Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Learning through Ratings under Endogeneous Product Quality

Fan Zhang

fan.zhang@novasbe.pt

Universidade Nova de Lisboa https://orcid.org/0000-0002-2744-9719

Przemyslaw Jeziorski

University of California, Berkeley

Research Article

Keywords: Rating Design, Asymmetric Information, Dynamic Games, Endogenous Quality

Posted Date: September 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10876732/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Learning through Ratings under Endogenous Product Quality*

Przemyslaw Jeziorski† Fan Zhang‡

August 25, 2026

Abstract

This paper examines the design of rating systems in markets where product quality is unobserved and dynamically managed by sellers. Ratings serve a dual purpose: they inform consumers about unobserved quality and, meanwhile, incentivize sellers' quality provision efforts. We model this interaction as a dynamic Bayesian game in which sellers can enhance quality through costly effort, while buyers update their beliefs about quality based on observed ratings. Using data from a short-term rental platform where we directly observe sellers' quality provision efforts, we empirically quantify sellers' intertemporal incentives and consumer learning. Using the structural estimates, we conduct counterfactual experiments to assess the impact of alternative rating systems, explicitly accounting for the endogenous provision of quality. We find that rating manipulation reduces equilibrium revenue, quality, and consumer welfare, while rating forgiveness improves these market outcomes.

Keywords: Rating Design, Asymmetric Information, Dynamic Games, Endogenous Quality

*We would like to thank Ganesh Iyer, Miguel Villas-Boas, Pedro Gardete, Thomas Otter, Samuel Goldberg, Robin Ng, and seminar participants at Nova SBE, University of Rochester Simon School of Business, and Bristol Empirical IO Workshop, 26th CEPR-JIE Conference on Applied Industrial Organization, Barcelona Summer Forum, 2026 Markets and Waves Conference, and 24th ZEW Conference on the Economics of Information and Communication Technologies for their valuable insights and feedback. Disclosure: Przemyslaw Jeziorski has received compensation exceeding $10,000 from the data-providing company. The findings and conclusions presented in this paper are independent and do not pose any conflict of interest. This work was funded by Fundação para a Ciência e a Tecnologia (UIDB/00124/2025, UID/PRR/124/2025, Nova School of Business and Economics) and LISBOA2030 (DataLab2030 - LISBOA2030-FEDER-01314200).

†przemekj@berkeley.edu, UC Berkeley, Haas School of Business

‡fan.zhang@novasbe.pt, Nova School of Business and Economics, Universidade NOVA de Lisboa, Campus de Carcavelos, 2775-405 Carcavelos, Portugal

1 Introduction

With the proliferation of digital platforms, reputation systems based on customer reviews have become the cornerstone of modern marketplaces. The primary function of online reviews is to improve market efficiency by addressing the problem of information asymmetry, including adverse selection (Akerlof, 1978), where exploitative sellers enter the market while conscientious ones exit, and moral hazard (Arrow, 1968), characterized by an insufficient provision of costly effort. In such circumstances, reviews serve a dual purpose: they act not only as an informational signal to other consumers but also as a motivational tool for the rated product and service providers. These two functions are not always aligned: fully transparent information does not necessarily lead to better outcomes, and rating uncertainty can outperform full-information outcomes by incentivizing quality provision. [1] The analysis can become even more complex in dynamic settings where product and service quality evolve over time. In such markets, frequent maintenance efforts and managerial investments are required and may cause great variation in quality over time. The primary rating design challenge in these settings lies in quantifying and balancing informing consumers and incentivizing sellers, accounting for sellers' intertemporal decisions.

This paper studies the rating design on digital platforms where quality evolves endogenously through sellers' costly effort choices. A central challenge in identifying moral hazard on digital platforms is that sellers' efforts are typically unobserved. For example, if we observe a sequence of ratings such as "5, 5, 3, 5, 3", it is difficult to determine whether the fluctuations reflect genuine changes in seller effort and underlying quality, or simply rating noise and idiosyncratic variation in individual guest experiences. To overcome this challenge, we utilize novel data where the sellers' efforts are directly observable, including their timing, frequency, and monetary expenses. These data allow us to quantify sellers' intertemporal reputation incentives and consumer learning in an equilibrium framework, and use the estimated model to study the counterfactual market outcomes of alternative rating designs.

We first present descriptive evidence on the key mechanisms underlying the dynamic equilibrium model. On the supply side, we show that hosts strategically increase maintenance efforts following negative ratings, suggesting the presence of moral hazard. On the demand side, property occupancy rates increase with average ratings, consistent with existing research on consumer learning from

[1] For example, Saeedi and Shourideh (2023) find that rating uncertainty is a fairly robust feature of optimal rating design under moral hazard. Recent empirical work (Vatter, 2023) also shows evidence on the benefit of coarsening information. Hörner and Lambert (2021) study the design of motivational ratings in a dynamic context.

reviews. Moreover, we find that additional ratings raise demand when average ratings are high but reduce demand when the average ratings are low, indicating active consumer learning from rating signals.

To quantify these mechanisms and conduct counterfactual analysis of alternative rating designs, we develop and estimate a dynamic structural model of quality provision and consumer learning. Property quality evolves over time, and hosts can improve quality through costly efforts. Meanwhile, consumers form beliefs about latent quality based on observed ratings and make booking decisions accordingly. The model captures the joint evolution of quality, reputation, consumer beliefs, bookings, and seller effort in equilibrium. To account for heterogeneity across hosts, we classify listings into two latent segments based on property characteristics, host behavior, and business performance, and allow each segment to have distinct underlying structural parameters. Our segmentation analysis reveals a value segment and a premium segment. Compared with the premium segment listings, the value segment listings tend to have smaller units, lower ratings, lower prices, less frequent maintenance, and more bookings. Further estimates reveal that relative to the premium segment listings, the value segment listings have slower organic quality depreciation, higher maintenance costs relative to nightly profits, lower effort effectiveness, greater idiosyncratic rating noise, and higher baseline demand. These heterogeneity estimates allow us to study how different listing segments respond to alternative counterfactual rating designs.

We use our estimated model to evaluate counterfactual rating policies, including rating manipulation and rating forgiveness. We study rating manipulation as a form of strategic practice that artificially inflates ratings without improving underlying quality, thereby introducing positive noise into the rating system. Beyond the misinformation effect documented in the existing literature(Gandhi et al., 2024), we identify an additional equilibrium mechanism: manipulated ratings weaken sellers' incentives to exert costly effort to maintain and improve true quality. As a result, the booking rate, average quality, and average booked quality all decline, even when consumers are sophisticated and can update their beliefs to account for the manipulation effect. Among the two segments, the value segment suffers a larger negative effect compared with the premium segment. We next study rating forgiveness policies that selectively remove negative noise in ratings. Unlike rating manipulation, rating forgiveness strengthens incentives for quality provision by reducing the likelihood that high-quality sellers are penalized by noisy negative reviews. We find that rating forgiveness increases maintenance effort, booking rates, average quality, and average booked quality relative to the baseline. These positive effects are especially strong in the premium segment relative

to the value segment. Overall, our results reveal an important asymmetry in rating system design: adding positive noise is harmful, while removing negative noise can improve market outcomes.

Related Literature This paper contributes to the literature on reputation systems in online marketplaces. [2] Specifically, we contribute to the growing literature on the supply-side effects of online reviews. [3] Existing studies document that reviews influence seller behavior through multiple channels, including review manipulation (Mayzlin et al., 2014; Luca and Zervas, 2016; He et al., 2022; Gandhi et al., 2024), manager responses to customers (Ma et al., 2015; Proserpio and Zervas, 2017; Chevalier et al., 2018; Wang and Chaudhry, 2018), quality improvement (Klein et al., 2016; Hollenbeck, 2018; Ananthakrishnan et al., 2023), and seller exit decisions (Cabral and Hortaçsu, 2010; Batikas and Kretschmer, 2018). Our paper focuses on the quality improvement channel and contributes to the existing literature in three dimensions.

First, most prior work indirectly infers unobserved seller effort from observed outcomes such as service-centric platform statistics (Rossi, 2024), excessive bunching around rating rounding thresholds (Hunter, 2022), and sentiment variation in textual reviews (Ananthakrishnan et al., 2023), etc. In contrast, we directly observe the timing, monetary cost, and subsequent outcomes of hosts' quality provision efforts over multiple years, which allows us to better measure the incentives and effectiveness of seller quality investment.

Second, existing studies adopt either an informational or incentive perspective on reviews. Under the informational perspective, firms learn from customer feedback and improve quality accordingly, in line with the quality function deployment (QFD) concept.[4] Under the incentive perspective, reviews incentivize firms' quality provision through reputation concerns, but existing work largely studies the *static* equilibrium effects of different policy environments (Vatter, 2023; Rossi, 2024; Leyden, 2025). Our paper extends this literature by studying the *intertemporal* seller incentives in environments with dynamically evolving quality, where forward-looking sellers strategically choose costly efforts in response to the reputation incentives generated by consumer learning.

[2] Tadelis (2016), Seiler et al. (2018), and Pocchiari et al. (2024) provide recent overviews of reputation systems in online platform marketplaces.

[3] Besides rating design based on supply side responses, there is also a large literature on the demand side response to reviews and the factors influencing review volume and valence (Chevalier et al., 2018; Fradkin et al., 2021; Filippas et al., 2022; Fradkin and Holtz, 2023) and the welfare analysis of review systems(Lewis and Zervas, 2016; Reimers and Waldfogel, 2021).

[4] For example, Ananthakrishnan et al. (2023) show that hotels improve dimensions of service quality that are frequently mentioned in reviews. Kaul et al. (2024) study the effect of managers eliciting feedback in the offline environment. When small business owners start seeking feedback from some of their customers, changes based on the feedback have positive spillovers to increase sales made by customers from whom feedback is not sought.

Third, we contribute to the literature on optimal information design under moral hazard. Recent theoretical work shows that fully transparent ratings may not maximize welfare when sellers strategically respond to ratings (Hörner and Lambert, 2021; Saeedi and Shourideh, 2023). We bring this literature to data by empirically estimating a structural model that jointly captures consumer learning and endogenous quality provision. This allows us to quantify how alternative rating systems, rating manipulation, and rating forgiveness policies affect equilibrium quality, consumer welfare, and seller incentives.

More broadly, our paper relates to the literature on moral hazard and mechanism design in various markets such as insurance and credit (Holmström, 1999; Einav et al., 2013; Jeziorski et al., 2017; Einav and Finkelstein, 2018; Kim et al., 2022; Jin, 2023; Xin, 2023; Vatter, 2023; Yang et al., 2024), as well as the literature on the dynamics of consumer learning under different rating systems (Papanastasiou et al., 2018; Besbes and Scarsini, 2018; Ifrach et al., 2019; Acemoglu et al., 2022). Our contribution is to integrate these two strands by modeling both endogenous seller effort and consumer learning dynamics in an equilibrium framework.

The remainder of the paper is organized as follows. Section 2 reports the industry background, data, and descriptive analysis. In Section 3, 4, and 5, we specify and estimate the model. We present our counterfactual analysis in Section 6 and conclude in Section 7.

2 Data and Descriptive Analysis

2.1 Industry Background

We study the short-term rental (STR) industry in this paper. STR is a business model in which sellers ("hosts") rent furnished properties for a short time, compared to long-term rentals - primarily unfurnished properties rented annually. The STR buyers ("guests") purchase vacation rentals as a relatively inexpensive alternative to hotels. The global STR was valued at $134.51 billion in 2024 and is expected to grow at a CAGR of 11.4% from 2025 to 2030. [5] STRs have a long history and were typically executed via informal markets. Today, most STRs occur via online platforms like Airbnb, VRBO, Expedia, and booking.com. These platforms facilitate the interaction between hosts and guests, from search, booking, and payment to check-out and feedback. STR platforms provide a contractual framework and arbitrate disputes. A significant difference between hotels and

[5] Source: "Short-term Vacation Rental market Size, Share & Trends Analysis Report", https://www.grandviewresearch.com/industry-analysis/short-term-vacation-rental-market-report

STRs is the concentration of ownership. The majority of hosts are individuals offering one or a few residences for rent. In comparison, a hotel typically offers multiple rooms and could be part of a larger chain. For this reason, the degree of asymmetric information about product quality can be more substantial for STR than for hotels. To mitigate the asymmetric information problem, the platforms introduce elaborate review systems. Ratings has become the leading driver of demand for STRs, as well as the search and ranking algorithms on the platforms.

We study the largest STR platform, Airbnb, which was founded in 2007, and has since grown to over 5 million hosts with over 8 million listings. [6] Airbnb's annual revenue was $9.9 billion in 2023, rivaling the revenue of large international hotel chains.[7] The platform review system allows double-blind guest and host ratings within two weeks after the stay, which are then published for other potential guests to see. On a scale from 1 to 5, guests can rate their stays in overall experience and multiple subcategories. [8] The current rating system shows consumers the full history of the overall experience ratings [9], the average rating and the total number of ratings.

Compared with long-term rentals, besides the more severe information asymmetry problem and thus the necessity of a rating system, short-term rentals also require significantly more resources to manage due to the high turnover rate. Regular extra requirements include replying to daily guest messages, calendar management across platforms, scheduling turn-overs, cleaning, and maintenance. Quality is not fixed over time and can be altered with hosts' costly efforts. These features of the STR market make it an ideal context for us to study the dynamic quality provision efforts under information asymmetry.

2.2 Data

We utilize data from an agency company specializing in the management of short-term rental (STR) properties. Property owners hire the agency to manage their accounts on STR platforms, and this project will focus on the Airbnb platform. The management agency does not engage in direct marketing or sales of accommodations but instead relies entirely on STR platforms to generate reservations. Acting as an intermediary, the agency assumes control over hosts' platform accounts and operates them on their behalf in exchange for a commission on rental revenue. These accounts

[6] Source: https://news.airbnb.com/about-us/.

[7] Source: https://www.statista.com/statistics/1193134/airbnb-revenue-worldwide/.

[8] Section A.1 provides more background information about the data and platform layout.

