Consumer Reviews and Dynamic Price Signaling*

Stepan Aleksenko† and Jacob Kohlhepp‡

[Click Here for Latest Version]

November 13, 2023

Abstract

Pricing decisions are crucial for managing a firm's reputation and maximizing profits. Consumer reviews reflect both the product quality and its price, with more favorable reviews being left when a product is priced lower. We study whether such review behavior can induce a firm to manipulate the review process by underpricing its product, or pricing it below current consumers' willingness to pay. We introduce an equilibrium model with a privately informed firm repeatedly selling its product to uninformed but rational consumers who learn about the quality of the product from past reviews and current prices. We show that underpricing can arise only when the firm reputation is low and then only under a specific condition on consumers' taste shock distribution, which we fully characterize. Rating manipulation unambiguously benefits consumers, because it operates via underpricing.

Keywords: Consumer Reviews, Reputation, Dynamic Pricing, Price Signaling
JEL Codes: D83, D21, (D82)

*We are grateful for insightful comments from Moritz Meyer-ter-Vehn, Simon Board, Alex Bloedel, William Zame, Sushil Bikhchandani, Jay Lu, Tomasz Sadzik, Ichiro Obara, Felipe Goncalves, and Luciano Pomatto. We thank participants in the Stony Brook Game Theory Conference; South-Western Economic Theory Conference; UCLA, UNC Chapel Hill, and SBU theory proseminars; and CalTech reading group where this work was presented.
†Department of Economics, UCLA. Contact: saleksenko@ucla.edu
‡Department of Economics, UNC Chapel Hill. Contact: kohlhepp@unc.edu.

1 Introduction

Consumer reviews often reflect not just product quality but also value for the price. Current consumers leave better reviews for higher-quality goods than for low-quality ones, but the reviews are worse if the same good is sold at a higher current price (Abrate, Quinton, and Pera 2021, Luca and Reshef 2021). Future consumers then make purchase decisions based on a firm's reputation for quality, which is shaped by these reviews (Chevalier and Mayzlin (2006)). In settings where consumers observe past reviews but not past prices of a product (such as Amazon, Google Maps, Airbnb, etc.), the firm selling the product may strategically manipulate its own reputation through prices, a phenomenon that is supported by empirical evidence (Sorokin 2021).

In this paper, we propose a novel model to analyze how firms make pricing decisions in the presence of reputational incentives driven by price-dependent 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, a phenomenon we call underpricing.

In our model, an infinite stream of short-lived consumers decide whether to purchase the product based on their expected utility from consuming it. After purchasing the product, a consumer can leave a review for it depending on the quality of the product and her realized utility, which additionally depends on the current price of the product and the consumer's idiosyncratic taste shock. Future consumers rationally learn about the quality of the product from past consumer reviews and the current price of the product, but they do not observe past prices. Our technical contribution is solving this reputation model with observable current actions (prices) and dynamic price-signaling.

To solve the model, we analyze the firm's trade-off between the reputational and the myopic pricing incentives. Consumers' ability to observe past reviews but not past prices of a product creates a reputation-management channel for the firm selling that product. Specifically, if the firm lowers the price of its product today, the firm will receive better consumer reviews today and build a better future reputation. Future consumers will have higher beliefs about the quality of the firm's product because they cannot distinguish whether the firm's better reviews are due to its higher product quality or lower 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 and lowering the current demand for that product.

We show that the consumer taste shock distribution, specifically, its *adjusted hazard rate*, determines whether underpricing occurs in equilibrium. For a wide range of primitives, including the case when consumer taste shocks are distributed uniformly, there is a unique equilibrium where the firm prices its product at consumers' willingness to pay at all reputation levels and never tries to build its reputation via underpricing. For another range of primitives, underpricing occurs only at lower reputation levels. Underpricing is unlikely to occur when the taste shock distribution is sufficiently dispersed. Additionally, we show that the reputation-management channel, and therefore underpricing, is more prominent when consumers arrive more frequently.

Although many insights can be extended to other review processes, we focus on the case when the reviews are modeled as perfect good news. Consumers leave only good reviews, and they can leave a review only for a high-quality product.[1] Each consumer's probability of leaving a good review is increasing in the utility delivered net of the price and is therefore decreasing in the price. We model this feature by introducing an explicit review rule: the consumer leaves a review if the overall ex-post utility depending on the price of the product and the consumer's idiosyncratic (ex-post) taste shock exceeds the given threshold.[2] The taste shocks are i.i.d. mean zero distributed and realized only after the decision to purchase the product. These taste shocks can be interpreted as after-purchase idiosyncratic experiences from consuming the product. Examples include faster or slower delivery of the product or service in the restaurant on a given date, or horizontal matching shocks unobserved prior to purchase.

To build intuition and show that underpricing need not occur, we first consider the case when ex-post taste shocks are distributed uniformly. We show that reputational incentives are always dominated by the static profit motive. At lower reputation levels, the firm is forced to charge low prices, whether via underpricing or not. Under the uniform shock distribution, the reputational incentives become insensitive to price reductions relative to the overall speed of reputation building. Therefore, future reputation-building crowds out any incentives to underprice a product today. This means that both the high-quality and low-quality types of firm have the same pricing incentives, so price signaling never occurs. Thus, both types of firm pool at consumers' maximum willingness to pay at any rating and underpricing does not occur.

More generally, for an arbitrary taste shock distribution, there do not exist equilibria where firms pool and price their products below consumers' willingness to pay. Either there is a unique equilibrium where both firm types pool at consumers' willingness to pay or the two types separate at low reputation levels and pool at high reputation lev-

1. There is empirical evidence for mostly positive reviews on online platforms including AirBnB (Carnehl et al. 2021) and Amazon (Hu, Zhang, and Pavlou 2009).

2. This implies only people with extreme experiences select into leaving a review, an idea with empirical support (Schoenmüller, Netzer, and Stahl 2019, Lafky 2014, Marinescu et al. 2021).

els. In the second case, for a given reputation level below a threshold, the high-quality type of firm ("high type") prices its product below that of the low-quality type of firm ("low type"). Then, prices are also fully informative about the quality of the product at low reputation levels. We derive necessary and sufficient conditions on the ex-post taste shock distribution and the consumer arrival rate that determine which type of equilibrium arises. Underpricing does not occur in equilibrium if the density of marginal consumers that the firm can win by using underpricing is low relative to the mass of consumers who will leave reviews regardless of whether underpricing is used.

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, makes consumers better off relative to a myopic benchmark for two main reasons. First, underpricing directly increases the consumer surplus. Second, underpricing by the high-quality firm increases the arrival rate of good news, which speeds up consumer learning and effectively transfers information rent from the low-quality firm to consumers. The welfare consequences for the high-quality firm are ambiguous because it cuts prices at low reputation levels but is sooner differentiated from the low-quality firm.

The rest of the paper is organized as follows. We discuss the related literature and our contribution to it in Section 2. We introduce the model in Section 3 and show the equilibrium characterization and main results in Section 4. The robustness of our results to a number of extensions and the welfare implications are discussed in Section 5. Section 6 concludes the paper.

