)\bar{\phi}^{-1})$. Note that, for states that we observe in the data, we can estimate exit rates non-parametrically. We can then use estimated exit rates $\hat{\chi}(x)$ to represent the value function before the operating cost realization as follows:

$$\hat{V}(x|\bar{\phi}) = -\ln \hat{\chi}(x)\bar{\phi} \tag{13}$$

Hence, the value of a listing in state x at the beginning of a period is identified by the empirical exit rate in that state up to the operating cost parameter. We can replace the value function in Equation (12) by its empirical equivalent:

$$c = p - \hat{m}(p,x) + \sum_{x'} \frac{\partial \hat{Pr}(x'|p,x)}{\partial p}(\hat{V}(x'|\bar{\phi}) - r_e(1-\hat{\chi}(x'))\bar{\phi}) \tag{14}$$

To be consistent with our state space, we parameterize marginal costs as follows: $c_{jt} = c_l \mathbb{1}(type_j = l) + \eta_{jt}$. Based on this specification, our moment conditions are $\mathbb{E}[Z_{jt}\eta_{jt}(\bar{\phi}_{l(j)})] = 0$. Hence we get marginal cost estimates as a function of the operating cost parameter guess: $\hat{c}(\bar{\phi})$.

Solving the Model Given a candidate for the operating cost and entry cost, the above step gives us the marginal costs. Next, for each candidate of the cost parameters $k = (\bar{\kappa}, \bar{\phi}, \bar{c}(\bar{\phi}))$, we solve the model in four steps. We start by formulating the initial guess of the pricing policy function $P_0(x,c)$, the number of listings $s_0(x)$ per state, and the host's value function $V_0(x)$.[25]

Step 1 – Based on the guess, we determine a host's best response $P_1(x,c)$ if all remaining hosts adhere to $P_0(x,c)$. We find $P_1(x,c)$ by combining numerical integration with the method of gradient ascent. We draw marginal costs based on Halton draws from the marginal cost distribution we estimated in the data. We then iterate over Equation (15), where k is the iteration step and $P_1^0 = P_0(x,c)$, until the change in the rental price, $P_1^{k+1} - P_1^k$ for any state x is less than \$0.1.

[25] We initialize the model with the empirical state distribution, value functions of 0 for all states, and stationary optimal prices.

$$P_1^{k+1}(x,c) = P_1^k(x,c) + w'(x,c|P_1^k, P_0) \tag{15}$$

where w' is the first-order derivative of the host's value after learning the cost shock given the host follows price rule P_1^k while competitors follow price rule P_0. Since hosts don't observe their rivals' marginal cost shocks, they have to form expectations about the rivals' prices based on the marginal cost distribution and the pricing function P_0.[26]

Step 2 – Assuming that all hosts set their prices according to $P_1(x,c)$, we compute the expected variable per-period profit $\pi_1(x)$ and occupancy rate $q_1(x)$ *before the cost shock realization*. The latter forms the baseline for our expected transition matrix $T_1(x)$. We use the expected profit, demand and transition matrix together with the current candidate for the operating cost to compute the value function $V_1(x)$ from Equation (9).

Step 3 – We use $V_1(x)$ to update the exit rate in state x to $\chi_1(x)$ and the entry rate for each type l to λ_{l1}. Together with the occupancy rates $\mathbf{q}_1$, $\boldsymbol{\chi}_1$ and $\boldsymbol{\lambda}_1$ allow us to compute an expanded transition matrix $\mathbf{F}_1$. We use $\mathbf{F}_1$ to solve for the new, stationary listing distribution $\mathbf{s}_1$, where $\mathbf{s}_1$ is given by:

$$\mathbf{s}_1 = \mathbf{s}_1 \mathbf{F}_1 \tag{16}$$

Step 4 – If the absolute difference between $V_1(x)$ and $V_0(x)$ or $\mathbb{E}_c[P_1(x,c)]$ and $\mathbb{E}_c[P_0(x,c)]$ or $s_1(x)$ and $s_0(x)$ for any x exceeds 0.000001, we update the guess to $P_1(x,c), s_1(x)$, and $V_1(x)$ for all x and repeat steps 1 to 3. Otherwise, $(\mathbf{P}_1, \mathbf{s}_1, \mathbf{V}_1)$ constitutes the model solution $(\mathbf{P}^*, \mathbf{s}^*, \mathbf{V}^*)$.

Doraszelski and Satterthwaite (2010) establish the existence of a symmetric equilibrium in pure strategies for a closely related model. Note that the model may have multiple equilibria. (Doraszelski and Satterthwaite, 2010; Weintraub et al., 2008).

7.2.2 Outer Maximum Likelihood estimation

Denote the set of states associated with type l by X_l. Equipped with the solution to the model from the previous step, the likelihood function is given by:

$$\begin{aligned} L(k|x_{jt},\theta) = \sum_t \Bigg\{ \sum_j \Bigg[\sum_{x \in X} \mathbb{1}(x_{jt} = x) \ln(s^*(x|k,\theta)) \Bigg] \\ + \sum_{l \in \{1,2,3,4\}} \left(J_l - \sum_{x \in X_l} \sum_j \mathbb{1}(x_{jt} = x) \right) \ln \left(\frac{J_l}{J} - \sum_{x \in X_l} s^*(x|k,\theta) \right) \Bigg\} \end{aligned} \tag{17}$$

The maximum likelihood estimation requires us to solve the model and determine $\mathbf{s}^*$ repeatedly for different cost parameter candidates.

7.3 Results

The estimation results are shown in Table 5. All estimated parameters are highly significant. Note that we estimate the mean parameters of the entry and operating cost distribution, which

[26] The exact expressions can be found in Appendix B.1.

		Type 1	Type 2	Type 3	Type 4
κ	Entry cost ($000s)	443.2	541.1	415.8	463.4
		(4.0)	(4.7)	(3.2)	(3.1)
φ	Operating cost ($00s)	36.09	61.56	73.88	89.87
		(0.06)	(0.08)	(0.08)	(0.09)
mc	Marginal cost ($)	65.72	63.72	77.73	117.00
		(0.26)	(0.29)	(0.35)	(0.52)

Table 5: Supply estimates.

is by many orders of magnitude larger than entry and operating costs that are actually *paid* upon entry as firms will only enter and incur the cost if the cost draw is low enough. Our estimates imply that hosts on average incur entry costs ranging from $1,231 to $3,892, which are monotone increasing in type and around $2,956 on average. This compares to an average expected lifetime profit of entrants between $2,467 and $7,815, with a total average of $5,933. Listings of any two types are by construction equally likely observed in the market although higher type listings are more profitable. Our estimates suggest that entry costs are consistently around 50% below the listing value upon market entry.

We estimate that on average hosts pay between $2,062 and $3,860 in operating costs every four weeks, depending on the type, with an average of $3,227. As with entry costs, higher-type listings incur higher operating costs. Intuitively, higher-type listings are more desirable to guests but also the hosts. The host's opportunity cost of renting out the apartment, rather than using it herself, is therefore higher. A type-4 listing, for example, costs almost twice as much as a type-1 listing to maintain.

While the average marginal costs are also broadly increasing in type, the increase is non-monotonic with type-2 listings having lower marginal costs on average than type 1 listings. Despite this non-monotonicity in marginal costs, variable weekly profits are monotonically increasing in type. The average variable weekly profit ranges from $255 to $481 with an average of $413. With average revenue per month ranging from $458 to $995, and a total average of $773, this implies an average variable profit margin of around 53.4% for active listings. The profit margin is quite similar across listing types.

8 Model Fit

To assess the accuracy of our model, we compare key empirical moments derived from the data with those arising in the model solution, when parameterizing it with the estimates obtained in the previous section. Table 6 presents a detailed comparison of these moments.

While the model's average number of listings and average prices are somewhat higher than those observed in the data, differences are small. The average occupancy rate in the simulated data aligns well with its empirical counterpart, underscoring the model's capacity to closely

match data moments. Similarly, the average rating differs only in the second decimal place. The average number of reviews is somewhat lower in the model (notably after winsorizing data reviews at 50), while the exit rate is slightly higher. We believe that this comparison demonstrates the model's strong ability to capture essential features of the data.

	# listings	Price	Occupancy	# reviews	Rating	Weekly exit	% unreviewed
Model	713	\$198	58%	16	4.27	4.0%	16%
Data	719	\$197	60%	18	4.69	2.9%	12%

Table 6: Model fit: simulated vs. observed moments.

Figure 9 compares the rating distribution in the model solution ("posterior (model)") to the rating distribution in the actual data ("posterior (data)"). The frequency of one-star, three-star, and four-star listings (i.e., ratings comprised of three good reviews and one bad review) are remarkably similar in the model and the actual data. The model predicts more five-star ratings than we see in the actual data. In the actual data, listings at the high end of the distribution typically have ratings between four and five stars. We believe that the difference can be explained by the fact that the number of reviews is censored at 50. Hence, the model allows for little variation in the ratings close to 5 stars. Ratings in the simulated data are based on at most 50 reviews but possibly on many more in the actual data. Note that Figure 9 illustrates the selection effect that underlies the model as well as the actual data. Recall that the mean quality of *entrant* listings is estimated to be 4.15 stars. This is lower than the average rating of *all active* listings, both in the model solution and actual data.

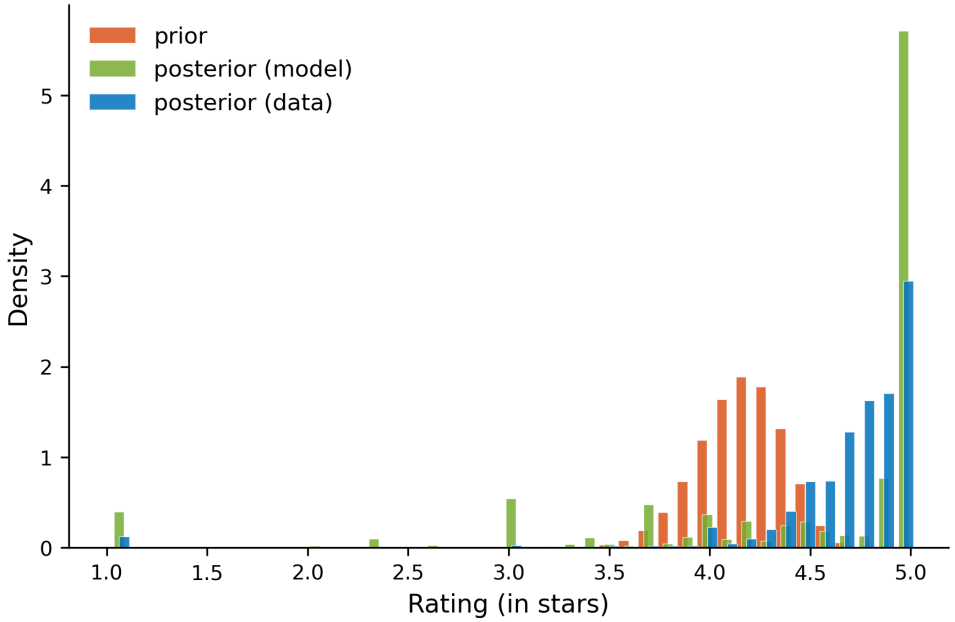


Figure 9: Rating distribution: model vs. data.