[9] By default, reviews appear in chronological order, with the most recent review shown first. Source: https://www.airbnb.com/help/article/13#section-heading-9.

are managed through programmatic access facilitated by official platform partnerships. When a listing is connected to the agency, its Airbnb presentation remains unchanged. Specifically, the agency's branding is not incorporated into the listing, which continues to display the host's name and photo. As a result, guests interacting with the platform—whether making reservations or providing feedback—remain unaware of the agency's involvement. Therefore, for the purpose of this study, the context is best understood as Airbnb's platform dynamics, with the agency serving solely as a data provider.

The management agency also offers web and mobile applications to property owners, enabling them to access various management services. Through this software, hosts can order services, including cleaning and property maintenance, and access business analytics. Routine cleaning and laundry services encompass standard cleaning, laundry, and turnover preparation, post-checkout restocking (upon request), and quality control monitoring. Additionally, maintenance and repair services include deep cleaning, recommended twice annually, as well as repairs such as plumbing, fixing leaking roofs and faucets, and patching damaged drywall, etc.

In our analysis, we differentiate between *extra maintenance services* and *basic turnover services*. Basic turnover services represent the essential minimum required for market participation, encompassing tasks such as cleaning, restocking supplies, and providing towel and linen services. The absence of these fundamental services is classified as an operational failure, often resulting in customer refunds and, in some cases, the suspension of the listing. Given these consequences, no rational host would deliberately forgo basic services, and such instances are exceedingly rare in our dataset.

The primary objective of our model is to capture the long-term evolution of quality, encompassing factors such as the condition of wall paint, the quality of furniture, and the overall state of towels, linens, and amenities. Unlike basic turnover services, which are essential for maintaining market presence, extra maintenance efforts involve more substantial and periodic upkeep activities, including repainting, furniture repairs, roofing maintenance, and the full replacement of towels and linens. A lack of maintenance services does not necessarily result in immediate guest dissatisfaction. Instead, temporary lapses in service may go unnoticed for extended periods, depending on the overall condition of the property and its amenities. However, over time, the absence of adequate maintenance leads to a gradual decline in quality, ultimately reflected in lower guest reviews. Given that extra maintenance activities often incur significant costs, hosts must strategically balance these

expenditures against the objective of sustaining high customer satisfaction and favorable ratings.[10]

The management agency provides comprehensive information on Airbnb listing characteristics, reservations, guest reviews, and maintenance requests. A novel key feature of this dataset is its ability to capture direct observations of the timing, frequency, and costs associated with hosts' quality provision efforts, particularly in the form of extra maintenance service requests. This granularity allows us to empirically assess the presence of moral hazard in quality provision within the framework of Airbnb's review system. By linking hosts' service efforts to their precedent review history and subsequent guest feedback, we can evaluate how the incentives created by the platform's reputation mechanisms influence maintenance decisions and overall service quality.

Table 1 presents the summary statistics of the data. The final sample contains 185 listings, 9,517 bookings, and 52,344 booked nights from March 2015 to November 2021. On average, listings have 1.69 bedrooms and 1.46 bathrooms, with a nightly price of $116.79. Maintenance requests occur in approximately 2% of stays. When maintenance occurs, the associated expenses are economically meaningful, averaging $91.85 per request. The average overall rating is 4.70 out of 5, while the median rating is 5, suggesting that hosts may perceive any rating below five stars as unfavorable. Meanwhile, the variation in ratings indicates cross-listing and intertemporal heterogeneity. A closer examination of within-listing intertemporal variation in ratings, together with observed maintenance efforts, provides insight into the dynamics of quality evolution and the effectiveness of hosts' efforts to improve service quality. Figure 1 shows the transitional probability of a rating increasing, remaining unchanged, or decreasing, conditional on whether a maintenance occurred between reviews. Relative to cases without maintenance, maintenance shifts probability mass upward in the rating transition distribution: it increases the probability of an upward transition and reduces the probability of a downward transition. This upward shift of probability mass provides evidence that maintenance efforts are effective in improving quality and subsequent ratings.

2.3 Descriptive Evidence

This section presents descriptive evidence on the two central mechanisms underlying the structural model. On the supply side, we show that hosts strategically invest in quality in response to repu-

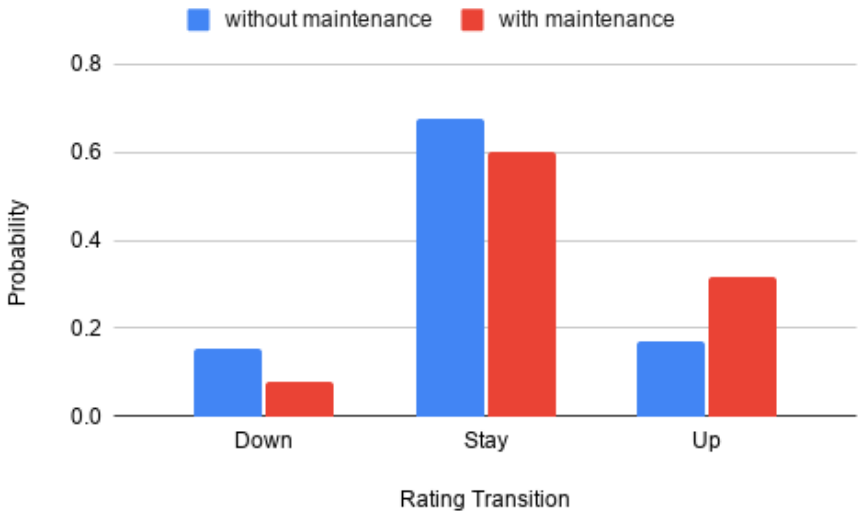
[10] Another perspective on quality depreciation is offered by Ericson and Pakes (1995), who allows for the possibility that a firm's relative quality deteriorates as the outside option improves. In their framework, improvements in the quality of the outside option impose a depreciation shock on all firms in the market, even if their absolute quality remains unchanged. This conceptualization extends our model to environments where natural (physical) quality erosion is less relevant, but where the outside option is itself dynamic.

Table 1: Summary Statistics

	Mean	Median	Std.	Min	Max
Number of Bedrooms	1.69	1	1.21	0	6
Number of Bathrooms	1.46	1	0.80	0	5
Maintenance Frequency	0.02	0	0.15	0	1
Maintenance Expenses ($)	91.85	47	124.21	6	643.75
Nightly Price ($)	116.79	88.56	113.86	8.8	1396.98
Number of Nights	4.97	3	9.93	1	186
Number of Guests	2.83	2	2.25	1	14
Occupancy Rate	0.63	0.67	0.33	0.01	1
Number of Ratings	42.33	29	42.82	1	266
Overall Experience Ratings	4.70	5	0.79	1	5
Subcategory Ratings					
Value	4.69	5	0.70	1	5
Check-in	4.85	5	0.54	1	5
Accuracy	4.79	5	0.63	1	5
Location	4.76	5	0.59	1	5
Communication	4.84	5	0.58	1	5
Cleanliness	4.73	5	0.78	1	5

Note: The table presents summary statistics of the data. The final sample contains 185 listings, 9,517 bookings, and 52,344 booked nights from March 2015 to November 2021.

Figure 1: Rating Transition and Maintenance Effort



Note: The figure reports the rating transition probability distribution, conditional on with and without a maintenance between reviews.

tation incentives generated by ratings. On the demand side, we show that consumers actively learn from rating histories when making booking decisions. Together, these findings motivate a dynamic equilibrium framework in which forward-looking hosts strategically manage quality while consumers update beliefs about latent listing quality from publicly observable reviews.

2.3.1 Ratings and Host Efforts

We first examine the relationship between ratings and hosts' maintenance efforts. We observe that listings with higher overall quality tend to submit more maintenance requests, with a raw correlation of 0.25 between ratings and maintenance activity. However, this relationship may be confounded by reverse causality and spurious correlation, potentially driven by unobserved heterogeneity in the cost of effort across hosts. To address these concerns, we exploit the timing feature of the data: the exact timing of guest reviews is stochastic, reflecting randomness in reservation arrivals and stay durations.[11] We implement a regression discontinuity (RD) design around the timing of review submissions. The key idea is that market beliefs about listing quality update discretely when a new review is posted. At the same time, quality depreciation is itself likely to be gradual and stable, assuming that guests do not deliberately sabotage the property upon checkout.[12] Importantly, hosts may respond to negative reviews even when the reviews themselves are noisy signals of true quality. Because future guests update their beliefs based on observed ratings, a low rating can reduce perceived listing quality regardless of whether the underlying service quality has actually changed. This creates a strategic incentive for hosts to exert additional effort following an unfavorable review in order to prevent further deterioration in reputation. In other words, even relatively high-quality hosts may "hedge" after receiving a negative review by investing in maintenance to avoid a sequence of poor ratings that could be particularly damaging.

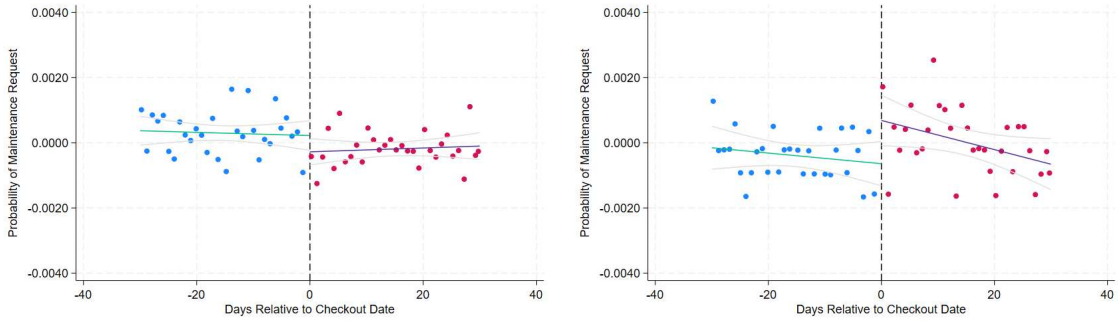
Figure 2 illustrates the probability of a maintenance request within a 30-day window before and after the checkout date of a reviewed stay. We classify a good review as a 5-star rating and a bad review as any rating below 5 stars. The figure shows a clear increase in maintenance activity following bad reviews, whereas maintenance probabilities remain relatively stable following good

[11] Similar identification has been employed by Vana and Lambrecht (2021) to identify the impact of the last review on demand.

[12] The identification strategy becomes increasingly valid as the time window around review submission narrows, since the likelihood of organic quality depreciation becomes minimal in a short time window. Even in cases when latent quality also changes, the observed behavioral response would still be consistent with moral hazard, since both quality and belief shocks generate incentives for costly efforts from forward-looking hosts. The structural model in the next section explicitly models and disentangles both channels.

reviews. The regression discontinuity results are reported in Table 2. Column (1) reports the results for stays with five-star ratings. Column (2) reports the results for stays with ratings of fewer than five stars. Consistent with the graph, we find that a bad rating increases the likelihood of placing a maintenance request, whereas a good rating does not. One alternative explanation is that hosts may have incomplete information about the quality, and learn about latent property problems from the negative reviews, and increase maintenance accordingly. Section A.2 provides evidence against this alternative host information mechanism. We analyze the informativeness of private feedback, which is unobservable to other consumers and find that hosts are responsive to bad ratings rather than informative private feedback. [13]

Figure 2: Timing of Hosts' Efforts



(a) Good Review (b) Bad Review

Note: The figure presents the maintenance probability around a 30-day window before and after the checkout date of a stay. 'Good Review' refers to a stay with a 5-star rating. 'Bad Review' is defined as a stay with a rating of less than 5 stars. The maintenance probability is demeaned by listing × checkout fixed effects.

2.3.2 Ratings and Guest Learning

We now provide evidence on the demand side and study whether consumers learn about listing quality from ratings. We define occupancy rate, d_{jt}, as the proportion of occupied nights before the end of the next stay after the end of the current stay, i.e., $d_{jt} = \frac{\text{occupied nights}_{j,t+1}}{\text{checkout}_{j,t+1} - \text{checkout}_{j,t}}$. To investigate the learning effect of ratings on consumers' posterior belief of quality, we run the following regression:

$$d_{jt} = \alpha + \beta R_{jt} + \gamma_1 \mathbf{1}(R_{jt} \geq \tau) \times N_{jt} + \gamma_2 \mathbf{1}(R_{jt} < \tau) \times N_{jt} + \epsilon_{jt}$$

[13] Furthermore, Section A.3 presents suggestive evidence of a "resting on the laurels" effect: hosts respond more strongly to bad reviews when their current reputation stock is relatively weak. In contrast, hosts with sufficiently strong reputations exhibit little behavioral response to negative reviews. Together, these findings are consistent with strategic maintenance behavior in response to reputation incentives.

Table 2: Timing of Hosts' Efforts

	Dependent Variable: Maintenance (m_{jt})	
	(1) Good Review	(2) Bad Review
Checkout (β)	-0.00047	0.00133***
	(0.00034)	(0.00046)
Observations	347,695	86,911

Note: The table reports results from regression discontinuity with triangular kernel results for a 30-day time window before and after the checkout date of a stay. Column (1) reports the results for stays with five-star ratings. Column (2) reports the results for stays with ratings of fewer than five stars. Standard errors are clustered at the reservation level and reported in parentheses. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

where R_{jt} denotes the average rating of listing j in period t, and N_{jt} denotes the total number of ratings. To allow for the number of ratings to have different effects on demand for good ratings and bad ratings, we allow the effect of review volume, γ_1 and γ_2, to be different for highly rated and poorly rated listings. The results are reported in Table 3. We report results using different average rating thresholds τ to define highly and poorly rated listings. Columns 1 and 2 show results where τ is set at the sample average rating of 4.70. Columns 3 and 4 report results where τ is set at 4. Consistent with findings in existing literature on the positive relationship between demand and ratings (Luca, 2016; Fan et al., 2016; Lewis and Zervas, 2019) and the theoretical framework of how informative rating systems help with the adverse selection problem, we observe that occupancy rate increases with average ratings. The total number of ratings exhibits heterogeneous effects for highly rated and poorly rated listings. When the threshold is set at four stars ($\tau = 4$), it increases demand for highly rated listings. For listings with low enough ratings, demand decreases with the total number of ratings. These patterns are consistent with consumer learning about unobserved listing quality through observed ratings.