2 Literature Review

Our paper is related to a set of papers that study how a firm sets prices over time in the presence of a consumer review system. In this literature, papers typically consider two types of reviews: those that do not depend on product prices (He and Chen (2018)) and those that do. Our model belongs to the second category. Many papers in this literature consider consumers which are behavioral rather than Bayesian. 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.

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 rational consumers that make Bayesian inferences from the current price and full history of reviews of a product. One important difference between our model and those in the literature is that we allow for static price signaling, and it occurs in equilibrium. Of the four papers mentioned in the previous paragraph, three explicitly assume price signaling does not occur.

A small literature (Huang, Li, and Zuo 2022, Martin and Shelegia 2021) considers pricing incentives and learning in the presence of both consumer reviews and Bayesian consumers, similar to our paper. One key difference between the models in these papers and ours is that they consider a single period of building reputation via reviews, while we consider an infinite number of periods. This is one reason underpricing is more common in the models in this literature: there is only one attempt at building reputation, while underpricing generically need not occur in our model, because the firms might simply wait for their future selves to build their reputation.

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. (2017) 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 into purchase impacts learning.

Our paper contributes to the literature on reputation management and is most related to Holmström (1999) and Board and Meyer-ter-Vehn (2013). We discuss these papers using the language of our model, for clarity. In Holmström (1999), the firm has a quality that is fixed but initially unknown by everyone. The firm makes costly effort choices that impact the utility delivered to consumers. Importantly, while the full history of utility delivered is observed by consumers, effort choices both today and in the past are unobserved. In Board and Meyer-ter-Vehn (2013), the firm can make costly effort choices that determine quality when it is redrawn. Importantly, while the firm observes its quality and its effort choices, consumers do not; instead, they observe signals about quality via a process that the firm cannot influence.

Our paper is also related to the strand of reputation literature with observable actions (such as the seminal paper by Fudenberg and Levine 1989, and most closely to Pei 2020). The main differences between our paper and this literature are that in our model (1) we do not have the commitment type, (2) current rather than past actions are observable, (3) the long-lived player is not necessarily patient, (4) we analyze a different class of stage games.

In our model, the firm cannot influence the quality of its product; instead, the firm makes costly effort choices (prices) that impact the utility delivered to consumers. Reviews are left when the firm is high quality, but the firm can influence the arrival rate by changing the prices of its product. Importantly, consumers observe the full history of reviews and the current price, but not the full history of prices. Our paper is like Holmström (1999) (and unlike Board and Meyer-ter-Vehn 2013) in that the firm can influence learning about quality but not quality itself. Our paper is like Board and Meyer-ter-Vehn (2013) (and unlike Holmström 1999) in that the firm knows its quality, but consumers do not. Our paper is different from both in that the firm takes an observable action (price). Similar to Board and Meyer-ter-Vehn (2013), we find that when incentives for underpricing exist, equilibria take a partition form with investment (underpricing) at low reputation levels and shirking (pricing at consumers' willingness to pay) at high reputation levels. Different from Board and Meyer-ter-Vehn (2013), we find that for a wide class of primitives, there is a unique equilibrium where investment never occurs.

Finally, our paper is related to a literature where individuals signal by choosing an information structure. A few examples are Rodríguez Barraquer and Tan (2022) (tasks on the job), Degan and Li (2021) (precision of information), and Daley and Green (2014) (grades across levels of education). In our case, the firm's choice of price influences the arrival rate of future information. In some cases, this gives rise to a form of endogenous single crossing where high-quality firms separate themselves from low-quality firms by choosing to price their products low, to invest in reputation. The biggest difference from our model is that we assume a repeated dynamic structure where the choice of signal structure is observed today, but its realization is observed in the future. For this reason, (semi-)separating equilibria are less common in our model.

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 $t \in \mathbb{R}_+$. Production is costless and the future is discounted at rate r. The quality of the firm's product θ_t 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 redrawn from the same prior distribution at a Poisson rate χ.[3] High quality

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

is normalized to $H = 1$, and low quality is assumed to be strictly positive ($L > 0$).

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.[4] This shock is realized only after the good is purchased, i.e., the consumer makes 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 only leaves a good review at t if (1) she purchases a high-quality product, $\theta_t = H$, and (2) her realized ex-post utility exceeds threshold $\bar{u}$: $u_t = H - p_t + \varepsilon_t \geq \bar{u}$. This implies that 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 consumption is higher. The firm's review history $h^{t-} = \{t, \tau_1, ..., \tau_n\}$ is a public history of good review arrival times 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 the $t = 0$ as

4. Normal, logistic, uniform, and type-1 extreme value random variables are examples.

an integral:

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

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.

The product quality in our model is exogenous and changing but highly persistent over time: $\chi > 0$ but small. This is mainly a technical assumption that guarantees the continuity of value functions at all points. It also guarantees that consumers do not learn the product quality perfectly and allows us to produce more realistic price dynamics. Intuitively, there may be some idiosyncratic changes in the firm's supply chain that impact its product quality that are observable by the firm but out of its control.

A binary type with perfect good news is a tractable benchmark model in the reputation literature. To model the idea that consumers are reciprocal and more likely to leave good reviews when the product is priced lower, we introduce an additional condition that a consumer leaves a good review only if the overall utility, including an idiosyncratic taste shock, is above some threshold. This model provides a micro-foundation and is also isomorphic to modeling the arrival rate function of perfect good news explicitly as a primitive. The ex-post taste shock ε_t captures unexpected differences in the consumer's individual experience of the product, such as faster or slower delivery of the product or service in the restaurant.

Consumers in our model are fully Bayesian in the way they update their beliefs based on the review history and the current price. However, they do not observe the full history of past prices. These ingredients imply that the firm can invest in its reputation by choosing off-path prices, and because the different types of firm may have different incentives, price signaling becomes a possibility. In this way, we allow for price signaling to occur in any given period instead of assuming it away ex-ante.

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, which we call the firm's reputation, as the Markovian states.

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. Under the Markov assumption, the firm's prices and consumers' beliefs depend on the review history only via the firm's reputation. Definition 2 below formalizes that the firm's reputation is Markovian and the full review history is not necessary for updating the firm's reputation in equilibrium.

To proceed with defining the equilibrium concept, we first 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 that the high-quality firm sells its product at price p if a consumer arrives:

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

Now we can formalize the equilibrium concept.

Definition 2 *A pure Markov Perfect Bayesian Equilibrium (MPBE)*[5] *consists of*

1. *Piecewise continuous*[6] *in q firm's pricing strategies: $p(\theta, q) : \{L, H\} \times [0, 1] \to [0, 1]$;*
2. *Value functions: $V(\theta, q) : \{L, H\} \times [0, 1] \to \mathbb{R}_+$;*
3. *Consumers' belief about prices: $\tilde{p}(\theta, q) : \{L, H\} \times [0, 1] \to [0, 1]$;*
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 the Hamilton-Jacobi-Bellman (HJB) equations (1) and (2).*[7]

5. Maskin and Tirole 2001
6. A condition the guarantees its integrability and differentiability of the value function.
7. 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.

