Online Reviews: Information Content, Biases, and Platform Design

Tommaso Bondi, Michelangelo Rossi*

October 8, 2025

Abstract

This review synthesizes recent research on biases in online reviews and their implications for digital platform design. It introduces a structured framework that distinguishes between strategic distortions (e.g., fake, incentivized, and reciprocal reviews) and non-strategic biases stemming from behavioral and aggregation mechanisms. The article highlights how a wide spectrum of root causes, ranging from misaligned incentives to user psychology, contextual factors, and self-selection can skew both individual evaluations and aggregate ratings.

Keywords: online reviews, digital platforms, behavioral biases, platform design.

*T. Bondi: Cornell Tech and SC Johnson School of Business, and CESifo. E-mail: tbondi@cornell.edu. M. Rossi: HEC Paris, and CESifo. E-mail: rossim@hec.fr.

1 Introduction

Online consumer reviews are a foundational element of modern digital marketplaces. A large and interdisciplinary literature has examined various aspects of reviews, including the characteristics of users who write them, how consumers interpret them, and, consequently, their influence on product sales and consumer welfare Prior review articles have provided valuable insights into several of these dimensions (Tadelis, 2016; Pocchiari et al., 2024).

The focus of the present review article is complementary: to summarize the growing and eclectic literature studying the drivers and consequences of biases in online review systems. Given the fast pace at which this literature is evolving, we place particular emphasis on recent contributions, focusing on papers published since 2015.

We begin by considering a stylized benchmark in which reviews are used to estimate a product's fixed, intrinsic quality. Let r_{ij} denote the rating left by user j for product i, and $\bar{r}_i$ the average rating for product i. We study deviations of $\bar{r}_i$ from Q_i. Are these empirically significant? Systematic? Predictable? Consequential for consumer choice? And from a purely normative standpoint, is $\bar{r}_i \to Q_i$ as the number of reviews grow desirable?

To answer these questions, we trace deviations of $\bar{r}_i$ from Q_i back to individual ratings and divide them into two conceptually distinct categories:

$$r_{ij} = Q_i + \eta_{ij} + \theta_{ij},$$

where:

- η_{ij} captures systematic distortions in individual ratings that are due to strategic, behavioral, or contextual factors,
- θ_{ij} reflects idiosyncratic user taste.

Under the η_{ij} term, we include a large range of empirically documented phenomena, ranging from fake and incentivized reviews (Section 2) to behavioral biases in consumers' quality perception and evaluation (Section 3.1). In Section 3.2, we turn our attention to θ_{ij}, and show that the aggregation of this term across consumers can introduce large and systematic biases in average ratings.

We conclude our review with an eye towards the future, by offering some speculation on how AI and future agentic technologies may affect reviews, their biases, and their use by platforms and consumers.

2 Strategic Biases

2.1 Fake Reviews

Fake reviewsare intentionally crafted user feedback designed to manipulate a product's or seller's perceived reputation. These reviews may be posted by the firm itself, hired agents, or even rival businesses, and typically aim either to inflate the review scores of the target product or to undermine competitors. Reflecting growing concern over this phenomenon, policymakers have recently initiated efforts to combat fake reviews and other deceptive endorsement tactics (see Federal Trade Commission, 2022b).

The prevalence of fake reviews have been well-documented empirically. Luca and Zervas (2016) show that the likelihood of promotional reviews increases under intensified local competition or reputation shocks. More recently, He et al. (2022b) uncover a market for fake reviews in which sellers, especially small and low-quality ones, procure fake reviews via social media in exchange for free products. These manipulative practices degrade the informational value of reviews, mislead consumers, and erode trust in the platform.

Fake reviews differ from traditional advertising in that they are not recognized by consumers as promotional content. While consumers typically discount explicit advertisements as biased, they often lack the ability to recognize fake reviews. This makes fake reviews particularly deceptive, and potentially more persuasive, than conventional ads.[1] As a result, fake reviews are detrimental to consumer welfare. Experimental evidence by Akesson et al. (2023) estimates significant losses for consumers, amounting to approximately \$0.12 for every dollar in consumer spending.