Recall that we estimate the cost parameters of the model by maximizing the likelihood of the equilibrium state distribution. According to our model, in equilibrium, there are on average 707

active listings, compared to 719 in the data. Figure 10 (left) shows that the number of listings generated by the model for each number of (good) reviews has a similar "pitched-fork" shaped distribution as the observed data, even though the pitchfork in the model might be somewhat more pronounced. Around 18% of active listings are unreviewed in the model, which compares to around 12% in the data. The model predicts that there are few to no poorly rated listings, especially for states associated with a relatively large number of reviews. Highly rated listings, on the other hand, tend to remain in the market and gather reviews. Figure 10 (right) shows that exit rates decrease with the number of reviews in both the model and the data, consistent with the selection mechanism: listings that accumulate good reviews are less likely to exit and more likely to get yet more reviews.

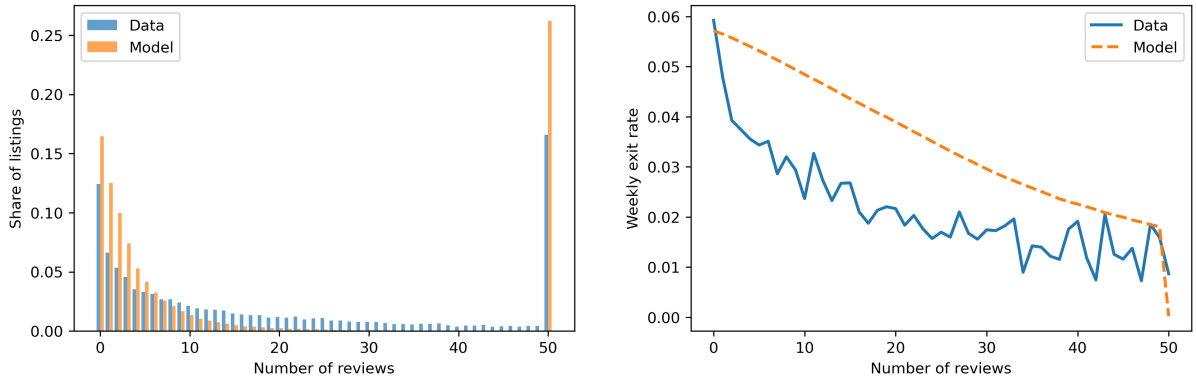


Figure 10: Model fit: listing distribution and exit rates by number of reviews.

Figure 11 compares occupancy rates and rental prices over a listing's lifecycle between model and data. The model reproduces the key patterns: new listings start with lower occupancy and prices, and both increase as listings age and accumulate reviews. The upward trajectories are broadly consistent, though the model predicts a somewhat flatter price profile than the data.

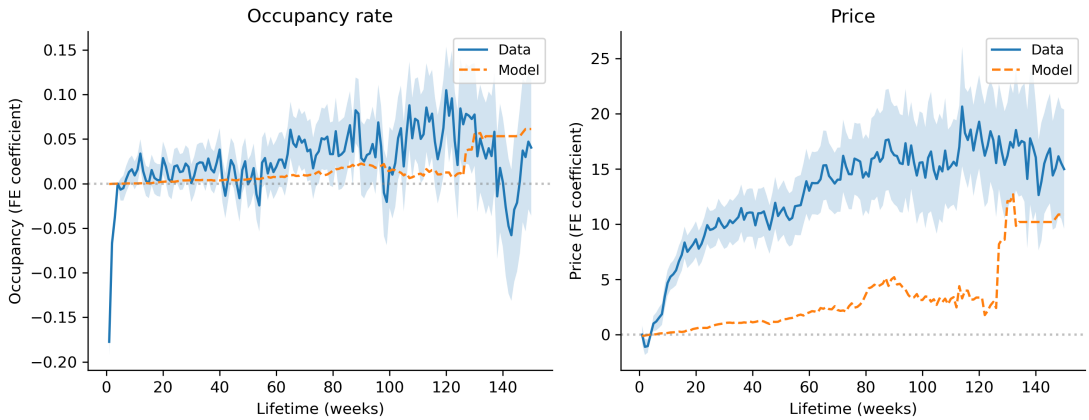


Figure 11: Model fit: occupancy rates and rental prices over a listing's lifetime.

Our model effectively captures many key aspects of the data. In particular, the model produces a state distribution similar to the empirical one. Furthermore, the direction and magni-

tude of the responses of hosts and guests to reviews in the model are broadly consistent with the data. Our counterfactual policies aim to increase the rate at which the quality of new listings is revealed by boosting their occupancy rates. According to our model, new listings underperform older ones economically. The occupancy rate of a new listing is typically 6.3% below the average, and its rental price is 4.1% lower than the average. Taken together, new listings generate 9.9% less revenue than the average listing. Consequently, our model predicts that new listings are more likely to exit the market compared to the average listing. We find a similar relationship in the data: new listings are twice as likely to exit in any given week. Moreover, new listings are booked 27% less frequently and generate 27% less revenue than the average listing.

9 Counterfactual analysis

We use the estimated model to quantify the informational externality inherent in the cold-start problem and to evaluate targeted policy instruments. We first measure the value of a first review, both in the short run, holding market structure fixed, and in the long run, allowing the full general equilibrium to adjust. We then evaluate two policies that target the externality from different sides of the market: a consumer transfer that compensates guests for booking unreviewed listings, and a host bonus that rewards hosts for obtaining their first review. For each policy, we consider three funding modes that differ in who bears the cost. Throughout, we calculate the compensating variation following Small and Rosen (1981) and McFadden (2012). We denote consumer surplus by CS, total host profits (variable per-period profits less operating and entry costs) by PS, Airbnb fee revenue by AR, and total surplus by $TS = CS + PS + AR$.

9.1 Measuring the informational externality

9.1.1 Short-run externality

We begin by asking: what is the immediate consumer-surplus value of entrants getting their first review? To answer this, we perturb the steady-state distribution by moving all unreviewed listings from state $(0, 0)$ to their expected post-review states and measure the resulting change in aggregate consumer surplus per review. Because this exercise abstracts from entry, exit, and re-optimization by other hosts, it isolates the short-run informational value of a review.

To understand how the supply side shapes this externality, we decompose the short-run consumer surplus change across three pricing regimes. First, under *uniform prices*, all listings charge the same price regardless of their review state; the resulting externality captures the pure information effect of a review. Second, under *static Nash prices* (setting the continuation value $V = 0$), markups vary across review states but hosts have no dynamic incentive to invest in reviews. Third, under *dynamic Nash prices* (P^*), hosts set prices accounting for the option value of future reviews, leading to lower prices early in a listing's lifetime.

Figure 12 presents the decomposition. Under uniform pricing, the pure information effect is large: a review partially resolves quality uncertainty, allowing consumers to reallocate demand toward higher-quality listings at no price premium.

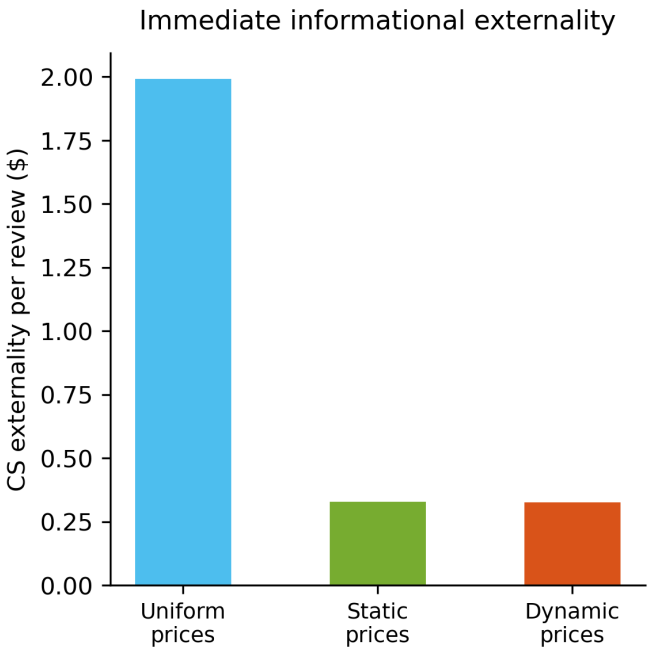


Figure 12: Short-run informational externality: consumer surplus gain per review under three pricing regimes.

Moving to static Nash prices reduces the externality by almost 90%. The mechanism is that prices increase in expected quality: once a listing receives a review, its host charges a higher markup that extracts much of the informational rent. Consumers still benefit from better information, but the host of the reviewed listing captures the lion's share of the value through higher prices.

The additional effect of dynamic pricing is less straightforward and could go in either direction. Dynamic considerations generally lead hosts to lower prices to attract bookings and accumulate reviews. However, whether this price reduction is larger for unreviewed listings than for listings with a single review depends on the shape of the value function: review accumulation does not end with the first review, so the option value of additional bookings may be similarly high for both states. If the price reduction is larger in the unreviewed state, then the value of a first review to consumers falls, because the reviewed listing now comes at a relatively higher price. If the price reduction is instead larger for the reviewed listing, dynamic pricing amplifies the externality. On average, these opposing forces approximately cancel: dynamic pricing neither adds to nor detracts from the short-run value of a review. However, this average masks interesting heterogeneity across listing types.

The by-type decomposition (Figure B.1 in Appendix B.4) reveals that the pure information effect is largest for high-type listings, whose quality uncertainty generates the greatest scope for demand reallocation. However, the net externality under dynamic pricing is most pronounced for low types, where investment pricing plays a smaller role. For the highest type, the net externality is in fact slightly negative: the price increase upon receiving a review is large enough that consumer surplus decreases in the short run.

