Reputational Underpricing*

Stepan Aleksenko† and Jacob Kohlhepp‡

December 28, 2025

Abstract

Consumer reviews reflect both product quality and price, with more favorable reviews for a lower-priced product. We study whether this review behavior induces a firm to manage its reputation by underpricing its product below consumers' willingness to pay. We introduce a model with a privately informed firm repeatedly selling its product to rational consumers who learn product quality from past value-based reviews and the current price. We characterize the necessary and sufficient condition for underpricing, which depends on the relative amount of vertical versus horizontal quality differentiation. This condition implies that underpricing need not occur even if the firm is perfectly patient. Reputation management via underpricing, when it occurs, unambiguously benefits consumers.

Keywords: Consumer Reviews, Reputation, Dynamic Pricing, Price Signaling

JEL Codes: D83, D21, (D82)

*First version: February 2023. Previous Version: November 2023. This paper supersedes an earlier version circulated as "Consumer Reviews and Dynamic Price Signaling." We are grateful for insightful comments from Sushil Bikhchandani, Alex Bloedel, Simon Board, Daniel Clark, Felipe Goncalves, Jay Lu, Moritz Meyer-ter-Vehn, Ichiro Obara, Luciano Pomatto, Tomasz Sadzik, Eran Shmaya, and William Zame. We thank participants in the ASSA 2024 Annual Meeting; SEA 2023 Annual Meeting; All-UC Theory Conference; 35th Stony Brook International Conference on Game Theory; South-Western Economic Theory Conference; UCLA, UNC Chapel Hill, and SBU theory proseminars; and Caltech reading group where this work was presented.

†Simon Business School, University of Rochester. Contact: saleksenko@ucla.edu

‡Department of Economics, UNC Chapel Hill. Contact: kohlhepp@unc.edu.

1 Introduction

Tens of millions of distinct sellers operate online, and in 2024, 20.1% of worldwide retail sales are expected to occur online (Snyder 2024, Kiniulis 2024). In this paper, we analyze how these sellers make pricing decisions in the presence of value-based consumer reviews. Value-based consumer reviews are a common feature in many online markets, where consumers tend to leave better reviews for the same product sold at a lower price.[1] The reviews, on the other hand, are crucial to the firms in these markets, determining profitability, demand, and ultimately the survival of these sellers.[2] An open question in such settings is whether or not sellers strategically manage their own reputation through prices.[3]

In this paper, we model how firms make pricing decisions in the presence of reputational incentives driven by value-based consumer reviews. We derive a necessary and sufficient condition for these reputational incentives to cause the firm to price its product below consumers' willingness to pay at the given reputation level, a phenomenon we call *underpricing*. Surprisingly, reputational incentives do not always cause underpricing, even for a highly patient firm. Our condition illustrates that the firm will only engage in underpricing in markets where vertical differentiation is large relative to horizontal differentiation.

Our theoretical contribution is introducing active pricing to the reputation literature as opposed to classic reputation models where passive prices are defined by the market expectations of the product value (e.g., Holmström 1982, Mailath and Samuelson 2001). The technical challenge of our approach is that the firm engages in both price signaling vis-à-vis current consumers (à la Spence signaling) and managing reputation via unobserved past prices vis-à-vis future consumers (à la Holmstrom signal jamming). To resolve this challenge, we introduce a novel Markov Perfect Bayesian equilibrium concept with a continuity refinement, which allows us to characterize equilibria in our model.

In our model, a single long-lived privately informed firm dynamically prices and sells a product of a binary quality (high or low) to an infinite stream of ex-ante identical short-lived consumers, who decide whether to purchase the product based on their expected

1. Luca and Reshef (2021) find among restaurants that a 1% price increase decreases average consumer ratings by 3–5%. Sunada 2024 and Abrate, Quinton, and Pera 2021 find similar evidence for hotel reviews.

2. Anderson and Magruder (2012) find that an increase from 3.5 to 4 stars on Yelp makes restaurants 21% more likely to be fully booked. Similarly, Chevalier and Mayzlin (2006) estimate that changing a single 5-star review to a 1-star review of a book on Amazon is expected to reduce sales by 20 books per week.

3. Fan, Ju, and Xiao 2016 show new sellers manage reputation via prices on Tao Bao. Hui et al. (2016) show that eBay sellers close to the margin reduce their prices to earn a Top Rated Seller badge. Sorokin 2021 shows similar behavior on Steam, where sellers reduce prices to push their rating above a rounding cutoff.

utility from consuming it. We adapt the Poisson perfect good news framework commonly used in the reputation literature to model consumer reviews in a simple and tractable way. Specifically, a consumer leaves a good review only for a high-quality firm AND if the realized net utility (which depends on the true quality of the product, its price, and the consumer's idiosyncratic taste shock) is greater than a given threshold. Future consumers rationally learn about the quality of the product from past reviews and the current price of the product, but they do not observe past prices. This allows us to capture the essence of value-based reviews, i.e., that lowering the price improves (in the FOSD sense) the distribution of perceived quality in the future.

To solve the model, we analyze the firm's trade-off between reputational and myopic pricing incentives. Consumers' ability to observe the past reviews but not the past prices of a product creates a reputation-management channel for the firm selling that product. Specifically, if the high-quality firm lowers the price of its product today, the firm will have a better distribution of consumer reviews today and the future consumers will believe the firm's product to be higher expected quality because they do not observe past prices. The downside of lowering the price today is a lower current profit, either directly, via price, or indirectly, via signaling a lower quality today.

We characterize a necessary and sufficient condition on the primitives under which there is underpricing in equilibrium. Specifically, underpricing occurs if and only if the *adjusted hazard rate* of the taste shock distribution is sufficiently high. If the adjusted hazard rate is low, the reputational incentives do not affect the firm's pricing decisions at all, and different types of the firm pool at the consumers' willingness to pay at all reputation levels in the unique equilibrium that survives the refinement. If the adjusted hazard rate is high, then there are multiple equilibria, and in all of them, both types of the firm underprice only at low reputation levels, and the high-quality firm prices its product *lower* than the low-quality firm. The adjusted hazard rate is higher and therefore underpricing is more likely to occur in product markets where vertical quality differentiation is large relative to horizontal taste differentiation. We illustrate this connection using a specific example in Section 5.

The adjusted hazard rate represents the reputational value of underpricing to the high-quality firm and is roughly equal to the ratio of the marginal reviewers the firm can obtain by underpricing to the inframarginal reviewers the firm has regardless of underpricing. Underpricing occurs if the adjusted hazard rate is sufficiently high, which is more likely if either the density of marginal reviewers is higher or the mass of inframarginal reviewers is lower. A higher density of marginal reviewers increases the value of underpricing since the same sacrifice of myopic profit leads to a larger increase in the arrival of good reviews. A lower mass of inframarginal reviewers increases the value of underpricing since the arrival rate of good reviews is low and it would take a long time to build a good reputation without underpricing.

Our model allows the price posted by firms to signal underlying product quality, and thus our model is a dynamic signaling game. Past work on consumer reviews has typically not confronted price signaling, either because the specific research question did not require doing so or because price signaling is believed to make some models intractable. Signaling is an unavoidable feature of our research question. Even though the assumption of privately informed firms builds asymmetric information into our model upfront, asymmetric information would arise organically in this setting due to unobserved past prices, even if the firms were initially uninformed. This makes signaling a crucial concern in our setting. Assuming ex-ante identical consumers helps us factor out signaling from the firm's problem, allowing us to analyze the trade-off between reputational and myopic incentives for any fixed beliefs first, and incorporate the pricing incentives into the signaling game afterwards.

Our condition for underpricing illustrates that reputation management via underpricing is not a generic feature of marketplaces with consumer reviews. Although a firm producing a high-quality product is always tempted to underprice to improve its reputation, whether the loss of profit today is worth the reputational gain in the future depends on whether the number of reviews gained by underpricing is large relative to the number of reviews the firm would obtain anyway. As a result, for many consumer taste distributions, the unique equilibrium features no underpricing even if the firm is very patient. This occurs because of the infinite time horizon of our model: because the firm always expects to operate tomorrow, there is always a chance it will receive a consumer review without underpricing today, lowering the reputational benefit of underpricing at any given point in time. Such a force does not exist in "static" or finite-horizon models.

With equilibrium pricing behavior in hand, we analyze the welfare implications. Contrary to the standard intuition that review manipulation harms consumers, underpricing, when it occurs, generally benefits consumers. If the firm is myopic and therefore has no incentives to underprice, the firm sets prices equal to consumers' willingness to pay, and the expected consumer surplus is zero. In contrast, when underpricing occurs, the firm sometimes sets prices below consumers' willingness to pay, and the expected consumer surplus is strictly positive.

This implies that a platform designer who wants to maximize consumer surplus would not necessarily want to post the full history of prices or ask for detailed component-wise reviews (i.e., quality and value separately) since either of those policies eliminates the firm's reputational incentives to underprice and thus reduces consumer surplus. Underpricing by the high-quality firm also increases the arrival rate of good news, which speeds up consumer learning[4] and effectively transfers information rent from the low-quality firm to consumers. The welfare consequences of underpricing for the high-quality firm are ambiguous because it sells at lower prices at low reputation levels but is sooner differentiated from the low-quality firm and therefore sooner charges higher prices.

In a broader sense, we develop a setting with a privately informed long-lived player (firm), which chooses short-term actions to strategically manage its reputation, and short-lived players (consumers), who learn about the type of the long-lived player from its current action and signals depending on its past actions and its type. Such a setting could be used for modeling applications other than online retailers, for example, the gig economy. On freelancing platforms such as Upwork, employers are uncertain about a worker's skill except for a performance rating. The worker also submits a proposal to the employer before beginning work, and the worker can choose to develop a more detailed and tailored proposal. Similar to underpricing, submitting a more detailed proposal is costly to the worker but beneficial to the employer and could lead to a better future rating of this worker.

The rest of the paper is organized as follows. We discuss the related literature and our contribution in Section 2. We introduce the model and the equilibrium notion in Section 3, and show the equilibrium characterization and main results in Section 4. We discuss comparative statics in vertical and horizontal product differentiation in Section 5. The robustness of our results and the welfare implications are discussed in Section 6. Section 7 concludes the paper.

2 Literature Review

In this section, we state our contribution relative to three main strands of literature. First, our paper contributes to a vast literature studying reputation effects on the long-run behavior of economic agents, starting with such seminal papers as Kreps and Wilson (1982), Milgrom and Roberts (1982), Holmström (1982), and Fudenberg and Levine (1989).[5] This literature has mainly focused on whether a patient long-lived player will achieve her Stackelberg payoff when she cannot commit and is motivated only by reputational concerns. Our paper is related to the applied papers in this literature, such as Holmström

4. Underpricing makes reviews more informative in the Blackwell sense.

5. There is a vast literature following these papers, including recent papers by Liu (2011), Cisternas (2018), Pei (2020), and Ekmekci et al. (2022). Surveying this literature is outside the scope of this paper.

(1982) and Mailath and Samuelson (2001). A key difference is that instead of assuming passive prices formed by market expectations, we focus on price formation by a strategic long-run player.[6]

It is informative to compare our paper to Holmström (1982). In our paper, the firm's quality plays the role of the manager's productivity in Holmström (1982), and the price plays the role of effort, where a lower price (higher effort) benefits consumers but is costly for the firm (manager). Endogenizing prices allows for positive consumer surplus, which is not possible with passive prices. Consumer reviews are similar to the manager's performance, which is a noisy signal that depends on both productivity and effort. Together with unobserved past actions, this leads to signal jamming incentives in both papers, yet in our paper, does not necessarily lead to underpricing (exerting effort). In our paper, if underpricing happens, the firm underprices (exerts effort) when the market is pessimistic.

The difference is that, in our paper, consumers learn about product quality from current prices and the firm knows its quality, which leads to repeated static signaling on top of dynamic signal jamming and the possibility of separating equilibria in the price-quality signaling game. Figure 1 shows key relationships within both models. Although our model shares aspects of Holmström (1982) and other reputation models, equilibrium dynamics can be starkly different: the firm in our model does not necessarily underprice (exert effort) early in the game, even if the firm is very patient and the state changes over time.

6. Additionally, a contemporary work by Knutsen (2023) analyzes a reputation model with endogenous prices. However, the focus in that paper is on achieving the Stackelberg payoff and effort mimicking, rather than the role of pricing in reputation formation.