Furthermore, a decline in review reliability can trigger broader equilibrium effects that undermine the usefulness of the review ecosystem. Glazer et al. (2021) show that when users grow uncertain about

[1] Supporting this distinction, Hollenbeck et al. (2019) finds that hotels treat TripAdvisor ratings and advertising spending as substitutes rather than complements.

the authenticity of reviews, the informativeness of the system as a whole diminishes. In their model, full transparency, (that is, reporting all reviews without filtering) emerges as the optimal strategy, because any attempt by the platform to suppress or moderate content can inadvertently encourage more extreme behavior from fake reviewers. This, in turn, distorts user beliefs and reduces the perceived credibility of honest reviews. Mostagir and Siderius (2023) examine how reviewers may actively respond to platform incentives and enforcement mechanisms. They find that stricter auditing policies may unintentionally backfire by raising sellers' willingness to pay for fake reviews, thereby increasing reviewers' incentives to engage in manipulation.

These dynamics highlight how platform design and reviewer incentives jointly shape the sustainability of trust in rating systems. In practice, the reputational risks associated with unchecked manipulation often compel platforms to intervene, even when such actions may weaken the perceived neutrality or informativeness of the system. Detection methods have traditionally relied on metadata and textual features (Wu et al., 2020), but as fraudulent practices evolve, platforms are developing more sophisticated tools. For instance, He et al. (2022a) propose a network-based detection strategy that revolves around the identification of suspicious users, not (just) suspicious reviews. However, even effective detection and removal may not fully eliminate the impact of fake reviews. Once incorporated into platform recommendation systems, early fake reviews can affect future exposure and learning trajectories. Adamopoulos (2024) highlight how biased early reviews, if used to train recommendation algorithms, can produce long-lasting distortions in consumer behavior and product visibility.

2.2 Incentivized and Influenced Reviews

Not all seller-driven biases in online reviews stem from fraudulent behavior. While *fake reviews* are typically the result of deliberate manipulation, *incentivized* and *influenced* reviews can distort either η_{ij} – the way users evaluate and report product quality – or θ_{ij} – the profile of users who choose to leave a review – in subtler but important ways.

In principle, incentivizing users to leave more ratings is a promising avenue to increase reviews informativeness, as limited participation of users in reviewing products has been widely documented. Brandes et al. (2022) and Brandes and Dover (2022) report that only 11% and 18.6% of hotel guests, respectively, leave a review after their stay on a traditional hotel booking platform. In contrast, review rates are substantially higher on peer-to-peer platforms: on Airbnb, for instance, approximately 70% of stays result in a review, according to Fradkin et al. (2021).

This low engagement rate creates problems for both platforms and sellers, especially for new entrants who suffer from the "cold-start" problem. That is, when a seller offers a high-quality product but lacks reviews, buyers may hesitate to purchase due to insufficient information. This can skew the classic exploration/exploitation trade-off (Papanastasiou et al., 2018; Che and Hörner, 2018) towards established products and create a barrier to entry that worsens equilibrium supply and, thus, welfare (Vellodi, 2018).

Low review rates present another, crucial issue: the subset of buyers deciding to rate a product is usually non-random (Schoenmueller et al., 2020). As buyers with strong opinions (whether positive or negative) are more likely to rate, the review distribution for many products is J-shaped, which impacts both the average and the variance of reviews. Encouraging a broader set of users to post reviews can therefore improve the informativeness of the signal by capturing a more diverse distribution of θ_{ij}. For example, Karaman (2021) finds that incentivizing reviews increases review volume and reduces selection bias. Li et al. (2020) develop a theoretical model grounded in signaling theory, predicting that only sellers of high-quality products will offer rewards for truthful feedback, particularly when they are not yet established or when their products lack prior reviews. Their empirical analysis of Taobao's reward-for-feedback mechanism supports this prediction: sellers are more likely to offer incentives for feedback on high-quality, low-review products, and such incentives lead to a 36% increase in sales.