This exercise abstracts entirely from entry and exit. All listings remain in their current states

after the perturbation, and no listing exits or accumulates further reviews. We analyze the contribution of these general-equilibrium channels below.

9.1.2 Long-run externality

We now turn to the long-run value of information by asking: how much welfare is lost because entrants start at $(0, 0)$ rather than having a review from the outset? In the long run, "faster learning" means that new listings obtain their first review sooner, and, crucially, that all agents incorporate this into their decision-making. The first-order consequences operate through the exit and entry margins. On the exit side, faster information revelation sharpens selection: high-quality listings are identified sooner and survive longer, while low-quality listings are revealed and exit more quickly. On the entry side, the prospect of faster learning changes the option value of entry, potentially altering the volume and composition of entrants. As the steady-state distribution shifts, prices adjust as a second-order response. Unlike the short-run exercise, these counterfactuals allow the full general equilibrium (prices, value functions, entry, exit, and the steady-state distribution) to adjust.

To isolate the effect of faster learning, we consider a seeded review experiment: the platform independently evaluates each new listing and provides a quality signal. Entrants start at a probabilistic mix of $(1, 1)$ and $(0, 1)$, weighted by the prior $\frac{a}{a+b}$, rather than at $(0, 0)$. Consumers play no role in generating the first review. One could imagine Airbnb implementing this by screening listings before they go live, for instance through professional inspections or algorithmic quality assessments. This is not designed primarily as a practical policy recommendation; rather, it provides a clean counterfactual for measuring the long-run informational externality.

	Seeded review
Δ CS	\$26,454 (+2.8%)
Δ PS	\$-694 (-2.5%)
Δ AR	\$2,103 (+2.9%)
Δ TS	\$27,863 (+2.6%)
Δ CS per add'l first review	\$988.50
Δ AR per add'l first review	\$78.60
Δ Entry value (wavg)	\$-180 (-4.0%)
Δ Reviews per period	+18 (+17.7%)
Δ Total listings	+4 (+0.7%)
Δ Avg. expected quality	+0.021 stars

Welfare in \$/period. Per-review: Δ welfare / Δ first reviews. Entrant value: expected PDV at entry.

Table 7: Long-run externality: welfare effects of providing entrants with a seeded first review.

Table 7 reports the results. The seeded review increases total surplus, but the gains accrue disproportionately to consumers: consumer surplus increases by \$4,511 per period, while total

host profits decline by $928 and Airbnb revenue increases modestly. These magnitudes confirm that the informational externality is economically meaningful and that an additional first review is worth substantially more to consumers than to hosts or the platform.

So why are the long-run consumer surplus gains so much larger than the short-run externality? To understand the contribution of faster selection, changes in entry incentives and price changes, we decompose ΔCS into three channels. Starting from the status quo, we first evaluate consumer surplus under counterfactual prices but the baseline state distribution (price effect). We then replace the baseline distribution with the counterfactual distribution rescaled to the same total number of listings, isolating the change in the quality composition of active listings (composition effect). Finally, the residual from the change in the total number of listings captures the entry effect.

	Seeded review
Δ CS	$26,454 (+2.8%)
price effect	$9,550 (+1.0%)
composition effect	$14,268 (+1.5%)
entry effect	$2,636 (+0.3%)

Decomposition of Δ CS ($/period). Composition: CF distribution rescaled to baseline total listings. Entry: residual from change in total listings.

Table 8: Decomposition of consumer surplus change into price, composition, and entry effects.

Table 8 reveals that, while all channels contribute positively, the composition effect is the dominant channel, dwarfing the price and entry effects. The mechanism is selection: when entrants receive a review at entry, high-quality listings are identified immediately and survive longer, while low-quality listings are revealed and exit the market faster. Figure A.2 in Appendix A.2 shows the resulting shift in the steady-state distribution: the mass at $N = 0$ disappears and redistributes toward higher review counts. Figure A.3 shows the corresponding shift in the quality distribution, with the density shifting rightward as the composition of active listings improves. The resulting improvement in quality composition is what generates the large consumer surplus gains. This stands in sharp contrast to the short-run exercise, which, by abstracting from entry and exit, captures only the immediate information value of a review and misses the selection channel entirely.

9.2 Policy interventions

Having established that the informational externality is economically meaningful, we evaluate two policies that target it from different sides of the market. Each policy is analyzed under three funding modes: (i) platform-funded, where Airbnb absorbs the cost; (ii) revenue-neutral with host financing, where entrant hosts pay a participation fee F at $N = 0$ calibrated so that total fees equal total policy transfers; and (iii) revenue-neutral with consumer financing, where all consumers pay a per-booking surcharge F_c calibrated analogously.

9.2.1 Consumer transfer

The first policy compensates consumers for the informational externality they create by booking unreviewed listings. Consumers who book a listing with $N = 0$ reviews receive a lump-sum transfer τ per booking. The transfer enters consumer utility additively, increasing the demand for unreviewed listings. More bookings in turn accelerate review accumulation, generating long-run information gains. We search over a grid of transfer values τ and, for each funding mode, select the transfer that maximizes total surplus.

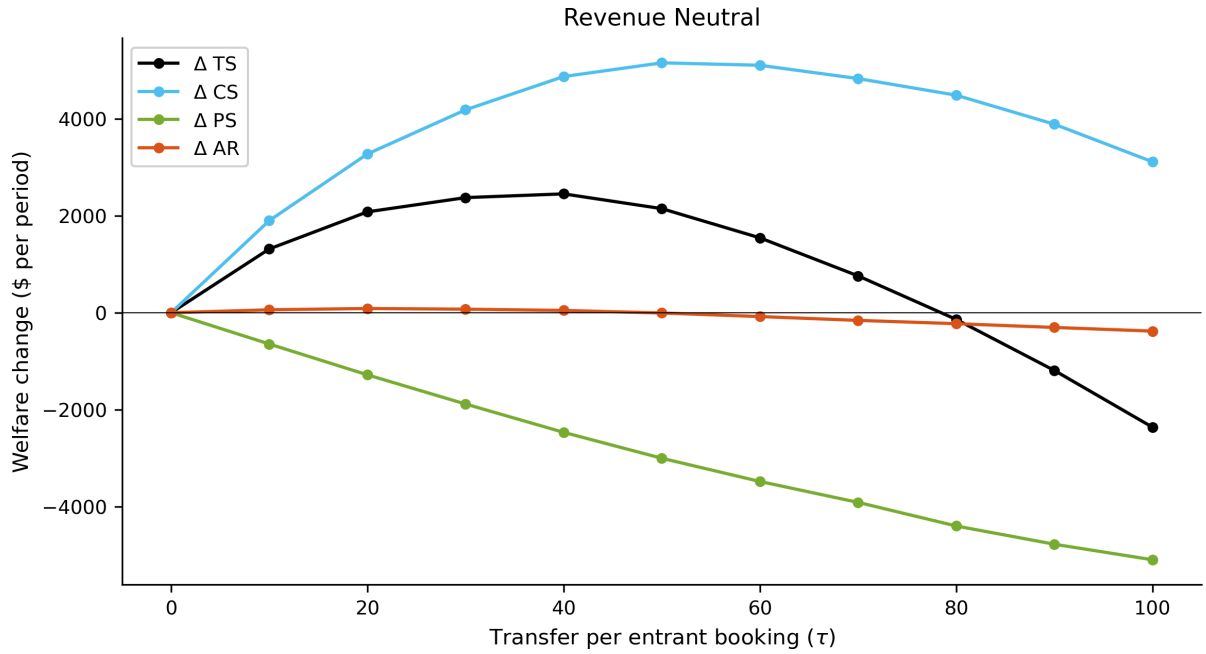


Figure 13: Consumer transfer policy (revenue-neutral, host-funded): welfare components as a function of the transfer τ. Left: consumer surplus, host profits, and Airbnb revenue. Right: total surplus decomposition.

Figure 13 illustrates how welfare components respond to the transfer under the revenue-neutral (host-funded) design.[27] As τ increases, consumer surplus rises while host profits decline, as the participation fee reduces entry incentives, partially offsetting the demand stimulus. The total surplus-maximizing transfer balances these forces. Table 9 compares the three funding modes at their respective optima.

The platform-funded mode generates the largest gains in consumer surplus ($22,793) and total surplus ($11,772), but at a substantial cost to Airbnb revenue (−$12,012). When hosts finance the policy, total surplus remains positive ($2,452) and Airbnb revenue is approximately unchanged, but the consumer surplus gain is considerably smaller ($4,873) as the participation fee screens out marginal entrants and reduces listing variety. Under consumer financing, the gains are more modest still ($366 in total surplus), because the surcharge is spread across all bookings and partially offsets the transfer's demand stimulus. Notably, the long-run externality

[27] The corresponding figures for the platform-funded and consumer-funded modes are in Figures A.4 and A.5 in Appendix A.3.

	Platform-funded	Revenue-neutral (hosts)	Revenue-neutral (consumers)
Δ CS	+22,793.00	+4,872.77	+825.02
Δ PS	+991.40	-2,467.97	-725.52
Δ AR	-12,011.94	+47.16	+266.12
Δ TS	+11,772.47	+2,451.96	+365.62
Δ CS (LR externality)	+66.56	+419.10	+22.37
Δ Total listings	+27	-23	+5
τ^*	\$40	\$40	\$10
Host fee F	—	\$160.58	—
Consumer surcharge F_c	—	—	\$1.2059

Each column evaluated at its own TS-maximizing τ^*. Changes relative to baseline, \$ per period.

Table 9: Consumer transfer: welfare effects under three funding modes, evaluated at the total surplus-maximizing transfer τ^*.

reduction (measured as the change in consumer surplus attributable to the information channel) is largest under host financing (\$419), reflecting the fact that the participation fee screens out low-value entrants and concentrates entry among hosts whose listings generate more information per booking.

9.2.2 First-review bonus

The second policy targets the supply side directly. Hosts receive a lump-sum bonus B upon transitioning from $(0,0)$ to $(1,1)$, i.e., upon obtaining their first good review. The bonus enters the host's pricing first-order condition, increasing the option value of being booked and thereby incentivizing hosts to lower prices at $N = 0$. Lower entrant prices attract more bookings, accelerating review accumulation.