Taken together, the evidence in this section suggests a dynamic interaction between ratings, consumer beliefs, and hosts' costly maintenance decisions. Ratings affect future demand through consumer learning while also shaping future quality by altering hosts' incentives to exert effort. These empirical patterns motivate our structural framework, in which forward-looking hosts strategically manage quality in response to evolving reputation incentives under the current review system.

Table 3: Demand Response to Ratings

	Dependent Variable: Occupancy Rate (d_{jt})			
	(1)	(2)	(3)	(4)
Average Rating (β)	0.068***	0.0308	0.078***	0.016
	(0.0135)	(0.0267)	(0.014)	(0.028)
Number of Ratings × High Average (γ_1)	0.001***	0.0003*	0.005**	0.005**
	(0.0001)	(0.0001)	(0.002)	(0.002)
Number of Ratings × Low Average (γ_2)	0.0002*	0.0002	-0.004*	-0.005*
	(0.0001)	(0.326)	(0.002)	(0.002)
Constant	0.602***	0.620***	0.608***	0.623***
	(0.006)	(0.007)	(0.006)	(0.007)
Listing FE	N	Y	N	Y
N	9,517	9,517	9,517	9,517
R^2	0.017	0.132	0.015	0.133

Note: The table reports demand response to ratings. One observation is one stay. R_{jt} is the average rating of listing j in period t. Columns 1 and 2 show results where τ is set at the sample average rating. Columns 3 and 4 report results where τ is set at 4. Standard errors are reported in parentheses. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

3 Model

We develop a model to incorporate the descriptive facts presented above into an equilibrium framework. The model rationalizes host effort choices and guest booking choices given the review system on the STR platform. We use the model to answer policy-relevant questions such as what the equilibrium effect of alternative rating designs will be in terms of consumer welfare, quality, and revenue.

3.1 Quality Evolvement and Host Effort

The model is a game between hosts and guests. Consider a continuum of hosts, each denoted by j, selling accommodations on the platform.[14] The modal host owns one short-term rental; thus, we use hosts, listings, and accommodations interchangeably. The framework analyses the incentives and choices of individual hosts exploring profit opportunities in a competitive marketplace under the current Airbnb review system. Each host is small and is a price taker. The model is in discrete time with infinite horizon $0 < t < \infty$.

In each period t, a host is characterized by discrete quality $q_t \in \{1, \ldots, Q\}$. While this quality level is privately observed by the host, it remains unobservable to guests. Each period, the host

[14]For ease of exposition, j subscript is omitted in this section.

observes the cost of extra maintenance effort c_t^E and the cost of basic turn-over service c_t^{NE}. The host then decides whether to exert extra maintenance effort and pay c_t^E ($e_t = 1$, hereafter referred to as simply *effort*), or provide a basic turn-over service and pay c_t^{NE} ($e_t = 0$, hereafter referred to as simply *no effort*). We set $c_t = (c_t^E, c_t^{NE})$.

The quality of a host's listing evolves stochastically according to the transition distribution $P^Q(q_{t+1}|q_t, e_t)$, where the future quality level depends on whether the host exerts extra maintenance effort. This distribution serves as a fundamental component of the model, governing both the effectiveness of maintenance efforts and the natural evolution of quality over time.

We allow for natural quality depreciation over time, such that $P^Q(q_{t+1} < q_t|e_t = 0)$ is generally positive. This reflects the reality that inadequate maintenance can lead to a gradual decline in quality. Importantly, our framework also permits exogenous quality improvements even in the absence of effort, meaning $P^Q(q_{t+1} > q_t|e_t = 0)$ can be positive. This feature captures external factors such as seasonality effects. For example, beachfront accommodations often experience an inherent increase in perceived quality during the summer months due to favorable weather conditions and higher demand. This distinction sets our model apart from prior literature, which typically infers deliberate maintenance effort based on observable indicators of higher quality, such as an increase in positive reviews. In contrast, our empirical setting enables a clear differentiation between seasonal fluctuations and intentional maintenance investment, as we observe extra maintenance effort directly.

This framework is sufficiently general to encompass a fixed-quality model, where quality remains constant over time, i.e., $P^Q(q_t = q_{t+1}) = 1$. Additionally, it allows for a setting in which quality evolves exogenously without moral hazard, meaning that maintenance effort does not influence quality transitions: $P^Q(q_{t+1}|q_t, e_t = 1) = P^Q(q_{t+1}|q_t, e_t = 0)$. In such cases, quality fluctuations arise solely from external factors, such as seasonality or broader market conditions, rather than from the host's strategic effort decisions.

3.2 Capacity

In addition to the quality level, in each period, the host is also characterized by the booking status $l_t \in \{E, 0, 1, \ldots, L\}$, which tracks the occupancy state of the property. Specifically:

- If $l_t = E$, the property is unoccupied and available for booking.
- If $l_t \neq E$, the property is occupied for the next l_t nights.

- If $l_t = 0$, a check-out occurs in the current period, meaning the property becomes available for booking today.

We distinguish between the states $l_t = E$ (unoccupied and available) and $l_t = 0$ (check-out today and available) because in the former case, no guest is present, and thus no review can be submitted, while in the latter case, a guest is checking out, so they may leave a review.

This formulation of capacity constraints simplifies the model by avoiding the need to track the full booking calendar, making the framework sharper and more elegant. In practice, hosts manage availability at a daily level, considering a potentially large booking window that begins from the current day. This reality introduces complexity into host decision-making, including strategies related to calendar and revenue management. However, our model abstracts from these aspects to maintain a focused analysis of quality evolution and review dynamics, as we consider sophisticated calendar management to be largely orthogonal to maintenance effort. For instance, hosts are likely to face similar incentives for maintenance effort provision regardless of whether their future calendar is fully booked or relatively empty. While hosts may have a slight incentive to shirk maintenance if their future calendar is already full, since negative reviews may have a more delayed impact on future bookings, this concern is mitigated by two key factors: (i) flexible cancellation policies, which are frequently observed in our data, allow guests to cancel bookings at no cost after observing a negative review. This reduces the host's ability to temporarily defer maintenance without consequences. And, (ii) long-term host incentives, as most hosts in our sample operate with a long-term perspective. Given this horizon, they are less likely to engage in short-term quality shirking, as the delay in negative review effects is unlikely to meaningfully alter their overall strategy.

3.3 Guest Demand

We consider a representative consumer i who arrives in the market each period. Consumers are short-lived, meaning they make a booking decision or exit after one period. Each consumer has an exogenously determined need for a stay for the duration of n_t nights, which is realized upon their arrival. The consumer's utility from booking the property is given by:

$$u_i = \theta_1^U + \theta_2^U \tilde{q}_t + \epsilon_{it}^B, \tag{1}$$

while the outside option utility (e.g., staying at a hotel or another accommodation) is: $u_i^0 = \epsilon_{it}^O$. The parameter θ_1^U represents the baseline utility of a stay at a reference quality level, net of

price. Consumers exhibit heterogeneity in their valuation of a stay, as captured by the idiosyncratic preference shock ϵ_{it}^B. Similarly, consumers differ in their valuation of the outside option, as reflected in ϵ_{it}^O, which accounts for their individual preferences for alternative accommodations, such as hotels or other short-term rentals. Hereafter we set $\epsilon_{it} = (\epsilon_{it}^B, \epsilon_{it}^O)$

The parameter θ_2^U represents the marginal utility of accommodation quality, capturing how an increase in quality influences a consumer's valuation of the stay. However, consumers do not directly observe the true quality of the accommodation. Instead, based on their information set H_t (specified later), they form beliefs, $\tilde{q}_t$, about the true quality level q_t, represented by the probability distribution, $G(\cdot|H_t)$, which defines their subjective expectations over the possible quality levels $\{1, \ldots, Q\}$. Given this setup, the consumer's decision-making process follows from their perceived distribution of quality rather than the true underlying state. Consumers then make a choice to book the property, $b_t = 1$, or choose the outside option, $b_t = 0$. Upon booking, the customer generates gross profit $\pi(n_t)$, reflecting nightly rates with all the extra fees, net of the marginal cost of the stay, which includes basic effort facilitating turnover.

After checking out, the guest may leave a rating r_t. For simplicity, we assume that guests submit ratings immediately after checkout, thereby avoiding the need to track a stock of pending ratings. Modeling such a stock would require accounting for both host and consumer beliefs about unobserved pending ratings. While this would add complexity to the model, it is likely to have only a second-order economic impact.

We further assume that guests observe the true quality level and its evolution during their stay. If a checkout occurs, the guest may leave a rating according to the conditional distribution $P^R(r_t|q_t)$, which captures how quality translates into ratings. This distribution includes the possibility of not leaving a review, represented by $r_t = 0$. Additionally, we assume that all submitted reviews are posted by the platform without modification. This assumption aligns with standard practices on Airbnb, where, based on our experience, it is exceedingly difficult to persuade the platform to alter the star rating, even if the textual content of the review violates the platform's terms of use.

In general, $P^R(r_t = q_t|q_t) \neq 1$, which introduces noise into the review process. This noise may arise from three primary sources:

1. Trembling-hand errors: These represent pure mistakes made by guests when submitting reviews, such as pressing the wrong button, misinterpreting the rating scale, or misunderstanding the review system.

2. Idiosyncratic experience factors: These include external, quality-unrelated shocks that affect a guest's stay but are not reflective of the true accommodation quality. Examples include bad weather, lockouts, traffic congestion, or difficulty finding the listing.

3. Horizontal preferences: Guests may have varying personal preferences that cause them to rate the same objective quality differently. For instance, a property in the city center might be highly rated by one guest who prefers transportation convenience but receive a lower rating from another who prefers a quiet neighbourhood.

The first two sources of noise, trembling-hand errors and idiosyncratic factors, are relatively straightforward to incorporate into the model, as they introduce exogenous randomness into the rating distribution P^R. These sources do not systematically bias the rating process but instead increase the variance in observed ratings around the true quality level. From the perspective of subsequent guests would simply lower the signal-to-noise ratio of reviews.

Empirically, horizontal preferences can also contribute to positive off-diagonal terms in P^R, reflecting disagreement among users regarding the overall quality of a listing. According to Equation (1), consumers are assumed to have full information about their own horizontal preferences, ϵ_{it}, and form inferences about the vertical quality measure q_t. As a result, we preclude segment-level learning, such as updates related to specific demographic factors (e.g., age, family-friendliness).[15] If horizontal preferences ϵ_{it} were unobserved, consumers might learn about it heterogeneously, depending on the identity of previous guests and the textual descriptions in their reviews. For example, a guest might adjust their expectations differently based on whether past reviewers were solo travelers, families, or business professionals. However, under our setup, all users who share the same initial belief about quality update their beliefs in the same manner upon observing a given rating. This assumption simplifies the learning dynamics, ensuring that observed ratings provide a common informational signal that influences consumer expectations about vertical quality in a consistent way within each segment across the market.

3.4 Information Structure and Strategies

The model is formulated as an incomplete information dynamic game between hosts and guests, where hosts act strategically and are forward-looking in their maintenance effort decisions. In contrast, guests make one-shot decisions upon arrival, without direct knowledge of the true quality

[15]We allow for heterogeneous learning across different segments, but preclude within-segment heterogeneity in consumer learning. Section 4.4 presents more details of how segments are classified and analyzed.

level. Instead, guests base their booking decisions on observed historical data, including past reservations and associated ratings. Hosts observe the same publicly available information as guests (i.e., past reviews and booking history) but additionally have private knowledge of the true current quality q_t before deciding on maintenance effort. Idiosyncratic preference shocks ϵ_{it} are realized at the moment the guest enters the market, and are observed by the guests but unobserved by the hosts.

In this paper, we restrict our analysis to Markov Perfect Equilibrium (MPE), where all players' strategies depend only on the smallest set of payoff-relevant state variables. Consequently, we do not define the entire general strategy space but instead focus on Markov strategies, which depend only on the current state. To formalize this equilibrium structure, we define the host's state at time t as: $s_t = (q_t, l_t, G_t)$, where: q_t is the true quality, l_t is a booking status, and G_t is the guest's belief about quality. As mentioned earlier, quality evolves as a controlled Markov process with transition kernel $P^Q(q_{t+1}|q_t, e_t)$. Conditional on s_t, the occupancy state is also a Markov process, such that

$$l_{t+1} = \begin{cases} l_t - 1 & \text{if } l_t \in \{1, \ldots, L\}, \\ E & \text{if } l_t \in \{E, 0\} \text{ and } b_t = 0, \\ n_t - 1 & \text{if } l_t \in \{E, 0\} \text{ and } b_t = 1, \end{cases}$$

where b_t indicates if a booking occurred, and n_t indicates the number of nights of the booking. The belief G_t can be expressed as a $Q \times 1$ dimensional continuous state, tracking probability distribution of the discrete quality levels on the set $\{1, \ldots, Q\}$. The belief also evolves as a Markov process (conditional on s_t), whose transition is explained later in this section.

Define a Markov strategy of the host as a function $\mathbf{e}(s_t, c_t) \in \{0, 1\}$ where c_t are the current costs of extra and basic efforts. Following Rust (1987), we assume that the cost of effort is conditionally independent of s_t, such that c_t is excluded from the state. Variation in c_t can help us to rationalize why observationally equivalent hosts make different maintenance decisions. Similarly, define a Markov strategy of the guest that is given by $\mathbf{b}(s_t \setminus q_t, \epsilon_t) \in \{0, 1\}$. The term $s_t \setminus q_t$ removes q_t from the state, signifying that guests use their beliefs G_t instead.

Fix the host and guest strategies $\mathbf{e}(\cdot)$ and $\mathbf{b}(\cdot)$, and state s_t. The game starts from the initial state s_0, when the listing enters the market and has no reservation history and no star rating. The term G_0 constitutes guests' prior beliefs about the new listing's quality. The timing of the game is as follows:

1. Three actions occur simultaneously:

   (a) If there is a check-out, i.e., if $l_t = 0$: previous guest leaves a review, r_t according to the current true quality q_t.