However, incentives to increase consumer participation come at a cost, as they also influence the *content* of reviews in several critical ways. Woolley and Sharif (2021) show that incentivized reviews tend to contain more positive emotional language than non-incentivized ones. By contrast, Fradkin and Holtz (2023) study an experiment on Airbnb in which guests were offered coupons in exchange for leaving reviews. The treatment increased review volume but led to more negative ratings and weakened the correlation between review content and transaction quality. Interestingly, despite these changes in review content, the treatment had no measurable effect on bookings or revenue—suggesting that incentivized reviews may still offer useful, albeit noisy, signals to future consumers.

Additional sources of potential bias arise in two-sided review systems, in which buyers and sellers rate each other. These settings often give rise to strategic feedback dynamics: users may condition their own reviews on the anticipated reactions from the counterparty. On platforms like eBay, which initially enabled mutual and publicly visible feedback, this structure led to a form of retaliatory equilibrium in which users refrained from leaving honest negative reviews to avoid receiving retaliatory ratings in return. The resulting bias reduced the informativeness of the review system and distorted market signals.

Empirical evidence supports the impact of platform rules on these dynamics. For instance, Hui et al. (2018) study eBay's 2008 policy change that prevented sellers from leaving negative feedback. Following the change, low-quality sellers experienced reduced market success or exited the platform altogether, while overall service quality improved. These outcomes suggest that institutional design choices can directly influence the magnitude and direction of η_{ij}, by shaping how users perceive the consequences of leaving truthful feedback. Fradkin et al. (2021) document similar reciprocity dynamics on Airbnb: users time their reviews strategically when these are revealed sequentially. In response, Airbnb redesigned its platform in 2014 to reveal reviews simultaneously. This intervention increased review rates and reduced both retaliatory and reciprocal rating behavior. Still, highly positive reviews remain prevalent on the platform. Proserpio et al. (2018) attribute this to the interpersonal nature of Airbnb transactions, which may discourage users from leaving critical feedback.

More broadly, Filippas et al. (2018) document a long-term trend of "reputation inflation" in online marketplaces. Using data from an online labor platform, they show that average ratings have become increasingly positive over time, even without corresponding improvements in actual performance.[2]

A final class of distortions stems from platform features that allow sellers to respond to user reviews. These managerial responses can affect both product quality (Chevalier et al., 2018; Ananthakrishnan et al., 2023) and user reviewing behavior. Proserpio and Zervas (2017) find that hotel managers are more likely to respond to negative reviews, and that such responses lead to a 0.12-star average increase in subsequent ratings and a 12% increase in review volume. Similarly, Wang and Chaudhry (2018) argue that responses to negative reviews can positively shape prospective buyers' beliefs by countering negative sentiment or demonstrating responsiveness. However, these effects are not uniform across user groups. Proserpio et al. (2021) show that after seller responses are introduced, female reviewers are significantly less likely to leave negative reviews, possibly due to a higher incidence of confrontational or dismissive replies. That is, this dynamic introduces both selection bias in θ_{ij} and distortions in η_{ij}.

3 Non-Strategic Biases

3.1 Behavioral Biases

A central challenge in interpreting online reviews is that they may not reflect a product's intrinsic quality Q_i, but rather consumers' subjective utility, which can be swayed by a range of contextual and psychological factors (on top of their idiosyncratic tastes, which we discuss more in Section 3.2).

One salient departure from reviews as quality signals involves prices and value-for-money considerations. Luca and Reshef (2021) show that a 1% increase in price leads to a 3–5% drop in average ratings on a food delivery platform, suggesting that more expensive experiences are rated more harshly. These findings are the result of two opposing channels: higher prices reduce value for money, worsening reviews, but they also deter low-valuation buyers, potentially improving reviews through favorable selection. Carnehl et al. (2024) formalize this trade-off and study its implications for the design of platform-level rating systems. In related empirical work, Carnehl et al. (2022) use Airbnb data to document a dominant value-for-money effect and show that strategic hosts adjust prices and effort in ways that exploit this bias to manage ratings and revenues.