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

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

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

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

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

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

(d) *Consumers' belief about price is correct $\tilde{p}(\theta,q) = p(\theta,q)$.*

We now explain each component of this definition:

- **Strategies and beliefs.** Definition 2 formalizes that the firm's price and continuation value depend only on its current quality and reputation: $p(\theta_t, q(h^{t-})) = p(\theta^t, h^{t-})$. Consumers' belief about prices for each type of 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 low- and high-quality firms' problems. The first line of (1) includes the revenue stream as well as the possibility of getting a value jump as the reputation jumps from q to 1 after receiving a good review. It is multiplied by an indicator function because 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 the 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 beliefs about those prices in order to update their belief about the firm's quality in the absence of good news. Intuitively, consumers "fill in" unobserved past prices 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 firm.

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 firm at any reputation level would only choose a price for its product among the acceptable prices $p \in \mathcal{P}_q$ (4), 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.[8] Thus the law of motion of the firm's reputation in equilibrium can be expressed as

$$\frac{dq}{dt} = -\lambda_g(p(H,q)) \cdot q(1-q) + \chi \cdot (q_0 - q) \tag{5}$$

There is a signaling game at every reputation level that endogenously defines both the firm's prices and consumers' expectations, which prohibits us from solving for full HJB equations paths for a given consumers' expectation function. Thus, to derive value functions and MPBE, we first need to define and solve an **auxiliary signaling game** as a static version of our model played at a single moment for a given reputation level, with the firm's payoffs derived from the HJB equations (1) and (2).

To be a signaling equilibrium at q for a given value function $V(H,q)$, the firm's prices

8. We show in online appendix Section 7.5 that $\forall q : \quad V(H,1) \geq V(H,q)$

must be optimal

$$p(H,q) \in \arg\max_{p \in \mathcal{P}_q} \{\lambda p + \lambda_g(p)(V(H,1) - V(H,q))\} \tag{6}$$
$$p(L,q) = \max \mathcal{P}_q$$

and the consumers' expectations function is correct for the equilibrium prices, i.e.,

1. $\tilde{\theta}(p(L,q),q) = L$ and $\tilde{\theta}(p(H,q),q) = H$, if $p(L,q) \neq p(H,q)$ OR
2. $\tilde{\theta}(p(L,q),q) = \tilde{\theta}(p(H,q),q) = qH + (1-q)L$, if $p(L,q) = p(H,q)$.

Equation (6) includes only the parts of the firm's HJB (1) that depend on the chosen price and therefore reflect the same pricing incentives for the firm. Equation (6) 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 payoff is increasing, and $p(L,q) = \max \mathcal{P}_q$. We continue the analysis of this trade-off for the high type and how it affects the signaling equilibrium and MPBE structure in Section 4.

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

4 Analysis

In this section, we derive necessary and sufficient conditions for *underpricing* in equilibria: pricing below the consumer's expected value of the product at the given reputation of the firm $\tilde{\theta}(q) := qH + (1-q)L$.

First, we analyze the case with a uniform taste shock distribution, which allows us to illustrate the main forces in the model and introduce a continuity equilibrium refinement. In this case, we show that underpricing will never arise in equilibrium and different firm types pool at consumers' willingness to pay $\tilde{\theta}(q)$.

Second, we solve the general case of the model and characterize the necessary and sufficient condition under which underpricing occurs in equilibrium. A key element of

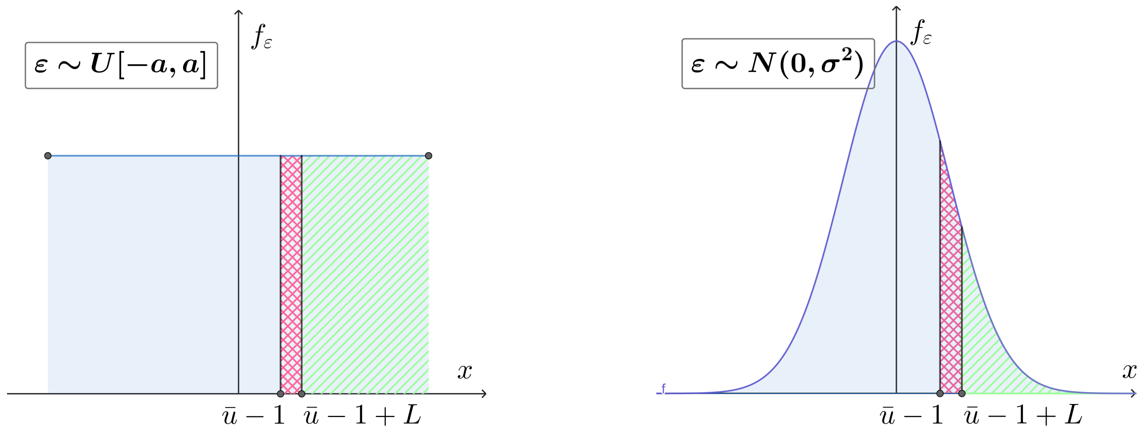

(a) Low Adjusted Hazard Rate

(b) High Adjusted Hazard Rate

Figure 1

the model is the taste shock CDF F_ε and a related primitive object that we call the *adjusted hazard rate*:

$$h_\varepsilon := \frac{\lambda(F_\varepsilon(\bar{u}-1+L) - F_\varepsilon(\bar{u}-1))/L}{\lambda(1-F_\varepsilon(\bar{u}-1+L)) + r} \tag{7}$$

The adjusted hazard rate represents the value of underpricing to the firm. The numerator of this expression is equal to the density of reviewers at $p = 0$ if $L \to 0$, and the average density of reviewers between $p = 0$ and $p = L$ if $L > 0$. It reflects the density of reviewers the firm can attract through underpricing when its reputation is lowest. The denominator includes the mass of consumers who leave reviews regardless of the firm's underpricing at its lowest reputation level. Thus, this expression is equal to the ratio of the marginal to the inframarginal reviewers (see Figure 1).

Additionally, the adjusted hazard rate is affected by the discount rate and the consumer arrival rate. A higher discount rate makes the current marginal reviewers less valuable. A higher consumer arrival rate makes the review channel, and therefore the marginal reviewers, more important.

In Theorem 1, we show that there is underpricing in (every) MPBE if and only if the adjusted hazard rate is high ($h_\varepsilon > h^*$). Otherwise, there is no underpricing in the unique MPBE.

4.1 Uniform Case

In this subsection, we analyze a version of our model where the taste shock ε is distributed uniformly. In this natural case, we show that the possibility of good news in the future crowds out any incentive of the firm to underprice its product today, and the reputational incentives do not lead to underpricing.

Assumption 1 *The taste shock has a uniform distribution with sufficiently large support ($a \geq \bar{u} \geq 1 - a$):*

$$\varepsilon \sim U[-a, a]$$

Under Assumption 1, we can show that λ_g is linear in price:

$$\lambda_g(p) = \lambda(1 - F_\varepsilon(\bar{u} - 1 + p)) = \lambda \cdot \frac{1 + a - \bar{u} - p}{2a} \tag{8}$$

We will show that underpricing does not occur, in two steps: First, we will solve the firm's problem and derive the pricing incentives for different types of firm in order to characterize all possible equilibria of the auxiliary game. Second, we will characterize a unique (in terms of prices) equilibrium in this model under a continuous belief refinement.