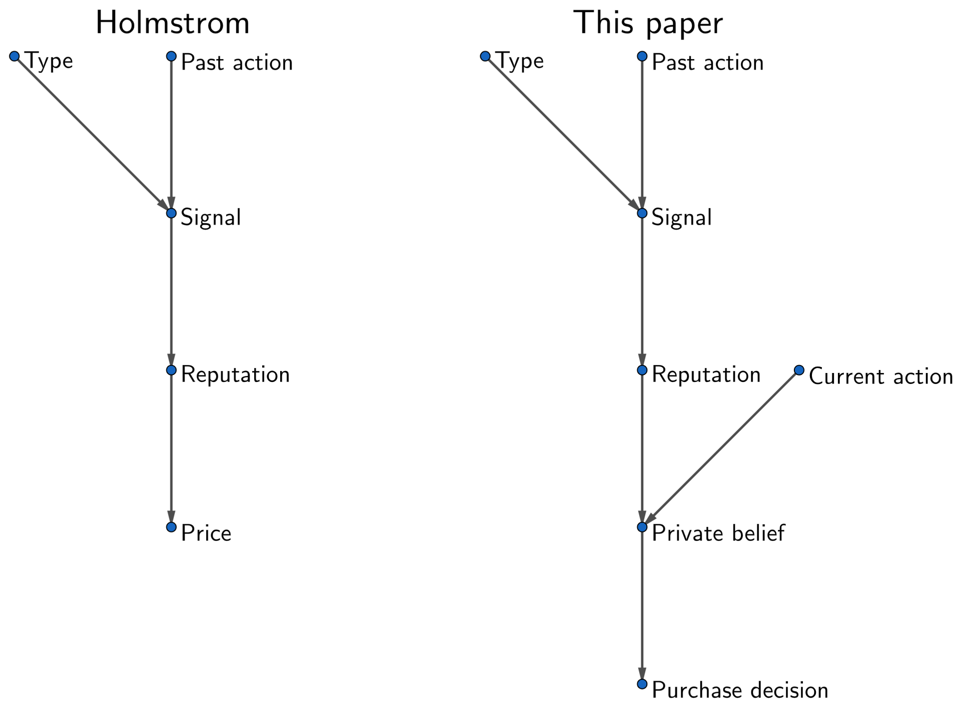

Figure 1: Comparison to Holmström 1982

Second, our paper contributes to a strand of the literature that studies how a firm sets prices over time in the presence of consumer reviews that depend on product prices. Many papers in this literature consider behavioral consumers that use heuristic learning instead of Bayesian updating.[7] Behavioral assumptions on consumers are typically used because "a fully rational consumer would have to solve a dynamic signaling game with rating systems, which is a highly complicated problem" (Carnehl, Stenzel, and Schmidt 2021). Our contribution to this literature is an analysis of fully Bayesian consumers who make inferences from the current price and the full history of product reviews. One important difference between our model and those in the literature is that we allow for static price signaling, and it indeed sometimes occurs in equilibrium.

A small literature considers pricing incentives in the presence of both consumer reviews and Bayesian consumers, similar to our paper. In Huang, Li, and Zuo (2024) and Martin and Shelegia (2021), a sufficiently patient firm underprices its product in the first period because it is the only opportunity to build a reputation. In our model, the firm

7. Shin, Vaccari, and Zeevi (2021) consider consumers who choose a single quality and a single price faced by past consumers that rationalize the observed current average rating. Carnehl, Stenzel, and Schmidt (2021) consider consumers who form beliefs about product quality that rationalize the current rating at the current price. Crapis et al. (2017) consider a firm that sets a price once and for all and consumers who assume all past consumers had the same information. He and Chen (2018) explicitly rule out price signaling.

can be perfectly patient yet never underprice because an infinite horizon allows the firm to postpone reputation building for its future self. In contemporaneous work by Chen, Chen, and Ishida (2025), heterogeneous buyers with coarse anonymous reviews lead to inevitable flash sales/underpricing at low reputation and pooling at any reputation. In contrast, in our paper, underpricing need not happen; both pooling and separation can occur.[8]

Finally, our paper is also related to work on sequential learning through review systems. In line with this literature, one goal of our paper is to understand how and what consumers learn from rating systems. However, we focus on how dynamic pricing undertaken by a forward-looking strategic firm interacts with consumer learning. In contrast, other work in this literature, including Acemoglu et al. (2022) and Koh and Li (2023), focuses on the case when prices are given but other questions are of interest, including how the selection of consumers impacts learning.

3 Model

Firm. A single long-lived firm repeatedly sells a single product. Time is continuous, and the firm posts a price $p_t \in [0, 1]$ at each moment of time $t \in \mathbb{R}_+$.[9] Production is costless, and the future is discounted at rate r. The quality of the firm's product is low or high: $\theta_t \in \{L, H\}$, with a prior probability $q_0 \in (0, 1)$ of being high at $t = 0$. Quality is exogenously and independently redrawn from the same prior distribution at a Poisson rate χ at any t; otherwise, it remains unchanged.[10,11] High quality is normalized to $H = 1$, and low quality is assumed to be strictly positive ($H > L > 0$).

Consumers. The market is composed of a stream of short-lived consumers that arrive at Poisson rate λ. When a consumer arrives, she decides whether to buy a single unit of the product. Consumer utility from purchasing a product of quality θ_t at price p_t is equal to $u_t = \theta_t - p_t + \varepsilon_t$, where the consumer's idiosyncratic ex-post taste shock ε_t is drawn i.i.d. from a symmetric, unimodal, and mean-zero distribution with CDF F_ε and PDF f_ε.[12] This shock is realized only after the good is purchased, i.e., the consumer makes

8. While the flash sales in Chen, Chen, and Ishida (2025) are self-defeating because they only add reviews on the extensive margin by attracting the type of shoppers with uninformative reviews, reputation building via underpricing in our paper is not self-defeating due to the perfect informativeness of extra reviews.

9. We assume non-negative prices, which is also equivalent to requiring the firm's profit to be non-negative. This assumption potentially comes with a loss of generality, but is a highly natural and common.

10. Some model elements, such as Poisson shocks and news structure, are based on previous work, e.g. Board and Meyer-ter-Vehn (2013), Halac and Prat (2016) and Board and Meyer-ter-Vehn (2022).

11. Almost all results in the paper will be for a positive but "small" χ.

12. Normal, logistic, uniform, and type-1 extreme value random variables with zero mean are examples.

the purchase decision based only on the expected quality net of the price. We normalize the consumer's outside option to 0, which implies each consumer purchases the good if her expected utility from consumption of that product is greater than 0, with indifference resolved in favor of purchasing.

Consumer Reviews. Consumer reviews are modeled by perfect good news. Specifically, a consumer leaves a good review at t only if (1) she purchases a high-quality product, $\theta_t = H$, and (2) her realized ex-post utility exceeds threshold $\bar{u}$: $u_t = H - p_t + \varepsilon_t \geq \bar{u}$. We assume $1 \leq \bar{u} \leq \sup\{x \in \mathbb{R} | f_\varepsilon(x) > 0\}$, which is equivalent to requiring that the fraction of consumers leaving reviews is always positive yet is never greater than $1/2$.[13]

Conditional on the product of high quality being purchased, the probability of a good review being generated is decreasing in the price of the product:

$$\Pr(\varepsilon_t > \bar{u} - (1 - p_t)) = 1 - F_\varepsilon(\bar{u} - 1 + p_t).$$

Consumers never leave good reviews for a low-quality product, but as long as the product is truly high quality, consumers are more likely to leave good reviews if their expected utility from the consumption of that product is higher. The firm's review history $h^{t-} = \langle t, \{\tau_1, ..., \tau_n\} \rangle$ is a public history of good review arrival times, τ_i, before time t ($\tau_i < t$) that also tracks the current calendar time t.

Information. A consumer at t observes the review history h^{t-} and the currently posted price p_t of a product, but not past prices, and forms an expectation about the firm's current quality $\tilde{\theta}(p_t, h^{t-}) \in [L, H]$. Then, she buys the product if her expected utility from the consumption of that product is weakly positive: $\tilde{\theta}(p_t, h^{t-}) - p_t \geq 0$. The firm is privately informed about the quality of its product, but consumers are initially uncertain of it. The firm also observes the review history prior to setting a price $p_t = p(\theta^t, h^{t-})$.

Firm's Problem. Even though consumers arrive at discrete times, the Poisson structure implies that expected discounted profit can be expressed as if consumers were arriving as a flow. From the firm's perspective, at any small interval of time (dt), a consumer arrives with probability λdt. Thus the firm's expected profit during dt is equal to $\mathbf{1}_{\{\tilde{\theta}(p_t, h^{t-}) \geq p_t)\}} \lambda p_t dt$, and we can write the firm's expected discounted value at $t = 0$ as an integral:

$$\max_{\{p(\theta^t, h^{t-})\}} \mathbb{E}\left[\int_0^{+\infty} e^{-rt} \mathbf{1}_{\{\tilde{\theta}(p(\theta^t, h^{t-}), h^{t-}) \geq p(\theta^t, h^{t-})\}} p(\theta^t, h^{t-}) \ \lambda dt \right].$$

13. This requirement is consistent with empirical evidence that a very small fraction (1 out of 1000) of consumers leave a review (Hu, Pavlou, and Zhang 2017).

3.1 Model Discussion

Several ingredients in our model allow us to highlight the key economic forces present in the environment without sacrificing tractability. We discuss these in turn.

First, we model the reviews as perfect good news about the product quality with a Poisson arrival rate. This assumption is crucial for tractability and allows us to solve for equilibrium in the presence of fully rational consumers and dynamic price signaling. The Poisson arrival rate implies that the time since the last good review is a sufficient statistic for the entire review history, allowing us to define strategies in continuous time. The perfect good news implies that the reputation jumps to one immediately after receiving a good review, which significantly simplifies the analysis. More general information structures in the Poisson family involve more complicated reputation dynamics that require solving a system of two nonlocal differential HJB equations.

The applied content of our assumption about the review process is that a consumer leaves a review if the product quality is truly good and her overall ex-post net utility, including the idiosyncratic taste shock, exceeds a threshold. This is isomorphic to explicitly modeling the arrival rate of perfect good news as a decreasing function of the price. This provides a tractable microfoundation for why lower prices lead to a "better" distribution of reviews and higher expected reputation in the future (for the high-quality firm). The limitation is that it does not easily map to many observed online review systems, which often feature good and bad reviews, and requires consumers to learn the quality of the product perfectly upon purchasing the product.[14]

Second, for tractability, we model consumers' idiosyncrasies with ex-post rather than ex-ante taste shocks. These taste shocks can be interpreted as after-purchase idiosyncratic experiences, like faster or slower delivery of the product, service in a restaurant on a given date, or horizontal match shocks. Because these shocks manifest after purchase, they do not impact purchase behavior, which yields a tractable bang-bang demand system, yet generates a well-behaved connection between prices and consumer reviews. Under perfect good reviews, the price-reviews connection we wish to model – i.e., that lower price today leads to higher expected consumers' beliefs in the future – would hold with either ex-ante or ex-post taste shocks.

Finally, we model the product quality as exogenous and changing but highly persistent over time, that is $\chi > 0$ but small. This is mainly a technical assumption that guarantees the continuity of value functions at all points. It also guarantees that consumers

14. There is also empirical evidence that consumers tend to leave positive reviews (Carnehl et al. 2022, Hu, Zhang, and Pavlou 2009) and that consumers with more extreme experiences are more likely to leave a review (Schoenmüller, Netzer, and Stahl 2019, Lafky 2014, Marinescu et al. 2021).

do not learn the product quality perfectly, and allows us to produce more realistic price dynamics. An example of such exogenous quality changes can be idiosyncratic changes in the firm's supply chain that impact its product quality, which are observable by the firm but out of its control.

3.2 Equilibrium Concept

We define a pure Markov Perfect Bayesian Equilibrium (MPBE) with the current firm's quality and the public belief about this quality as the Markovian states. We call the public belief about the firm's quality *the firm's reputation* throughout the paper.

Definition 1 *The* ***firm's reputation*** q *is the public belief that the firm's quality is high:*

$$q(h^{t-}) := \frac{\tilde{\theta}(h^{t-}) - L}{H - L} \in [0, 1].$$

The firm's reputation is a single sufficient statistic that summarizes the whole review history and describes the quality distribution in the market for this review history. With a slight abuse of notation, we use $\tilde{\theta}(h^{t-})$ as the consumer's expectation about the firm's quality based only on the public review history h^{t-}, i.e., prior to observing the price of the product p_t. Under the Markov assumption, the firm's prices and consumers' beliefs depend on the review history only via the firm's reputation. Definition 2 formalizes that the firm's reputation is Markovian and the full review history is not necessary for updating the firm's reputation in equilibrium.

For notational clarity, we also introduce an auxiliary function, the **good news arrival rate**, which is the arrival rate of good news for the high-quality product sold at price p, conditional on the high-quality firm selling its product at price p:

$$\lambda_g(p) := \lambda \cdot (1 - F_\varepsilon(\bar{u} - 1 + p)).$$

Now we can formalize the equilibrium concept.

Definition 2 *A pure Markov Perfect Bayesian Equilibrium (MPBE) consists of*

1. *Piecewise continuous*[15] *in q firm's pricing strategies*[16]*:* $p(\theta, q) : \{L, H\} \times [0, 1] \to [0, 1]$*;*

2. *Value functions:* $V(\theta, q) : \{L, H\} \times [0, 1] \to \mathbb{R}_+$*;*

15. This condition guarantees integrability of the price function and differentiability of the value function.
16. We consider only pure strategies, although adding mixed strategies would not change the results significantly.