   (b) If there is a check-out, or if the property is empty, i.e., if $l_t \in \{E, 0\}$: Consumers arrive in the market and learn their horizontal preferences. Booking decision and reservation length are determined according to the current beliefs G_t and booking strategy $\mathbf{b}(\cdot)$

   (c) Hosts make effort decisions $e(s_t)$.

2. Reservation payout π_t and optimal effort costs are paid out to the host.

3. True quality updates according to the distribution $P^Q(q_{t+1}|q_t, e_t)$ and is revealed to the hosts.

4. Guests observe a review r_t, if any, and update their prior according to Bayes rule:

$$\bar{G}_t(\tilde{q}_t) = \frac{P^R(r_t|\tilde{q}_t)G_t(\tilde{q}_t)}{\sum_q P^R(r_t|q)G_t(q)},$$

   where $\bar{G}_t$ is the interim posterior. If there was no possibility to submit a review, i.e., if $l_t \neq 0$, then $\bar{G}_t = G_t$.

5. Guests update their interim posterior for the quality update process, consistently with the host's strategy:

$$G_{t+1}(\tilde{q}_{t+1}) = \sum_{\bar{q}_t} E_{c_t} P^Q\Big(\tilde{q}_{t+1}|\bar{q}_t, \mathbf{e}(s_t(\bar{q}_t), c_t)\Big)\bar{G}_t(\bar{q}_t),$$

   where $\bar{G}_t$ is the interim posterior, and $s_t(\bar{q}_t)$ is the current state of the market, with the true quality value replaced by a belief $\bar{q}_t$.

In Step 4, guests form interim posteriors about quality based on available reviews. If a review is submitted for every checkout, guests do not need to observe the booking status l_t, as the presence or absence of bookings does not affect the information available for updating beliefs. If review submission is stochastic but uncorrelated with true quality q_t, then the absence of a review does not carry additional informational value regardless of whether a check-out occurred. In this case, guests update their beliefs solely based on observed reviews, without drawing inferences from missing ones. If the reviews are missing selectively, for instance, in case q_t is particularly low, then the absence of the review carries information. In such a case, consumers may update according to $P^R(\cdot|r_t = 0)$,

provided they know whether a checkout occurred.[16]

In this model, guests do not infer quality from the mere fact that a property was booked. This is because all potential guests share the same information set at the time of booking, meaning the booking action itself does not convey new information about quality. This follows from the assumption that guests lack private information about q_t at the time of booking. This assumption would be violated if guests received and selected on private signals about vertical quality beyond past ratings. In such a scenario, herding behavior could emerge, where bookings themselves signal high quality to future guests. However, this is unlikely in our setting for several reasons. First, guests typically cannot inspect the listing before booking, as the management agency's policies explicitly prohibit pre-booking visits for security reasons. Second, repeat guests are extremely rare, further limiting the possibility of private information about vertical quality. Conversely, guests possess private information about their own horizontal preferences, which may influence their reviews; however, as we ruled out segment-level learning, this private information does not generate information revelation from mere booking decisions.

In Step 5, guests update their interim posterior to account for quality evolution, which occurs both organically and as a result of deliberate maintenance effort. We assume that guests have knowledge of the parameters governing quality evolution—specifically, they understand the likelihood of quality deterioration in the absence of maintenance and the effectiveness of maintenance effort in mitigating this decline. However, guests do not observe the host's effort choice directly and must therefore infer it based on available information. The belief formation process follows three steps:

(i) Beliefs about the host's maintenance strategy: Guests form expectations about the host's effort policy, denoted as $\mathbf{e}(\cdot)$, which maps the state to an effort decision.

(ii) Beliefs about effort distribution: Since guests do not observe the true quality q_t, while hosts do, guests must infer the distribution of effort choices using their prior belief over quality, G_t. The result is a random variable $\mathbf{e}(s_t(\tilde{q}_t))$, where $\tilde{q}_t$ is distributed according to G_t.

(iii) Posterior update: Guests update their interim beliefs, $\bar{G}_t$, about quality by integrating out

[16] Alternatively, if consumers are uncertain about check-outs, they may form beliefs over the likelihood of a check-out, denoted as $\tilde{l}_t$. While the case of uncertain check-outs is conceptually possible, it is not pursued in this study, as incorporating such uncertainty would significantly expand the belief state space, adding complexity to the model. Moreover, this extension would only affect outcomes if the rating submission were highly skewed. Even in that case, modeling uncertainty in check-outs yields only numerically negligible differences, as these beliefs collapse to certainty every time a guest observes a review.

the unobserved maintenance action, incorporating their inferred distribution of effort into the overall quality transition process.

Overall, for a given guest's and hosts' strategies, this structure defines a Markov process on (s_t, b_t, c_t, e_t). The structure also ensures that guests update their beliefs rationally, incorporating their knowledge of quality dynamics and host incentives, while accounting for the uncertainty surrounding individual maintenance decisions. However, this process introduces complex strategic incentives for hosts when deciding whether to exert effort.

Since effort depends on true quality q_t, hosts experiencing quality depreciation may choose to undertake maintenance to restore the property's condition. Importantly, because guests infer quality changes through reviews, such maintenance can be viewed as preemptive, aimed at maintaining a favorable reputation before negative reviews occur and accumulate. However, hosts may also delay maintenance, compared to a setting where guests directly observe quality. This is because market learning about quality is gradual, meaning that a decline in quality may not immediately translate into lower demand. As a result, hosts may time their maintenance decisions strategically, responding to information releases in the form of reviews, which shape the updated prior G_t. Thus, even if true quality remains unchanged, the host's effort decision will likely depend on G_t, reflecting shifts in consumer expectations.

Furthermore, since both guests and hosts have common knowledge of the likelihood of quality depreciation, there exists an implicit expectation regarding the minimum level of maintenance necessary to sustain a reasonable quality standard. Conversely, as both sides know about favorable quality shocks, a certain level of shirking is expected in high-demand periods. In equilibrium, defined in the next section, hosts will exert effort at a level consistent with these expectations, balancing the costs of maintenance with the reputation incentives driven by consumer learning.

3.5 Markov Perfect Equilibrium

The value function of the host is defined for each feasible state and the current costs of effort, (s_t, c_t), given the Markov strategies, $\mathbf{b}(\cdot)$ and $\mathbf{e}(\cdot)$, adopted by guests and hosts, respectively. Formally, the value function is expressed as:

$$V(s_t, c_t; \mathbf{e}, \mathbf{b}) = E\left[\sum_{s=t}^{\infty} \beta^{s-t}\pi(n_t)\mathbf{1}(b_t = 1) - c_t^E\mathbf{1}(e_t = 1) - c_t^{NE}\mathbf{1}(e_t = 0)\Big| s_t\right], \tag{2}$$

where the expectation is taken with respect to the unique Markov process over (s_t, b_t, c_t, e_t) induced by the strategies.

Before we define an equilibrium, we define a few more objects that are useful when formalizing an equilibrium. To match the standard discrete choice formulation, introduce some normalizations. The cost of extra maintenance effort can be written as $c_t^E = c^E + \sigma^C \epsilon_t^E$, whereas the cost of basic maintenance, without extra effort, is $c^{NE} + \sigma^C \epsilon_t^{NE}$. In this formulation, c^E represents the baseline monetary cost of additional effort, while c^{NE} denotes the baseline cost of basic maintenance. The errors are orthogonal shocks, representing time-varying innovations to the cost of effort. In practice, only the cost difference $\Delta^C = c^E - c^{NE}$ matters. This is because our model does not allow for exit, and we assume that the company always exerts at least the basic level of effort. Under this assumption, the model yields identical outcomes if we set $\Delta^C = c^E$ and normalize $c^{NE} = 0$. This normalization does not affect other model estimates or counterfactual analyses, as long as both c^E and c^{NE} are not simultaneously varied.

Under the above specification, the ex-ante value function is given by

$$V(s_t; \mathbf{e}, \mathbf{b}) = E_{c_t} V(s_t, c_t; \mathbf{e}, \mathbf{b}),$$

while the choice-specific value function is defined as

$$\bar{V}(s_t, e_t; \mathbf{e}, \mathbf{b}) = E[\pi(n_t)\mathbf{1}(b_t = 1)] - \Delta^C \mathbf{1}(e_t = 1) + \beta E[V(s_{t+1})|s_t]$$

which represents one-shot deviation from $\mathbf{e}$ and following $\mathbf{e}$ thereafter. For notational convenience, the choice-specific value function does not include the current shock to effort costs.

Definition 1 (Markov Perfect Equilibrium) *We define Markov Perfect Equilibrium (MPE) as a pair of Markov strategies* $(\mathbf{b}(\cdot), \mathbf{e}(\cdot))$ *that satisfy the following conditions:*

- *Host Optimization: Fixing* $\mathbf{b}(\cdot)$*, for each* (s_t, c_t)*, the strategy* $\mathbf{e}(\cdot)$ *maximizes host's value function.*
- *Guest Optimization: Fixing* $\mathbf{e}(\cdot)$*, for each* $(s_t \backslash q_t, \epsilon_t)$*, the strategy* $\mathbf{b}(\cdot)$ *maximizes guest's utility.*
- *Consistency of Beliefs: The beliefs of hosts and guests about actions and the evolution of the state space are consistent with the equilibrium strategy profile.*

The MPE defined above constitutes a Nash equilibrium because neither hosts nor guests can

improve their payoffs by deviating to a non-Markovian strategy. For hosts, this result follows from the standard arguments based on the optimality principle in dynamic programming. For guests, the result holds because their belief G_t serves as a sufficient statistic for their booking decision, meaning no additional information from the game's history influences their choice beyond what is captured in the belief state.

Under standard conditions (such as bounded profits), MPE-implied value functions must satisfy a Bellman equation

$$V(s_t) = E \max\{\bar{V}(s_t, e_t = 0) + \epsilon_t^{NE}, \bar{V}(s_t, e_t = 1) + \epsilon_t^{E}\}.$$

We remove the dependence on strategies to simplify the notation. In such cases, we presume equilibrium strategies. In the next section, we present the details of the estimation of the model.

4 Estimation

In this section, we present the empirical parametrization of the model to help with identification with a finite sample. We then detail our GMM estimator and discuss the identification.

4.1 Parametrization

We start with parametrization of the utility function u_{it} given by equation (1). We assume that horizontal preference shocks ϵ_{it}^{B} and ϵ_{it}^{O} follow the type one value distribution. In such a case, given a prior belief on quality, G_t, the likelihood of receiving the booking equals

$$P(b_t = 1|G_t) = \frac{\exp(\theta_1^U + \theta_2^U E[q_t|G_t])}{1 + \exp(\theta_1^U + \theta_2^U E[q_t|G_t])}.$$

The preference for the number of nights is assumed to follow a Poisson process with parameter λ. In addition to idiosyncratic shocks, we allow for persistent heterogeneity in both the cost of effort and gross profits across hosts, indexed by j, to capture variation in local market conditions and property characteristics. For example, hosts operating in high-cost areas such as California or New York City typically generate higher rental revenue due to elevated demand and price levels, but they also face substantially higher maintenance costs. Conversely, hosts located in less expensive regions experience lower revenues and reduced service costs. Beyond geographic variation, we also account for differences in property size and type. Our sample includes a diverse range of listings, from small

studios to five-bedroom homes. Larger properties generally command higher nightly rates, leading to greater gross profits, but they also incur higher turnover and maintenance expenses. In contrast, studios tend to yield lower revenues and are less costly to maintain.

The management agency's pricing strategy involves passing through standard turnover costs (e.g., cleaning and restocking) directly to guests as *cleaning fee*, allowing to model gross profits using nightly rates. This feature enables us to capture cross-host heterogeneity in a tractable manner by adopting a linear functional form for gross profits, given by: $\pi_j(n_t) = \theta_j^\pi n_t$, where θ_j^π denotes the host-specific average nightly profit rate, and n_t is the number of booked nights in period t. We then model the persistent cost difference associated with exerting extra maintenance effort as a fixed share of gross profits: $\Delta_j^C = \theta^C \theta_j^\pi$, where θ^C is a parameter that maps nightly earnings into effort cost. This formulation ensures that both the incentives to provide effort and the costs of doing so are endogenously tied to the economic environment in which the host operates. Importantly, this structure allows us to accommodate persistent, host-level heterogeneity in a reduced-form yet economically meaningful way. As a result, θ_j^π drops out of the value function (2), and the model does not need to be solved separately for each host during the estimation.

We parametrize the quality transition kernel using a classical *quality ladder specification*, originally introduced by Ericson and Pakes (1995) and adapted here to discrete effort rather than continuous investment. Specifically, we model the evolution of quality as:

$$q_{t+1} = q_t + \iota^E \mathbf{1}(e_t = 1) + \iota^O, \tag{3}$$

where: ι^E represents the innovation to quality resulting from costly effort, and ι^O represents organic quality innovation that occurs independently of effort. For numerical tractability, q_{t+1} is truncated to remain within the set $(1, 2, \ldots, Q)$.

The term ι^E is modeled as a Bernoulli random variable, taking the value:

$$\iota^E = \begin{cases} 1 & \text{with probability } p^E, \\ 0 & \text{with probability } 1 - p^E. \end{cases}$$

This specification allows for the possibility that effort is unsuccessful with probability $1 - p^E$. If $p^E = 0$, then no quality improvements occur due to effort, effectively eliminating moral hazard from

the model. The term ι^O takes three values with the following probabilities:

$$\iota^O = \begin{cases} -1 & \text{with probability } \delta, \\ 0 & \text{with probability } 1-\delta-\alpha, \\ 1 & \text{with probability } \alpha. \end{cases}$$

If $\delta = 0$, quality never depreciates, whereas if $\alpha = 0$, it cannot improve without costly effort. When p^E, δ, and α are all zero, quality is fixed at a constant level q^*, encompassing a case of fixed quality previously studied in the learning and feedback literature. In such a case, guests will eventually learn the true quality through reviews. However, when quality evolves dynamically, whether due to moral hazard or external factors, consumer learning becomes an ongoing process. In such cases, new quality innovations create persistent uncertainty, requiring guests to continuously update their beliefs.