We can see three different components of the pricing incentives in the high-quality firm's objective function in equation (6) and (4). The reputational incentive is the combination of the probability of the good news with the value jump $\lambda_g(p) \cdot (V(H, 1) - V(H, q))$. The myopic profit maximization incentive is reflected by λp. The signaling incentive and the demand confound the other two incentives, because firms must choose a price at which the consumer's belief is sufficiently high to purchase the good ($p \in \mathcal{P}_q$).

We will analyze the first-order condition (FOC) of the firm's problem formulated in (6) subject to the selling-price constraint $\mathcal{P}_q$ to determine the optimal price $p(H, q)$. We determine the pricing incentives by differentiating the firm's objective. Because the firm's problem is linear in p, the sign of the FOC will determine if H chooses the maximum or the minimum price in $\mathcal{P}_q$:

$$\frac{\partial}{\partial p}[\lambda p + \lambda_g(p) \cdot (V(H, 1) - V(H, q))] = \lambda - \frac{\lambda}{2a} \cdot (V(H, 1) - V(H, q)) \tag{9}$$

Lemma 1 [9] *For small χ*[10]*, the high-type firm always prefers the highest possible price, i.e., (9) is positive and $p(H, q) = \max \mathcal{P}_q$ (4).*

This lemma implies that for a low redrawing rate χ, the high-quality firm has no incentives to underprice its product. A high-quality firm gains the most from receiving a

9. The proof is provided in the online appendix, Section 7.1.
10. There exists $\chi^* > 0$ such that for any $\chi < \chi^*$

good review, $(V(H,1) - V(H,q))$, when its reputation q and the current value $V(H,q)$ are low. However, when the firm's reputation is low, the product sells only at low prices. Because the firm always sells in any profit-maximizing strategy, the only deviations we need to consider are those from already low prices to even lower prices.

The marginal benefit of reducing the price today is decreasing the expected time until getting a good review and jumping to a high reputation level. Without good reviews at low reputation levels, the firm also has to charge low prices in the future. For a uniform distribution of taste shocks, the density of marginal consumers that the firm can win by underpricing, $\frac{\lambda}{2a}$, is low relative to the mass of consumers who will leave reviews regardless of whether the firm uses underpricing today or will use it in the near future (because the density in the tails of the uniform distribution is high). Thus, underpricing does not significantly reduce the expected time until getting a good review, and the benefit of it is lower than the cost of sacrificing the profit today.

From an algebraic perspective, one could try to break this result by making the $\lambda_g(p)$ function steeper, i.e., increasing the density of marginal consumers $\lambda_g'(p) = \frac{\lambda}{2a}$ in order to make the continuation value more sensitive to price. However, this effect is completely counteracted by also increasing the reviews arrival rate at $q \approx 0$, when the firm is forced to charge $p \approx 0$ no matter what: $\lambda_g(0) \geq \frac{\lambda}{2a}$. Increasing this rate decreases the value jump $V(H,1) - V(H,0)$, because a review will arrive very soon when $q = 0$ and q will jump to 1 almost immediately. For this reason, the firm has no incentives to decrease the price when the firm's reputation level is low because the good news will arrive very soon regardless of whether the firm uses underpricing.[11]

Thus, the high and low types have the same pricing incentives and, as a result, all equilibria under uniform taste shocks will be pooling equilibria. The problem is that both types of firm are choosing the same price for each equilibrium belief, so on-path behavior pins down only one point of the equilibrium belief function. Thus there is a continuum of equilibria with pooling at any given price point and fully pessimistic off-path beliefs at all other prices ($\tilde{\theta}(p,q) = \mathbf{1}_{\{p=p(H,q)\}}\tilde{\theta}(q)$). Further, many common refinements do not help select an equilibrium because both types of firm have exactly the same preferences over actions for any consumers' belief functions. To remedy this problem, we introduce a continuity refinement.[12]

Assumption 2 *For all q, the expectation function $\tilde{\theta}(p,q)$ is continuous in p.*

11. This and the previous arguments can be thought of as a form of the one-shot deviation principle, where we consider a price path that is equal to consumers' willingness to pay at every reputation level of the firm, and then show one-shot price cuts are not profitable.

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

This 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. Under this refinement, there is a unique equilibrium.

Proposition 1 *Under Assumptions 1 and 2, and for small χ, no underpricing is the unique MPBE: the high and low types pool at the willingness to pay, $p(L, q) = p(H, q) = \tilde{\theta}(q)$, at any reputation level q.*

Proof. We show that any pooling equilibrium with the pooling price strictly below consumers' belief ($p^* < \tilde{\theta}(q)$) cannot be an equilibrium of the signaling game if the belief function is continuous 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^*, this consumers' expectations about the quality of the product is still above the price. Thus, consumers would also buy the good for a price a little higher than the equilibrium price and both types of firm would prefer to deviate to that higher price (see Figure 2). This contradicts optimality. Therefore, in the unique equilibrium of the auxiliary game at any q, both types of firm always charge the consumers' willingness to pay: $\tilde{\theta}(q)$. Thus, the unique MPBE features no underpricing. ■

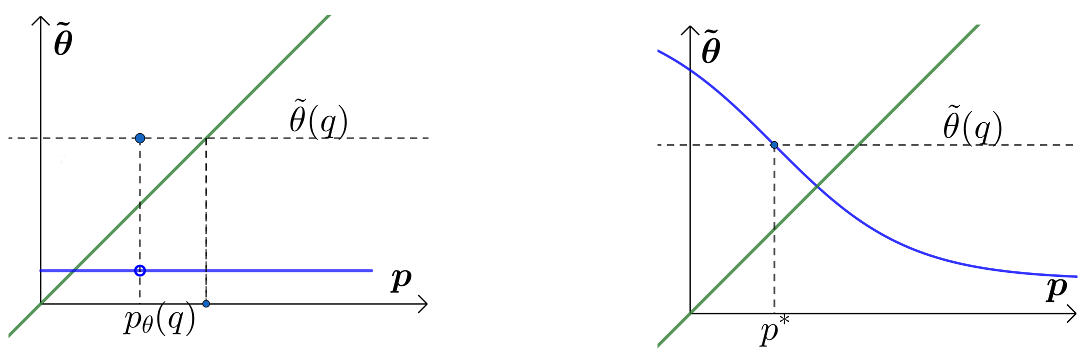

(a) Unreasonable Underpricing

(b) NUP Proof by Contradiction

Figure 2

4.2 General Case

In this section, we derive the main results of the paper for a general class of unimodal taste shock distributions and show that underpricing occurs if and only if the adjusted hazard rate, h_ε, is sufficiently high. We show that any equilibrium is defined as a partition of the firm reputation interval, with no underpricing at higher reputation levels and possibly full underpricing at lower reputation levels.