3. *Consumers' price conjectures:* $\tilde{p}(\theta, q) : \{L, H\} \times [0, 1] \to [0, 1]$;

4. *Consumers' expectations about the firm's quality:* $\tilde{\theta}(p, q) : [0, 1]^2 \to [L, H]$

such that:

(a) *The value functions $V(\theta, q)$ solve Hamilton–Jacobi–Bellman (HJB) equations (1) and (2):*[17]

$$rV(H,q) = \max_p \Big\{ \mathbf{1}_{\{\tilde{\theta}(p,q) \geq p\}} \cdot [\lambda p + \lambda_g(p) \cdot (V(H,1) - V(H,q))] \\ + V_q(H,q) \cdot \frac{dq}{dt} + \chi(1-q_0)(V(L,q) - V(H,q)) \Big\} \tag{1}$$

$$rV(L,q) = \max_p \Big\{ \mathbf{1}_{\{\tilde{\theta}(p,q) \geq p\}} \cdot \lambda p + V_q(L,q) \cdot \frac{dq}{dt} + \chi q_0 (V(H,q) - V(L,q)) \Big\}, \tag{2}$$

where q jumps to 1 at rate $\mathbf{1}_{\{\tilde{\theta}(p(H,q),q) \geq p(H,q)\}} \cdot \lambda_g(p(H,q))$, and otherwise drifts as

$$\frac{dq}{dt} = -\mathbf{1}_{\{\tilde{\theta}(\tilde{p}(H,q),q) \geq \tilde{p}(H,q)\}} \cdot \lambda_g(\tilde{p}(H,q)) \cdot q(1-q) + \chi \cdot (q_0 - q). \tag{3}$$

(b) *Prices $p(\theta, q)$ maximize the right-hand sides of HJB equations (1) and (2).*

(c) *Consumers' expectations about the firm's quality $\tilde{\theta}(p, q)$ are Bayesian for the on-equilibrium-path prices $\{p(L, q), p(H, q)\}$.*

(d) *Consumers' price conjectures are correct, $\tilde{p}(\theta, q) = p(\theta, q)$.*

We explain each component of this definition in turn:

- **Strategies and beliefs.** Definition 2 formalizes that the firm's price and continuation value depend only on the firm's current quality and reputation: $p(\theta_t, q(h^{t-})) = p(\theta^t, h^{t-})$. Consumers' price conjecture for each type of the firm depends only on the current reputation of the firm, and their expectations about the firm's quality depend only on the current reputation of the firm and the price of the product: $\tilde{\theta}(p_t, q(h^{t-})) = \tilde{\theta}(p_t, h^{t-})$.

- **HJB.** Equations (1) and (2) are recursive formulations of the high- and low-quality firms' problems, respectively. The first line of (1) includes the revenue stream as well as the possibility of getting a value jump as the firm's reputation jumps from q to 1 after receiving a good review. It is multiplied by an indicator function because

17. Where $V_q(\theta, q)$ is a left or right derivative $\frac{\partial V(\theta,q)}{\partial q}$ depending on whether $\frac{dq}{dt}$ is negative or positive.

the firm gets the revenue and reviews only if a consumer buys the firm's product at the chosen price. The second line reflects how the future continuation value drifts down without good reviews and might also jump to a different type's value if the quality is redrawn.

- **Law of motion of reputation.** To derive the law of motion for the firm's reputation (3), we need to understand how consumers form their beliefs about the firm's quality based on the firm's review history. The review history process is governed by the prices chosen by the firm. Consumers do not observe past prices, so they use their conjectures about those prices to update their belief about the firm's quality in the absence of good news. Intuitively, consumers form those price conjectures using their understanding of equilibrium and the review history.

  Consumers believe that the high-quality firm charges price $\tilde{p}(H, q)$ and receives a good review at arrival rate $\lambda_g(\tilde{p}(H, q)) \cdot \mathbf{1}\{\tilde{\theta}(\tilde{p}(H, q), q) \geq \tilde{p}(H, q)\}$. We derive the law of motion for the firm's reputation q using the fact that without redrawing the state, it is a martingale and it jumps to 1 immediately after a good review.

  Otherwise, since the time of the last review, the reputation drifts down. An additional term $\chi \cdot (q_0 - q_t)$ represents the mean reversion of the firm's quality because the quality is stochastically redrawn at rate χ. The HJB equations include all these events to calculate the expected continuation value of each type of the firm.

  The discrete-time intuition behind this process is that consumers at period $t = \tau$ recursively reconstruct past price and reputation conjectures, starting from $t = 0$ and incorporating available reviews along the way. They begin by conjecturing the unobserved price at $t = 0$, use it along with any review to update the reputation at $t = 1$, then form a new price conjecture at $t = 1$, incorporate new reviews, and update the reputation at $t = 2$. This recursive process continues until reaching $t = \tau$.

We further introduce an equilibrium refinement to deal with the multiplicity of equilibria. Without refinement, there always is a non-empty interval of reputation levels at which there are multiple pooling equilibria. For instance, there can be a continuum of equilibria with pooling at any given price point supported by pessimistic off-path beliefs at all other prices.[18] Further, many common refinements (e.g., intuitive criterion, D1) are satisfied for all those equilibria and do not help select among them because both types of

18. We show later that this refinement only eliminates unreasonable underpricing and helps us reduce unnecessary equilibria multiplicity, but it does not affect the condition which identifies where underpricing must always happen with or without this refinement.

the firm can have the same preferences over actions for many consumers' belief functions. To remedy this problem, we introduce a continuity refinement.[19]

Assumption 1 ***(Continuity Refinement):*** *For all q, the expectation function $\tilde{\theta}(p, q)$ is continuous in p.*

The continuity refinement requires that small differences in price do not cause large jumps in perceived quality. This requirement is reasonable in the context of online marketplaces such as Amazon: we do not expect consumers to believe a product priced at $99.99 has a significantly different quality than a product priced at $100. *We further impose this assumption throughout the paper and simply call any pure MPBE satisfying the continuity refinement an equilibrium*.

4 Analysis

In this section, we formally define underpricing, characterize equilibria, and derive the necessary and sufficient condition for *underpricing* in equilibria.

To characterize MPBE and the firm's value functions, we must first analyze the static signaling game that occurs at every possible Markov state, i.e. reputation level. Both the firm's prices and consumers' expectations at a given reputation level are endogenous to this signaling game. To analyze this game, we introduce the notion of acceptable prices, which allows us to rewrite the firm's pricing incentives more concisely, given the binary demand function. A set of *acceptable prices* is a set of prices at which consumers purchase the good:

$$\mathcal{P}_q := \{p \in [0,1] | \tilde{\theta}(p, q) \geq p\} \tag{4}$$

Any type of the firm at any reputation level would only choose a price for its product among the acceptable prices $p \in \mathcal{P}_q$, because selling today adds weakly positive revenue to a stream of payoffs and allows a possibility of getting good news, which increases revenue in the future. Thus, for any MPBE, reputation level q, and the value function $V(H, q)$ at this reputation, the firm's pricing incentives, which are captured by the first line of the HJB equation (1) and the first term of (2), can be also expressed as:

$$p(H, q) \in \arg\max_{p \in \mathcal{P}_q} \{\lambda p + \lambda_g(p)(V(H, 1) - V(H, q))\} \tag{5}$$

$$p(L, q) = \max \mathcal{P}_q, \tag{6}$$

19. A similar continuity refinement is used in Gertz (2014).

The right-hand side of (5) only includes the parts of the firm's HJB (1) that depend on the chosen price, because the rest of (1) is exogenous to the signaling game. Equation (5) illustrates the main trade-off in the model. Conditional on choosing a price that signals a quality high enough to sell today, a higher price increases the revenue today but decreases the arrival rate of good news (for the high type), which decreases future payoffs.

The low type's arrival rate of good news is always zero, and therefore the low type's payoff is increasing in the price and $p(L,q) = \max \mathcal{P}_q$. In Section 4, we continue the analysis of this trade-off for the high type and how it affects the signaling equilibrium and MPBE structure.

Finally, the consumers' expectations function must be correct for the equilibrium prices within the signaling game, i.e.,

$$\tilde{\theta}(p,q) = \begin{cases} L, & \text{if } p = p(L,q) \neq p(H,q) \\ H, & \text{if } p = p(H,q) \neq p(L,q) \\ qH + (1-q)L, & \text{if } p = p(L,q) = p(H,q) \end{cases} \tag{7}$$

and the law of motion 3 of the firm's reputation in MPBE can be expressed as

$$\frac{dq}{dt} = -\,\lambda_g(p(H,q)) \cdot q(1-q) + \chi \cdot (q_0 - q). \tag{8}$$

Therefore, MPBE conditions (a)–(d) are satisfied if and only if (A) the firm's prices and consumers' expectations satisfy the equilibrium conditions (5), (6), and (7) of the signaling games at (almost) every q for the given value functions, and (B) the value functions are derived from HJB equations (1) and (2) for given price functions with the law of motion for q given by equation (8).

We define consumers' willingness to pay and underpricing next. Consumers' willingness to pay for a product of the firm with a given reputation q is the maximum price a consumer is ready to pay for this product prior to observing the firm's price choice. Given the consumer's outside option and utility function, the consumer's willingness to pay is equal to the expected value of the product to the consumer, i.e., $\tilde{\theta}(q) = qH + (1-q)L$.

Definition 3

1. *No-underpricing is pricing at the consumers' willingness to pay for a product of the firm with a given reputation q: $p = \tilde{\theta}(q)$.*

2. *Underpricing is pricing below the consumers' willingness to pay for a product of the firm with a given reputation* q*:* $p < \tilde{\theta}(q)$*.*

Before we characterize all possible equilibria, we begin by examining the firm H's objective function at a given q, $\lambda p + \lambda_g(p)(V(H,1) - V(H,q))$ (5), which the firm maximizes within the fixed set of acceptable prices for the given consumer's beliefs: $p \in \mathcal{P}_q$ (4).[20] The firm has three distinct but intertwined incentives: First, there is a *reputational incentive* captured by the combination of the arrival rate of good news with the following continuation value jump: $\lambda_g(p) \cdot (V(H,1) - V(H,q))$. Second, there is the *myopic profit incentive* captured by λp. Third, there is the *signaling incentive* combined with demand considerations which confounds the other two incentives because the firm must choose a price at which the consumer's belief is sufficiently high to purchase the good ($p \in \mathcal{P}_q$).

The following lemma characterizes the high-quality firm's set of possible optimal prices for a fixed set of acceptable prices, which is either the maximal or minimal acceptable price. This follows from the fact that only consumers with extreme tastes leave reviews, and thus further underpricing allows the firm to gain more typical reviewers per $1 of underpricing. Thus, if the high-quality firm decides to underprice, it lowers the price to zero.

Lemma 1 *At any reputation level* q*, the high-quality firm's optimal price is either* 0 *or the maximal acceptable price:* $p(H,q) \in \{0, \max \mathcal{P}_q\}$*.*

Proof sketch. Because the signaling incentive confounds the other two, we temporarily set it aside and consider first how the firm navigates the tension between the myopic and reputational incentives. For any given $(V(H,1) - V(H,q))$,[21] the high-quality firm's objective function (5) is convex in p because the arrival rate of good reviews $\lambda_g(p)$ is convex due to unimodality of the idiosyncratic taste shock ε distribution and given the fact that $\bar{u} \geq 1$. Thus, the optimal solution to (5) is bang-bang: $p(H,q) \in \{0, \max \mathcal{P}_q\}$. ■

The condition under which underpricing occurs in equilibrium crucially depends on a function of primitives that we call the *adjusted hazard rate* of the taste shock distribution F_ε:

$$h_\varepsilon := \frac{(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{1 - F_\varepsilon(\bar{u} - 1 + L) + r/\lambda}. \tag{9}$$

20. Static signaling concerns are captured by the set of acceptable prices, which is evolving with reputation and is endogenous to the exact equilibrium being considered.

21. Given all the future strategies of the firm and consumers fixed, which determines $V(H,q)$ and $V(H,1)$.

The adjusted hazard rate represents the reputational value of underpricing to the high-quality firm. This expression is roughly equal to the ratio of the marginal reviewers the firm can obtain by underpricing (magenta areas in Figures 2a and b) to the inframarginal reviewers the firm already has without underpricing (green areas in Figures 2a and b). More specifically, the numerator is equal to the average density of the marginal reviewers the firm gains per $\$1$ of underpricing at the lowest reputation level, $q = 0$ (underpricing from $p = L$ to $p = 0$). The denominator includes the mass of the inframarginal reviewers the firm has without underpricing at $q = 0$ ($p = L$).

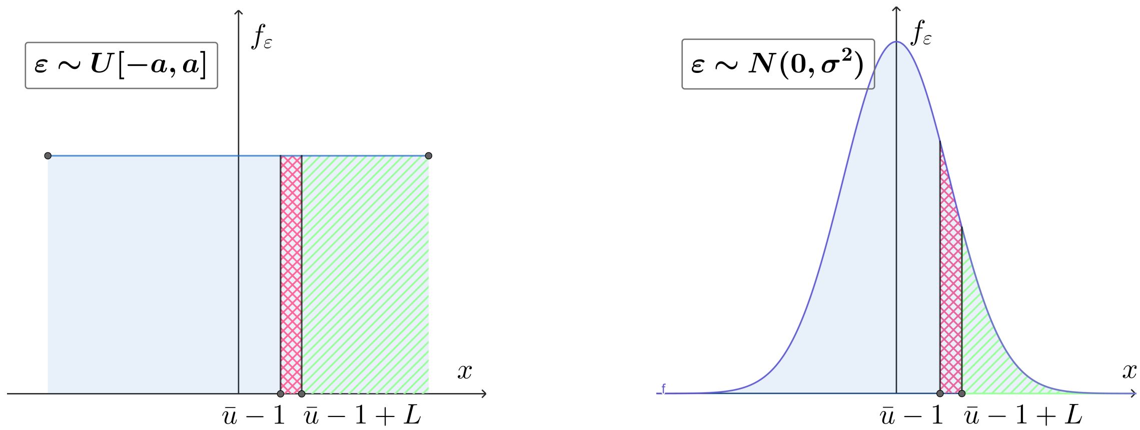

(a) Low adjusted hazard rate

(b) High adjusted hazard rate

Figure 2: Marginal vs. Inframarginal Reviewers

Now, we present the main results of the paper. In Theorem 1, we show that there is underpricing in (every) equilibrium if and only if the adjusted hazard rate is sufficiently high ($h_\varepsilon > \frac{1}{1-L}$). If the adjusted hazard rate is high, then any equilibrium is defined as a partition of the firm reputation interval, with no underpricing at higher reputation levels and underpricing at lower reputation levels. If the adjusted hazard rate is low, there is no underpricing in the unique equilibrium.