To simplify the computation of the optimal effort, we set error terms ϵ_t^E and ϵ_t^{NE}, that capture time-varying fluctuations in costs, to follow Type-I extreme value distributions, with a scaling parameter σ. In such a case, the Bellman equation can be written as

$$V(s_t) = \sigma \log[\exp(\sigma^{-1}\bar{V}(s_t, e_t = 0)) + \exp(\sigma^{-1}\bar{V}(s_t, e_t = 1))] + \text{CONSTANT}$$

We disregard the constant in the remainder of the paper, as it does not affect estimation or counterfactuals. Further, Conditional Choice Probability (CCP) for optimal effort may be obtained as

$$P(e_t = 1|s_t) = \frac{\exp(\sigma^{-1}\bar{V}(s_t, e_t = 1))}{\exp(\sigma^{-1}\bar{V}(s_t, e_t = 0)) + \exp(\sigma^{-1}\bar{V}(s_t, e_t = 1))}.$$

Next, we parametrize the review quality and review process. We assume that review can take the same values as quality, from 1 to Q stars. We set $Q = 5$ to match our empirical setting. A review equal to quality is an unbiased review, whereas a review different from quality is biased. As a first approximation, we consider a process where some reviews are omitted exogenously, where the rating probability is $P^R(r_t > 0|q_t) = p^R$. Also, we allow for noise in the review process given by $P^R(r_t \neq q_t|q_t) = \eta$. If $\eta = 0$, the reviews are a fully informative signal, and one review fully reveals the current quality. This does not mean that the learning process is over, or that there is no moral hazard, as quality may further change organically or as a result of maintenance. Notably, even in the case of $\eta = 0$, there is a delay in quality revelation, since reviews typically would not

arrive immediately after innovation to quality.

Lastly, we parametrize the initial distribution of quality for new properties, denoted by q_0. We assume that q_0 is drawn from a probability mass function defined as

$$P_0(q_0; \mu^Q, \sigma^Q) = \frac{\phi(q_0; \mu^Q, \sigma^Q)}{\sum_{\tilde{q}_0 \in \{1,\ldots,5\}} \phi(\tilde{q}_0; \mu^Q, \sigma^Q)},$$

where $\phi(\tilde{q}_0; \mu^Q, \sigma^Q)$ is the density of a normal distribution with mean μ^Q and standard deviation σ^Q evaluated at $\tilde{q}_0$. This specification allows for variation in both the central tendency and dispersion of the initial quality distribution. We further impose rational expectations by setting the prior G_0 equal to the actual initial distribution of initial quality.

4.2 Method of Moments

We aim to estimate the following vector of structural parameters:

$$(\mu^Q, \sigma^Q, \theta_1^U, \theta_2^U, \lambda, \theta^C, \sigma^C, p^E, \delta, \alpha, p^R, \eta).$$

Some of these parameters can be estimated directly from the data prior to solving the structural model. In particular, the demand for the number of nights, λ, and the rating probability, p^R, can be inferred from the raw data. Note that, θ^C and σ^C, the mean and dispersion of the cost of exerting effort, cannot be directly identified from observed maintenance cost data. This is due to a selection issue, as we only observe costs conditional on effort being exerted. That is, the data reflect the conditional expectation $E[\theta^C + \sigma^C \epsilon_t^E \mid e_t = 1]$. As a result, both θ^C and the dispersion parameter σ^C must be estimated jointly with the remaining structural parameters within the model.

We implement a Simulated Generalized Method of Moments (Pakes and Pollard, 1989) with the full solution nested fixed-point approach. Define τ_{jk} as stopping times at which the k-th rating for the property j is submitted, defined recursively as $\tau_{jk} = \min\{t > \tau_{jk-1} : r_{jt} > 0\}$ and $\tau_{j0} = 0$. We match the following moments:

(M1) The likelihood of an increase and decrease in ratings, conditional on placing/not placing at least one maintenance request in between:

$$E\left[\mathbf{1}(r_{\tau_{k+1}} > r_{\tau_k}) \Big| \max_{\tau_k \le t < \tau_{k+1}} e_{jt} = 1\right] \qquad E\left[\mathbf{1}(r_{\tau_{k+1}} < r_{\tau_k}) \Big| \max_{\tau_k \le t < \tau_{k+1}} e_{jt} = 1\right],$$

$$E\left[\mathbf{1}(r_{\tau_{k+1}} > r_{\tau_k})\Big|\max_{\tau_k \le t < \tau_{k+1}} e_{jt} = 0\right] \qquad E\left[\mathbf{1}(r_{\tau_{k+1}} < r_{\tau_k})\Big|\max_{\tau_k \le t < \tau_{k+1}} e_{jt} = 0\right].$$

(M2) Distribution of the first 3-ratings: $E[r_{\tau_k}|k \in \{1,2,3\}]$.

(M3) Indirect inference matching regression 4 coefficients of Model (3) in Table 3: average impact of current average rating, and the interaction of good and bad cumulative rating with the number of ratings, on the occupancy rate.

(M4) Probability of submitting at least one maintenance request within 30 days, given the last rating:

$$E\left[\max_{t \le k \le t+30} \mathbf{1}(e_k = 1)|r_t = r\right], \text{ for } r \in \{1,2,3,4,5\}.$$

In the estimation, we set the discount factor to $\beta = 0.95$, the number of quality levels $Q = 5$, and the maximum reservation length to $L = 30$.

4.3 Identification

A common challenge in the identification of moral hazard versus adverse selection arises because effort is not directly observed. In such settings, the insurance literature has developed panel data estimators that enable the identification of moral hazard under asymmetric information. These approaches typically rely on structural assumptions to distinguish true state dependence from unobserved heterogeneity. For example, it is often assumed that asymmetric information is limited at the start of the contractual relationship, or that the transition process (see Abbring et al., 2003) for unobserved quality is sufficiently restricted (see Jeziorski et al., 2017). However, these frameworks are less well-suited to dynamic product quality settings. In our context, asymmetric information is likely to be most pronounced at the beginning of the relationship, when guests have no prior experience with the seller. Additionally, seasonality in demand may induce complex and non-stationary patterns in organic quality transitions, making it difficult to disentangle unobserved heterogeneity from endogenous dynamics.

To solve this problem, the literature on consumer reviews has often relied on firm responses, such as replies to reviews, as a proxy for effort. We extend this approach by leveraging direct data on actual effort, allowing for more precise identification of the relationship between effort and quality dynamics. The central challenge in our identification strategy lies in disentangling the marginal utility of quality from the signal-to-noise ratio of reviews. This distinction is difficult to infer using

demand data alone, even if effort is observed. For example, if reviews are uninformative, we would expect to observe only a weak demand response to changes in ratings. However, a similarly muted demand response could also arise if the marginal utility of quality is low, even when reviews are highly informative. In the extreme case, the absence of a demand reaction to review updates may reflect either that consumers place little value on quality or that they fail to learn about quality from the reviews themselves.

The intuition of our approach is that we use a combination of demand and supply moments for identification. Suppose for now that the econometrician knows the informativeness of the review signal, in our case η. In such case, the first group of moments, denoted (M1), directly identifies the productivity of effort, p^E, as well as the parameters governing, organic quality transitions, α and δ, as one can invert the review transitions into latent quality transitions using Bayes rule for a given value of η.[17]

Similarly, the second group of moments (M2) identifies the prior distribution of quality, as the distribution of observed reviews, combined with the signal precision parameter η, jointly reveals the mean and dispersion of initial quality, denoted (μ^Q, σ^Q). The third group of moments (M3), capturing the relationship between average reviews and demand, identifies the marginal utility of quality. Given the prior, the quality transition process, and the signal-to-noise ratio η, the econometrician can compute the expected average quality for each listing conditional of the observed ratings, enabling identification of how changes in quality affect booking behavior.

The remaining parameters to be identified are the signal precision, η, and the cost parameters (θ^C, σ^C) governing the supply side. The interaction between the number of reviews and average review score is particularly informative about the speed of learning, as it captures the relative value of an additional review and the informativeness of the existing review distribution. This is best demonstrated by the negative coefficient on the number of ratings in Table 3, if the current average of ratings is inferior. Our estimation chooses η that rationalizes this coefficient. [18]

In addition, conditional on the marginal utility of quality, the supply-side moments (M4) are

[17]Depending on data availability, this identification strategy can be extended to non-parametrically identify non-Markovian seasonality in quality. In such cases, moment (M1) could be augmented to include conditioning on a longer history of past reviews, such as $r_{\tau_{k-s}}$.

[18]This effect is most transparent under the assumption of risk neutrality. In that case, if consumers do not learn from reviews, or if learning is extremely fast, the current number of reviews should not affect demand. Conversely, if consumers were risk-averse, they would benefit from additional reviews because they reduce uncertainty. Under risk aversion, holding the average review constant, an increase in the number of reviews should raise demand by decreasing posterior variance. Our framework could accommodate the identification of risk preferences, as we can examine how the impact of an additional review differs when the average review score is low versus high.

also informative about the signal precision parameter η. Specifically, our estimation procedure uses variation in hosts' responses to negative reviews to infer η, imposing a previously identified relationship between η and the marginal utility of quality from demand-side moments. Intuitively, fixing guests' value for quality, making reviews more informative (low η), makes hosts more responsive to unfavorable review signals. Conversely, if reviews are less informative (higher η), hosts will have weaker incentives to react to rating changes. Our framework allows η to vary across the rating distribution, acknowledging that lower ratings may be more informative than higher ratings. However, the estimated functional form of this relationship depends on the quantity and richness of the data available.

Finally, the cost of effort is identified from the frequency and intensity of maintenance activities. The intuition here is similar to regression discontinuity presented earlier. In general, the challenge is that multiple combinations of (θ^C, σ^C) can generate similar aggregate levels of effort. For example, a high mean cost paired with high variance may produce the same average effort level as a lower mean with lower variance. Yet, such parameter combinations carry different strategic implications. In particular, high variance in costs would imply that effort decisions are less strategic, more random, and as a result, less dependent of the current review stock. In the extreme, where effort is entirely random, we would expect it to be uncorrelated with recent reviews. The supply-side moments (M4), therefore, play a crucial role in separately identifying θ^C and σ^C by exploiting the relationship between recent reviews and observed maintenance behavior. Specifically, if effort is responsive to the latest review, this indicates that idiosyncratic cost realizations are not a dominating factor when timing effort, allowing us to pin down θ^C and σ^C separately. [19]

4.4 Listing Heterogeneity

Listings generally differ substantially in their supply and demand conditions. For example, a large beach house may follow a different quality evolution process, face different maintenance costs, and attract different customers than a small house near a train station. To account for such heterogeneity, we use the k-means method to classify listings into latent segments based on observed property characteristics, host behavior, and business performance. Specifically, we use the number of bedrooms, the number of bathrooms, average nightly rate, location score, average rating, average

[19] Section B.2 presents simulation tests on the identification of structural parameters based on simulated data of similar size.

Table 4: Listing Segments

Listing Features	Segment 1 (73%)	Segment 2 (27%)
Number of Bedrooms	1.13	3.24
Number of Bathrooms	1.09	2.49
Location	4.74	4.87
Rating	4.69	4.84
Nightly Rate($)	95.82	197.09
Maintenance Frequency	2.2%	3.8%
Nights per Booking	5.46	3.45
Occupancy Rate	64.9%	57.0%

Note: The table reports the key heterogeneous features of the two segments of listings. We use the k-means method to cluster listings into latent segments based on observed features of the listings, including the number of bedrooms, the number of bathrooms, average nightly rates, location score, average rating, average maintenance frequency, and average booking rate.

maintenance frequency, and average booking rate. [20] We then construct the estimation moments separately for each segment, and allow distinct structural parameters across segments.

Table 4 reports the resulting two listing segments. Segment 1 accounts for 73% of listings. These listings, on average, have one bedroom and one bathroom, booked 5.46 nights per stay, receive moderate ratings, are in moderate locations, are maintained less often, have lower nightly rates, and have more bookings. We therefore refer to this segment as the value segment. Segment 2 accounts for 27% of listings. These listings, on average, have three bedrooms and two bathrooms, booked 3.45 nights per stay, receive high ratings, are in better locations, are maintained more often, have higher nightly rates, and have fewer bookings. We henceforth refer to this segment as the premium segment. Table 5 presents the regression discontinuity results separately by segment. In both segments, hosts increase maintenance effort following stays with bad ratings. Additionally, the magnitude of the response differs across segments: the premium segment exhibits a stronger maintenance response after a bad review. These results suggest that hosts' reputation incentives, maintenance costs, and effort effectiveness vary across listing segments, and indicate different responses across segments in the counterfactual analysis with alternative rating designs.

[20] Similar segmentation practice can be found in the salesforce management literature, see Kim et al. (2022) and Yang et al. (2024) for recent examples.

Table 5: Timing of Effort in Different Segments

	Dependent Variable: Maintenance (m_{jt})			
	Segment 1		Segment 2	
	(1) Good Review	(2) Bad Review	(3) Good Review	(4) Bad Review
Checkout (β)	-0.000365	0.00102***	-0.000976	0.00335**
	(0.000262)	(0.000347)	(0.000829)	(0.00151)
Observations	274,439	75,024	73,256	11,887

Note: The table reports results from regression discontinuity with triangular kernel results for a 30-day time window before and after the checkout date of a stay. Columns (1) and (2) report the results for stays in segment 1 listings, with Column (1) for stays with five-star ratings and Column (2) for stays with ratings below five stars. Columns (3) and (4) report the results for stays in segment 2, with Column (3) for stays with five-star ratings and Column (4) for stays with ratings below five stars. All standard errors are clustered at the reservation level and reported in parentheses. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

5 Results

Table 6 reports the structural estimates for the two listing segments.[21] The initial quality estimates distinguish the two segments. The premium segment starts from a higher and more concentrated quality distribution, while the value segment contains a wider range of listing quality. The estimated quality transition parameters show that listing quality evolves dynamically in the absence of maintenance. In the value segment, the daily probability of organic quality appreciation is 0.012, while the probability of organic depreciation is 0.015. In the premium segment, the corresponding probabilities are 0.008 and 0.013. Thus, absent maintenance, quality is slightly more likely to deteriorate than improve in both segments. Meanwhile, strategic maintenance is effective; i.e., it offsets 260 days of organic quality depreciation in the value segment and 160 days of organic quality depreciation in the premium segment. These estimates indicate that observed maintenance activities are quantitatively important inputs into the evolution of listing quality.