Another key behavioral distortion arises from user expectations shaped by external signals. Platforms or third parties quality certifications – such as Michelin stars, Airbnb Superhost badges, or Academy Award nominations – affect consumers' expectations, which can in turn influence their ratings. High expectations often lead to disappointment, and thus to systematically lower ratings. Li et al. (2024) find that restaurants losing their Michelin stars see an uptick in the valence of their reviews. Meister and Reinholtz (2025) show that Airbnb's Superhost designation leads to lower subsequent ratings, possibly due to a combination of inflated expectations, adjustments in host behavior and changes in the reviewing population. Rossi and Schleef (2024) isolate a disappointment-driven penalty in movie ratings following Academy Award nominations, especially among inexperienced users. These mechanisms are supported

[2] Platforms may not always be opposed to inflated ratings, as overly positive reviews can reassure buyers and boost transactions. A case in point is Fashion Nova, which was fined by the FTC for systematically suppressing reviews below four stars to maintain an illusion of quality (Federal Trade Commission, 2022a).

by belief-elicitation evidence from Aridor et al. (2024), who document how prior expectations, shaped by recommendation systems, strongly influence satisfaction and post-consumption evaluations.

Worse, behavioral biases can also emerge from contextual and situational factors that are completely unrelated to product attributes. Brandes and Dover (2022) show that users are more likely to post reviews on rainy days and that these reviews tend to be more negative, suggesting that emotional state at the time of reviewing, not consumption, can skew evaluation. In addition to transient moods, social influence can systematically distort review behavior. Jacobsen (2015) and Sunder et al. (2019) document herding dynamics in online platforms, showing that ratings are shaped by both crowd and peer opinions, undermining the wisdom of the crowd. These effects vary with user experience and can lead to divergence or convergence in evaluations, depending on the reference group and the reviewer's own history.

Finally, reviewer identity and broader societal biases can influence both the content and likelihood of reviews. Recent studies document gender- and race-based disparities – both on the consumer and the product side – in evaluation outcomes. Aguiar (2024) show that female-led movies receive disproportionately lower ratings from male crowd reviewers, despite receiving similar assessments from professional critics. Bayerl et al. (2024) find that women leave more favorable online reviews than men, potentially due to heightened concerns about social backlash. Bairathi et al. (2023) provide evidence that buyers give higher public ratings to male freelancers than to female ones, even when private satisfaction is the same, likely reflecting stereotype-driven concerns about confrontation. Aneja et al. (2025) find that labeling restaurants as Black-owned increases engagement and alters the composition of reviewers, suggesting that identity disclosures systematically affect both participation and perceived quality.

3.2 Aggregation Biases

Thus far, we have examined sources of bias that originate at the level of individual reviews, either through strategic manipulation or behavioral distortions in how users perceive and report their experiences. We now turn our attention to a different (and subtler) family of distortions: aggregation biases. We argue that, even in an ideal scenario in which individual reviews do not suffer from the several biases discussed so far, the process through which they are accumulated and summarized, typically via simple averages, can lead to systematic biases in aggregate ratings.

Hu et al. (2017) provide a useful conceptual foundation for this discussion by distinguishing between two forms of self-selection: selection into reviewing (or underreporting bias – discussed in Section 2.2) and selection into consumption (or acquisition bias). Simply put, even if all buyers of a product reviewed it (that is, no underreporting bias), these are likely to differ in fundamental ways from non-buyers – especially as it pertains to their taste for the product itself.

This framing also helps illuminate one key fact, which we alluded to in our intro: whenever reviews reflect both quality and idiosyncratic fit (as they arguably do in virtually all empirical applications – be it movies, restaurants or online courses), $\bar{r}_i \to Q_i$ is not a realistic benchmark. Buyers generally have a more favorable opinion of a product than non-buyers: $\mathbb{E}(\theta_{ij}|\text{Purchase}) > 0$, so that $\bar{r}_i > Q_i$.

This is not necessarily problematic. The key observation is that, whenever consumers employ reviews to choose between products, their choices are shaped by *relative ratings* (and *rankings*) rather than by the absolute ratings themselves. Thus, the central issue is that $\mathbb{E}(\theta_{ij}|\text{Purchase})$ varies widely across products, so that both *differences in ratings* and, therefore, *rankings* can be misleading.

Although consumers partially recognize these distortions and may use distributional cues to correct for them, such adjustments are extremely complex (and product-specific) and thus partial at best.