If underpricing occurs, the high-quality firm charges lower prices than the low-quality firm, and prices are fully informative about quality. This happens because the high-quality firm values establishing its reputation, while the low-quality firm does not. The high-quality firm knows that its quality is high and it only needs to "convince" consumers that this is so, to start collecting high profits. Thus, the high-quality firm prefers to underprice its product heavily when its current reputation is low.

In contrast, the low-quality firm at the lower reputation level cannot "convince" consumers that the quality of its product is high, not even by charging low prices. Thus, the best the firm can do is collect as much profit as possible given its reputation, which entails selling the product at face value to consumers, whereas the high-quality firm prefers to price its product even lower.

Throughout this section, we maintain two important assumptions. First, we retain our continuity belief refinement introduced in the linear case (Assumption 2). Second, we assume $\bar{u} \geq 1$, which implies that reviews are sufficiently selected, and therefore the review arrival rate $\lambda_g(u)$ is convex (Lemma 2). This assumption is equivalent to requiring that the fraction of consumers leaving reviews be never above $1/2$. 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).

Lemma 2 *If $\bar{u} \geq 1$, $\lambda_g(p)$ is decreasing and convex.*

Proof. 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. ■

We show that convex good news arrival rate functions can induce underpricing and separating equilibria at lower reputation levels. Intuitively, a convex arrival rate means that the good reviews arrival rate function is more sensitive to price when the price is low.

In Theorem 1, we will show precisely which distributions F_ε lead to underpricing for small χ. The necessary and sufficient condition for underpricing is that h_ε (7) be sufficiently high or F_ε be convex "enough" around the marginal consumer point, or the taste shock be sufficiently concentrated around its mean. If F_ε is linear or nearly linear, then there is no underpricing. On the other hand, making F_ε steeper between $\bar{u}-1$ and $\bar{u}-1+L$ generates underpricing.

Theorem 1 *An equilibrium exists.*

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

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

2. *If $h_\varepsilon > \frac{1}{1-L}$, then for small χ, there must be underpricing at low reputation levels and no underpricing at high reputation levels in any MPBE, i.e., $\exists\, 0 < q^* < q^{**} < 1$:*

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

High h_ε corresponds to a case where there is a large density of consumers who can be convinced to leave a review after the firm cuts the price of its product from L to 0 and a small mass of consumers who leave reviews even if the 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 the firm's reputation and the profit stream are already high.

In contrast, the low-quality firm, at any reputation level, knows that regardless of how low 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. In this way, the two types of firm engage in separate pricing strategies at low reputation levels, but pool and engage in the same pricing strategy at high reputation levels.

Proof. We prove Theorem 1 in three steps: First, we discuss how the high-quality firm's pricing incentives determine the signaling-game equilibria at different reputation levels and characterize all possible MPBEs as a partition of the reputation interval. Second, we prove equilibrium existence and show that more generally there is a dichotomy: either there is a unique equilibrium with no underpricing or all equilibria have underpricing. Third, we derive the condition on λ_g along which the dichotomy occurs and translate it into a condition on the primitives.[13]

First, similarly to 4.1, we first need to analyze the high-quality firm's pricing incentives for a given reputation q. Let us fix all the future strategies of the firm and consumers,

13. The full proof is provided in the online appendix, Section 7.2.

which determines $V(H,q)$ and $V(H,1)$, and analyze an auxiliary signaling game at q. The high-type (H) firm's problem is characterized by (6):

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

Because $\lambda_g(p)$ is convex, the whole objective function in (6) is also convex in p, for any given $(V(H,1) - V(H,q))$. This implies that the optimal solution to (6) is bang-bang: $p(H,q) \in \{0, \max \mathcal{P}_q\}$.

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 auxiliary game: separating or pooling. Knowing the consumer's preference, we can characterize the two equilibria. 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. In the pooling equilibrium $p(L,q) = p(H,q) = \tilde{\theta}(q)$ because of the continuity refinement.

Because all these equilibrium prices are within $[0, \tilde{\theta}(q)]$ interval, we need to determine the high type's preferences over prices on this interval, or, more specifically, its preferences over 0, L, and $\tilde{\theta}(q)$. Both the continuation value and the public belief about the quality q (3) are continuous functions of time, so the value function $V(\theta,q)$ is also a continuous function in any MPBE. To simplify the next argument, we assume just for now that $V(H,q)$ is monotone increasing in q.[14]

$$\frac{\partial}{\partial p}\bigg(\lambda p + \lambda_g(p)(V(H,1) - V(H,q))\bigg) = \lambda + \lambda_g'(p)(V(H,1) - V(H,q)) \tag{10}$$

We start by showing that pooling is a unique equilibrium of the auxiliary 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 (10) is positive for any p. Therefore, the objective in (6) is monotone increasing in p and H always prefers higher prices to lower prices (see Figure 3a), thus $p(H,q) = \max \mathcal{P}_q$ and pooling at $\tilde{\theta}(q)$ is a unique equilibrium of the auxiliary 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 auxiliary game.

Next, we show that separating is a unique equilibrium of the auxiliary game for low reputation levels. When reputation q is low and the value gap $V(H,1) - V(H,q)$ is large, the high type prefers $p = 0$ to $p = \tilde{\theta}(q)$ and therefore to all prices in $[0, \tilde{\theta}(q)]$ (see Figure 3b).

14. We relax this assumption in the full proof in the online appendix, Section 7.2.

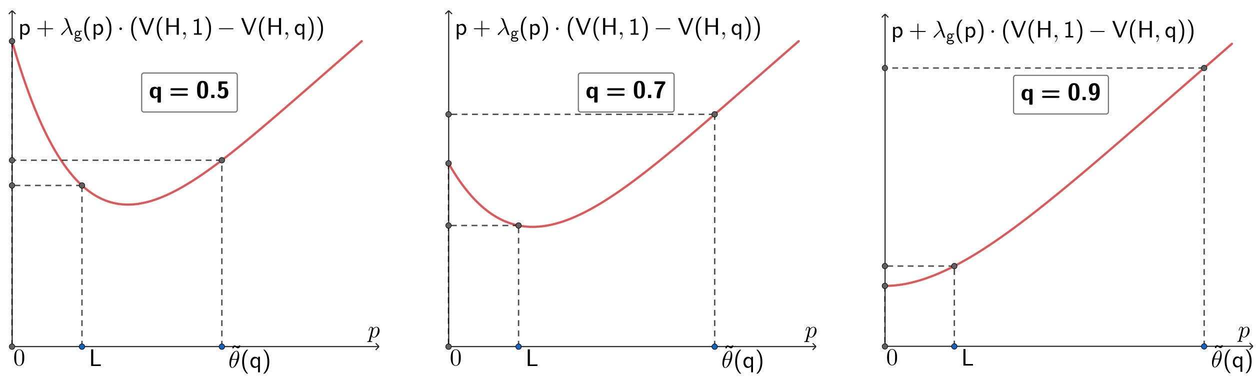

(a) high reputation **(b)** low reputation **(c)** intermediate reputation

Figure 3: Pricing Incentives

This happens because the reputational incentive becomes more significant than the static motive in (10). 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 auxiliary 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 3c). In this case, both pooling and separating equilibria are possible in the auxiliary 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 $p(L,q) = L$.