Figure 14 displays the welfare effects under the revenue-neutral (host-funded) design.[28] The qualitative pattern mirrors the consumer transfer: consumer surplus increases with B, host profits decline as the participation fee rises, and the total surplus-maximizing bonus balances these margins. Table 10 compares the three funding modes.

The host bonus is somewhat more effective than the consumer transfer across all funding modes. Under platform funding, total surplus increases by \$12,659 (compared to \$11,772 for the transfer), and consumer surplus gains are approximately 10% larger. Under host financing, total surplus rises by \$4,243, nearly double the transfer's \$2,452. The bonus's advantage stems from its direct effect on the host's pricing problem: by increasing the value of the transition to $(1,1)$, the bonus lowers equilibrium entrant prices, which both attracts more bookings and accelerates information diffusion simultaneously. The consumer transfer, by contrast, increases

[28] The corresponding figures for the platform-funded and consumer-funded modes are in Figures A.6 and A.7 in Appendix A.3.

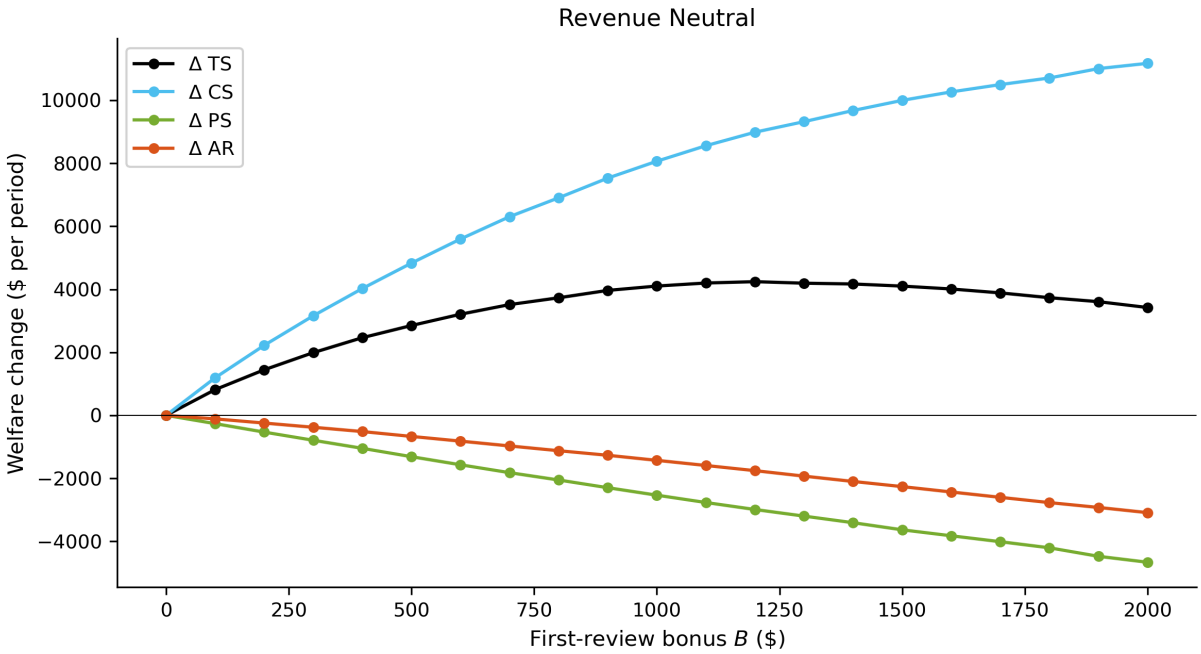


Figure 14: First-review bonus policy (revenue-neutral, host-funded): welfare components as a function of the bonus B. Left: consumer surplus, host profits, and Airbnb revenue. Right: total surplus decomposition.

	Platform-funded	Revenue-neutral (hosts)	Revenue-neutral (consumers)
Δ CS	+25,119.07	+8,986.36	+2,236.12
Δ PS	+1,179.03	-2,988.30	-1,094.66
Δ AR	-13,638.69	-1,755.46	-189.63
Δ TS	+12,659.41	+4,242.61	+951.83
Δ CS (LR externality)	+75.12	+520.82	+36.52
Δ Total listings	+30	-23	+9
B^*	$1,000	$1,200	$400
Host fee F	—	$195.20	—
Consumer surcharge F_c	—	—	$1.9203

Each column evaluated at its own TS-maximizing B^*. Changes relative to baseline, $ per period.

Table 10: First-review bonus: welfare effects under three funding modes, evaluated at the total surplus-maximizing bonus B^*.

demand for unreviewed listings without directly affecting host pricing incentives, so hosts may partially offset the demand stimulus by raising prices.

9.2.3 Comparing the two policies

The consumer transfer targets the demand side of the externality by compensating consumers for the information they generate, while the first-review bonus targets the supply side by rewarding hosts for obtaining reviews. Although the bonus is somewhat more effective, generating 7–73% larger total surplus gains depending on the funding mode, the two policies are not dramatically different in magnitude. The more consequential dimension is the choice of funding mode.

Under platform funding, both CS and TS are highest, because the platform effectively subsidizes information production without distorting entry incentives. However, this comes at a significant cost to Airbnb revenue, making implementation unlikely absent regulatory pressure. Under host financing, the participation fee screens out low-value entrants, which reduces listing variety but concentrates the market among hosts who generate more information per booking. Accordingly, the cold-start effectiveness (measured by the change in consumer surplus attributable to the information channel) is largest under host financing. Under consumer financing, the per-booking surcharge dilutes the policy's demand-side stimulus and yields the smallest welfare gains.

These results underscore a common theme: the supply-side response, through pricing, entry, and exit, mediates the welfare effects of any policy targeting the cold-start problem. Who bears the cost of the policy matters at least as much as which side of the market it targets.

10 Conclusion

In this study, we have examined the informational externality underlying the cold-start problem on Airbnb, taking into account the endogenous supply-side responses of hosts through pricing, entry, and exit decisions. We developed and estimated a structural model in which consumers learn from peer reviews while hosts use reviews to inform their dynamic decisions. Our estimates indicate that reviews significantly influence both booking decisions and host profitability.

We quantified the informational externality through two complementary exercises. In the short run, we gave every unreviewed listing its first review and decomposed the resulting consumer surplus gain per review across pricing regimes, finding that investment pricing by hosts partially substitutes for the information channel but does not fully internalize the social value of reviews. Entry and exit dynamics amplify the externality beyond what the short-run exercise captures. In the long run, we showed that providing entrants with a first review generates consumer surplus gains that substantially exceed the costs borne by hosts and the platform, confirming that the externality is economically meaningful.

We then evaluated two targeted policies, a consumer transfer for booking unreviewed listings and a host bonus for obtaining a first review, each under three funding modes. The host bonus is somewhat more effective, as it directly alters hosts' pricing incentives and accelerates review accumulation through the supply side. However, the choice of funding mode matters at least as much as the choice of policy instrument. Platform-funded designs generate the largest welfare gains but are costly for Airbnb. Revenue-neutral designs funded by host participation

fees screen out low-value entrants and are most effective at addressing the cold-start problem specifically, though they reduce overall listing variety.

These findings highlight the importance of accounting for the supply-side response when designing policies to promote information aggregation and diffusion on two-sided platforms. Hosts' pricing, entry, and exit decisions mediate the welfare effects of any intervention, and the alignment of private incentives with the policy objective is a first-order consideration. Our results generalize to any platform where peer reviews drive demand and sellers make dynamic decisions.

References

Aguirregabiria, Victor, Allan Collard-Wexler, and Stephen P Ryan, "Dynamic games in empirical industrial organization," in "Handbook of industrial organization," Vol. 4, Elsevier, 2021, pp. 225–343.

AirDNA, "Airbnbs in New York, Newark and Jersey City," 2019.

Almagro, Milena and Tomás Domínguez-Iino, "Location sorting and endogenous amenities: Evidence from amsterdam," *Econometrica*, 2025, *93* (3), 1031–1071.

Anderson, Michael and Jeremy Magruder, "Learning from the crowd: Regression discontinuity estimates of the effects of an online review database," *The Economic Journal*, 2012, *122* (563), 957–989.

Bao, Ying, Limin Fang, and Matthew Osborne, "The Effect of Quality Disclosure on Firm Entry and Exit Dynamics: Evidence from Online Review Platforms," *Available at SSRN 4804928*, 2024.

Barwick, Panle Jia and Parag A Pathak, "The costs of free entry: an empirical study of real estate agents in Greater Boston," *The RAND Journal of Economics*, 2015, *46* (1), 103–145.

Bergemann, Dirk and Juuso Välimäki, "Learning and strategic pricing," *Econometrica*, 1996, *64* (5), 1125–1149.

_ and _, "Market diffusion with two-sided learning," *The RAND Journal of Economics*, 1997, pp. 773–795.

_ and _, "Experimentation in markets," *The Review of Economic Studies*, 2000, *67* (2), 213–234.

Besanko, David, Ulrich Doraszelski, and Yaroslav Kryukov, "The economics of predation: What drives pricing when there is learning-by-doing?," *American Economic Review*, 2014, *104* (3), 868–97.

_ , _ , and _, "How efficient is dynamic competition? the case of price as investment," *American Economic Review*, 2019, *109* (9), 3339–64.

_ , _ , _ , and Mark Satterthwaite, "Learning-by-doing, organizational forgetting, and industry dynamics," *Econometrica*, 2010, *78* (2), 453–508.

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

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou, "Geography, transportation, and endogenous trade costs," *Econometrica*, 2020, *88* (2), 657–691.

Butler, Jeffrey V., Enrica Carbone, Pierluigi Conzo, and Giancarlo Spagnolo, "Past performance and entry in procurement: An experimental investigation," *Journal of Economic Behavior & Organization*, 2020, *173*, 179–195.

Cabral, Luis and Ali Hortacsu, "The dynamics of seller reputation: Evidence from eBay," *The Journal of Industrial Economics*, 2010, *58* (1), 54–78.

_ **and Michael H Riordan**, "The learning curve, market dominance, and predatory pricing," *Econometrica*, 1994, *62* (5), 1115–1140.

_ **and** _ , "The learning curve, predation, antitrust, and welfare," *The Journal of Industrial Economics*, 1997, *45* (2), 155–169.

Calder-Wang, Sophie, "The distributional impact of the sharing economy on the housing market," *Working Paper*, 2021.