Acemoglu et al. (2022) develop a Bayesian model of online review dynamics and show how the structure of the review system – specifically, whether it reports full histories or summary statistics – alters the consumption (and reviewing) decisions of future users. They derive conditions under which learning is biased, and show that, due to consumers' idiosyncratic tastes, *i*) learning from simple averages lead to biased outcome and *ii*) more detailed information does not always lead to better learning outcomes.

Bondi (2025) builds a model of naïve social learning in which consumers' beliefs, choices, and ratings dynamically feed into each other, and characterizes its long-run behavior. He shows that equilibrium ratings systematically advantage lower quality and more polarizing products, since these induce stronger consumer self-selection: $\mathbb{E}(\theta_{ij}|\text{Purchase})$. Thus, excessive choice fragmentation and systematic long-run biases persist even without any selection into reviewing, or other behavioral or strategic biases.

These theoretical results are mirrored in empirical studies that document how reviewer heterogeneity affects aggregate ratings. Bondi et al. (2024) show that experienced users tend to simultaneously consume higher quality products and to post harsher evaluations for any level of quality. When ratings from experienced and novice users are pooled without adjustment, the result is a compression of scores, or equivalently a penalty for high-quality products. The authors propose a debiasing algorithm that recovers more accurate ratings and rankings and document *ranking reversals* for roughly 8% of movie pairs in their MovieLens dataset. Similarly, Dai et al. (2018) and Carayol and Jackson (2024) argue that heterogeneity in reviewer stringency and accuracy should be accounted for when aggregating ratings. Using Yelp and Bordeaux wine ratings data respectively, both papers demonstrate that simple structural corrections yield large gains in informativeness relative to the arithmetic mean.

Finally, the design of the review system itself can influence the extent to which users trust aggregate scores and whether they consult individual reviews. Lee et al. (2021) find that when rating dispersion is low, users are more likely to treat the average rating as a sufficient summary and skip reading individual reviews. In contrast, high dispersion encourages deeper review reading, especially of extreme opinions, due to their clarity and decisiveness. These behavioral responses may partially mitigate or exacerbate aggregation bias, depending on the context.

Taken together, these findings suggest that the informativeness of aggregated ratings cannot be taken for granted, even in a first best scenario in which everyone reviews and does so honestly (but, crucially, subjectively).

4 AI-Driven Biases, Detection and Consumer (Mis)Learning

The rise of AI and agentic technologies introduces both new challenges and novel opportunities for online review systems (Lee et al., 2023). These tools may amplify existing distortions or offer new strategies for mitigating them. In this section, we outline several key ways in which AI is likely to affect review bias in the coming years.

The arms race in fake reviews. One of the most pressing developments is the use of AI to generate fake reviews. Large language models (LLMs) can now produce persuasive, context-specific fake reviews at scale. Agentic AI, systems capable of autonomous goal pursuit, may further automate the entire manipulation pipeline, from identifying products to generating and posting fake content. In response, platforms may deploy sophisticated detection systems (Shukla et al., 2019), but these often lag behind the generative capabilities. Detection relies on pattern recognition and behavioral metadata, which adversarial systems can quickly adapt to. This asymmetry raises concerns that generative tools will outpace efforts to detect and moderate fake content.

AI-generated review responses. Businesses increasingly rely on AI to produce real-time, personalized responses to user reviews (Tsekouras et al., 2020). While this can enhance engagement and customer service, it also risks reinforcing biases in both the extensive (*who*) and intensive (*how*) margin of reviewing behavior.

Summarization bias. AI is frequently used to condense large sets of reviews into textual or visual summaries. These summaries, if not carefully designed, can introduce new forms of bias, especially when the underlying data reflect the self-selection of reviewers. Summarization algorithms may overemphasize extreme or unrepresentative opinions and fail to account for structural distortions in who leaves reviews.

We note that this issue is not limited to review summaries: similar concerns apply to recommendation systems that rely on user ratings to generate personalized suggestions, as these systems often assume that observed ratings are unbiased signals of user preferences. Without proper correction, recommendation algorithms may lead to systematically skewed suggestions and reinforcing feedback loops. As with review summarization, addressing these biases in recommendations requires greater awareness of the data-generating process and explicit mechanisms for bias correction.