Theorem 1 *An equilibrium exists.*

1. *If $h_\varepsilon < \frac{1}{1-L}$, then for small χ,*[22] *no-underpricing is the unique equilibrium:*

$$\forall\, q: \quad p(L, q) = p(H, q) = \tilde{\theta}(q).$$

2. *If $h_\varepsilon > \frac{1}{1-L}$, then for small χ, there must be underpricing at low reputation levels and no*

22. $\exists \varepsilon > 0$, s.t. $\forall\, 0 < \chi < \varepsilon$.

underpricing at high reputation levels in every equilibrium, i.e., $\exists\, 0 < q^* < q^{**} < 1$*:*

$$\forall q < q^* :\ p(L, q) = L,\ p(H, q) = 0,$$
$$\forall q > q^{**} :\ p(L, q) = p(H, q) = \tilde{\theta}(q).$$

Low h_ε corresponds to a case where the density of marginal consumers that the firm can get by underpricing at low reputation levels is low relative to the mass of inframarginal consumers who will leave reviews regardless of underpricing. That implies that the firm will likely get a good review soon even without underpricing, and underpricing does not significantly reduce the expected time until a good review. Therefore, the benefit of underpricing is less than the cost of sacrificing profit today even when the current reputation of the firm is low. Thus, under the continuity refinement, pooling at consumers' willingness to pay is the unique equilibrium (at any reputation level).

High h_ε corresponds to a case where there is a large density of marginal consumers who can be convinced to leave a review after the high-quality firm cuts the price of its product from L to 0 (maximal underpricing at $q = 0$) and a small mass of inframarginal consumers who leave reviews even if the high-quality firm does not cut the price. Whenever this is the case, the high-quality firm with a low reputation exploits the opportunity because it knows that if a customer leaves a review, the product will be revealed for what it truly is: high quality. The high-quality firm with a good reputation does not cut the price of its product, because the gains it obtains from a good review are small when its reputation and profit stream are already high. The high-quality firm with an intermediate reputation may prefer to underprice or not to, depending on consumers' equilibrium beliefs.

In contrast, the low-quality firm, at any reputation level, knows that regardless of how much it cuts the price of its product, no customer will leave a good review. Therefore, it sets the price of its product as high as it can to exploit its current reputation as much as possible. However, at low reputation levels, it cannot price its product at the consumers' willingness to pay, and has to underprice to L, since the consumer understands that positive prices signal low quality. In this way, the two types of the firm engage in separate pricing strategies. Prices are therefore fully informative about quality at low reputation levels. At high reputation levels, both types of the firm pool on the same pricing strategy, and prices are completely uninformative (conditional on reviews).

The proof sketch of Theorem 1 is given next. We omit select technical details in the proof and include them in Appendix Section 8.1. This sketch provides an intuition for the results in the limit ($\chi = 0$, $q = 0$), whereas the full proof proves equilibrium existence and its particular structure ($\forall q$) for positive χ, and discusses the roles of $\chi > 0$ and value functions' continuity in q for our results.

Proof sketch.

1. (No-)Underpricing Dichotomy.

We begin by showing that conditions from Theorem 1 determine whether the high-quality firm prefers to underprice at the lowest reputation level ($q = 0$). First, we rewrite the condition from part 2 of Theorem 1, $h_\varepsilon > \frac{1}{1-L}$, as follows:

$$\lambda \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1)) \cdot \frac{\lambda \cdot (1 - L)}{\lambda \cdot (1 - F_\varepsilon(\bar{u} - 1 + L)) + r} > \lambda L. \qquad (10)$$

We can notice that the equivalent condition is exactly the myopic-reputational trade-off that the high-quality firm faces at $q = 0$ (when considering underpricing). We can decompose the left-hand side of inequality (10) into the following components, which together represent the reputational gain of the continuation value when switching from no-underpricing ($p = \tilde{\theta}(0) = L$) to underpricing ($p = 0$):

1. $\lambda \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))$ is the increase of good news arrival rate;
2. $\frac{\lambda \cdot (1-L)}{\lambda \cdot (1-F_\varepsilon(\bar{u}-1+L))+r}$ is the value of a good review at $q = 0$ (the numerator reflects the profit increase when the price increases from $p = L$ to $p = 1$ after getting a good review; the denominator is equal to the effective discount rate – the sum of the discount rate and the "probability" of the "end of the game", which is getting a good review).

The right-hand side of inequality (10), λL, is the myopic loss of profit when switching from underpricing ($p = 0$) to no-underpricing ($p = L$) at $q = 0$. Since the left-hand side is greater than the right-hand side, the high-quality firm has incentives to underprice and prefers $p = 0$ to $p = L$ at $q = 0$ in every equilibrium.

If the opposite condition from part 1 of Theorem 1, $h_\varepsilon < \frac{1}{1-L}$, is satisfied, then the high-quality firm has no incentives to underprice and prefers $p = L$ to $p = 0$ at $q = 0$ in every equilibrium. This implies that at any $q > 0$, the high-quality firm has no incentive to underprice its product because the reputational incentives get progressively smaller as the firm's reputation increases and the value of a good review decreases $(V(H, 1) - V(H, q))$. Finally, since choosing $p = 0$ is not optimal, the high-quality firm chooses from the remaining set of optimal prices outlined in Lemma 1, which is the maximal acceptable price $p = \mathcal{P}_q$ (at any q).

2. Unique No-Underpricing Equilibrium.

Given that the low type always plays $p(L,q) = \max \mathcal{P}_q$ and the high type plays 0 or $\max \mathcal{P}_q$, there can be two possible equilibria in the signaling game: separating or pooling. If $h_\varepsilon > \frac{1}{1-L}$, then at every q the signaling game equilibrium is pooling ($p(H,q) = p(L,q) = \max \mathcal{P}_q$). With this in hand, we can show that no-underpricing is the unique equilibrium; that is, the high- and low-quality firms pool at the consumers' willingness to pay, $p(L,q) = p(H,q) = \tilde{\theta}(q)$, at any reputation level q.

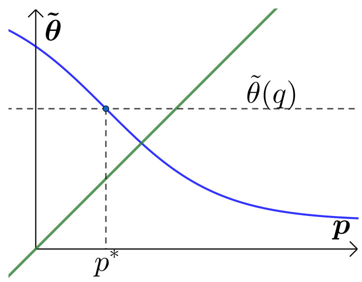

Figure 3: Deviation to a Higher Price

We prove this by contradiction. Assume that at some q, the signaling equilibrium is pooling at $p(H,q) = p(L,q) = p^*(q) < \tilde{\theta}(q)$. By definition of the equilibrium, the expectations $\tilde{\theta}(p,q)$ should be correct on-path (for p^*) and strictly above the price $\tilde{\theta}(p^*,q) > p^*$. By continuity of $\tilde{\theta}(p,q)$, for a small range of prices around p^*, the consumers' expectations about the quality of the product are strictly above the price. Thus, consumers would also buy the good for a price a little higher than the equilibrium price and both types of the firm would prefer to deviate to that higher price (see Figure 3). This contradicts optimality.

Alternatively, the pooling price cannot be above the consumers' willingness to pay either ($p(H,q) = p(L,q) = p^*(q) > \tilde{\theta}(q)$), since consumers would not buy at such price. Thus, at every q, pooling at $\tilde{\theta}(q)$ is the unique equilibrium of the signaling game.

3. Underpricing Equilibrium Structure.

Lastly, we characterize the structure of equilibria when the underpricing condition (12) is satisfied. There can be at most two possible types of equilibria at any q: no-underpricing pooling at $p(L,q) = p(H,q) = \tilde{\theta}(q)$ and underpricing separating $p(L,q) = L, p(H,q) = 0$ (because the consumers are not ready to buy a low-quality good at any price above L, and they are ready to buy it at any price weakly below L). Thus, to determine the possible types of equilibria (underpricing or no-underpricing) at the given reputation level, it suffices to determine the high-quality firm's preferences over three prices: 0, L, and $\tilde{\theta}(q)$.

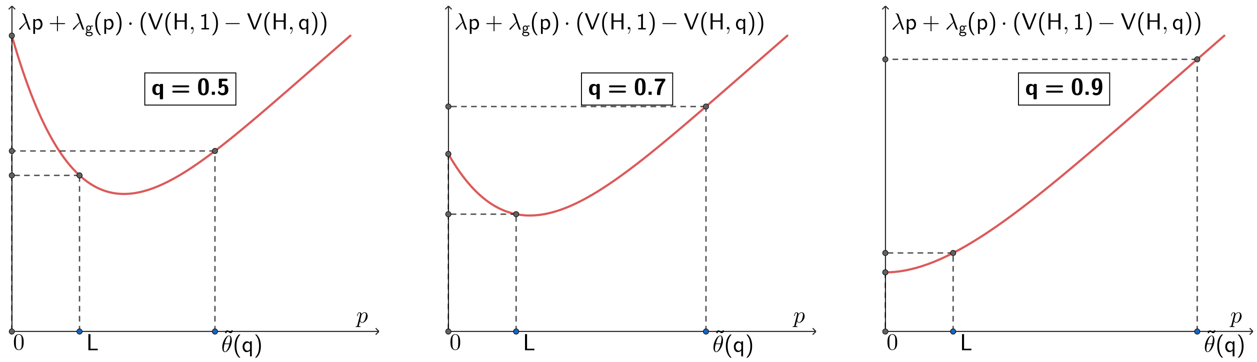

(a) low reputation **(b)** intermediate reputation **(c)** high reputation

Figure 4: Pricing Incentives

First, when the reputation q is high and $V(H,1) - V(H,q)$ is small (since $V(H,q)$ is continuous and increasing in q), the static profit motive dominates the reputational incentive and the right-hand side (5) is monotone increasing in p (Figure 4c). Thus no-underpricing ($p(L,q) = p(H,q) = \tilde{\theta}(q)$) is a unique equilibrium of the signaling game at high q ($q > q^{**}$). Second, when the reputation q is low and the value gap $V(H,1) - V(H,q)$ is large, then the high-quality firm prefers $p = 0$ to all prices in $[0, \tilde{\theta}(q)]$ and underpricing ($p(L,q) = L, p(H,q) = 0$) is the unique equilibrium for all $q < q^*$ (Figure 4a). This happens because the reputational incentive becomes more significant than the static motive in (5). Finally, if q is intermediate, such that both $\tilde{\theta}(q)$ and $V(H,1) - V(H,q)$ are sufficiently large for the high type to prefer $p = \tilde{\theta}(q)$ to $p = 0$ to $p = L$ (Figure 4b), then both no-underpricing and underpricing equilibria are possible in the signaling game.[23] ■

4.1 Discussion

The results derived thus far have implications for consumer welfare, price dynamics, and the occurrence of price signaling. We discuss each in turn.

Welfare Implications of Underpricing. Underpricing benefits consumers at the expense of the low-quality firm. To illustrate this, we compare our model to a myopic benchmark. In this myopic model, the first-order conditions of the low- and high-quality firm are the same, with both desiring to price their products at consumers' willingness to pay, which implies zero consumer surplus.

Whenever the adjusted hazard rate is high, consumer surplus is unambiguously higher under our baseline model than in the myopic benchmark via the direct benefit of paying

23. The full characterization of the possible equilibria for the intermediate reputation levels is neither clear nor economically insightful.

a strictly lower price for a range of reputation levels. Moreover, the high-quality firm sets a lower price to speed up the arrival rate of good reviews. When this happens, the low-quality firm is revealed to be low quality and forced to charge exactly its market value.

Price Dynamics. In our model, a combination of review arrivals, reputation evolution, and (in some cases) strategic underpricing generates nontrivial price dynamics over time. We show these dynamics in Figure 5 using a representative simulated draw of reviews.

1. In the two left panels, we plot price paths when the adjusted hazard rate is high enough so that there is underpricing. We select an equilibrium where underpricing occurs if and only if the reputation is below $\tilde{\theta}(q^*)$. Before a review arrives, the price moves downward over time. Then, when the reputation reaches the threshold $\tilde{\theta}(q^*)$, the price discontinuously jumps downward. For the low-quality firm, the price jumps down to L and remains there. For the high-quality firm, the price jumps down even further to 0 and remains there until a good review arrives.

2. In the two right panels, we plot price paths when the adjusted hazard rate is low enough so that underpricing does not occur. Before a review arrives, the price moves downward over time. Then, for the firm with a high-quality product only, a good review arrives and the price jumps to H. For the low-quality firm, the price gradually converges to a price level that is slightly above L. In this way, the true quality of the firm's product matters for pricing only indirectly via the evolution of reputation.

The taste shock distribution is the primitive of the model which determines whether underpricing occurs. Because preference distributions are typically unobserved, we may want an alternative empirical test for underpricing. Price dynamics provide such a test. When underpricing is occurring, we should observe both upward jumps in prices when good reviews arrive and downward jumps in prices when the firm's reputation is low. When underpricing is not occurring, we should observe only upward jumps in prices at all levels of a firm's reputation.