Che, Yeon-Koo and Johannes Hörner, "Recommender Systems as Mechanisms for Social Learning," *The Quarterly Journal of Economics*, 2018, *133* (2), 871–925.

Chen, Jiawei, "How do switching costs affect market concentration and prices in network industries?," *The Journal of Industrial Economics*, 2016, *64* (2), 226–254.

Chen, Yanyou and Daniel Yi Xu, "A Structural Empirical Model of R&D, Firm Heterogeneity, and Industry Evolution," *The Journal of Industrial Economics*, 2023, *71* (2), 323–353.

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

Ching, Andrew T, "A dynamic oligopoly structural model for the prescription drug market after patent expiration," *International Economic Review*, 2010, *51* (4), 1175–1207.

Collard-Wexler, Allan, "Demand fluctuations in the ready-mix concrete industry," *Econometrica*, 2013, *81* (3), 1003–1037.

Conlon, Christopher T and Julie Holland Mortimer, "Demand estimation under incomplete product availability," *American Economic Journal: Microeconomics*, 2013, *5* (4), 1–30.

Dickstein, Michael J, "Efficient provision of experience goods: Evidence from antidepressant choice," *Working Paper*, 2021.

Doraszelski, Ulrich and Mark Satterthwaite, "Computable Markov-perfect industry dynamics," *The RAND Journal of Economics*, 2010, *41* (2), 215–243.

Dubé, Jean-Pierre H, Günter J Hitsch, and Pradeep K Chintagunta, "Tipping and concentration in markets with indirect network effects," *Marketing Science*, 2010, *29* (2), 216–249.

Fang, Limin, "The effects of online review platforms on restaurant revenue, consumer learning, and welfare," *Management Science*, 2022, *68* (11), 8116–8143.

Farronato, Chiara and Andrey Fradkin, "The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry," *American Economic Review*, 2022, *112* (6), 1782–1817.

Ferrari, Silvia and Francisco Cribari-Neto, "Beta regression for modelling rates and proportions," *Journal of Applied Statistics*, 2004, *31* (7), 799–815.

Fradkin, Andrey and David Holtz, "Do incentives to review help the market? Evidence from a field experiment on Airbnb," *American Economic Review: Insights*, 2022.

_ , Elena Grewal, and David Holtz, "Reciprocity and unveiling in two-sided reputation systems: Evidence from an experiment on Airbnb," *Marketing Science*, 2021, *40* (6), 1013–1029.

Frechette, Guillaume R, Alessandro Lizzeri, and Tobias Salz, "Frictions in a competitive, regulated market: Evidence from taxis," *American Economic Review*, 2019, *109* (8), 2954–92.

Hagiu, Andrei and Julian Wright, "Platforms and the exploration of new products," *Management Science*, 2020, *66* (4), 1527–1543.

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams, "Organizational Structure and Pricing: Evidence from a Large U.S. Airline," *The Quarterly Journal of Economics*, 09 2023, *139* (2), 1149–1199.

Huang, Yufeng, "Pricing Frictions and Platform Remedies: The Case of Airbnb," *Available at SSRN 3767103*, 2022.

Igami, Mitsuru and Kosuke Uetake, "Mergers, innovation, and entry-exit dynamics: Consolidation of the hard disk drive industry, 1996–2016," *The Review of Economic Studies*, 2020, *87* (6), 2672–2702.

Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica*, 1982, *50* (3), 649–670.

Kellogg, Ryan, "The effect of uncertainty on investment: Evidence from Texas oil drilling," *American Economic Review*, 2014, *104* (6), 1698–1734.

Kovbasyuk, Sergey and Giancarlo Spagnolo, "Memory and Markets," *The Review of Economic Studies*, 06 2023, *91* (3), 1775–1806.

Kremer, Ilan, Yishay Mansour, and Motty Perry, "Implementing the "wisdom of the crowd"," *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Lewis, Gregory and Georgios Zervas, "The welfare impact of consumer reviews: A case study of the hotel industry," *Working Paper*, 2016.

Li, Lingfang, Steven Tadelis, and Xiaolan Zhou, "Buying reputation as a signal of quality: Evidence from an online marketplace," *The RAND Journal of Economics*, 2020, *51* (4), 965–988.

Luca, Michael, "Reviews, reputation, and revenue: The case of Yelp. com," *Working Paper*, 2016, (12-016).

McFadden, Daniel, "Computing willingness–to–pay in random utility models," in "Trade, theory and econometrics" 2012, pp. 275–296.

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

Proserpio, Davide, Wendy Xu, and Georgios Zervas, "You get what you give: theory and evidence of reciprocity in the sharing economy," *Quantitative Marketing and Economics*, 2018, *16*, 371–407.

Reimers, Imke and Joel Waldfogel, "Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings," *American Economic Review*, 2021, *111* (6), 1944–71.

Rossi, Michelangelo, "Competition and Reputation in an Online Marketplace: Evidence from Airbnb," *Management Science*, 2024, *70* (3), 1357–1373.

Small, Kenneth A and Harvey S Rosen, "Applied welfare economics with discrete choice models," *Econometrica*, 1981, pp. 105–130.

Sweeting, Andrew, Dun Jia, Shen Hui, and Xinlu Yao, "Dynamic price competition, learning-by-doing, and strategic buyers," *American Economic Review*, 2022, *112* (4), 1311–1333.

_ , James W Roberts, and Chris Gedge, "A model of dynamic limit pricing with an application to the airline industry," *Journal of Political Economy*, 2020, *128* (3), 1148–1193.

Takahashi, Yuya, "Estimating a war of attrition: The case of the us movie theater industry," *American Economic Review*, 2015, *105* (7), 2204–2241.

Vellodi, Nikhil, "Ratings design and barriers to entry," *Working Paper*, 2024.

Weintraub, Gabriel Y, C Lanier Benkard, and Benjamin Van Roy, "Markov perfect industry dynamics with many firms," *Econometrica*, 2008, *76* (6), 1375–1411.

_ , _ , and _ , "Computational methods for oblivious equilibrium," *Operations Research*, 2010, *58* (4-part-2), 1247–1265.

Williams, Kevin R, "The welfare effects of dynamic pricing: Evidence from airline markets," *Econometrica*, 2022, *90* (2), 831–858.

Wollmann, Thomas G, "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles," *American Economic Review*, 2018, *108* (6), 1364–1406.

Wu, Chunhua, Hai Che, Tat Y Chan, and Xianghua Lu, "The economic value of online reviews," *Marketing Science*, 2015, *34* (5), 739–754.

Zervas, Georgios, Davide Proserpio, and John W Byers, "A first look at online reputation on Airbnb, where every stay is above average," *Marketing Letters*, 2021, *32*, 1–16.

Appendix

A.1 The informational externality

This appendix provides derivations and proofs for the results in Section 3. Under the Gumbel distributional assumption on taste shocks, the consumer chooses product j with probability $q_j = \exp(\nu_j) / (1 + \exp(\nu_I) + \exp(\nu_E))$ and the outside option with probability $q_0 = 1 - q_E - q_I$, where $\nu_j = \omega_j - p_j$. The expected consumer surplus is $CS = \ln(1 + \exp(\nu_E) + \exp(\nu_I)) + \gamma$, where γ is the Euler-Mascheroni constant.

Proof of Lemma 1 (Uniform prices)

Under uniform prices $p_E = p_I = \bar{p}$, consumer surplus $CS^U(\omega_E^e) = \ln(1 + \exp(\omega_I - \bar{p}) + \exp(\omega_E^e - \bar{p})) + \gamma$ is convex in ω_E^e, since $\frac{d^2 CS^U}{d(\omega_E^e)^2} = q_E(1 - q_E) > 0$. Since the expected posterior equals the prior ($\mathbb{E}[\omega_E^e] = \omega_E^0$), $\Delta CS^U > 0$ follows from Jensen's inequality. □

Proof of Lemma 2 (Static Nash prices)

We prove the two parts of the lemma in turn.

$\Delta CS^* > 0$**:** We show that $CS^*(\omega_E^e)$ is increasing and convex in ω_E^e; the result then follows from Jensen's inequality and $\mathbb{E}[\omega_E^e] = \omega_E^0$.

Step 1 (Equilibrium and comparative statics). The first-order condition for firm j yields $p_j^* = \frac{1}{1-q_j}$. The logit choice probability satisfies $q_j = e^{\delta_j} \cdot q_0$, where $\delta_j = \omega_j - p_j$ and $q_0 = 1 - q_E - q_I$. Taking logs and substituting the Nash markup:

$$\ln q_j = \omega_j - p_j^* + \ln q_0 = \omega_j - \tfrac{1}{1-q_j} + \ln q_0.$$

Rearranging, the equilibrium shares satisfy

$$\underbrace{\ln q_j + \tfrac{1}{1-q_j}}_{\equiv h(q_j)} = \omega_j + \ln q_0, \quad j \in \{E, I\}, \tag{18}$$

where h is strictly increasing ($h'(q) = \frac{1}{q} + \frac{1}{(1-q)^2} > 0$). Totally differentiating (18) with respect to ω_E gives the system

$$\begin{aligned} h'(q_E)\frac{dq_E}{d\omega_E} &= 1 - \frac{1}{q_0}\left(\frac{dq_E}{d\omega_E} + \frac{dq_I}{d\omega_E}\right), \\ h'(q_I)\frac{dq_I}{d\omega_E} &= -\frac{1}{q_0}\left(\frac{dq_E}{d\omega_E} + \frac{dq_I}{d\omega_E}\right), \end{aligned} \tag{19}$$

Solving yields $\frac{dq_E}{d\omega_E} > 0$ and $\frac{dq_I}{d\omega_E} < 0$. Since $\frac{dp_j}{d\omega_E} = \frac{1}{(1-q_j)^2}\frac{dq_j}{d\omega_E}$, the entrant's price is increasing and the incumbent's price is decreasing in ω_E. Moreover, $(1 - q_E)^2 h'(q_E) > 1$ implies $\frac{dp_E}{d\omega_E} < 1$: the entrant's price increase is smaller than the quality improvement.

Step 2 (Consumer surplus is increasing). Consumer surplus is $CS^* = \ln(1 + e^{\delta_E} + e^{\delta_I}) + \gamma$, where $\delta_j = \omega_j - p_j^*$ and $\delta_0 = 0$. The total derivative is

$$\frac{dCS^*}{d\omega_E} = q_E \frac{d\delta_E}{d\omega_E} + q_I \frac{d\delta_I}{d\omega_E} > 0, \tag{20}$$

since $\frac{d\delta_E}{d\omega_E} = 1 - \frac{dp_E}{d\omega_E} \in (0,1)$ and $\frac{d\delta_I}{d\omega_E} = -\frac{dp_I}{d\omega_E} > 0$ by Step 1.