Personalization and aggregation bias. A particularly promising use of AI is in addressing *aggregation bias*. Traditional review systems assume that all users value the same aspects of a product. In contrast, AI systems with access to detailed preference data can personalize the review signal. Instead of a single average score, platforms could show ratings and reviews from users with similar tastes and priorities. For example, a user who consistently values durability might be shown different reviews than one who prioritizes aesthetics. This personalization can make review signals more relevant and informative. However, it also introduces risks of over-filtering, privacy violations, and reduced exposure to diverse perspectives. Transparency around how such personalization is implemented will be essential to maintaining trust and accountability.

Overall, it is easy to imagine that AI will bring both opportunities and challenges to online WoM. While its impact on the quantity of information seems clear, drawing conclusions about *quality* appears more complex. Crucially, these will depend not only on technical capabilities but also on how platforms deploy these tools.

References

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

Adamopoulos, P. (2024). The spillover effect of fraudulent reviews on product recommendations. *Management Science*, 70(12):8818–8832.

Aguiar, L. (2024). Bad apples on rotten tomatoes: critics, crowds, and gender bias in product ratings. *Marketing Science*.

Akesson, J., Hahn, R. W., Metcalfe, R. D., and Monti-Nussbaum, M. (2023). The impact of fake reviews on demand and welfare. Technical report, National Bureau of Economic Research.

Ananthakrishnan, U., Proserpio, D., and Sharma, S. (2023). I hear you: does quality improve with customer voice? *Marketing Science*, 42(6):1143–1161.

Aneja, A., Luca, M., and Reshef, O. (2025). The benefits of revealing race: Evidence from minority-owned local businesses. *American Economic Review*, 115(2):660–689.

Aridor, G., Gonçalves, D., Kong, R., Kluver, D., and Konstan, J. (2024). The movielens beliefs dataset: Collecting pre-choice data for online recommender systems. In *Proceedings of the 18th ACM Conference on Recommender Systems*, pages 1–1.

Bairathi, M., Lambrecht, A., and Zhang, X. (2023). Gender disparity in online reputation: Evidence from an online freelance platform. *Available at SSRN 4635329*.

Bayerl, A., Dover, Y., Riemer, H., and Shapira, D. (2024). Gender rating gap in online reviews. *Nature Human Behaviour*, pages 1–14.

Bondi, T. (2025). Alone, together: A model of social (mis) learning from consumer reviews. *Marketing Science*.

Bondi, T., Rossi, M., and Stevens, R. L. (2024). The good, the bad and the picky: Consumer heterogeneity and the reversal of product ratings. *Management Science*.

Brandes, L. and Dover, Y. (2022). Offline context affects online reviews: The effect of post-consumption weather. *Journal of Consumer Research*, 49(4):595–615.

Brandes, L., Godes, D., and Mayzlin, D. (2022). Extremity bias in online reviews: The role of attrition. *Journal of Marketing Research*, 59(4):675–695.

Carayol, N. and Jackson, M. O. (2024). Finding the wise and the wisdom in a crowd: Estimating underlying qualities of reviewers and items. *The Economic Journal*, 134(663):2712–2745.

Carnehl, C., Schaefer, M., Stenzel, A., and Tran, K. D. (2022). Value for money and selection: How pricing affects airbnb ratings. *Innocenzo Gasparini Institute for Economic Research Working Paper Series*.

Carnehl, C., Stenzel, A., and Schmidt, P. (2024). Pricing for the stars: Dynamic pricing in the presence of rating systems. *Management Science*, 70(3):1755–1772.

Che, Y.-K. and Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925.

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

Dai, W., Jin, G., Lee, J., and Luca, M. (2018). Aggregation of consumer ratings: an application to yelp. com. *Quantitative Marketing and Economics*, 16:289–339.

Federal Trade Commission (2022a). Fashion Nova will Pay $4.2 Million as part of Settlement of FTC Allegations it Blocked Negative Reviews of Products.

Federal Trade Commission (2022b). FTC to Explore Rulemaking to Combat Fake Reviews and Other Deceptive Endorsements.