The cost estimates suggest that hosts face meaningful but not prohibitive maintenance costs. The estimated ratio of the maintenance cost relative to nightly profits is 6.179 in Segment 1 and 5.295 in Segment 2. Thus, maintenance is relatively more cost-efficient for the premium listings than for the value listings. These estimates help explain why maintenance behavior differs across segments: premium hosts face higher maintenance cost efficiency from maintenance effort, making them more responsive to reputation incentives.

On the demand side, the value segment has a higher baseline demand level (after accounting

[21] Section B.3 presents results on the out-of-sample fit of our model. We show that the observed maintenance effort can be rationalized by the tractable model of endogenous quality provision.

Table 6: Estimation Results

	Segment 1		Segment 2	
	Estimate	Std.	Estimate	Std.
Quality Evolvement and Effort Effectiveness				
Quality Appreciation Prob. (α)	0.012	0.005	0.008	0.003
Default Depreciation Prob. (δ)	0.015	0.006	0.013	0.004
Effort Success Prob. (p^E)	0.781	0.015	0.814	0.017
Consumer Demand				
Constant (θ_1^U)	-3.101	0.191	-3.689	0.215
Quality (θ_2^U)	0.578	0.013	0.461	0.041
Preference for Nights (λ)	5.461	0.010	3.448	0.038
Review System				
Review Informativeness (η)	0.025	0.009	0.014	0.005
Prob. of Submitting a Rating (P^R)	0.769	0.024	0.826	0.029
Initial Quality				
Central Tendency (μ^Q)	4.404	0.017	4.651	0.020
Dispersion (σ^Q)	1.129	0.059	0.840	0.064
Cost of Effort				
Central Tendency (θ^C)	6.179	0.037	5.295	0.042
Dispersion (σ^C)	1.102	0.067	0.987	0.065

Note: The table reports the estimated parameters of the model for the two segments of listings. Standard errors are estimated through bootstrapping.

for the price effect). These patterns are consistent with the descriptive segment differences: value listings have lower prices and higher occupancy, while premium listings serve a smaller but higher-quality segment of the market. The estimates also show that ratings are informative but noisy signals of underlying quality. The estimated rating noise parameter is 0.025 in Segment 1 and 0.014 in Segment 2, implying that ratings are noisier in the value segment and consumers learn more precisely from ratings in the premium segment.

The estimates imply that the two segments operate under different business conditions. Segment 1, the value segment, has lower initial quality, noisier ratings, lower relative maintenance cost efficiency, but stronger baseline demand. These differences imply that rating designs can generate different equilibrium responses across segments. We examine these heterogeneous responses in the counterfactual analysis in the next section.

6 Counterfactual Analysis

This section illustrates how our model can be used for counterfactual analysis. We use the estimated model to evaluate how alternative rating system designs affect equilibrium quality, demand, and welfare. We focus on two counterfactual policies: rating manipulation and rating forgiveness. Rating manipulation is the strategic practice that biases ratings upwards regardless of true quality. Rating forgiveness reduces negative noise in ratings by eliminating some observable negative noise, such as bad weather and traffic conditions. Comparing the two policies allows us to quantify how different rating policies affect both consumer learning and seller incentives.

6.1 Rating Manipulation

We study rating manipulation in the form of fake reviews (Mayzlin et al., 2014; He et al., 2022). Other forms of rating manipulation, such as reciprocal practices like reviews for rewards (Cabral and Li, 2015), have a similar effect to the rating system. Specifically, we define a fake review as one that assigns a five-star rating with probability ψ, irrespective of the listing's underlying quality. Formally, we modify the original review arrival process P^R to a counterfactual process $\tilde{P}^R$, under which each incoming review is five stars with probability ψ, regardless of the host's true quality level.

Table 7 presents the results of counterfactual simulations under the assumption that 20% of submitted reviews are fake(ψ =20%), defined as always assigning a five-star rating regardless of true quality. We first consider a setting in which hosts maintain their original maintenance strategies from the baseline (i.e., no fake reviews), and consumers are unaware of the presence of fake reviews. That is, consumers continue to update their beliefs using the original review process. In this case, the booking rate increases by 2.1%, as properties appear better than they truly are. However, this perception bias leads to a 4% standard deviation decrease in the quality of booked properties, even though average quality remains essentially unchanged. This case captures the direct misinformation effect of fake reviews: consumers book more, but they are less able to distinguish high-quality from low-quality listings.

When we maintain the assumption that hosts do not adjust their maintenance strategy but allows consumers to be sophisticated, meaning that they correctly incorporate the probability of fake reviews into their posterior beliefs. In this case, consumers are able to partially mitigate the misinformation effects of fake reviews. However, both the booking rate and the quality of booked

properties decline, reflecting a reduction in the signal-to-noise ratio of the review system. As reviews become less informative, consumers make more mistakes in assessing listing quality, resulting in lower match efficiency.

In the setting where consumers remain unaware, but hosts are allowed to reoptimize their effort decisions in response to the changed information environment, we observe a significant reduction in maintenance effort, leading to an 7% standard deviation decline in the overall quality of listings. Interestingly, the booking rate rises slightly, which is due to two opposing forces: (1) unsophisticated consumers become less selective, mistakenly interpreting fake reviews as truthful (as in the first case), which increases booking rates; and (2) the endogenous decline in equilibrium quality reduces demand. These two effects approximately offset each other. This result is intuitive, as hosts strategically balance the gains from misleading consumers through fake reviews with the losses from declining product quality, choosing a level of shirking that equates these marginal effects.

Finally, in the case with sophisticated consumers and host reoptimization, hosts still reduce their maintenance effort, although to a lesser extent, knowing that consumers correctly infer the likelihood of fake reviews. Nevertheless, shirking remains substantial. Even with accurate updating, fake reviews reduce the informativeness of the rating system, thereby weakening the learning process. While consumers mitigate part of the welfare loss through better belief formation, the decline in equilibrium quality still passes through to booked quality, ultimately lowering consumer surplus. The segment-level results show that the value segment is more vulnerable to rating manipulation, with a higher loss in booking rate, average quality, and consumer surplus.

Taken together, these results highlight the quality substitution effect of fake reviews. When consumers are unaware of the manipulation, fake reviews distort demand and degrade allocation efficiency, even in the absence of endogenous host response. Sophisticated consumers can avert most of this loss through proper inference. However, when hosts endogenously adjust quality provision in response to the information environment, even sophisticated consumers are unable to fully prevent welfare losses, as they have limited influence over equilibrium quality. This extends the results from Gandhi et al. (2024) where quality and effort provision are assumed to be fixed.

6.2 Rating Forgiveness

We next study the effect of rating forgiveness, which platforms often offer in the form of a rating removal policy. For example, Instacart, a delivery platform, allows shoppers to apply for removal of

Table 7: Counterfactual Analysis on Rating Manipulation

20% Chance of Fake Review				
		Booking Rate Revenue	Average Quality	Average Booked Quality Consumer Surplus
Unaware Guests Hosts do not Re-optimize Effort	Segment 1	+2.3%	4.72 (+0%)	4.70 (-5% std.)
	Segment 2	+1.5%	4.81 (+0%)	4.82 (-3% std.)
	All	+2.1%	4.74 (+0%)	4.73 (-4% std.)
Sophisticated Guests Hosts do not Re-optimize Effort	Segment 1	-0.1%	4.72 (+0%)	4.74 (-1% std.)
	Segment 2	-0.1%	4.81 (+0%)	4.84 (-1% std.)
	All	-0.1%	4.74 (+0%)	4.77 (-1% std.)
Unaware Guests Hosts Re-optimize Effort	Segment 1	+0.2%	4.64 (-8% std.)	4.66 (-9% std.)
	Segment 2	+0.1%	4.75 (-6% std.)	4.79 (-6% std.)
	All	+0.1%	4.67 (-7% std.)	4.70 (-8% std.)
Sophisticated Guests Hosts Re-optimize Effort	Segment 1	-0.8%	4.65 (-7% std.)	4.68 (-7% std.)
	Segment 2	-0.7%	4.76 (-5% std.)	4.81 (-4% std.)
	All	-0.8%	4.68 (-6% std.)	4.72 (-6% std.)

Note: The table reports the counterfactual analysis on rating manipulation with 20% rate of fake reviews, including overall level and segment level results.

their previous low ratings. [22] Airbnb also deals with complaints from hosts in the community who get accidental low ratings and allows for rating removal. [23] We study rating forgiveness as one that partially reduces negative noise in the rating system. In reality, this corresponds to practices like Instacart, conditional removal of bad ratings based on the observed factors such as harshness of the consumer, the weather, and traffic conditions, etc. Formally, we modify the original review arrival process P^R to a counterfactual process $\hat{P}^R$, under which the negative noise ($r < q$) is removed with probability ϕ.

Similar to the previous section, we examine the four scenarios of guest knowledge and host reoptimization. Table 8 presents the results of counterfactual simulations under the assumption that 30% of negative noise is removed(ϕ=30%). When consumers are unaware and hosts do not reoptimize, the booking rate increases by 0.8%, as properties appear better due to the removal of bad ratings. What's more, this perception bias leads to a 1% standard deviation increase in the average quality of all properties and booked properties, because the high-quality properties suffer from negative noise more, and the removal of negative noise benefits high-quality properties disproportionately.

[22]Source: https://www.instacart.com/company/shopper-community/changing-how-ratings-work/.

[23]Source: Airbnb rating removal policy: https://www.airbnb.com/help/article/13. The policy is much stricter than Instacart's version. Moreover, hosts can't easily relist and reset their rating record because Airbnb has a detection and alarm system for duplicate listings.

When consumers are sophisticated and hosts do not reoptimize effort, the booking rate drops relative to the previous case due to the correction of misbelief, while the quality of booked properties increases further, reflecting an increase in the signal-to-noise ratio of the review system. As reviews become more informative, consumers make fewer mistakes in assessing listing quality, resulting in higher match efficiency.

When consumers remain unsophisticated, but hosts reoptimize their effort decisions, we observe an increase in maintenance effort, leading to a 3% standard deviation increase in the overall quality. The booking rate rises by 1.2%, which results from two reasons: unsophisticated consumers overestimate the proportion of negative noise and over-optimistically update their belief upon seeing a positive signal; and the endogenous increase in equilibrium quality increases demand.

With sophisticated consumers and hosts' re-optimization, hosts further increase their maintenance effort, knowing that consumers correctly infer the likelihood of removing bad noise. With accurate updating, reducing negative noise increases the informativeness of the rating system, especially for the top quality listings, which suffer from negative noise more. Although relative to the previous case, the booking rates decrease due to consumers' correction of misbelief, the average quality and booked quality increase further by 4% standard deviation. The segment-level results show that the gains from rating forgiveness are larger in the premium segment.

In summary, these results reveal the positive consequences of rating forgiveness. When consumers are unaware of the manipulation, reducing negative noise from reviews increases demand and improves allocation efficiency, even in the absence of endogenous host response. Sophisticated consumers can book less but better through proper inference. Additionally, when hosts endogenously adjust quality provision in response to the information environment, the positive effect on booking rates and average quality in equilibrium is reinforced. [24]

7 Conclusion

This paper studies the design of rating systems in online marketplaces where product quality is unobserved and dynamically managed by sellers. Using unique data from a STR management agency on Airbnb, we show that ratings serve the dual role of informing consumers about unobserved quality

[24] Section B.4 presents more counterfactual results on other rating systems that are of theoretical values, including the full transparency and pure noise benchmark, and reducing noise asymmetrically at different levels of quality, etc. Note that the current rational belief updating model has already captured the learning property that more recent reviews matter more in the posterior belief. Therefore, we focus on investigating different noise structures and do not consider reweighting of ratings with different vintages as in the career concern literature(Hörner and Lambert, 2021).

Table 8: Counterfactual Analysis on Rating Forgiveness

Removing 30% of Negative Noise				
		Booking Rate Revenue	Average quality	Average Booked Quality Consumer Surplus
Unaware Guests Hosts do not Re-optimize Effort	Segment 1	+0.7%	4.73 (+1% std.)	4.76 (+1% std.)
	Segment 2	+0.9%	4.82 (+1% std.)	4.86 (+1% std.)
	All	+0.8%	4.75 (+1% std.)	4.79 (+1% std.)
Sophisticated Guests Hosts do not Re-optimize Effort	Segment 1	+0.1%	4.73 (+1% std.)	4.77 (+2% std.)
	Segment 2	+0.2%	4.82 (+1% std.)	4.88 (+3% std.)
	All	+0.1%	4.75 (+1% std.)	4.80 (+2% std.)
Unaware Guests Hosts Re-optimize Effort	Segment 1	+1.1%	4.74 (+2% std.)	4.78 (+3% std.)
	Segment 2	+1.3%	4.85 (+4% std.)	4.89 (+4% std.)
	All	+1.2%	4.77 (+3% std.)	4.81 (+3% std.)
Sophisticated Guests Hosts Re-optimize Effort	Segment 1	+0.4%	4.75 (+3% std.)	4.79 (+4% std.)
	Segment 2	+0.5%	4.87 (+6% std.)	4.91 (+7% std.)
	All	+0.4%	4.78 (+4% std.)	4.82 (+4% std.)

Note: The table reports the counterfactual analysis of rating forgiveness removing 30% of negative noise, including overall level and segment level results.

and incentivizing sellers to provide quality. We develop and estimate an infinite-horizon dynamic model of consumer learning and host quality investment. Leveraging directly observable seller efforts, we quantify the informational and incentive effects of ratings on the platform and study the implications of alternative rating designs, including rating manipulation and rating forgiveness.

Our results provide evidence that sellers strategically invest in quality and that ratings play an important role in determining the timing and intensity of these investments. Consumers also actively learn from observed ratings when forming beliefs about latent quality and making booking decisions. The findings highlight the importance of rating systems in mitigating asymmetric information not only through information provision but also through their impact on dynamic quality provision incentives. To account for heterogeneity across hosts, we classify listings into two segments and obtain distinct estimates for each segment. The heterogeneity implies that alternative rating system designs can generate distinct equilibrium responses across market segments.