In many models with uncertain quality, firms price their products low early on in order to build their reputation later, a strategy called introductory pricing.[24] Although there are conceptual differences between our model and many of those in the literature, underpricing in our model can be viewed as a form of introductory pricing. To see this, consider when high-quality products are uncommon but nevertheless a firm begins its life with a high-quality product. Then the firm's reputation will start out low, and in some

24. See for example, Shapiro (1983).

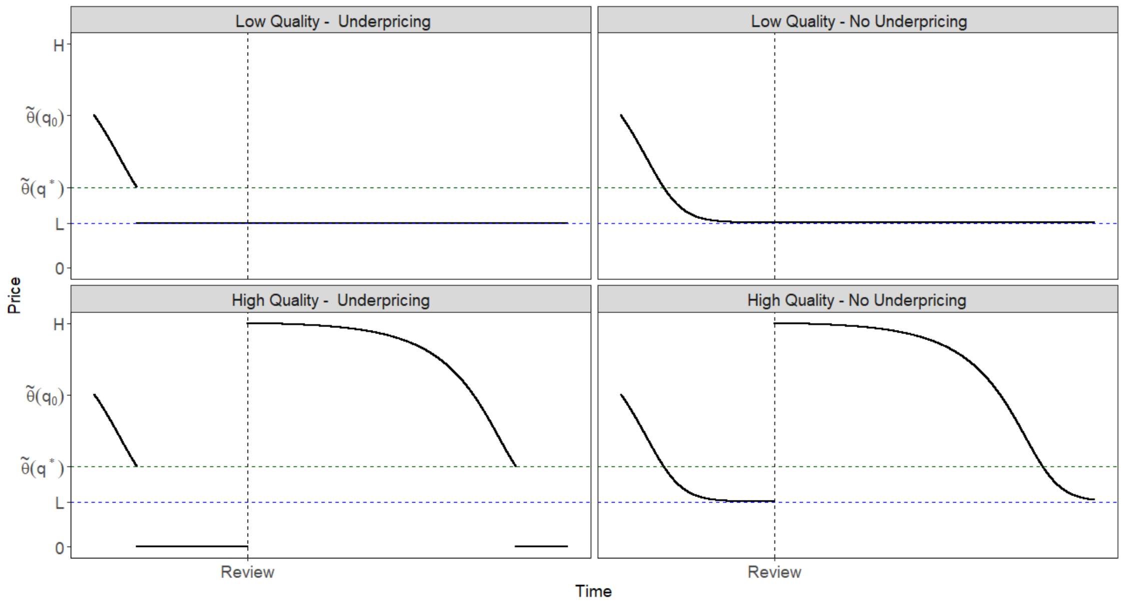

Figure 5: Price Dynamics

cases, the firm will underprice its product initially in order to build reputation, which it will exploit via higher prices when a good review arrives. Crucially, introductory pricing of this type occurs only for certain consumer taste shock distributions.

In contrast to other papers studying consumer reviews manipulation, like Huang, Li, and Zuo (2024) and Martin and Shelegia (2021), introductory underpricing does not always occur in equilibrium, even if the firm is very patient, due to the fully dynamic nature of our model. In Appendix Section 8.4, we show that in a two-period model, whenever the firm is sufficiently patient, underpricing must occur in the first period. This is because a two-period model lacks the procrastination motive that only arises when there is an infinite horizon. In our model, even if the firm is perfectly patient, as long as the $\frac{f_\varepsilon(\bar{u}-1)}{1-F_\varepsilon(\bar{u}-1)} < 1$, underpricing will not occur (for any $r > 0, L \in (0,1)$).

Price Signaling. In our model, price signaling occurs when underpricing occurs, and as a result, the firm with a high-quality product signals it with a low price. This may seem to clash with past work, where high-type firms typically signal with higher prices (Bagwell and Riordan 1991, Milgrom and Roberts 1986). However, to better understand our result, recall that the high-quality firm prices its product lower than the low-quality firm only conditional on reputation. From an unconditional perspective, high-quality firms are more likely to have high reputations at any given moment of time because they have some probability of receiving good reviews, while low-quality firms do not. As a result, there is still a sense in which high-quality firms set higher prices due to their higher average reputation.

Even with this qualification, it is still true that when price signaling occurs, a price of 0 signals high quality, while a price of $L > 0$ signals low quality. To understand why this occurs, we can ask what single crossing supports separation. Recall that there is no difference in the costs of producing high- and low-quality products, and conditional on a firm's reputation, there is no difference in demand. The only source that could generate single-crossing is the review process. This is why the high-quality firm has an incentive to underprice its product, while the low-quality firm does not: using underpricing today can increase the arrival rate of good reviews in the future.

5 Quality and Taste Differentiation

In this section, we show that the strength of horizontal differentiation relative to vertical differentiation determines whether underpricing occurs in equilibrium. Vertical differentiation is captured in the model by the value difference between high- and low-quality products ($H - L$). Horizontal differentiation is captured in the model by the spread of the idiosyncratic utility shock distribution. Greater vertical differentiation increases the total informativeness of reviews for consumers, making underpricing more likely. Greater horizontal differentiation causes consumer review behavior to be less sensitive to price changes, making underpricing less likely.

To establish these results formally, we introduce an additional parameter, $\alpha > 0$, which reflects horizontal differentiation. For any set of model primitives, we will be considering comparative statics over models with the taste shock $\varepsilon' = \alpha\varepsilon$, where higher (lower) α corresponds to higher (lower) horizontal taste differentiation. The amount of vertical quality differentiation is modeled by the difference $H - L$, which is equal to $1 - L$, given the normalization $H = 1$. Thus, various levels of L on the interval $(0, 1)$ correspond to different levels of vertical differentiation, with $L = 0$ leading to the maximum quality differentiation and $L \to 1$ leading to no quality differentiation.

Given the variation in horizontal taste differentiation, we will also assume $\bar{u} = 1$ for this section, which is a necessary and sufficient condition for the high-quality firm to be able to obtain good reviews with a non-vanishing probability if α becomes small. The next proposition establishes the formal results.

Proposition 1 *For any ε and λ/r and $\bar{u} = 1$, there exists a weakly increasing function $\alpha^*(H - L) : (0, 1) \to [0, +\infty)$, s.t.*

1. *If $\alpha < \alpha^*(H - L)$, there is underpricing in equilibrium.*

2. *If $\alpha > \alpha^*(H - L)$, there is no underpricing in equilibrium.*

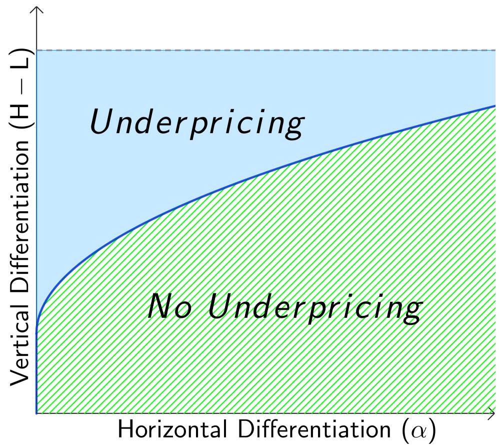

Figure 6: Underpricing: Vertical and Horizontal Differentiation

The proof of Proposition 1 relies on simply showing that the (no-)underpricing condition from Theorem 1 is single-crossing in both α and L.[25] Proposition 1 states that (1) increasing (decreasing) the consumer taste shock spread above (below) some threshold (by multiplying it by α) guarantees that there is no underpricing (underpricing) in equilibrium; and (2) increasing (decreasing) the vertical quality differentiation above (below) some threshold (by changing L) guarantees that there is underpricing (no underpricing) in equilibrium. This is illustrated in Figure 6.

Proposition 1 implies that the adjusted hazard rate is higher and therefore underpricing is more likely to occur in product markets where vertical quality differentiation is more important relative to horizontal taste differentiation. Consider two online markets, as an example: one for probiotic drinks and another for coffee grinders. Probiotic drinks differ in the amount of viable bacteria and therefore their effectiveness. The effectiveness of a probiotic drink is its vertical characteristic, and most consumers care about it much more than about the product's horizontal characteristics, e.g., the taste of the drink. In contrast, horizontal characteristics dominate for coffee grinders: a grinder that is good for a customer who makes cold-brew coffee (coarse grinding) may be almost useless for a customer who makes espresso (fine grinding). While (vertical) quality impacts consumers' utility and reviews, in this case, it is secondary to horizontal aspects. Our results imply that underpricing is more likely to occur in the market for probiotic drinks than in the market for coffee grinders. A firm is less prone to underprice the coffee grinder be-

25. See appendix Section 8.3 for detailed proof.

cause the reviews of it heavily depend on consumer tastes and are less elastic with respect to the price.

Additionally, we can see that the no-underpricing condition is single-crossing in $\frac{r}{\lambda}$. This has two interpretations. First, the more patient firm (lower r) is more likely to underprice. Second, higher λ makes the review channel more prominent and underpricing more likely.

6 Extensions

In this section, we analyze two natural extensions of our baseline model: perfect-bad-news consumer reviews and review-dependent demand.

6.1 Perfect Good and Bad News

In this section, we show that the analysis can be extended to the case with both good and bad reviews, although it is not nearly as tractable as the good-reviews-only setting. Assume that the model is the same as outlined in Section 3, except that consumers can also leave perfect bad reviews. Formally, suppose consumers leave good reviews if the quality is high and the overall utility is above $\bar{u}$, and bad reviews if the quality is low and the overall utility is below $\underline{u} < \bar{u}$.

We further analyze a specific parametric example in which the taste shocks are distributed uniformly between $-a$ and a. A consumer leaves perfect good reviews for a high-quality product if her utility is above $\bar{u} = 1$, or a perfect bad review for a low-quality product if her utility is below $\underline{u} = 0.5$. Additionally, assume that $\lambda = r = 1$, $a = 1$, and $L = 0.5$. In this example, it is guaranteed that any product has a positive probability of getting a review upon consumption for any price between 0 and 1. We continue formulating results for small χ.

In the benchmark model with only good reviews, there would be no underpricing in the unique equilibrium in this example. We further show that in this example with both good and bad reviews there is an equilibrium without underpricing.[26]

First, the condition from part one of Theorem 1 is satisfied and the condition remains sufficient for the high-quality firm to choose the highest acceptable price $\mathcal{P}_q$ at every reputation level.

Second, we need to analyze the low-quality firm's preference as it becomes non-trivial after adding bad reviews. The low-quality firm's value function is now given by:

26. The detailed proof of it is provided in the Appendix Section 8.5.

$$rV(L,q) = \max_{p \in \mathcal{P}_q} \left\{ \lambda p - \lambda_b(p) \cdot (V(L,q) - V(L,0)) + \frac{dq}{dt} \cdot V_q(L,q) \right\},$$

where $\lambda_b(p) = \lambda \cdot \frac{\underline{u}-L+a+p}{2a}$ is linear in the price. Moreover, the low-quality firm has an option of not selling the product at any given moment, which would give it zero profit but also eliminate the possibility of a bad review. Given the convexity of the low-quality firm's objective function $\max_p\{\lambda p - \lambda_b(p) \cdot (V(L,q) - V(L,0))\}^+$, it would never choose to underprice its product but would rather choose to not sell it at all. This would require $0 > \lambda \max \mathcal{P}_q - \lambda_b(\max \mathcal{P}_q) \cdot (V(L,q) - V(L,0))$ at some reputation level q, which we show is not possible in the no-underpricing equilibrium in this example (see Appendix Section 8.5).

Therefore, both the high- and the low-quality firms always prefer $p = \max \mathcal{P}_q$, and pooling at the consumers' willingness to pay (no-underpricing) at any q is an equilibrium for this specific set of parameters. More generally, this illustrates that no-underpricing can be an equilibrium even when the review system is expanded to include both perfect good and perfect bad news.

In our baseline model, perfect good news makes it such that interesting strategic choices and forces operate when the firm has a high-quality product. When underpricing occurs, it is used by a firm with a high-quality product attempting to improve its low reputation. In contrast, under perfect bad news, interesting strategic choices and forces operate when the firm has a low-quality product. In terms of price dynamics, our baseline model produces a smooth downward trend in price punctuated by sudden upward jumps when good reviews are left. Perfect bad news produces the opposite: a smooth upward trend in price punctuated by sudden downward jumps when bad reviews are left.

6.2 Review-Dependent Demand

Assume that the model is the same as outlined in Section 3, except that consumers are more likely to arrive at any given instant when a firm has a higher reputation (i.e., $\lambda(q)$ is increasing). We show that the analysis remains analogous to the baseline model, and we can characterize the condition for underpricing to occur in equilibrium in a similar fashion.

Notice that $\lambda(q)$ does not directly affect pricing incentives and signaling equilibria since $\lambda(q)$ multiplies both the myopic and the reputational parts of the high-quality firm's objective. However, $\lambda(q)$ affects the pricing incentives indirectly through the solution of HJB equation (1), $V(H,q)$. Therefore, the analog of the underpricing condition from

Theorem 1 depends both on the $\lambda(1)$ and $\lambda(0)$ since the reputational incentive at $q = 0$ depends on the value gap $V(H, 1) - V(H, 0)$, which depends on both $\lambda(1)$ and $\lambda(0)$:

$$\frac{(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{1 - F_\varepsilon(\bar{u} - 1 + L) + r/\lambda(0)} > \frac{1}{\frac{\lambda(1)}{\lambda(0)} - L}.$$

The left-hand side of this inequality, h_ε, determines the underpricing/reviews intertemporal trade-off at $q = 0$ and thus depends on $\lambda(0)$. The right-hand side is inversely proportional to the static value of a single good review, which is equal to $\lambda(0) \cdot (\frac{\lambda(1)}{\lambda(0)} - L)$ in this extension instead of $\lambda \cdot (1 - L)$ as in the benchmark model.

Thus, underpricing is more likely under review-dependent demand compared to the benchmark model (which corresponds to $\lambda(0) = \lambda$) because the value of a good reputation now is relatively higher due to high demand and not just high prices.