This auxiliary-game equilibria characterization suggests that equilibria can be described as partitions with separating, or underpricing, at low reputation levels, pooling at high reputation levels, and multiple equilibria at intermediate reputation levels (we formalize how this partition and the thresholds q^* and q^{**} are defined in the appendix).

Second, we prove equilibrium existence by constructing a partition equilibrium with underpricing below some reputation threshold and no underpricing above it. We start with a no-underpricing strategy profile and keep increasing the underpricing interval of the partition until it becomes an equilibrium (a formal proof is provided in the appendix and relies on the intermediate value theorem).

Then we show generally that if there exists an equilibrium with complete pooling and no underpricing, then it is unique. Otherwise, there must be some non-empty interval of low reputations where the auxiliary game exhibits separating equilibria. We show it by proving that a complete pooling equilibrium generates the highest values for the high type at any q, including $V(H,1)$ (the full proof relies on comparing equilibria values by imitating the MPBE with underpricing as a long deviation from the MPBE without underpricing). It also generates the largest gain from a positive review, which is the value gap $V(H,1) - V(H,0)$ (this follows from the 1). Underpricing cannot occur in any equilibrium, because the high type has no incentive to use underpricing even when the value gap is largest.

Finally, the condition separating between cases (1) and (2) of the theorem follows from the HJB equation (1). If χ is small, then approximately $V(H,1) = \lambda/r$. Then $V(H,0) = (\lambda_g(L) \cdot V(H,1) + \lambda L)/(\lambda_g(L) + r)$. Then for the high type to prefer $p = L$ 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 as equation 11 in terms of the model primitives:

$$\frac{F_\varepsilon(1 - \bar{u}) - F_\varepsilon(1 - L - \bar{u})}{L} < \frac{F_\varepsilon(1 - L - \bar{u}) + r/\lambda}{1 - L} \tag{11}$$

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

4.3 Comparative Statics

In this section, we discuss how the spread of the taste shock distribution, consumer arrival rate, and discount factor affect the occurrence of underpricing in equilibrium. The following proposition gives a partial characterization of the set of situations when underpricing does not occur.

Corollary 1.1 [15] *Take a set of primitives $L, q_0, \lambda, r, F_\varepsilon$. Then*

1. *There exists $\alpha^* < +\infty$, such that for any $\alpha > \alpha^*$ and $\varepsilon' = \alpha\varepsilon$ no underpricing is the unique MPBE (for small χ).*

2. *There exists $\lambda^* > 0$, such that for any $\lambda < \lambda^*$ no underpricing is the unique MPBE (for small χ).*

3. *There exists $r^* < \infty$, such that for any $r > r^*$ no underpricing is the unique MPBE (for small χ).*

4. *There exists $L^* < 1$, such that for any $L > L^*$ no underpricing is the unique MPBE (for small χ).*

15. The full proof is in the online appendix.

The first statement of Proposition 1.1 states that increasing the consumer taste shock spread above some threshold (by multiplying it by a constant) guarantees that there is no underpricing in equilibrium. The same is true if the discount rate is sufficiently large, the consumer arrival rate is sufficiently small, or the low quality is sufficiently high.

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

To prove show these comparative statics, we rewrite the no-underpricing condition from Theorem 1 as (12). Increasing (resp., decreasing) any of the parameters from 1–4 above (resp., below) some threshold decreases the left-hand side of (12) below its right-hand side $\frac{1}{1-L}$ and thus guarantees the equilibrium from the first case of Theorem 1[16].

5 Discussion and Future Work

5.1 Welfare Analysis

Underpricing benefits consumers at the expense of the low-quality firm. To illustrate this, compare our model to a myopic benchmark. Consider a model where the firm ignores its ability to influence future reviews through prices. 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 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.

Although the welfare benefits of underpricing seem straightforward ex-post, they are not obvious ex-ante. This is because once a good review arrives, the price increases to 1, and consumer surplus is 0. Thus revealing the quality of the product does not directly lead to high consumer surplus.

5.2 Price Dynamics

In our model, a combination of quality changes, review arrivals, and (in some cases) strategic underpricing generates nontrivial price dynamics over time. We highlight sev-

16. See online appendix Section 7.3 for a detailed proof.

eral examples in turn.

Consider the case with uniform consumer taste shocks so that there is no underpricing. Then the price moves downward over time until a good review arrives, at which point the price jumps to 1. The true quality of the firm's product matters only indirectly via the reputation dynamics.

Next, consider the case when the adjusted hazard rate is high enough so that there is underpricing. If a good review arrives, the price jumps immediately to 1 and then moves downward until it reaches a critical threshold. At this point, what occurs depends on the quality of the firm's product. If it is high quality, the firm will engage in underpricing and the price will suddenly drop to 0 and stay there until a good review arrives. If the product is low quality, the price will still drop, but only to $L > 0$. The price will remain at L until the quality becomes high, at which point the firm will drop the price further to 0 and engage in underpricing.

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.[17] 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 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.

5.3 Price Signaling

In our model, when price signaling occurs, the firm with a high-quality product signals with a low price. This may seem to clash with past work, where high-type firms typically signal with higher prices (Milgrom and Roberts 1986). However, to better understand

17. See for example, Shapiro (1983).

our result, recall that the high-quality firm prices its product lower than the low-quality firm only conditional on reputation. However, 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 firms with high-quality products set generally higher prices: they enjoy a high reputation for longer.

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 is supporting 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.

5.4 Perfect Bad News

In our main model, we consider a review process with perfect good news. Under this model, the 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. One could also consider a perfect-bad-news review process, where consumers leave reviews only for low-quality products. Our framework can be easily adjusted to accommodate such an extension.

In this alternative model, the interesting strategic choices and forces operate when the firm has a low-quality product. If underpricing occurs, it will be used by the low-quality firm with a good reputation attempting to preserve its reputation. In this case underpricing can harm consumers, because it allows the low-quality firm to slow down consumer learning.

In general, the forces in this model will differ from those in the perfect-good-news model. However, when idiosyncratic taste shocks follow a uniform distribution, it continues to be true that no underpricing occurs.[18] In this sense, one of the main results of the paper is robust to some alternative review processes.

In terms of price dynamics, our benchmark model produces a smooth downward trend in price punctuated by sudden upward jumps when good reviews are left. An alternative model with perfect bad news produces the opposite: a smooth upward trend

18. See online appendix Section 7.4 for formal results.

in price punctuated by sudden downward jumps when bad reviews are left.

6 Conclusion

This paper proposes a model of dynamic pricing, where a firm privately informed about the quality of its product faces a market of rational consumers and a rating system that depends on both product quality and net utility delivered. The model allows us to fully account for three important economic forces at play in many online markets: static profit maximization, price signaling, and ratings-based reputation building. We characterize a simple necessary and sufficient condition for when such a review system leads to underpricing in equilibrium.