Step 3 (Consumer surplus is convex). Since CS^* is the log-sum-exp of $(\delta_0, \delta_E, \delta_I)$, the chain rule and the positive semi-definiteness of the log-sum-exp Hessian yield the decomposition

$$\frac{d^2CS^*}{d\omega_E^2} = \underbrace{\mathrm{Var}_q\left[\frac{d\delta_j}{d\omega_E}\right]}_{\text{information effect } (\geq 0)} + \underbrace{\sum_{j\in\{E,I\}} q_j \frac{d^2\delta_j}{d\omega_E^2}}_{\text{curvature correction}}, \tag{21}$$

where the variance is over $j \in \{0, E, I\}$ with weights (q_0, q_E, q_I) and values $\left(0, \frac{d\delta_E}{d\omega_E}, \frac{d\delta_I}{d\omega_E}\right)$. The information effect is strictly positive: the outside option anchors $\frac{d\delta_0}{d\omega_E} = 0$, while $\frac{d\delta_E}{d\omega_E}, \frac{d\delta_I}{d\omega_E} > 0$ by Step 2. The curvature correction captures second-order price adjustments; for $\omega_j \in [0,1]$, the information effect strictly dominates the curvature correction, ensuring $\frac{d^2CS^*}{d\omega_E^2} > 0$.[29]

The outside option is essential for this result. Without it, the two products form a duopoly with $q_E + q_I = 1$ and consumer surplus becomes concave in ω_E under Nash pricing ($\frac{dCS}{d\omega_E} = 1 - q_E$ is strictly decreasing), so that reviews would reduce expected consumer surplus.

$\Delta CS^* < \Delta CS^U$**:** Under Nash pricing, the entrant raises its price after a good review and lowers it after a bad review, extracting part of the informational rent. Since the rent extraction $CS^U(\omega_E^e) - CS^*(\omega_E^e)$ is convex in ω_E^e (the Nash curvature $\frac{d^2CS^*}{d\omega_E^2}$ is strictly smaller than the uniform curvature $q_E(1-q_E)$), Jensen's inequality gives $\mathbb{E}[CS^U - CS^*] > CS^U(\omega_E^0) - CS^*(\omega_E^0)$, which rearranges to $\Delta CS^U > \Delta CS^*$.[30] □

Proof of Proposition 1 (Dynamic pricing)

We first characterize the dynamic equilibrium, then prove the consumer surplus result.

Equilibrium characterization. The game has three states: no review (0), good review (g), and bad review (b). States g and b are absorbing with per-period profits $\pi_j^i = q_j^i/(1-q_j^i)$ and present values $\Pi_j^i = \pi_j^i/(1-\delta)$. In state 0, a purchase of E generates a review with probability v_r. Firm j in state 0 solves:

$$V_j^0 = \max_{p_j} \; q_j p_j \; + \; \delta\left[q_E v_r\left(\omega_E^0 \Pi_j^g + (1-\omega_E^0)\Pi_j^b\right) + (1 - q_E v_r)\, V_j^0\right]. \tag{22}$$

Denoting $\mathbb{E}[\Pi_j] = \omega_E^0 \Pi_j^g + (1-\omega_E^0)\Pi_j^b$ and $\Gamma_j \equiv \mathbb{E}[\Pi_j] - V_j^0$, and rearranging (22):

$$V_j^0(1-\delta) = \max_{p_j} \; q_j p_j + \delta\, q_E v_r\, \Gamma_j.$$

The first-order condition is:

$$\frac{\partial q_j}{\partial p_j}\, p_j + q_j + \delta v_r \frac{\partial q_E}{\partial p_j}\, \Gamma_j = 0. \tag{23}$$

[29] For all $\omega_E, \omega_I \in [0,1]$, the ratio of the information effect to the absolute value of the curvature correction exceeds 2.

[30] The convexity of the rent extraction function is verified for all $\omega_E, \omega_I \in [0,1]$.

Using the logit derivatives $\frac{\partial q_E}{\partial p_E} = -q_E(1-q_E)$ and $\frac{\partial q_E}{\partial p_I} = q_E q_I$, the equilibrium prices are:

$$p_E^0 = \frac{1}{1-q_E} - \delta v_r \Gamma_E, \qquad p_I^0 = \frac{1}{1-q_I} + \frac{\delta v_r q_E}{1-q_I} \Gamma_I.$$

Sign of Γ_j. Solving (22) for V_j^0:

$$\Gamma_j = \frac{\mathbb{E}[\pi_j] - \pi_j^0}{1 - \delta(1 - q_E^0 v_r)}, \tag{24}$$

where $\pi_j^0 = q_j^0 p_j^0$ and $\mathbb{E}[\pi_j] = \omega_E^0 \pi_j^g + (1-\omega_E^0)\pi_j^b$. Hence $\Gamma_j > 0$ iff $\mathbb{E}[\pi_j] > \pi_j^0$.

Entrant: $\pi_E^*(\omega_E^e) = q_E/(1-q_E)$ is convex in ω_E^e (at fixed prices, $\frac{\partial^2 \pi_E}{\partial(\omega_E^e)^2} = \pi_E > 0$; this extends to equilibrium prices for $\omega_j \in [0,1]$). By Jensen's inequality, $\mathbb{E}[\pi_E] > \pi_E^*(\omega_E^0) \geq \pi_E^0$,[31] so $\Gamma_E > 0$.

Incumbent: $\pi_I^*(\omega_E^e) = q_I/(1-q_I)$ is concave in ω_E^e at Nash prices (with the outside option, $q_I \neq 1 - q_E$, so the convexity of $1/x$ that holds in a duopoly does not carry over).[32] Jensen's inequality gives $\mathbb{E}[\pi_I] < \pi_I^*(\omega_E^0)$. For small δ, $\pi_I^0 \approx \pi_I^*(\omega_E^0)$, so $\Gamma_I < 0$.

Consumer surplus comparison. Let $W_g = CS^*(\omega_E^g)/(1-\delta)$, $W_b = CS^*(\omega_E^b)/(1-\delta)$, and $\bar{W} = \omega_E^0 W_g + (1-\omega_E^0) W_b = \mathbb{E}[CS^*]/(1-\delta)$. Consumer welfare in state 0 satisfies:

$$W_0 = CS^0 + \delta\left[q_E^0 v_r \bar{W} + (1 - q_E^0 v_r) W_0\right]. \tag{25}$$

Solving and computing $\bar{W} - W_0$:

$$\bar{W} - W_0 = \frac{\mathbb{E}[CS^*] - CS^0}{1 - \delta(1 - q_E^0 v_r)}.$$

Since the denominator is positive, $\bar{W} > W_0$ iff $\Delta CS^0 = \mathbb{E}[CS^*] - CS^0 > 0$. The decomposition follows from adding and subtracting $CS^*(\omega_E^0)$.

The introductory pricing effect $CS^0 - CS^*(\omega_E^0) \geq 0$ is weakly positive: the entrant always cuts its price ($\Gamma_E > 0$), and for moderate δ the incumbent also cuts its price ($\Gamma_I < 0$).[33] Since the introductory pricing effect is increasing in δ while ΔCS^* is independent of δ, ΔCS^0 is decreasing in δ. At $\delta = 0$, $\Delta CS^0 = \Delta CS^* > 0$. □

Simulations of firm's value gain from a review and consumer surplus. Figure A.1 illustrates the key comparative statics of the dynamic equilibrium by solving the fixed-point problem numerically for a grid of discount factors. Panel (a) shows that the entrant's value gain from a review (Γ_E) is positive and increasing in δ, confirming that the entrant always engages in introductory pricing. The incumbent's value gain (Γ_I) starts negative (the incumbent is hurt by reviews in expectation, due to the concavity of its profit) but may turn positive for large δ, as the dynamic equilibrium deviation becomes so large that even the reduced expected post-review profit exceeds the heavily distorted pre-review profit. Panel (b) shows that the consumer surplus value of a review (ΔCS^0) is decreasing in δ: the introductory pricing effect ($CS^0 - CS^*$, increasing in δ) gradually erodes the information value (ΔCS^*, independent of δ). For sufficiently

[31] $\pi_E^*(\omega_E^0) \geq \pi_E^0$ holds when firm E's best response to the other firm's dynamic price yields weakly lower per-period profit than the Nash profit. This is satisfied in the cold-start setting and verified for all $\omega_j \in [0,1]$.

[32] At fixed prices, $\frac{\partial^2 \pi_I}{\partial(\omega_E^e)^2} < 0$ for all $\omega_j \in [0,1]$.

[33] $CS^0 \geq CS^*(\omega_E^0)$ is verified for all $\omega_E^0, \omega_I \in [0,1]$ and all $\delta \in [0,1)$.

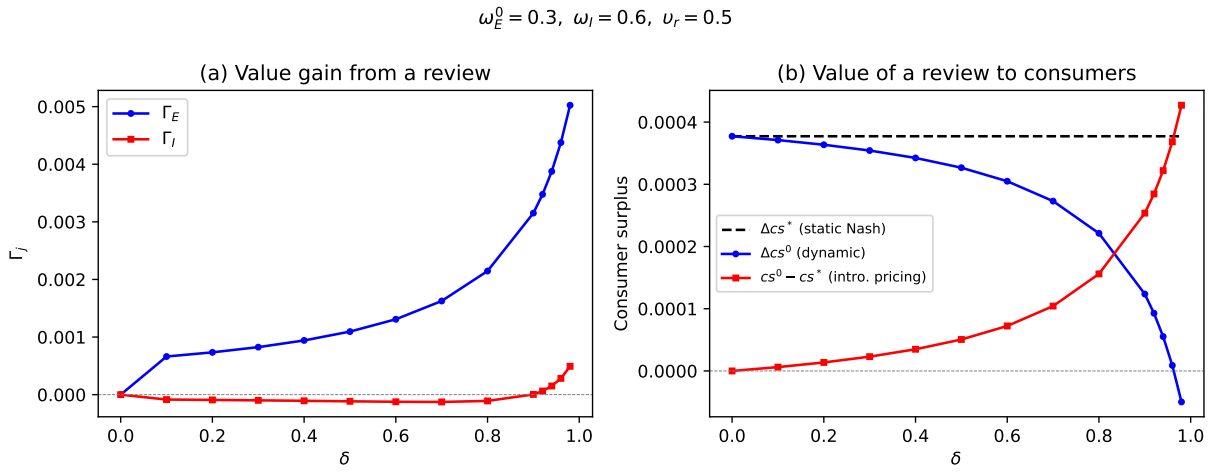


Figure A.1: Equilibrium objects as a function of δ for $\omega_E^0 = 0.3$, $\omega_I = 0.6$, $\upsilon_r = 0.5$. Panel (a): the entrant's value gain from a review (Γ_E) is always positive, while the incumbent's (Γ_I) starts negative and may turn positive for large δ. Panel (b): the consumer surplus value of a review (ΔCS^0) is decreasing in δ and may turn negative; the introductory pricing effect ($CS^0 - CS^*$) is increasing in δ, while the information value (ΔCS^*) is independent of δ.

large δ, ΔCS^0 turns negative, meaning consumers prefer the pre-review state with its introductory discounts. The threshold δ at which ΔCS^0 turns negative is lower when υ_r is higher, since a larger review probability amplifies the dynamic pricing incentive.