7 Conclusion

Online marketplaces are populated by firms selling products of uncertain quality to consumers who must make purchase decisions based on the price today and reviews left by past consumers. In this paper, we ask how a firm makes dynamic pricing decisions in such an environment. To answer this question, we develop a model where a long-lived privately informed firm sets its prices over time to sell to Bayesian consumers. Consumers leave reviews based on true quality, the price at purchase, and an idiosyncratic utility shock. The firm engages in both Spencian signaling vis-à-vis current consumers and Holmstromian signal jamming vis-à-vis future consumers in order to manage its reputation. We say a firm underprices if it sets a price below consumers' willingness to pay at a given reputation level.

We characterize equilibria and show that there are two mutually exclusive cases. In the first case, there is a unique equilibrium where underpricing does not occur and both the high- and low-quality firm pool at consumers' willingness to pay at every reputation level. In the second case, a firm of high quality builds up its low reputation by underpricing. Which case occurs depends on the adjusted hazard rate of the consumer taste shock distribution. We show via comparative statics that when horizontal differentiation is large (relative to vertical differentiation), the adjusted hazard rate is low and underpricing does not occur, but when horizontal differentiation is small (relative to vertical differentiation), the adjusted hazard is high and underpricing does occur.

Although firms engage in underpricing in order to bolster their own reputation and increase their profit, the practice actually increases the expected consumer surplus. Modifications to the environment that remove the incentive to underprice reduce consumer surplus to zero. For example, if the history of past prices is made public, the reputational incentive to underprice is eliminated and both types of the firm pool and price at consumers' willingness to pay. Our results therefore imply that a platform designer wishing to maximize consumer surplus should not make prices public.

We propose several directions for future research. First, our paper focuses on perfect good news reviews, but consumer reviews are often best thought of as continuous or multi-valued. It is interesting and practically relevant to understand firm pricing behavior under these alternative review systems. Second, bang-bang demand arises in our model because consumer heterogeneity is ex-post. Allowing for ex-ante heterogeneity is interesting both in terms of realism but also because it changes the firm's pricing problem significantly. In particular, the firm will want to choose interior prices, and the price elasticity of demand will play a key role. Third, consumer review behavior is exogenous in our model, but in many cases, those who leave reviews may have specific motivations. Modeling these motivations explicitly would help us understand how consumers select into the review process, which is important for information aggregation. Finally, we study a single firm, but competition could potentially play a large role both in how consumers leave reviews and in how firms value their reputation. We believe that in many of these settings, reputational concerns will impact firm pricing decisions in a similar fashion to that outlined in this paper, and our model provides a first step at modeling this force.

References

Abrate, Graziano, Sarah Quinton, and Rebecca Pera. 2021. "The relationship between price paid and hotel review ratings: Expectancy-disconfirmation or placebo effect?" *Tourism Management* 85 (August): 104314. ISSN: 0261-5177.

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

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

Bagwell, Kyle, and Michael H Riordan. 1991. "High and declining prices signal product quality." *The American Economic Review,* 224–239.

Board, Simon, and Moritz Meyer-ter-Vehn. 2013. "Reputation for Quality." *Econometrica* 81, no. 6 (November): 2381–2462. ISSN: 1468-0262.

———. 2022. "A reputational theory of firm dynamics." *American Economic Journal: Microeconomics* 14 (2): 44–80.

Carnehl, Christoph, Maximilian Schaefer, André Stenzel, and Kevin Ducbao Tran. 2022. "Value for money and selection: How pricing affects airbnb ratings." *Innocenzo Gasparini Institute for Economic Research Working Paper Series.*

Carnehl, Christoph, André Stenzel, and Peter Schmidt. 2021. "Pricing for the Stars." *SSRN Electronic Journal* (November).

Chen, Chia-Hui, Kong-Pin Chen, and Junichiro Ishida. 2025. "Social learning and strategic pricing with rating systems." *American Economic Journal: Microeconomics* 17 (4): 147–180.

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

Cisternas, Gonzalo. 2018. "Career concerns and the nature of skills." *American Economic Journal: Microeconomics* 10 (2): 152–189.

Crapis, Davide, Bar Ifrach, Costis Maglaras, and Marco Scarsini. 2017. "Monopoly pricing in the presence of social learning." *Management Science* 63 (11): 3586–3608.

Ekmekci, Mehmet, Leandro Gorno, Lucas Maestri, Jian Sun, and Dong Wei. 2022. "Learning from manipulable signals." *American Economic Review* 112 (12): 3995–4040.

Fan, Ying, Jiandong Ju, and Mo Xiao. 2016. "Reputation premium and reputation management: Evidence from the largest e-commerce platform in China." *International Journal of Industrial Organization* 46:63–76.

Fudenberg, Drew, and David K. Levine. 1989. "Reputation and Equilibrium Selection in Games with a Patient Player." *Econometrica* 57 (4): 759–778.

Gertz, Christopher. 2014. "Quality uncertainty with imperfect information acquisition." *Institute of Mathematical Economics Working Paper*, no. 487.

Halac, Marina, and Andrea Prat. 2016. "Managerial attention and worker performance." *American Economic Review* 106 (10): 3104–3132.

He, Qiao Chu, and Ying Ju Chen. 2018. "Dynamic pricing of electronic products with consumer reviews." *Omega (United Kingdom)* 80 (October): 123–134. ISSN: 03050483.

Holmström, Bengt. 1982. "Managerial Incentive Problems: A Dynamic Perspective." Reprinted in Review of Economic Studies 66: 169-182, *Essays in Economics and Management in Honor of Lars Wahlbeck.*

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

Hu, Nan, Jie Zhang, and Paul A Pavlou. 2009. "Overcoming the J-shaped distribution of product reviews." *Communications of the ACM* 52 (10): 144–147.

Huang, Yangguang, Chenyang Li, and Si Zuo. 2024. "Price Signaling and Reputation Building: Evidence from a Consulting Platform." *Available at SSRN 4080448.*

Hui, Xiang, Maryam Saeedi, Zeqian Shen, and Neel Sundaresan. 2016. "Reputation and regulations: Evidence from eBay." *Management Science* 62 (12): 3604–3616.

Kiniulis, Marius. 2024. *How many Ecommerce sites Are There in 2024?* Accessed: 2024-11-21. Markinblog.

Knutsen, Magnus Vage. 2023. "Endogenous Prices in Markets with Reputational Concerns."

Koh, Andrew, and Ricky Li. 2023. "Competition and Social Learning." *Available at SSRN.*

Kreps, David M, and Robert Wilson. 1982. "Reputation and imperfect information." *Journal of economic theory* 27 (2): 253–279.

Lafky, Jonathan. 2014. "Why do people rate? Theory and evidence on online ratings." *Games and Economic Behavior* 87:554–570.

Liu, Qingmin. 2011. "Information acquisition and reputation dynamics." *The Review of Economic Studies* 78 (4): 1400–1425.

Luca, Michael, and Oren Reshef. 2021. "The Effect of Price on Firm Reputation." *SSRN Electronic Journal* (April). ISSN: 1556-5068.

Mailath, George J, and Larry Samuelson. 2001. "Who wants a good reputation?" *The Review of Economic Studies* 68 (2): 415–441.

Marinescu, Ioana, Andrew Chamberlain, Morgan Smart, and Nadav Klein. 2021. "Incentives can reduce bias in online employer reviews." *Journal of Experimental Psychology: Applied.*

Martin, Simon, and Sandro Shelegia. 2021. "Underpromise and overdeliver?-Online product reviews and firm pricing." *International Journal of Industrial Organization* 79:102775.

Milgrom, Paul, and John Roberts. 1982. "Predation, reputation, and entry deterrence." *Journal of economic theory* 27 (2): 280–312.

———. 1986. "Price and advertising signals of product quality." *Journal of Political Economy* 94 (4): 796–821.

Pei, Harry. 2020. "Reputation effects under interdependent values." *Econometrica* 88 (5): 2175–2202.

Schoenmüller, Verena, Oded Netzer, and Florian Stahl. 2019. "The extreme distribution of online reviews: Prevalence, drivers and implications." *Columbia Business School Research Paper,* nos. 18-10.

Shapiro, Carl. 1983. "Optimal pricing of experience goods." *The Bell Journal of Economics,* 497–507.

Shin, Dongwook, Stefano Vaccari, and Assaf Zeevi. 2021. "Dynamic Pricing with Online Reviews." *SSRN Electronic Journal* (November). ISSN: 1556-5068.

Snyder, Kristy. 2024. *E-commerce Statistics for 2024.* Accessed: 2024-08-26. Forbes Advisor.

Sorokin, D. 2021. *Reaching for The Stars: Discounts and Review Tier Transitions in the Video Games Market.* Technical report. working paper.

Sunada, Takeaki. 2024. “Why Do Product Prices Affect Review Ratings? Evidence from a Japanese Online Travel Agency.” *Evidence from a Japanese Online Travel Agency (February 24, 2024).*

8 Appendix [Latest Version of Online Appendix]

8.1 Section 3.2 Derivations

Law of motion of q without redrawing the state:

$$q_t = q_t \cdot \lambda_g(\tilde{p}(q_t))dt \cdot 1 + (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)q_{t+dt}$$

$$\begin{aligned} q_t &= q_t \cdot \lambda_g(\tilde{p}(q_t))dt \cdot 1 + (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)(q_t + dq_t) \\ dq_t \cdot (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt) &= q_t(1-q_t)\lambda_g(\tilde{p}(q_t))dt \\ \frac{dq_t}{dt} &= q_t(1-q_t)\lambda_g(\tilde{p}(q_t)) \cdot \frac{1}{(1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)} \\ \frac{dq_t}{dt} = \lim_{dt\to 0} q_t(1-q_t)\lambda_g(\tilde{p}(q_t)) \cdot \frac{1}{(1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)} &= q_t(1-q_t)\lambda_g(\tilde{p}(q_t)) \end{aligned}$$

Adding mean reversion and a possibility of selling or not selling at $\tilde{p}(q)$ implies the law of motion in equation (3).

HJB for high type (1). We show an intuitive way of deriving this HJB, simplifying some aspects, such as demand, for ease of notation. This HJB can also be derived more formally from the continuation value of the firm introduced in Section 3.1:

$$\begin{aligned} V(H,q) =& \lambda p dt + (1 - rdt)[(1 - \chi(1-q_0)dt)\lambda_g(p)dtV(H,1) \\ &+ (1 - \chi(1-q_0)dt)(1 - \lambda_g(p)dt)V(H, q+dq) + \chi(1-q_0)dtV(L, q+dq)] \end{aligned}$$

$$\begin{aligned} &rV(H,q)dt = \lambda p dt + (1 - rdt)[(1 - \chi(1-q_0)\lambda_g(p)dt(V(H,1) - V(H,q)) \\ &+ (1 - \chi(1-q_0)dt)(1 - \lambda_g(p)dt)dV(H,q) + \chi(1-q_0)dt(V(L,q) + dV(L,q) - V(H,q))] \end{aligned} \tag{11}$$

$$\begin{aligned} rV(H,q) &= \lim_{dt\to 0} \text{RHS}(11) \\ &= \lambda p + \lambda_g(p)(V(H,1) - V(H,q)) + \frac{dV(H,q)}{dt} + \chi(1-q_0)(V(L,q) - V(H,q)). \end{aligned}$$

Incorporating the demand $\mathbf{1}_{\{\tilde{\theta}(p,q)\geq p\}}$ into this equation gives us equation (1).

8.2 Section 4 Proofs

Proof of Lemma 1

We first show that $\lambda_g(p)$ is decreasing and convex. Ex-ante utility, u_t, is bounded by 1 from above in any equilibrium. Thus $u_t - \bar{u}$ is below zero if $\bar{u} \geq 1$. A random variable that is unimodal has a CDF that is concave above its mode (zero for ε); thus F_ε is concave over the relevant domain of $\bar{u} - 1 + p$. Finally, $1 - F_\varepsilon(\bar{u} - 1 + p)$ is a linear increasing function composed within a convex decreasing function and is thus decreasing and convex.

Now, we examine the firm H's objective function, $\lambda p + \lambda_g(p) \cdot (V(H,1) - V(H,q))$ (5), which the firm maximizes subject to choosing an acceptable price: $p \in \mathcal{P}_q = \{p \in [0,1] | \tilde{\theta}(p,q) \geq p\}$ (4). The firm has three distinct but intertwined incentives: First, there is a *reputational incentive* captured by the combination of the arrival rate of good news with the following continuation value jump: $\lambda_g(p) \cdot (V(H,1) - V(H,q))$. Second, there is the *myopic profit incentive* captured by λp. Third, there is the *signaling incentive* combined with demand considerations which confounds the other two incentives because firms must choose a price at which the consumer's belief is sufficiently high to purchase the good ($p \in \mathcal{P}_q$).

Because the signaling incentive confounds the other two, we temporarily set it aside and consider first how the firm navigates the tension between the myopic and reputational incentives. For any given $(V(H,1) - V(H,q))$ (note that $V(H,q)$ and $V(H,1)$ are fixed and determined by the strategies of the firm and consumers at all reputation levels other than q), the high-quality firm's objective function (5) is convex in p because the arrival rate of good reviews $\lambda_g(p)$ is convex due to unimodality of the idiosyncratic taste shock ε distribution and given the fact that $\bar{u} \geq 1$. This intuitively follows from the fact that only consumers with extreme tastes leave reviews and every further one cent of underpricing allows the firm to gain even more typical reviewers. Thus, the optimal solution to (5) is bang-bang: $p(H,q) \in \{0, \max \mathcal{P}_q\}$. ■

Proof of Theorem 1

To proceed with the characterization of the set of all possible equilibria, we require additional notation. First, we need to determine what happens to the firm's reputation on path in any given equilibrium.