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. https://doi.org/10.1016/J.TOURMAN.2021.104314.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman E. Ozdaglar. 2017. "Fast and Slow Learning From Reviews." *SSRN* (November). ISSN: 1556-5068. https://doi.org/10.3386/W24046. https://www.nber.org/papers/w24046.

Bardi, Martino, Italo Capuzzo Dolcetta, et al. 1997. *Optimal control and viscosity solutions of Hamilton-Jacobi-Bellman equations.* Vol. 12. Springer.

Board, Simon, and Moritz Meyer-ter-Vehn. 2013. "Reputation for Quality." *Econometrica* 81, no. 6 (November): 2381–2462. ISSN: 1468-0262. https://doi.org/10.3982/ECTA9039. https://onlinelibrary.wiley.com/doi/full/10.3982/ECTA9039%20https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9039%20https://onlinelibrary.wiley.com/doi/10.3982/ECTA9039.

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

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

Carnehl, Christoph, Maximilian Schaefer, André Stenzel, and Kevin Ducbao Tran. 2021. "Value for Money and Selection: How Pricing Affects Airbnb Ratings."

Carnehl, Christoph, André Stenzel, and Peter Schmidt. 2021. "Pricing for the Stars." *SSRN Electronic Journal* (November). https://doi.org/10.2139/SSRN.3305217. https://papers.ssrn.com/abstract=3305217.

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.

Crandall, Michael G, and Pierre-Louis Lions. 1983. "Viscosity solutions of Hamilton-Jacobi equations." *Transactions of the American mathematical society* 277 (1): 1–42.

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

Daley, Brendan, and Brett Green. 2014. "Market signaling with grades." *Journal of Economic Theory* 151, no. 1 (May): 114–145. ISSN: 0022-0531. https://doi.org/10.1016/J.JET.2013.10.009.

Degan, Arianna, and Ming Li. 2021. "Persuasion with costly precision." *Economic Theory* 72 (3): 869–908.

Dynamic Pricing Under Customer Review Effects. https://conservancy.umn.edu/handle/11299/206378.

Feng, Juan, Xin Li, and Xiaoquan Zhang. 2019. "Online Product Reviews-Triggered Dynamic Pricing: Theory and Evidence." *Information Systems Research* 30 (4): 1107–1123. ISSN: 1526-5536. https://doi.org/10.1287/isre.2019.0852. http://pubsonline.informs.org.https://doi.org/10.1287/isre.2019.0852http://www.informs.orghttp://pubsonline.informs.org/journal/isrehttp://orcid.org/0000-0002-0548-1531.

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.

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. https://doi.org/10.1016/J.OMEGA.2017.08.014.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *The Review of Economic Studies* 66, no. 1 (January): 169–182. ISSN: 0034-6527. https://doi.org/10.1111/1467-937X.00083. https://academic.oup.com/restud/article/66/1/169/1666374.

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. 2022. "Price Signaling and Reputation Building: Evidence from a Consulting Platform." *Available at SSRN 4080448.*

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

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

Luca, Michael. 2011. "Reviews, Reputation, and Revenue: The Case of Yelp.com."

Luca, Michael, and Oren Reshef. 2021. "The Effect of Price on Firm Reputation." *SSRN Electronic Journal* (April). ISSN: 1556-5068. https://doi.org/10.2139/SSRN.3625175. https://papers.ssrn.com/abstract=3625175.

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.

Maskin, Eric, and Jean Tirole. 2001. "Markov perfect equilibrium: I. Observable actions." *Journal of Economic Theory* 100 (2): 191–219.

Milgrom, Paul, and John Roberts. 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.

Rodríguez Barraquer, Tomás, and Xu Tan. 2022. "A model of competitive signaling with rich message spaces." *Review of Economic Design* (January): 1–43. ISSN: 14344750. https://doi.org/10.1007/S10058-021-00279-Y/FIGURES/2. https://link.springer.com/article/10.1007/s10058-021-00279-y.

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. https://doi.org/10.2139/SSRN.3439836. https://papers.ssrn.com/abstract=3439836.

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

Sutherland, Ian, Youngseok Sim, and Seul Ki Lee. 2021. "Impacts of quality certification on online reviews and pricing strategies in the hospitality industry." *International Journal of Hospitality Management* 93 (February): 102776. ISSN: 0278-4319. https://doi.org/10.1016/J.IJHM.2020.102776.

Welch, Ivo. 1992. "Sequential sales, learning, and cascades." *The Journal of finance* 47 (2): 695–732.

7 Appendix [Latest Version of Online Appendix]

7.1 Section 4.1 Proofs

Proof of Lemma 1. We prove this lemma by contradiction. First, we notice that the high-quality firm's value $V(H, q)$ at any q cannot be lower than that of charging $p = 0$ until receiving a good review. The high type can always charge zero price at any reputation level since the consumers are ready to buy a product of any quality at any price weakly below L.

Moreover, $V(H, 1) \leq \lambda/r$, that is selling the product to every consumer arriving at the maximum price consumers are possibly ready to pay.

Then for small χ, rearranging (1) together with (3) under Assumption 1 implies

$$\frac{\lambda}{2a} \cdot (V(H,1) - V(H,q)) \leq \lambda_g(0) \cdot (V(H,1) - V(H,q)) \leq rV(H,0) \leq rV(H,1) \leq r * (\lambda/r) = \lambda$$

Therefore

$$\lambda - \frac{\lambda}{2a} \cdot (V(H,1) - V(H,q)) \geq 0$$

This inequality immediately implies that 9 is positive for any q. ■

7.2 Section 4.2 Proofs

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 3 *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 3 *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 new. Thus, all reputation levels $q < \underline{q}$ are off-path. ■

Proof of Theorem 1. 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 auxiliary-game 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 $p(L, q) = L,\ p(H, q) = 0$ 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 MPBE existence and characterize the equilibrium dichotomy in terms of these thresholds.

Lemma 4 *MPBE exists.*

Proof. We prove existence of MPBE by constructing a bi-partition equilibrium with underpricing below some reputation threshold and no underpricing above it: $q^* = q^{**}$. We start with a no-underpricing strategy profile $\underline{q} = q^* = q^{**}$ and check if it is an equilibrium from the firm's optimality perspective. If it is not, then the high type wants to underprice at low ratings, and we start increasing the underpricing region by increasing $q^* = q^{**}$. Our goal is to find a threshold which the H's pricing incentives are consistent with (there are multiple such thresholds), for instance, at which H is indifferent between $p = 0$ and $p = \tilde{\theta}(q)$ (this would be just one possible equilibrium). We know that at $q^* = q^{**} = \underline{q}$ H strictly prefers 0. Also, if $q^* = q^{**} = 1$, then $V(H, q) = 0$ for any q and H prefers any positive price to 0. By the continuity of the values in $q^* = q^{**}$, there exists a threshold at which H is indifferent between $p = 0$ and $p = \tilde{\theta}(q)$, and this will be a threshold for which the strategy profile is an equilibrium. Thus, an equilibrium exists. ■

Then we show generally that if there exists an equilibrium with complete pooling and no underpricing, then it is unique: $q^* = q^{**} = \underline{q}$. Otherwise, there must be some nonempty interval of low reputations where the auxiliary game exhibits separating equilibria: $q^* > \sup \underline{q} > 0$. We show it by proving that a no-underpricing equilibrium generates the highest values for the high type at any q, including $V(H, 1)$ and therefore the largest gap $V(H, 1) - V(H, \underline{q})$ (this follows from the 1). Then, underpricing cannot occur in any equilibrium, because the high type has no incentive to underprice in the NUP MPBE, where the value gap is largest.