A.2 Long-run externality: seeded review

A.3 Policy interventions: additional funding modes

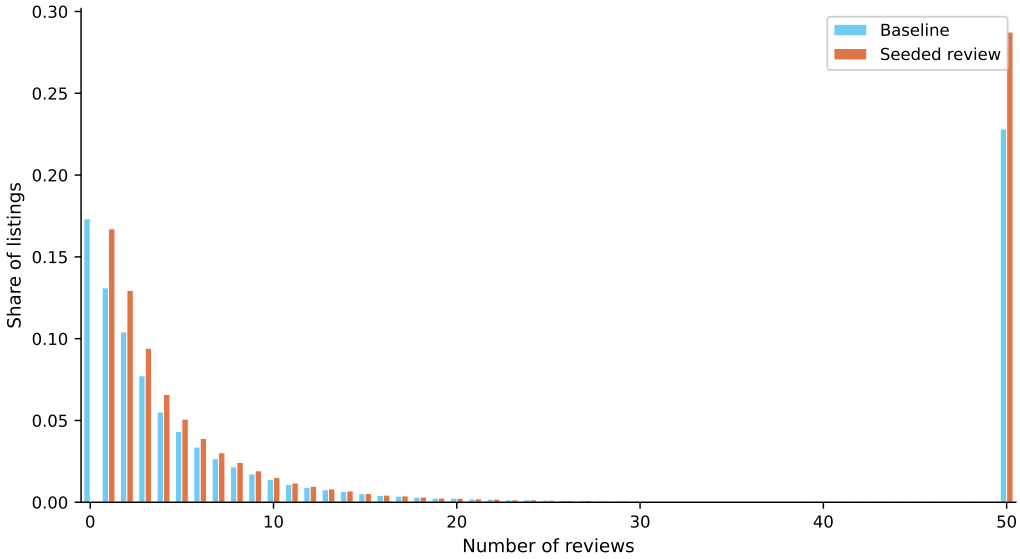


Figure A.2: Steady-state distribution of listings by number of reviews (N), baseline vs. seeded review. The seeded review eliminates the mass at $N = 0$ (entrants receive a review at entry) and shifts the distribution toward higher review counts.

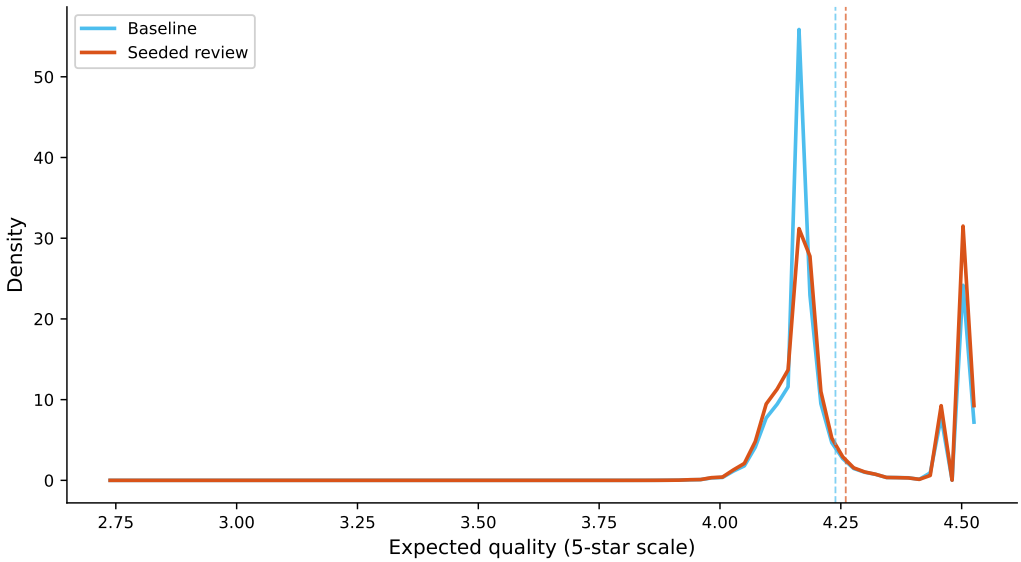


Figure A.3: Density of expected listing quality (5-star scale), baseline vs. seeded review. Dashed lines indicate the average expected quality in each regime. The seeded review sharpens selection: the distribution shifts rightward as high-quality listings survive longer and low-quality listings exit faster.

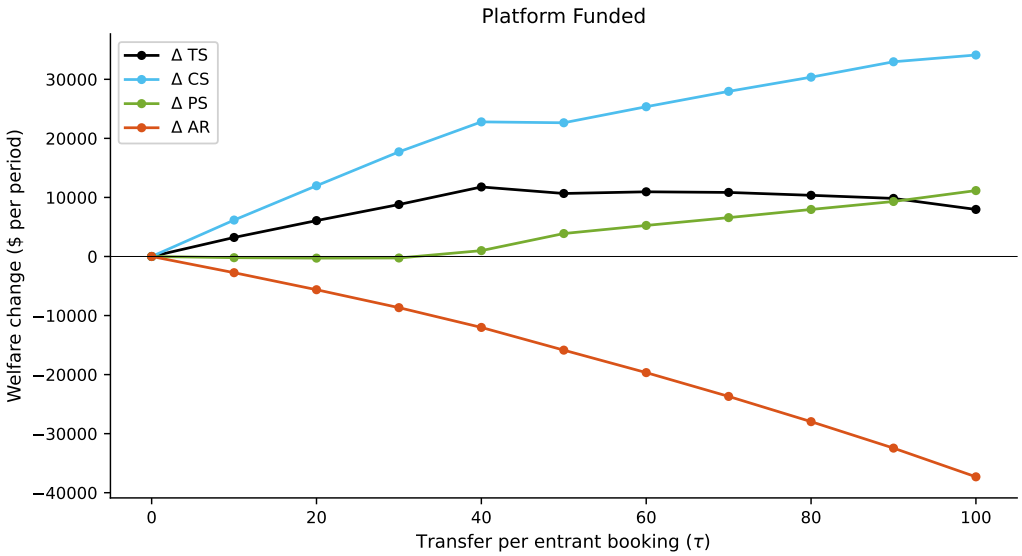


Figure A.4: Consumer transfer policy (platform-funded): welfare components as a function of the transfer τ.

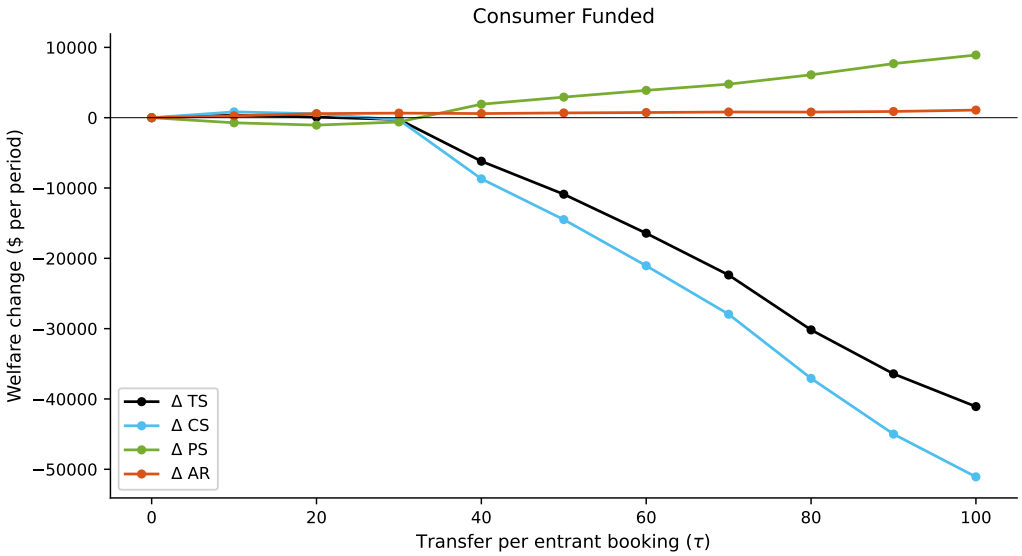


Figure A.5: Consumer transfer policy (revenue-neutral, consumer-funded): welfare components as a function of the transfer τ.

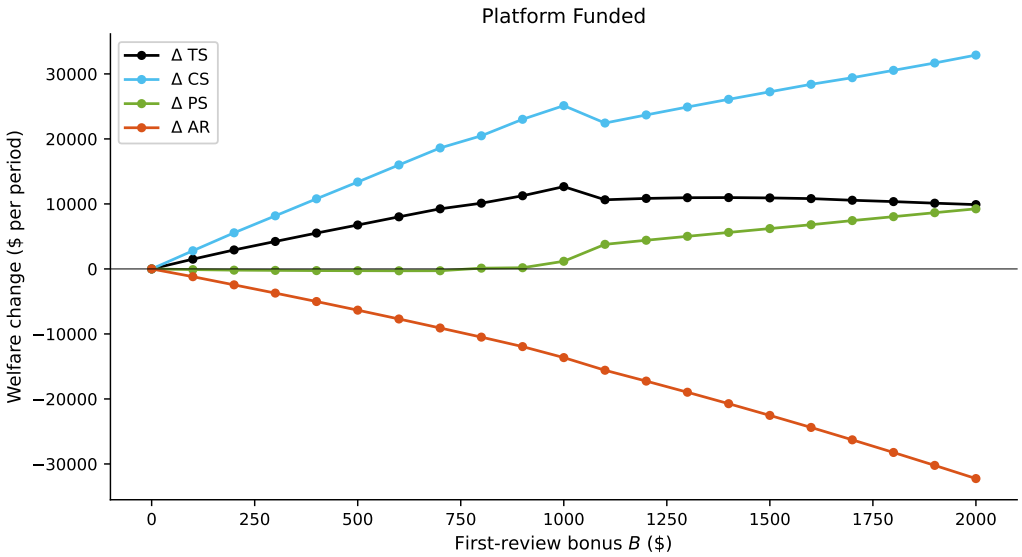


Figure A.6: First-review bonus policy (platform-funded): welfare components as a function of the bonus B.