Definition 4 *The lowest rating is* $\underline{q} := \sup\{q \in (0,1) | dq/dt \geq 0\}$ *for a given equilibrium price* $\tilde{p}(H,q)$.

Lemma 2 *The lowest rating is well defined, and* $\underline{q} \in (0, q_0)$. *Without good reviews, the firm's rating drifts down until it reaches* $\underline{q}$, *where it stays forever.*

Proof. By definition of $\underline{q}$, $dq/dt < 0$ for $q > \underline{q}$ and q drifts down without good reviews. When q drifts down to $\underline{q}$, it does not drift up or down anymore and stays at $\underline{q}$ without good news. Thus, all reputation levels $q < \underline{q}$ are off-path. ■

Lemma 3 *An MPBE exists.*

Proof. In equilibrium, H's price $\tilde{p}(H,q)$ determines $\underline{q}$, and given $\underline{q}$, equilibrium prices should be defined on $[\underline{q}, 1]$. We are now ready to define equilibrium partition thresholds for a given equilibrium and unify those partitions across all equilibria as q^* and q^{**}:

$$q^* = \inf_{\text{all equilibria}} \inf\{q \in [\underline{q}, 1] | p(L,q) > L\}$$
$$q^{**} = \sup_{\text{all equilibria}} \sup\left\{\{q \in [\underline{q}, 1] | p(L,q) < \tilde{\theta}(q)\} \cup \{\underline{q}\}\right\}.$$

The signaling equilibria characterization suggests that equilibria can be described as a partition with separating, or underpricing, at low reputation levels, pooling at high reputation levels, and multiple equilibria at intermediate reputation levels. By the definition of thresholds q^* and q^{**}, $[0, q^*)$ and $(q^{**}, 1]$ are separating and pooling regions, respectively; that is, in any equilibrium, high and low types always pool at $\tilde{\theta}(q)$ for any $q \in (q^{**}, 1]$ and always separate at $p(L,q) = L,\ p(H,q) = 0$, respectively, for any on-path $q \in [0, q^*)$. For a range of reputations $[q^*, q^{**}]$, the prices can vary across multiple equilibria. Now we need to show the existence of MPBE and characterize the equilibrium dichotomy in terms of these thresholds.

We prove the existence of MPBE by constructing a bi-partition $\hat{q}$-equilibrium with underpricing below some reputation threshold and no underpricing above it:

$$\inf\{q \in [\underline{q}, 1] | p(L,q) > L\} = \sup\left\{\{q \in [\underline{q}, 1] | p(L,q) < \tilde{\theta}(q)\} \cup \{\underline{q}\}\right\} = \hat{q}.$$

We start with a no-underpricing strategy profile $\hat{q} = \underline{q}$ and check if it is an equilibrium from the firm's optimality perspective. If it is, then we found an equilibrium. If it is not, then the high type wants to underprice at low ratings, and we start increasing the underpricing region by increasing $\hat{q}$. Our goal is to find a threshold with which the H's pricing incentives are consistent (there are multiple such thresholds), for instance, at which H is indifferent between $p = 0$ and $p = \tilde{\theta}(\hat{q})$ (this would be just one possible $\hat{q}$-equilibrium). We know that at $\hat{q} = \underline{q}$, H strictly prefers 0. Also, if $\hat{q} = 1$, there are no incentives to underprice at $q = \hat{q} - \varepsilon$ due to the continuity of $V(H,q)$ in q (because of $\chi > 0$), and therefore $V(H,1) - V(H, \hat{q} - \varepsilon) \approx 0$. By the continuity of the value functions in $\hat{q}$ and underpricing incentives in q, there exists a threshold at which H is indifferent between $p = 0$

and $p = \tilde{\theta}(\hat{q})$, and this will be a threshold for which the strategy profile is an equilibrium. Thus, a $\hat{q}$-equilibrium exists. ■

Given that the low type always plays $p(L, q) = \max \mathcal{P}_q$ and the high type plays 0 or $\max \mathcal{P}_q$, there can be two possible equilibria in the signaling game: separating or pooling. Consumer preferences further pin down the nature of each type of equilibrium. In any separating equilibrium, $p(H, q) = 0$ and $p(L, q) = L$ because the consumers are not ready to buy an obviously low-quality good for any price above L, and they are ready to buy any kind of good for any price weakly below L. Thus, the separating equilibrium features underpricing.

1. No-Underpricing ("only if" direction).

The right-hand side of the inequality (10), λL, is the myopic gain of profit when switching from underpricing ($p = 0$) to no-underpricing ($p = L$) at $q = 0$. Since the LHS is smaller than the RHS, underpricing at $q = 0$ cannot be supported in any equilibrium.

The next step is to show that if no-underpricing is an MPBE, it is unique. Here we explicitly rely on small χ. Specifically, when χ is small, $\underline{q}$ is small and $\tilde{\theta}(\underline{q}) \approx L$. Thus the choice between $p = 0$ and $p = L$ is nearly the same as between $p = 0$ and $p = \tilde{\theta}(\underline{q})$. Then by the continuity of the problem in χ, if χ is small enough, H prefers $p = L$ to $p = 0$ because it prefers $p = \tilde{\theta}(\underline{q})$ to $p = 0$.

The previous point implies that the condition separating between cases (1) and (2) of the theorem can be characterized as an explicit condition for when no-underpricing is an MPBE. This condition follows directly from the HJB.

If χ is small, then $V(H, 1) \approx \lambda/r$ and $V(H, 0) \approx (\lambda_g(L) \cdot V(H, 1) + \lambda L)/(\lambda_g(L) + r)$. Then for the high type to prefer $p = \tilde{\theta}(\underline{q}) \approx L$ to $p = 0$, the absolute value of the average slope of λ_g between them $(\lambda_g(1) - \lambda_g(L))/L$ should be smaller than $\lambda/(V(H, 1) - V(H, 0))$. We can rewrite this condition as

$$\frac{\lambda_g(1) - \lambda_g(L)}{L} < \frac{\lambda_g(L) + r}{1 - L}$$

or in terms of the model primitives:

$$\frac{F_\varepsilon(1 - \bar{u}) - F_\varepsilon(1 - L - \bar{u})}{L} < \frac{F_\varepsilon(1 - L - \bar{u}) + r/\lambda}{1 - L},$$

which is equivalent to $h_\varepsilon < \frac{1}{1-L}$.

Finally, when relaxing the monotonicity of $V(H, q)$ assumption, we need to show that the underpricing and no-underpricing regions in case (2) of Theorem 1 are non-empty. This follows from the continuity of $V(H, q)$ in q (even at $q = 1$ due to χ being positive). Specifically, if $h_\varepsilon > \frac{1}{1-L}$, then H prefers to underprice at $\underline{q}$ and at least some small interval of q's around it.

We further show that at any $q > 0$, the high-quality firm has no incentive to underprice its product because the reputational incentives get progressively smaller as the firm's reputation increases. At any reputation level q, switching from $p = 0$ to $p = L$ increases the high-quality firm's objective by $-\lambda \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))(V(H, 1) - V(H, q)) + L$. A high-quality firm gains the most from receiving a good review, $(V(H, 1) - V(H, q))$, when the current value $V(H, q)$ is low. The current value for the high-quality firm $V(H, q)$ is lowest when its reputation q is lowest since the firm can get at least as high of a continuation value at any other reputation level $q' > 0$ by imitating its behavior at $q = 0$ ($p(H, q) = p(H, 0)\ \forall q < q'$). Therefore, if the expression is positive for the lowest reputation ($q = 0$), it is positive for all reputation levels, and deviating from $p = 0$ to $p = L$ remains profitable for the high-quality firm.

Finally, according to Lemma 1, since choosing $p = 0$ is not optimal, the high-quality firm chooses the maximal acceptable price $p = \mathcal{P}_q$.

We further prove that under the continuity refinement any underpricing pooling equilibrium ($p^* < \tilde{\theta}(q)$) cannot be a signaling equilibrium by contradiction. If it is, then by definition of the equilibrium, the expectations $\tilde{\theta}(p, q)$ should be correct on-path (for p^*) and strictly above the price $\tilde{\theta}(p^*, q) > p^*$. By continuity of $\tilde{\theta}(p, q)$, for a small range of prices around p^*, the consumers' expectations about the quality of the product are strictly above the price. Thus, consumers would also buy the good for a price a little higher than the equilibrium price and both types of the firm would prefer to deviate to that higher price (see Figure 3b). This contradicts optimality. Therefore, in the unique equilibrium of the signaling game at any q, both types of the firm always charge the consumers' willingness to pay: $\tilde{\theta}(q)$. Thus, the unique MPBE features no underpricing.

2. Underpricing ("only if" direction).

Similarly to the no-underpricing case, it is sufficient to check whether the high-quality firm prefers to underprice at the lowest reputation $q = 0$, from $p = L$ to $p = 0$, because that is where its incentives to underprice are the strongest due to both the highest value of a good review and the fact the good news arrival rate is the most sensitive with respect to the price between $p = L$ and $p = 0$.

The condition from Theorem 1, part 2 can be rewritten as follows (which is the opposite of the no-underpricing condition):

$$\lambda \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1)) \cdot \frac{1 - L}{\lambda \cdot (1 - F_\varepsilon(\bar{u} - 1 + L)) + r} > L \tag{12}$$

Similarly to the underpricing case, we can notice that this condition reflects the same myopic-reputational trade-off at $q = 0$ under a different reference point (no-underpricing). Specifically, we can decompose the left-hand side of this inequality (12) into components which reflect the reputational gain of the continuation value when switching from no-underpricing ($p = L$) to underpricing ($p = \tilde{\theta}(0) = L$):

1. $\lambda \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))$ is the increase of the good news arrival rate
2. $\frac{1}{\lambda \cdot (1 - F_\varepsilon(\bar{u} - 1 + L)) + r}$ is the value of a good review at $q = 0$ (the numerator reflects the price increase from $p = 0$ to $p = 1$ after getting a good review, and the denominator reflects the effective discount rate – the sum of the discount rate and the possibility of the "end of the game", which is getting a good review without underpricing).

The right-hand side of the inequality (12) L reflects the myopic loss of profit when switching from no-underpricing ($p = L$) to underpricing ($p = 0$) at $q = 0$. Since the LHS of (12) is bigger than the RHS, then no-underpricing at $q = 0$ cannot be supported in any equilibrium. Thus there must be underpricing at some levels of q in every equilibrium.

Lastly, we characterize the structure of equilibria when the underpricing condition (12) is satisfied. We show that all equilibria can be described as a partition with underpricing, at low reputation levels ($q < q^*$) in every equilibrium, no underpricing at high reputation levels ($q > q^{**}$) in every equilibrium, and multiple equilibria at intermediate reputation levels.

To determine the possible types of equilibria (underpricing or no-underpricing) at the given reputation level, we need to determine the high-quality firm's preference over prices. Since we determined that there are only two possible types of equilibria with $\max \mathcal{P}_q \in \{L, \tilde{\theta}(q)\}$, and the high-quality firm's pricing incentives are bang-bang: $p(H, q) \in \{0, \max \mathcal{P}_q\}$, it suffices to determine the high-quality firm's preferences over 0, L, and $\tilde{\theta}(q)$. To simplify the next argument, we assume just for now that $V(H, q)$ is monotone increasing in q.

We start by showing that pooling is a unique equilibrium of the signaling game at high reputation levels. When q is high and $V(H, 1) - V(H, q)$ is small (since $V(H, q)$ is continuous), the static profit motive dominates the reputational incentive and (5) is positive for any p. Therefore, the objective in (5) is monotone increasing in p and H always prefers higher prices to lower prices (see Figure 4c), thus $p(H, q) = \max \mathcal{P}_q$ and pooling at $\tilde{\theta}(q)$ is a unique equilibrium of the signaling game. Because of the continuity of $V(H, q)$, there must be a non-empty interval of high reputations at which this is a unique equilibrium of the signaling game.

Next, we show that separating is a unique equilibrium of the signaling game for low reputation levels. When reputation q is low, the value gap $V(H, 1) - V(H, q)$ is large and the maximum possible equilibrium price $\tilde{\theta}(q)$ is low. Thus the high type prefers $p = 0$ to $p = \tilde{\theta}(q)$ and therefore to all prices in $[0, \tilde{\theta}(q)]$ (see Figure 4a). This happens because the reputational incentive becomes more significant than the static motive in (5). In this case, the high type unambiguously chooses 0 in any equilibrium $\mathcal{P}_q$, and separating $p(H,q) = 0,\ p(L,q) = L$ is a unique equilibrium of the signaling game.