Using the estimated model, we investigate counterfactual rating policies, including rating manipulation and rating forgiveness. Rating manipulation artificially inflates ratings without improving quality, thereby introducing positive noise into the rating system. We highlight an additional mechanism that manipulated ratings weaken sellers' incentives to improve quality. As a result, booking rates, average quality, and average booked quality all decline, even when consumers are sophisticated and update their beliefs to account for rating manipulation. These negative effects are especially

strong in the value segment. We also study rating forgiveness policies that selectively remove negative rating noise. Unlike rating manipulation, rating forgiveness strengthens incentives for quality provision by reducing the likelihood that high-quality sellers are penalized by noisy negative reviews. We find that rating forgiveness increases maintenance effort, booking rates, average quality, and average booked quality relative to the baseline. These positive effects are stronger in the premium segment. Overall, our findings reveal an important asymmetry in rating system design: adding positive noise is harmful, while removing negative noise can improve market outcomes. The counterfactual analysis demonstrates that rating systems affect welfare not only through information provision but also through their impact on sellers' dynamic incentives to provide quality.

We conclude by highlighting several promising avenues for future research. Our study focuses on costly investments with persistent effects on quality, such as deep cleaning and repairs, due to our setting and data. In practice, sellers also exert more transient and perishable forms of effort, such as communication quality and responsiveness. Because these activities are delegated to the management agency in our data, they are largely standardized and non-strategic in our context. Future work could study environments with multidimensional effort choices and jointly determine how different forms of effort interact with reputation incentives. Extending the framework to incorporate endogenous pricing decisions, advertising investments, platform competition, and entry-exit dynamics would also provide important directions for future research.

Funding and Competing Interest

The authors declare no funding or competing interests during the review process.

References

Abbring, J. H., J. J. Heckman, P.-A. Chiappori, and J. Pinquet (2003). Adverse selection and moral hazard in insurance: Can dynamic data help to distinguish? *Journal of the European Economic Association 1*(2-3), 512–521.

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Learning from reviews: The selection effect and the speed of learning. *Econometrica 90*(6), 2857–2899.

Akerlof, G. A. (1978). The market for "lemons": Quality uncertainty and the market mechanism. In *Uncertainty in Economics*, pp. 235–251. Elsevier.

Ananthakrishnan, U., D. Proserpio, and S. Sharma (2023). I hear you: Does quality improve with customer voice? *Marketing Science*.

Arrow, K. J. (1968). The economics of moral hazard: further comment. *The American Economic Review 58*(3), 537–539.

Batikas, M. and T. Kretschmer (2018). Entrepreneurs on the darknet: Reaction to negative feedback. *Available at SSRN 3238141*.

Besbes, O. and M. Scarsini (2018). On information distortions in online ratings. *Operations Research 66*(3), 597–610.

Cabral, L. and A. Hortaçsu (2010). The dynamics of seller reputation: Evidence from ebay. *The Journal of Industrial Economics 58*(1), 54–78.

Cabral, L. and L. Li (2015). A dollar for your thoughts: Feedback-conditional rebates on ebay. *Management Science 61*(9), 2052–2063.

Chevalier, J. A., Y. Dover, and D. Mayzlin (2018). Channels of impact: User reviews when quality is dynamic and managers respond. *Marketing Science 37*(5), 688–709.

Einav, L. and A. Finkelstein (2018). Moral hazard in health insurance: what we know and how we know it. *Journal of the European Economic Association 16*(4), 957–982.

Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013). Selection on moral hazard in health insurance. *American Economic Review 103*(1), 178–219.

Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies 62*(1), 53–82.

Fan, Y., J. Ju, and M. Xiao (2016). Reputation premium and reputation management: Evidence from the largest e-commerce platform in china. *International Journal of Industrial Organization 46*, 63–76.

Filippas, A., J. J. Horton, and J. M. Golden (2022). Reputation inflation. *Marketing Science 41*(4), 733–745.

Fradkin, A., E. Grewal, and D. Holtz (2021). Reciprocity and unveiling in two-sided reputation systems: Evidence from an experiment on airbnb. *Marketing Science 40*(6), 1013–1029.

Fradkin, A. and D. Holtz (2023). Do incentives to review help the market? evidence from a field experiment on airbnb. *Marketing Science 42*(5), 853–865.

Gandhi, A., B. Hollenbeck, and Z. Li (2024). Misinformation and mistrust: The equilibrium effects of fake reviews on amazon. com.

He, S., B. Hollenbeck, and D. Proserpio (2022). The market for fake reviews. *Marketing Science 41*(5), 896–921.

Hollenbeck, B. (2018). Online reputation mechanisms and the decreasing value of chain affiliation. *Journal of Marketing Research 55*(5), 636–654.

Holmström, B. (1999). Managerial incentive problems: A dynamic perspective. *The Review of Economic Studies 66*(1), 169–182.

Hörner, J. and N. S. Lambert (2021). Motivational ratings. *The Review of Economic Studies 88*(4), 1892–1935.

Hunter, M. (2022). Chasing stars: Firms' strategic responses to online consumer ratings. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 298–299.

Ifrach, B., C. Maglaras, M. Scarsini, and A. Zseleva (2019). Bayesian social learning from consumer reviews. *Operations Research 67*(5), 1209–1221.

Jeziorski, P., E. Krasnokutskaya, and O. Ceccarini (2017). Adverse selection and moral hazard in the dynamic model of auto insurance. *UC Berkeley, Haas School of Business working paper*.

Jin, Y. (2023). Re-examining moral hazard under inattention: New evidence from behavioral data in auto insurance. *Available at SSRN 4453814*.

Kaul, R., S. J. Anderson, P. K. Chintagunta, and N. Vilcassim (2024). Call me maybe: Does customer feedback seeking impact nonsolicited customers? *Marketing Science 0*(0), null.

Kim, M., K. Sudhir, and K. Uetake (2022). A structural model of a multitasking salesforce: Incentives, private information, and job design. *Management Science 68*(6), 4602–4630.

Klein, T. J., C. Lambertz, and K. O. Stahl (2016). Market transparency, adverse selection, and moral hazard. *Journal of Political Economy 124*(6), 1677–1713.

Lewis, G. and G. Zervas (2016). The welfare impact of consumer reviews: A case study of the hotel industry.

Lewis, G. and G. Zervas (2019). The supply and demand effects of review platforms. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 197–197.

Leyden, B. T. (2025). Platform design and innovation incentives: Evidence from the product rating system on apple's app store. *International Journal of Industrial Organization 99*, 103133.

Luca, M. (2016). Reviews, reputation, and revenue: The case of yelp. com. *Com (March 15, 2016). Harvard Business School NOM Unit Working Paper* (12-016).

Luca, M. and G. Zervas (2016). Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science 62*(12), 3412–3427.

Ma, L., B. Sun, and S. Kekre (2015). The squeaky wheel gets the grease—an empirical analysis of customer voice and firm intervention on twitter. *Marketing Science 34*(5), 627–645.

Mayzlin, D., Y. Dover, and J. Chevalier (2014). Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review 104*(8), 2421–2455.

Pakes, A. and D. Pollard (1989). Simulation and the asymptotics of optimization estimators. *Econometrica: Journal of the Econometric Society*, 1027–1057.

Papanastasiou, Y., K. Bimpikis, and N. Savva (2018). Crowdsourcing exploration. *Management Science 64*(4), 1727–1746.

Pocchiari, M., D. Proserpio, and Y. Dover (2024). Online reviews: A literature review and roadmap for future research. *International Journal of Research in Marketing*.

Proserpio, D. and G. Zervas (2017). Online reputation management: Estimating the impact of management responses on consumer reviews. *Marketing Science 36*(5), 645–665.

Reimers, I. and J. Waldfogel (2021). Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review 111*(6), 1944–1971.

Rossi, M. (2024). Competition and reputation in an online marketplace: Evidence from airbnb. *Management Science 70*(3), 1357–1373.

Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 999–1033.

Saeedi, M. and A. Shourideh (2023). Optimal rating design under moral hazard. Technical report, arXiv. org.

Seiler, S., S. Yao, and G. Zervas (2018). Causal inference in word-of-mouth research: Methods and results. Technical report, Working paper.

Tadelis, S. (2016). Reputation and feedback systems in online platform markets. *Annual Review of Economics 8*(1), 321–340.

Vana, P. and A. Lambrecht (2021). The effect of individual online reviews on purchase likelihood. *Marketing Science 40*(4), 708–730.

Vatter, B. (2023). *Quality disclosure and regulation: Scoring design in medicare advantage*. SSRN.

Wang, Y. and A. Chaudhry (2018). When and how managers' responses to online reviews affect subsequent reviews. *Journal of Marketing Research 55*(2), 163–177.

Xin, Y. (2023). Asymmetric information, reputation, and welfare in online credit markets. *SSRN*.

Yang, B., T. Chan, H. Owan, and T. Tsuru (2024). Incentives from career concerns in a contract package: An empirical investigation. *Management Science 70*(9), 6093–6116.

A Background and Descriptive Evidence

A.1 Background

This section provides more background details of the review system on the Airbnb platform. On the Airbnb platform, on a scale from 1 to 5, guests can rate their stays in overall experience. As is shown in Figure A.1 (a)-(c), the average overall rating and the cumulative number of reviews are prominently displayed with the listings on both the search page and the product page. Besides the full history of overall experience rating, the platform also presents the average rating on six detailed subcategories: value, cleanness, check-in, accuracy, location, and communication (Figure A.1(c)). In the detailed reviews section, Airbnb also presents reviewer and trip details in 'review highlights' (Figure A.1(d)) where information includes reviewers' city, country, continent or region, date of

stay, trip type, and length of stay. [25] These rating designs suggest that ratings are important for the platform design and consumer decision-making.

A.2 Host Learning from Reviews

Our model assumes that hosts have full information on their current quality, and ratings affect their incentive to invest in quality. One potential violation of this assumption is that hosts are not fully informed about the quality of their property, and learn from the ratings and guest feedback about the current quality problem with the listing. Neglecting this host-learning effect would overestimate the function of ratings as an incentive tool. To study this effect, we leverage a unique feature of our dataset, which is the private feedback that guests sent to the hosts. On the Airbnb platform, besides the public ratings and reviews, guests can also leave a private feedback to the hosts about their experience in the STR property. This private feedback is unobservable to other consumers, and therefore only serve as a potential information tool, which enables us to evaluate the responsiveness of hosts' efforts to actionable information.

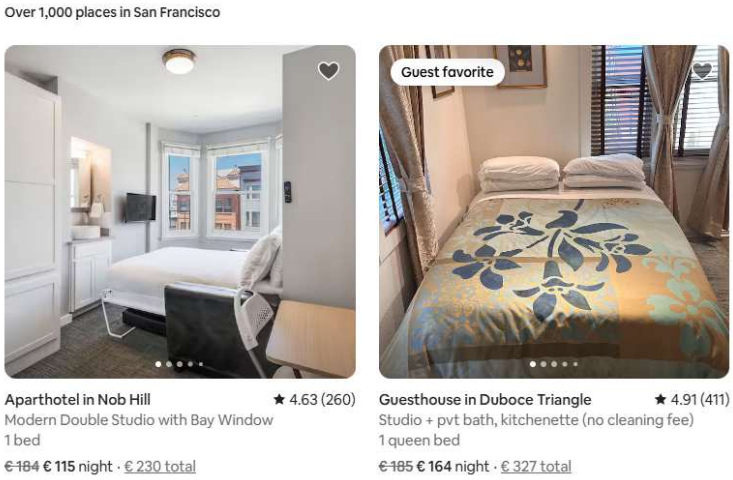
We conduct a textual analysis of the private feedback to the hosts on how actionable and informative they are about quality improvement, and create a continuous variable between 0 and 1 measuring how much actionable information is contained in the private feedback.[26] We then conduct the regression discontinuity analysis around checkout date by splitting the sample into stays with actionable private feedback and those without. We use the average actionability score as the cutoff of actionable v.s. not actionable[27]. To isolate the effect of bad rating and actionable private feedback, we further split the good stay sample and bad stay sample in Section 2.3 into stays with actionable and not actionable private messages. Table A.1 reports the regression results. Columns (1) and (2) report the results for stays with actionable and not actionable private feedback using only the sample of five-star stays. Columns (3) and (4) report the results for stays with

[25] Besides ratings, Airbnb also has badges for hosts and listings, including Superhosts, Guest Favorites, and top home highlights. These badges aggregate information on past ratings, as well as other statistics such as response rates, host cancellation rates, and quality-related incidents reported to Airbnb customer service: https://hospitable.com/airbnb-badges/. These designs can be seen as aiming to reduce noise asymmetrically in a local interval of quality.

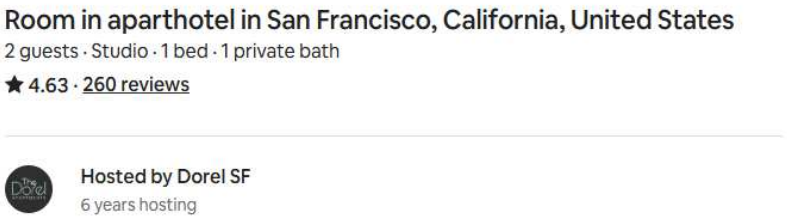
[26] Not actionable private feedback includes both empty feedback and those with little information, such as basic comments like "Thank you!" The mean and standard deviation of the actionability score of the private feedback are 0.129 and 0.258, respectively.

[27] We also tried the median actionability score as the cutoff of actionable vs not actionable, and results are very similar.