Lemma 5 *If there are multiple MPBE and no-underpricing is one of them, then it generates the largest $V(H, 1)$ among all MPBE.*

Proof. Consider any other possible equilibrium with underpricing at some reputation levels. Let us recreate it as a long off-path deviation of H from the NUP equilibrium (we can do it because pricing at 0 is always allowed). This deviation is not profitable for H, since we assumed that complete pooling is an equilibrium and no single-shot or longer deviations are profitable. Thus the pooling equilibrium value is higher than the off-path deviation one at any q. The off-path deviation value on the other hand is higher than the value in the underpricing equilibrium we picked because the prices are the same for every q, but the public belief q drifts down slower in the off-path deviation as the consumers do not expect lower prices and a higher rate of the good news arrival in the past, which benefits H. This implies that the value in the pooling equilibrium is higher than in the underpricing one and this concludes the proof.■

Thus if no-underpricing is an equilibrium, it generates the largest $V(H, 1)$ and thus the largest gap $V(H, 1) - V(H, \underline{q})$ because

$$V(H,1) - V(H,\underline{q}) = \frac{(r^2 + r\chi)V(H,1) - (r+\chi)\lambda L}{r^2 + r\chi + (r + \chi q_0)\lambda_g(L)}$$

which is increasing in $V(H, 1)$ (the expression follows from equations (1) and (2)).

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 $p = 0$. Therefore, there can be no underpricing signaling equilibria with the largest gap $V(H, 1) - V(H, \underline{q})$, which implies that there cannot be any underpricing signaling equilibria (since all other value gaps $V(H, 1) - V(H, q)$ are smaller).

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 as equation 11 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. 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. ■

7.3 Section 4.3 Proofs

Proof of Comparative Statics. Let us rewrite h_ε (7) in an alternative form:

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

We want to show that we can make this expression below $1 < \frac{1}{1-L}$ in either part of Corollary 1.1.

1. Increasing α above large threshold increases $1 - F_{\alpha\varepsilon}(\bar{u}-1+L)$ and 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. 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. 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. $h_\varepsilon \leq \lim\limits_{L\to 0} h_\varepsilon < \infty$. By making $L < L^*$ and $\frac{1}{1-L}$ above this limit, we guarantee no underpricing.

7.4 Section 5 Proofs

Perfect Bad News with Uniform Taste Shock.

Assume model modification where consumers leave only bad reviews if the quality is low and the overall utility is below $\underline{u}$ and the taste shocks are distributed i.i.d. uniformly between $-a$ and q with $a > \max\{L-\underline{u}, 1-(L-\underline{u})\}$. Then the bad news arrival rate is linear in the price

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

Proposition 2 *For small χ, no underpricing is the unique MPBE.*

Proof. L's pricing incentives are given by

$$\frac{\partial V}{\partial p} = \lambda - \frac{\lambda}{2a} \cdot (V(L,q) - V(L,0))$$

Given that $V(L,0) \approx 0$ and $rV(L,1) \approx \max_p\{\lambda p + \lambda_b(p)(V(L,0) - V(L,1)\}$, we can show that

$$V(L,1) - V(L,0) \leq \frac{\lambda}{r + \lambda_b(1)}$$

and

$$\frac{\partial V}{\partial p} \geq \lambda - \frac{\lambda}{2a} \cdot \frac{\lambda}{r + \frac{\lambda}{2a}} > 0$$

Thererfore, L always prefers $\max \mathcal{P}_q$, and pooling at the consumers' willingness to pay is a unique signaling equilibrium at any q. Thus, no underpricing is the unique MPBE.

Consumer Arrival Rate Depending on q.

Assume a model modification, where consumers are more likely to arrive for a firm with a higher reputation, i.e. $\lambda(q)$ is increasing. We want to show that the analysis remains the same as in the main model and we can easily characterize the condition for when there is underpricing in equilibrium.

Notice that for any given value function $V(H,q)$ and q, $\lambda(q)$ multiplies both the static and the dynamic parts of H's objective. Therefore, signaling equilibrium characterization remains the same.

However, $\lambda(q)$ affects the overall solution of the HJB equation (1), $V(H,q)$. Therefore, we need to verify the condition on the taste shock distribution for the lowest q and $\lambda(0)$:

$$h_\varepsilon = \frac{\lambda(0) \cdot (F_\varepsilon(\bar{u} - 1 + L) - F_\varepsilon(\bar{u} - 1))/L}{\lambda(0) \cdot (1 - F_\varepsilon(\bar{u} - 1 + L)) + r}$$

But for the threshold following from $V(H,1) \approx \frac{\lambda(1)}{r}$, which is equal to $\frac{1}{\frac{\lambda(1)}{\lambda(0)} - L}$ instead of $\frac{1}{1-L}$ in the benchmark model. Thus, if $\lambda(0) = \lambda$ from the benchmark model, then the reputation-based consumer arrival rate (popularity-based demand) increases the possibility of underpricing, which is quite intuitive given that the higher reputation increases the profit stream both via the price and demand. ■

7.5 Section 3 Proofs

Law of motion of q without redrawing the state:

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

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

$$dq_t \cdot (1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt) = q_t(1 - q_t)\lambda_g(\tilde{p}(q_t))dt$$

$$\frac{dq_t}{dt} = q_t(1 - q_t)\lambda_g(\tilde{p}(q_t)) \cdot \frac{1}{(1 - q_t \cdot \lambda_g(\tilde{p}(q_t))dt)}$$

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

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 simplicity of notation). It can also be derived more formally from the continuation value of the firm introduced at Section 3.1.

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

$$\begin{aligned} rV(H,q)dt = &\lambda p dt + (1 - r dt)[(1 - \chi(1 - q_0)\lambda_g(p)dt(V(H,1) - V(H,q)) \\ &+ (1 - \chi(1 - q_0)dt)(1 - \lambda_g(p)dt)dV(H,q) + \chi(1 - q_0)dt(V(L,q) + dV(L,q) - V(H,q))] \end{aligned} \tag{13}$$

$$\begin{aligned} rV(H,q) &= \lim_{dt \to 0} \text{RHS(13)} \\ &= \lambda p + \lambda_g(p)(V(H,1) - V(H,q)) + \frac{dV(H,q)}{dt} + \chi(1 - q_0)(V(L,q) - V(H,q)) \end{aligned}$$

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

Finally, we want to show that, for small χ, $V(H,1) \geq V(H,q)\ \forall q$. That follows (almost) immediately from the fact that $\lim_{\chi \to 0} V(H,1) = \lambda/r$ and $\forall q: \quad V(H,q) < \lambda/r$.