Finally, if q is intermediate, such that both $\tilde{\theta}(q)$ and $V(H,1) - V(H,q)$ are sufficiently large for the high type to prefer $p = \tilde{\theta}(q)$ to $p = 0$ to $p = L$ (see Figure 4b). In this case, both pooling and separating equilibria are possible in the signaling game. In the pooling equilibrium, both types charge $p = \tilde{\theta}(q)$ and have no profitable deviations. In the separating equilibrium, the high type would like to deviate from 0 to $\tilde{\theta}(q)$, but it is not available because $\max \mathcal{P}_q = L$. Thus, neither type has a profitable deviation from $p(H,q) = 0$ and $p(L,q) = L$. ■

8.3 Section 5 Proofs

Proof of Proposition 1. Proposition 1 states that (1) increasing (decreasing) the consumer taste shock spread above (below) some threshold (by multiplying it by α) guarantees that there is no underpricing (underpricing) in equilibrium; and (2) increasing (decreasing) the vertical quality differentiation above (below) some threshold (by changing L) guarantees that there is underpricing (no underpricing) in equilibrium (see Figure 6). To prove these comparative statics, we rewrite the no-underpricing condition ($h_\varepsilon < \frac{1}{1-L}$) from Theorem 1 as

$$\frac{(F_\varepsilon(L/\alpha) - F_\varepsilon(0))/L}{1 - F_\varepsilon(L/\alpha) + r/\lambda} < \frac{1}{1-L} \tag{13}$$

1. Increasing α above a large threshold increases $1 - F_{\alpha\varepsilon}(\bar{u} - 1 + L)$ and decreases density $f_{\alpha\varepsilon}(x)$ for any $x \in [\bar{u} - 1, \bar{u} - 1 + L]$. By decreasing the density at any point of this interval below $\frac{r}{\lambda(1-L)}$, we make $h_{\alpha\varepsilon}$ below $\frac{1}{1-L}$.

2. The adjusted hazard rate h_ε is increasing in λ and $\lim\limits_{\lambda \to 0} h_\varepsilon = 0$. Thus there is a threshold λ^* above which there is no underpricing.

3. The adjusted hazard rate h_ε is decreasing in r and $\lim\limits_{r \to \infty} h_\varepsilon = 0$. Thus there is a threshold r^* above which there is no underpricing.

4. $h_\varepsilon \le \lim\limits_{L \to 0} h_\varepsilon < \infty$. By making $L < L^*$ and $\frac{1}{1-L}$ above this limit, we guarantee no underpricing.

8.4 Two-Period Model

Consider a discrete-time variation of our model with only two periods: $t = 1, 2$, where the first period is the introductory period and the second period is the continuation period. Now $\lambda \in (0, 1]$ is the probability that a single consumer arrives at each period. The discount factor on the second period, $\beta \in (0, +\infty)$, can be larger than one because the continuation period serves as an abstraction of all future payoffs after the introduction period. In the first period, the firm maximizes the expected discounted payoff: $\Pi = \pi_1 + \beta\pi_2$. The rest of the model is the same as in the main benchmark.

Proposition 2 *If the firm is sufficiently patient, i.e.,* $\beta^* > 0$ *s.t. if* $\beta > \beta^*$*, then there is underpricing in the first period in the unique MPBE:* $p_1(L) = L,\ p_1(H) = 0.$

Proof. There is no underpricing in the second period. Thus, after a good review, the high-quality firm charges $p_2(G) = 1$. A high-quality firm selling its product at price $\tilde{p}_1(H)$ in the first period, gets a good review with probability $F_\varepsilon(1 - \bar{u} - p)$. Therefore, without a good review, both types of the firm pool at $p_2(\varnothing) = \frac{q_0 + (1-q_0)L - \lambda q_0 F\varepsilon(1-\tilde{p}_1(H)-\bar{u})}{1-\lambda q_0 F\varepsilon(1-\tilde{p}_1(H)-\bar{u})}$. This is a unique equilibrium in the second period.

In the first period, the signaling equilibrium is either pooling at $(1 - q_0)L$ (features no underpricing) or separating with $p_1(L) = L$ and $p_1(H) = 0$ (features underpricing). We next show that (1) the underpricing equilibrium exists under the condition from the proposition, (2) in the conjectured no-underpricing equilibrium, the high-quality firm has a profitable deviation. The high-quality firm's problem in the first period boils down to

$$\max_{p \in \mathcal{P}_1} \left\{ \lambda p + \beta\lambda^2 F_\varepsilon(1 - p - \bar{u})(p_2(G) - p_2(\varnothing)) + \beta\lambda p_2(\varnothing) \right\}$$

(1) Next, we show that there is an underpricing equilibrium. We construct an equilibrium with $\tilde{p}_1(H) = 0$ and $\max \mathcal{P}_1 = L$. Given the convexity of H's objective function, we next show that a sufficiently patient H prefers $p = 0$ to $p = L$, and therefore to all other

acceptable prices:

$$\begin{aligned}\beta\lambda^2 F_\varepsilon(1-\bar{u})(p_2(G)-p_2(\varnothing)) &> \lambda L + \beta\lambda^2 F_\varepsilon(1-L-\bar{u})(p_2(G)-p_2(\varnothing))\\ &\Leftrightarrow \lambda L < \beta\lambda^2(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-L-\bar{u}))(p_2(G)-p_2(\varnothing))\\ &\Leftrightarrow \beta > \frac{L}{\lambda(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-L-\bar{u}))(p_2(G)-p_2(\varnothing))}\\ &\Leftrightarrow \beta > \frac{L}{\lambda(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-L-\bar{u}))\cdot\frac{(1-q_0)(1-L)}{1-\lambda q_0 F_\varepsilon(1-\bar{u})}}\\ &\Leftrightarrow \beta > \frac{L(1-\lambda q_0 F_\varepsilon(1-\bar{u}))}{\lambda(1-q_0)(1-L)(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-L-\bar{u}))} =: \beta_1\end{aligned}$$

(2) Finally, we show by contradiction that a no-underpricing equilibrium does not exist. We try to construct an equilibrium with $\tilde{p}_1(H)=\tilde{\theta}(q_0)$ and $\max\mathcal{P}_1=\tilde{\theta}(q_0)$. Given the convexity of H's objective function, we next show that a sufficiently patient H prefers $p=0$ to $p=\tilde{\theta}(q_0)$, and therefore H has a profitable deviation:

$$\begin{aligned}\beta\lambda^2 F_\varepsilon(1-\bar{u})(p_2(G)-p_2(\varnothing)) &> \lambda\tilde{\theta}(q_0) + \beta\lambda^2 F_\varepsilon(1-\tilde{\theta}(q_0)-\bar{u})(p_2(G)-p_2(\varnothing))\\ &\Leftrightarrow \lambda\tilde{\theta}(q_0) < \beta\lambda^2(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-\tilde{\theta}(q_0)-\bar{u}))(p_2(G)-p_2(\varnothing))\\ &\Leftrightarrow \beta > \frac{\tilde{\theta}(q_0)}{\lambda(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-L-\bar{u}))(p_2(G)-p_2(\varnothing))}\\ &\Leftrightarrow \beta > \frac{\tilde{\theta}(q_0)}{\lambda(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-\tilde{\theta}(q_0)-\bar{u}))\cdot\frac{(1-q_0)(1-L)}{1-\lambda q_0 F_\varepsilon(1-\tilde{\theta}(q_0)-\bar{u})}}\\ &\Leftrightarrow \beta > \frac{(q+L(1-q))(1-\lambda q_0 F_\varepsilon(1-(q+L(1-q))-\bar{u}))}{\lambda(1-q_0)(1-L)(F_\varepsilon(1-\bar{u})-F_\varepsilon(1-(q+L(1-q))-\bar{u}))} =: \beta_2\end{aligned}$$

Therefore, if $\beta > \beta^* := \max\{\beta_1,\beta_2\}$, then underpricing is a unique equilibrium in the first period. ■

8.5 Section 6.1 Proofs

Lemma 4 *In general, without redrawing the state, the reputation drifts as*

$$\frac{dq_t}{dt} = q_t(1-q_t)(\lambda_b(p(L,t))-\lambda_g(p(H,t))).$$

Short proof via infinitesimal hazards.
Let $q_t=\Pr(\theta=1 \mid \text{no signal by time } t)$ and fix a small $\Delta t>0$. Conditional on no signal by t, the chance of no signal in $(t,t+\Delta t]$ is $1-\lambda_g(p(H,t)\Delta t+o(\Delta t)$ under $\theta=1$ and

$1 - \lambda_b(p(L,t)\Delta t + o(\Delta t)$ under $\theta = 0$. Bayes' rule gives

$$q_{t+\Delta t} = \frac{q_t - q_t\lambda_g(p(H,t))\Delta t}{1 - [q_t\lambda_g(p(H,t))\Delta t + (1-q_t)\lambda_b(p(L,t))\Delta t]} + o(\Delta t).$$

To first order,

$$q_{t+\Delta t} = (q_t - q_t\lambda_g(p(H,t))\Delta t) * (1 + [q_t\lambda_g(p(H,t))\Delta t + (1-q_t)\lambda_b(p(L,t))\Delta t]) + o(\Delta t).$$

$$q_{t+\Delta t} = q_t(1 + [q_t\lambda_g(p(H,t))\Delta t + (1-q_t)\lambda_b(p(L,t))\Delta t]) - q_t\lambda_g(p(H,t))\Delta t + o(\Delta t).$$

$$q_{t+\Delta t} - q_t = q_t(1-q_t)(\lambda_b(p(L,t)) - \lambda_g(p(H,t)))\Delta t + o(\Delta t).$$

Divide by Δt and let $\Delta t \to 0$:

$$\frac{dq_t}{dt} = q_t(1-q_t)(\lambda_b(p(L,t)) - \lambda_g(p(H,t))). \blacksquare$$

Lemma 5 *In the parametric example from Section 6.1, if there is no underpricing, the reputation drift is given by:*

$$\frac{dq}{dt} = \tfrac{1}{2}\lambda q(1-q^2) + \chi(q_0 - q).$$

Proof.

According to Lemma 4, the general updating rule without redrawing the state:

$$\frac{dq}{dt} = q(1-q)\big(\lambda_b(p(L,q)) - \lambda_g(p(H,q))\big).$$

Next, we compute $\lambda_b - \lambda_g$ if there is no underpricing, $p(L,q) = p(H,q) = \tilde{\theta}(q)$. Given $\tilde{\theta}(q) = \frac{1}{2} + \frac{1}{2}q(t)$ and $\varepsilon \sim \text{Unif}[-1,1]$,

$$\lambda_b(q) = \lambda \Pr(1/2 - p(L,q) + \varepsilon < 1/2) = \lambda \Pr(\varepsilon < p(L,q)) = \lambda \tfrac{\tilde{\theta}(q)+1}{2}$$

$$\lambda_g(q) = \lambda \Pr(1 - p(H,q) + \varepsilon < 1) = \lambda \Pr(\varepsilon > p(H,q)) = \lambda \tfrac{1-\tilde{\theta}(q)}{2},$$

so $\lambda_b(q) - \lambda_g(q) = \lambda\tilde{\theta}(q) = \lambda(\frac{1}{2} + \frac{1}{2}q)$.

Then

$$\frac{dq}{dt} = \lambda q(1-q)\Big(\tfrac{1}{2} + \tfrac{1}{2}q\Big) = \tfrac{1}{2}\lambda q(1-q^2).$$

Finally, redrawing the state will have a regular additive effect on the reputation dynamics: $+\chi(q_0 - q)$. $\blacksquare$

Proposition 3 *In the parametric example from Section 6.1, there is a no-underpricing equilibrium.*

Proof.

As a reminder: $\bar{u} = 1,\ \underline{u} = 0.5,\ \lambda = r = 1, a = 1,\ L = 0.5.$

First, we verify the NUP condition from Theorem 1 for the high-quality firm: $h_\varepsilon < \frac{1}{1-L}$. Substituting in the parametric values, we have that:

$$h_\varepsilon = \frac{(F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{1 - F_\varepsilon(\bar{u} - 1 + L) + r/\lambda} = \frac{(F_\varepsilon(0.5) - F_\varepsilon(0)]/0.5}{1 - F_\varepsilon(0.5) + 1} = \frac{(0.75 - 0.5)/0.5}{1 - 3/4 + 1} = \frac{2}{5} < 2 = \frac{1}{1 - 0.5} = \frac{1}{1 - L}$$

$$\lambda_b(p) = \lambda \cdot \frac{\underline{u} - L + a + p}{2a} = \frac{1 + p}{2}$$

$$\lambda_g(p) = \lambda \cdot \frac{a + p - \bar{u} + H}{2a} = \frac{1 - p}{2}$$

Second, we prove $\forall q$ that the low-quality firm never underprices its product. The convexity of the low-quality firm's pricing problem, $\max_p\{\lambda p - \lambda_b(p) \cdot (V(L, q) - V(L, 0))\}$, implies that it has a bang-bang solution. Furthermore, the low-quality firm will never underprice to $p = 0$ since it is strictly worse than not selling the product at all (same revenue, with no bad news). This guarantees that it would never choose to underprice its product but would rather choose not to sell it at all. Thus, we need to show that the low-quality firm prefers selling at $\tilde{\theta}(q)$ to not selling at all. Assuming small χ, this would require:

$$\lambda\tilde{\theta}(q) - \lambda_b(\tilde{\theta}(q)) \cdot (V(L, q) - V(L, 0)) \geq 0$$

or equivalently

$$V(L, q) \leq V(L, 0) + \frac{\lambda\theta(q)}{\lambda_b(\tilde{\theta}(q))} = \frac{\lambda L}{r} + \frac{\lambda(q + (1 - q)L)}{\lambda \cdot \frac{a + \underline{u} - L + (q + (1 - q)L)}{2a}} = 0.5 + \frac{0.5 + 0.5q}{(1.5 + 0.5q)/2} \tag{14}$$
$$= 0.5 + 2 \cdot \frac{0.5 + 0.5q + 1 - 1}{(1.5 + 0.5q)} = 2.5 - \frac{4}{3 + q}.$$

$$V(L, q) \leq V(H, 1) \leq 1 < \frac{7}{6} = 2.5 - \frac{4}{3} \leq 2.5 - \frac{4}{3 + q}$$

■