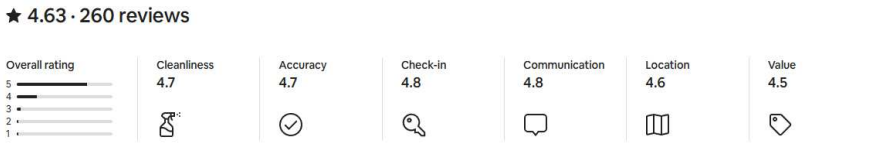
Figure A.1: Rating Display on Airbnb Platform



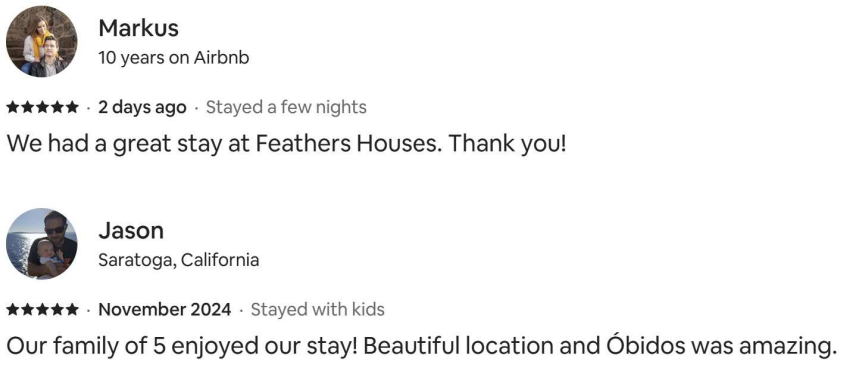
(a) Search Page



(b) Product Page Title Section



(c) Product Page Rating Section



(d) Product Page Rating Section

Note: The figure presents the layout of the platform and the prevalence of ratings throughout the customer journey.

actionable and not actionable private feedback using the sample of stays with ratings less than five stars. We do not find a significant difference between actionable vs not actionable private feedback among good stays and bad stays. They are always either both insignificant (all stays and good stays) or both significant (bad stays), which suggest the informativeness of private feedback does not drive the difference in host efforts.

Table A.1: Robustness Check: Host Learning from Guest Feedback

	Dependent Variable: Maintenance (m_{jt})			
	Good Review		Bad Review	
	(1)Actionable	(2)Not Actionable	(3)Actionable	(4)Not Actionable
β	-0.00021	-0.00056	0.00175*	0.00110**
	(0.00059)	(0.00049)	(0.00098)	(0.00047)
Obs	64,215	283,480	30,469	56,442

Note: The table reports results from regression discontinuity with triangular kernel results for a 30-day window before and after the checkout date of a stay. Columns (1) and (2) report the results for stays with actionable and not actionable private feedback using the sample of five-star stays. Columns (3) and (4) report the results for stays with actionable and not actionable private feedback using the sample of stays with ratings less than five stars. Standard errors are clustered at the reservation level and reported in parentheses. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

A.3 Resting on the Laurels Effect

Depending on the hosts' current reputation stock, measured by their average rating, the strength of their incentive to improve quality following a bad rating differs. To further investigate their response to different levels of incentive, we conduct the regression discontinuity analysis by splitting the sample into the cases when the listings' current average rating is above their overall average, vs when they are below their overall average. Table A.2 reports results from regression discontinuity analysis. Columns (1) and (2) report the results for stays with above-average and below-average ratings using the sample of five-star stays. Columns (3) and (4) report the results for stays with above-average and below-average ratings using the sample of stays with ratings less than five stars. We find that hosts are most responsive to bad ratings when their current average ratings are low, and we do not find a significant response to bad ratings when they are high. This provides suggestive evidence of a resting on the laurels effect: hosts are more responsive to bad ratings when their current reputation stock is lower.

Table A.2: Resting on the Laurels Effect

	Good Review		Bad Review	
	(1) High Average	(2) Low Average	(3) High Average	(4) Low Average
β	-0.00030	-0.00052	0.00045	0.00196***
	(0.00029)	(0.00047)	(0.00050)	(0.00073)
Obs	186,005	141,189	34,127	49,563

Note: The table reports results from regression discontinuity with triangular kernel results for a 30-day time window before and after the checkout date of a stay. Columns (1) and (2) report the results for stays with above-average and below-average ratings using the sample of good five-star stays. Columns (3) and (4) report the results for stays with above-average and below-average ratings using the sample of bad stays with ratings less than five stars. Standard errors are clustered at the reservation level and reported in parentheses. *** denotes significance at the 1% level, ** at 5% level, and * at 10%.

B Estimation and Results

B.1 Computing the Equilibrium

This section contains the description of the algorithm to compute the equilibrium of the game used in the nested loop of our GMM estimator.

The algorithm is a version of best response dynamics that solves the game backward over a large time horizon. At each time step r, we discretize the state space, excluding quality, using a fixed grid:

$$\Omega = [\underline{\omega}_p, \ldots, \bar{\omega}_p]^{Q-1} \times [\underline{\omega}_n, \ldots, \bar{\omega}_n].$$

The first part of the grid represents the belief parameters, while the second part discretizes the remaining days of the current reservation. For each point $\omega \in \Omega$ (and off the grid), we obtain implied prior probabilities for $\tilde{q} < Q$ using a logistic transformation.[28]

We define our algorithm recursively. Consider ex-ante value function V_{r-1}. We obtain V_r in a following way.

1. For each $\omega \in \Omega$, and each $q \leq Q$:

   (a) Construct the state vector s corresponding to a pair (ω, q).

[28] For $\tilde{q} < Q$, the prior is given by

$$G_t(\tilde{q}) = \frac{\exp(\omega(\tilde{q}))}{1 + \sum_{q<Q} \exp(\omega(q))}, \quad \text{and} \quad G_t(Q) = \frac{1}{1 + \sum_{q<Q} \exp(\omega(q))},$$

where $\omega_p(q)$ denotes the q-th dimension of the grid point ω.

Table B.1: Monte Carlo Test

	True Parameter	Estimate	Std.
Quality Evolvement and Effort Effectiveness			
Quality Appreciation Prob. (α)	0.050	0.041	0.014
Default Depreciation Prob. (δ)	0.060	0.070	0.015
Effort Success Prob. (p^E)	0.900	0.850	0.052
Consumer Demand			
Constant (θ_1^U)	-2.000	-1.947	0.048
Quality (θ_2^U)	0.500	0.558	0.050
Preference for Nights (λ)	5.000	4.962	0.037
Review System			
Review Informativeness (η)	0.020	0.026	0.005
Prob. of Submitting a Rating (P^R)	0.800	0.725	0.021
Initial Quality			
Central Tendency (μ^Q)	4.500	4.533	0.029
Dispersion (σ^Q)	1.000	0.967	0.031
Cost of Effort			
Central Tendency (θ^C)	5.000	4.972	0.025
Dispersion (σ^C)	1.000	0.948	0.048

Note: The table reports results from a Monte Carlo simulation test. We simulate data with a similar size of our sample. We consider 200 listings over five years. The true parameters and estimated parameters are reported.

(b) Guess the CCP of effort at $c_r(s)$ (to make guess efficient, set the guess to CCP at $r-1$).

(c) Compute best response booking CCP $b_r(s)$, given $e_r(s)$.

(d) Compute best response booking CCP $e_r(s)$, given $b_r(s)$.

(e) Iterate until convergence

2. Obtain value function V_r by using quadratic spline interpolation.

B.2 Monte Carlo Test

We verify the identification property of our dynamic model with a Monte Carlo simulation test on whether we are able to recover the structural parameters. We simulate data with a similar size of our sample. We consider 200 listings over five years. Table B.1 presents the results. Overall, coefficients are precisely estimated; starndard errors are small, and all estimates lie within two standard errors from the truth. We conclude that we are able to successfully recover structural parameters in the simulated sample.

Table B.2: Model Fit

Moment	Predicted	Data
Rating Transition		
Without a Maintenance Request		
Prob. Up ($r_{\tau_{k+1}} > r_{\tau_k}$)	0.10	0.17
Prob. Stay ($r_{\tau_{k+1}} = r_{\tau_k}$)	0.78	0.68
Prob. Down ($r_{\tau_{k+1}} < r_{\tau_k}$)	0.12	0.15
With a Maintenance Request		
Prob. Up ($r_{\tau_{k+1}} > r_{\tau_k}$)	0.32	0.32
Prob. Stay ($r_{\tau_{k+1}} = r_{\tau_k}$)	0.61	0.60
Prob. Down ($r_{\tau_{k+1}} < r_{\tau_k}$)	0.07	0.08
Average First Rating	4.78	4.70
Average Occupancy Rate	0.62	0.64
Maintenance Frequency	0.01	0.02

Note: The table reports results from an out-of-sample validation exercise. We split the sample randomly into an 80% training sample and 20% holdout sample. We estimate the baseline homogeneous model on the training sample and use the estimated model to predict several key moments in the holdout sample.

B.3 Model Fit

We examine model fit by comparing several key predictions to their empirical counterparts. We split the sample randomly into an 80% training sample and 20% holdout sample. We estimate the model on the training sample and use the estimated model to predict several key moments in the holdout sample. Table B.2 presents the simulated predictions and the empirical counterparts from the holdout sample. We evaluate the model's out-of-sample fit in terms of booking rate, initial ratings, rating transition probability conditional on with and without maintenance, maintenance probability conditional on last rating, and demand response to average rating and number of ratings. We find that the model predicts the holdout sample data reasonably well in terms of the key moments.

B.4 More Counterfactual Analysis

This section reports additional counterfactual exercises that complement the main analysis in Section 6. These counterfactuals are designed to clarify the mechanisms through which rating systems affect equilibrium outcomes. Some of the designs condition directly on latent true quality and therefore should be interpreted as theoretical benchmarks rather than immediately implementable

platform policies.[29]

We consider four additional classes of rating systems:

No Rating We study a no-rating benchmark, in which guests do not observe ratings when making booking decisions. This benchmark isolates the overall value of the rating system relative to an environment in which the rating system is absent. Besides ratings, Airbnb also provides channels for guests to report listings that are misleading or scams to maintain quality provision[30], therefore, in this case the marketplace will not collapse in the traditional adverse selection sense.

Full Transparency We consider a full-transparency benchmark, in which ratings perfectly reveal true quality. This case is useful for evaluating the tradeoff between information provision and effort incentives. In static environments with adverse selection, full transparency is typically desirable because it allows consumers to sort efficiently across products. In dynamic environments with moral hazard, however, full transparency may weaken incentives for costly quality provision by eliminating the reputation uncertainty that motivates sellers to maintain quality (Hörner and Lambert, 2021; Saeedi and Shourideh, 2023).

Missing Ratings This counterfactual removes a fraction of ratings at random and therefore lowers the amount of information available to consumers without systematically biasing ratings upward or downward. Comparing missing ratings with fake reviews helps distinguish the effect of lower information amount from the effect of biased positive information.

Asymmetric Noise We examine asymmetric changes in rating noise. Motivated by theoretical results showing that optimal rating systems under moral hazard may involve different degrees of informativeness at different signal realizations, we consider three cases: eliminating positive noise, eliminating negative noise, and a rating system that fully reveals whether a listing is of the highest quality while pooling all lower-quality listings together. These exercises clarify whether welfare gains come primarily from making all ratings more transparent or from reducing particular types of rating noise.

Table B.3 presents the results of different rating counterfactuals. The no-rating benchmark substantially worsens market outcomes as expected, which confirms that ratings play an important

[29]Because the platform does not directly observe true quality, rating rules conditional on true quality are best interpreted as theoretical benchmarks. In practice, platforms may approximate these rules using observable signals, dispute procedures, anomaly detection, or auxiliary quality reports.

[30]Source:https://www.airbnb.com/help/article/3598

role in mitigating asymmetric information. Without ratings, consumers are less able to direct demand toward high-quality listings, and hosts face weaker incentives to maintain quality.

Randomly missing ratings generates much smaller losses than rating manipulation. When 20% of ratings are missing at random, the booking rate decreases by only 0.2 percent, average quality declines by 3 percent of a standard deviation, and average booked quality declines by 1 percent of a standard deviation. In contrast, fake reviews are more harmful because they not only reduce available information but also increase bias in the signals. This comparison shows that biased positive noise is more damaging than a comparable reduction in rating availability.

The asymmetric-noise counterfactuals reveal that the direction of rating noise matters. Removing positive noise improves outcomes modestly, increasing the booking rate by 0.3 percent and raising both average quality and average booked quality. Removing negative noise generates substantially larger gains: the booking rate increases by 1.4 percent, average quality rises by 11 percent of a standard deviation, and average booked quality rises by 12 percent of a standard deviation. These results suggest that negative noise is especially costly in this environment because it may distort consumer beliefs about high-quality listings and weaken hosts' incentives to maintain top quality.

The transparency benchmarks further illustrate the tradeoff between information and incentives. Full transparency improves outcomes relative to the status quo, increasing the booking rate by 1 percent and average booked quality by 10 percent of a standard deviation. However, full transparency does not dominate all alternative designs. In particular, removing negative noise and revealing only top quality while pooling lower qualities both perform at least as well as full transparency on some dimensions. This finding is consistent with the idea that, when sellers strategically choose effort, the optimal rating system need not be fully transparent.

Overall, these additional counterfactual analyses reinforce the main message of the paper: effective rating design in dynamic-quality markets requires balancing consumer learning and sellers' intertemporal incentives to provide quality.

Table B.3: More Counterfactual Results

Rating System	Booking rate Revenue	Average quality	Average booked quality Consumer surplus
No Rating	-2.2%	4.57 (-17% std.)	4.52 (-26% std.)
Missing 20% Ratings	-0.2%	4.71 (-3% std.)	4.77 (-1% std.)
Fake Reviews 10%	-0.2%	4.72 (-2% std.)	4.75 (-3% std.)
Fake Reviews 20%	-0.8%	4.68 (-6% std.)	4.72 (-6% std.)
Rating Forgiveness	+0.4%	4.78 (+4% std.)	4.82 (+4% std.)
Remove Positive Noise	+0.3%	4.77 (+3% std.)	4.82 (+4% std.)
Remove Negative Noise	+1.4%	4.82 (+11% std.)	4.90 (+12% std.)
Transparency at 5 and Pooling Below 5	+0.8%	4.77 (+10% std.)	4.86 (+8% std.)
Full Transparency	+1%	4.74 (+7% std.)	4.88 (+10% std.)

Note: The table reports more counterfactual results of alternative rating systems. All scenarios are compared relative to the status quo as the benchmark.