Filippas, A., Horton, J. J., and Golden, J. (2018). Reputation inflation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 483–484.

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

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

Glazer, J., Herrera, H., and Perry, M. (2021). Fake reviews. *The Economic Journal*, 131(636):1772–1787.

He, S., Hollenbeck, B., Overgoor, G., Proserpio, D., and Tosyali, A. (2022a). Detecting fake-review buyers using network structure: Direct evidence from amazon. *Proceedings of the National Academy of Sciences*, 119(47):e2211932119.

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

Hollenbeck, B., Moorthy, S., and Proserpio, D. (2019). Advertising strategy in the presence of reviews: An empirical analysis. *Marketing Science*, 38(5):793–811.

Hu, N., Pavlou, P. A., and Zhang, J. (2017). On self-selection biases in online product reviews. *MIS quarterly*, 41(2):449–475.

Hui, X., Saeedi, M., and Sundaresan, N. (2018). Adverse selection or moral hazard, an empirical study. *The Journal of Industrial Economics*, 66(3):610–649.

Jacobsen, G. D. (2015). Consumers, experts, and online product evaluations: Evidence from the brewing industry. *Journal of Public Economics*, 126:114–123.

Karaman, H. (2021). Online review solicitations reduce extremity bias in online review distributions and increase their representativeness. *Management Science*, 67(7):4420–4445.

Lee, P. S., Chakraborty, I., and Banerjee, S. (2023). Artificial intelligence applications to customer feedback research: A review. *Artificial Intelligence in Marketing*, pages 169–190.

Lee, S., Lee, S., and Baek, H. (2021). Does the dispersion of online review ratings affect review helpfulness? *Computers in Human Behavior*, 117:106670.

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

Li, X., Deng, Y., Manchanda, P., and De Reyck, B. (2024). Can lower expert opinions lead to better consumer ratings?: The case of michelin stars. *Available at SSRN 4275944*.

Luca, M. and Reshef, O. (2021). The effect of price on firm reputation. *Management Science*, 67(7):4408–4419.

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

Meister, M. and Reinholtz, N. (2025). Quality certifications influence user-generated ratings. *Journal of Consumer Research*, page ucaf008.

Mostagir, M. and Siderius, J. (2023). Strategic reviews. *Management Science*, 69(2):904–921.

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

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

Proserpio, D., Troncoso, I., and Valsesia, F. (2021). Does gender matter? the effect of management responses on reviewing behavior. *Marketing Science*, 40(6):1199–1213.

Proserpio, D., Xu, W., and Zervas, G. (2018). You get what you give: theory and evidence of reciprocity in the sharing economy. *Quantitative Marketing and Economics*, 16:371–407.

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

Rossi, M. and Schleef, F. (2024). Quality disclosures and disappointment: Evidence from the academy nominations. *CESifo Working Paper Series*.

Schoenmueller, V., Netzer, O., and Stahl, F. (2020). The polarity of online reviews: Prevalence, drivers and implications. *Journal of Marketing Research*, 57(5):853–877.

Shukla, A., Wang, W., Gao, G. G., and Agarwal, R. (2019). Catch me if you can—detecting fraudulent online reviews of doctors using deep learning. *Available at SSRN: 3320258*.

Sunder, S., Kim, K. H., and Yorkston, E. A. (2019). What drives herding behavior in online ratings? the role of rater experience, product portfolio, and diverging opinions. *Journal of Marketing*, 83(6):93–112.

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

Tsekouras, D., Gutt, D., and Heimbach, I. (2020). The rise of robo-reviews: The effects of chatbot-mediated review elicitation on online reviews. *Available at SSRN 3754200*.

Vellodi, N. (2018). Ratings design and barriers to entry. *Available at SSRN 3267061*.

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

Woolley, K. and Sharif, M. A. (2021). Incentives increase relative positivity of review content and enjoyment of review writing. *Journal of Marketing Research*, 58(3):539–558.

Wu, Y., Ngai, E. W., Wu, P., and Wu, C. (2020). Fake online reviews: Literature review, synthesis, and directions for future research. *Decision Support Systems*, 132:113280.