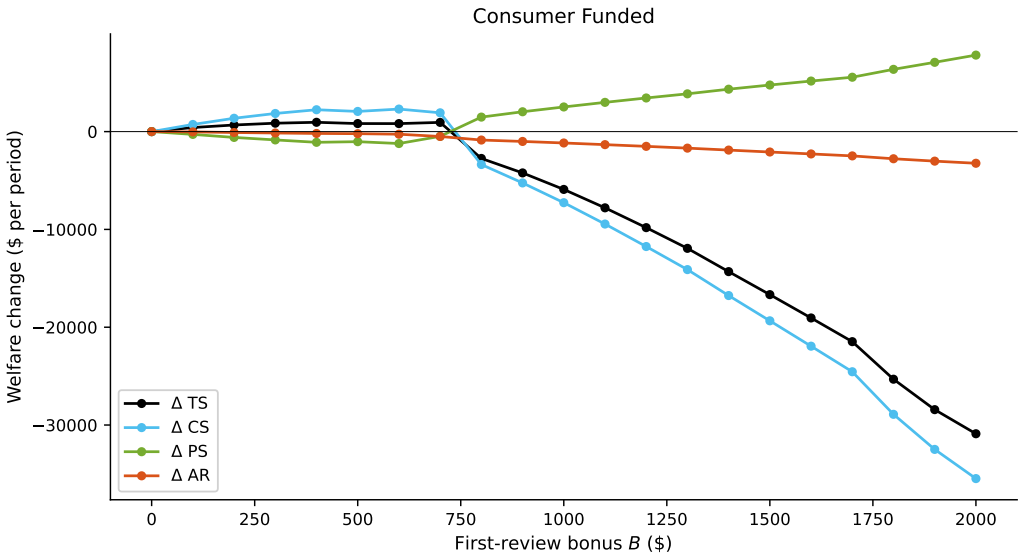


Figure A.7: First-review bonus policy (revenue-neutral, consumer-funded): welfare components as a function of the bonus B.

Online Appendix

B.1 First-order derivative of the pricing stage

Remember that hosts set their prices after they learn about their marginal cost shock, but they do not know about their rival's cost shock. To decrease the computational burden, we assume that hosts react to average prices by competitors instead of the full price distribution induces by $P_0(x, c)$. Hence, the FOC is given by:

$$w'(x, c|P_1^k, P_0) = \mathbb{E}\left[q(P_1^k, \mathbb{E}_0[p_{-j}], x_j) + \frac{\partial q(P_1^k, \mathbb{E}_0[p_{-j}], x)}{\partial p}(P_1^k - c) + \delta \frac{\partial \mathbb{E}_{x'}[V(x') - (1-\chi(x'))\bar{\phi}|\mathbb{E}_0[p_{-j}]]}{\partial P_1^k}\right]$$

where $\mathbb{E}_0[p_{-j}] = \mathbb{E}_c[p_{-j}|p_{-j} = P_0(x, c)]$.

B.2 Expected operating cost

The expected operating cost is given by the following expression.

$$\begin{aligned}\mathbb{E}[\phi_l|\phi_l \leq V(x)](1-\chi(x)) =& \bar{\phi}_l - \chi(x)\mathbb{E}[\phi_l|\phi_l > V(x)] \\ =& \bar{\phi}_l - \chi(x)(\bar{\phi}_l + V(x)) \\ =& (1-\chi(x))\bar{\phi}_l - \chi(x)V(x)\end{aligned}$$

B.3 Expected entry cost

The expected cost of entry is given by the following expression.

$$\begin{aligned}\lambda \mathbf{E}_\kappa[\kappa|\kappa \leq \delta\mathbb{E}_l[V((0,0,l))]] =& \bar{\kappa} - (1-\lambda)\mathbb{E}_\kappa[\kappa|\kappa > \delta\mathbb{E}_l[V((0,0,l))]] \\ =& \bar{\kappa} - (1-\lambda)\left(\kappa + \delta\mathbb{E}_l[V((0,0,l))]\right) \\ =& \lambda\bar{\kappa} - (1-\lambda)\delta\mathbb{E}_l[V((0,0,l))] \qquad (26)\end{aligned}$$

B.4 Additional Tables and Figures

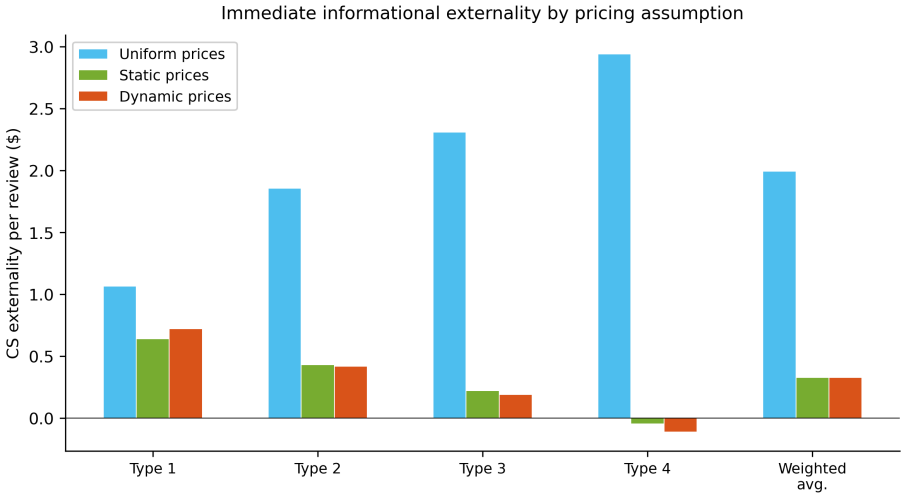


Figure B.1: Short-run informational externality by listing type: consumer surplus gain from one additional review under three pricing regimes.

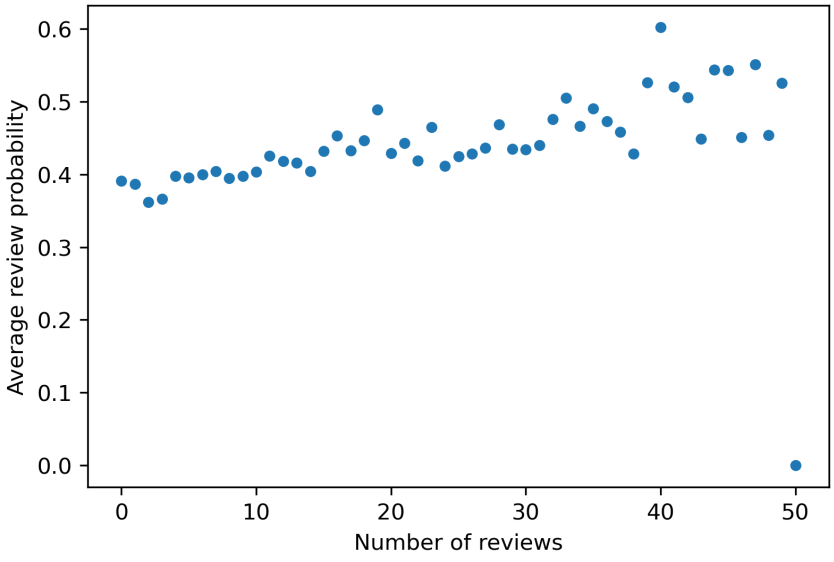


Figure B.2: Empirical review probabilities conditional on the number of reviews.

Dependent variable:	p	B	exit
	(1)	(2)	(3)
K	0.26***	0.0024***	-0.0008***
	(0.052)	(0.0004)	(0.0001)
N-K	-1.28***	-0.0066*	0.0003
	(0.38)	(0.0035)	(0.0003)
Fixed effects			
Property ID	Yes	Yes	
Time	Yes	Yes	Yes
Observations	149,533	149,533	149,533
R^2	0.002	0.001	0.007
Within R^2	0.009	0.002	-0.005

Significance: ***p<0.01, **p<0.05, *p<0.1

Table B.1: Reduced-form regressions of price, occupancy, and exit on good and bad reviews.

	Price	N	K	Within-group share
	(1)	(2)	(3)	(4)
z_1	-0.37	-5.01***	-4.68***	0.0000***
	(0.24)	(0.14)	(0.13)	(0.0000)
B_{t-3}	-2.35*	0.24	0.19	0.0000*
	(1.42)	(0.80)	(0.74)	(0.0000)
B_{t-4}	0.022	-0.62	-0.69	-0.0000
	(1.53)	(0.84)	(0.78)	(0.0000)
B_{t-5}	6.51***	12.3***	11.4***	-0.0001***
	(1.33)	(0.75)	(0.70)	(0.0000)
K/N	29.5***	25.0***	28.6***	0.0007***
	(4.11)	(1.41)	(1.18)	(0.0000)
Age IV	2.45***	-1.04***	-0.77***	0.0000
	(0.32)	(0.088)	(0.074)	(0.0000)
$K/N \times$ Age IV	-2.56***	0.055	-0.14*	0.0000
	(0.34)	(0.094)	(0.081)	(0.0000)
n	0.0072**	-0.031***	-0.030***	-0.0000***
	(0.0028)	(0.0019)	(0.0018)	(0.0000)
Constant	168.9***	75.0***	65.3***	0.0017***
	(4.72)	(2.37)	(2.09)	(0.0001)
Observations	24,012	24,012	24,012	24,012
R^2	0.006	0.114	0.117	0.072
F-statistic	12.7	421.2	464.0	265.9

Robust (HC1) standard errors in parentheses. ***p<0.01, **p<0.05, *p<0.1.

Table B.2: First-stage regression results.