

Pricing for the Stars

Dynamic Pricing in the Presence of Rating Systems*

Christoph Carnehl[†] André Stenzel[‡] Peter Schmidt[§]

February 21, 2023

Maintaining good ratings increases the profits of sellers on online platforms. We analyze the role of strategic pricing for ratings management in a setting where a monopolist sells a good of unknown quality. Higher prices reduce the value for money, which on average worsens reviews. However, higher prices also induce only those consumers with a strong taste for the product to purchase, which on average improves reviews. Our model flexibly parametrizes the two effects. This parametrization can rationalize the observed heterogeneity in the relationship between reviews and prices. Based on an analytic characterization of the optimal dynamic pricing strategy, we study a platform’s choice of the sensitivity of its rating system to incoming reviews. The optimal sensitivity depends on the effect of prices on reviews and on how the platform weighs consumers and sellers in its objective. While sellers always benefit from more sensitivity, consumers may suffer from higher prices and from slower learning from reviews due to endogenously emerging price and rating cycles.

Keywords: Rating Systems, Dynamic Pricing, Asymmetric Information

The system will learn what reviews are most helpful to customers...and it improves over time. It’s all meant to make customer reviews more useful.

- Amazon spokeswoman Julie Law, *Interview with cnet.com*, 2015

1. Introduction

Most online platforms feature rating systems to mitigate asymmetric information about product quality. These exert a substantial impact on consumer demand. According to EPRS (2017), 82% of European consumers read reviews before shopping online. In Britain alone, online reviews affect consumer spending worth GBP 23 billion each year (see CMA 2015). The effect on

*This paper supersedes an earlier version titled *Consumer Rating Dynamics*. We are grateful to Jérôme Adda, Heski Bar-Isaac, Luis Cabral, Justin Johnson, Xavier Lambin, Volker Nocke, Marco Ottaviani, Martin Peitz, Martin Pollrich, Johannes Schneider, Nicolas Schutz, Anton Sobolev, Konrad Stahl and Stefan Weiergräber for helpful comments and suggestions, and appreciate input from seminar audiences at the Bank of Canada, Berlin, Bocconi, Genova, Leicester, Mannheim, OsloMet, EARIE 2019, ESEM 2019, MaCCI Summer Institute 2019, IIOC 2018, EIEF-UNIBO-Igier Workshop 2018, MaCCI Annual Conference 2016, and VfS Jahrestagung 2016. Special thanks go to Johannes Dittrich for making the Steam data available to us.

[†]Bocconi University and IGIER. E-Mail: christoph.carnehl[at]unibocconi.it.

[‡]Bank of Canada. E-Mail: astenzel[at]bankofcanada.ca. Funding by the German Research Foundation (DFG) via CRC TR 224, Project C03, during previous employment at the University of Mannheim is gratefully acknowledged. The views expressed in this paper are those of the author and no responsibility for them should be attributed to the Bank of Canada.

[§]QuantCo. E-Mail: peter.schmidt[at]quantco.com.

individual firms is significant; for example, Luca (2016) estimates that a one-star rating increase on Yelp.com raises restaurant revenues by 5 to 9 percent.¹

Aggregate ratings are computed from individual reviews that reflect not only quality but also value for money and heterogeneous tastes.² Importantly, these factors determining the review are typically unobservable to future consumers. Sellers, therefore, can use current prices to affect reviews and thus future consumers’ quality perceptions and profits. Our paper focuses on strategic pricing as a means for online reputation management. We complement the existing literature on strategic pricing and consumer ratings by emphasizing the non-trivial interplay between the design of rating systems and sellers’ dynamic pricing incentives, which reflect product and environmental characteristics.

We provide a tractable model of dynamic pricing in the presence of rating systems. Our framework flexibly captures the heterogeneity to which reviews reflect purchase prices. A key ingredient of our model is the sensitivity of aggregate ratings to incoming reviews which is an important rating design parameter.³ We analytically characterize the optimal dynamic pricing strategy. Based on this characterization, we derive implications for sellers’ optimal price reactions to rating changes. In addition, we discuss the interaction of the optimal pricing strategy with the rating system’s sensitivity and its impact on the speed of consumer learning, seller profits, and consumer surplus. Building on these insights, we highlight novel considerations for platforms in the optimal design of rating systems.

Higher prices lead to better ratings when reviews reflect consumers’ overall enjoyment of the product, but they lead to worse ratings when consumers primarily judge the value for money. When the rating system becomes more sensitive to recent reviews, sellers benefit, as it facilitates the management of ratings via their pricing decision. Consequently, long-run prices and profits are higher for products where higher prices positively affect reviews, while consumer surplus is lower. Our results indicate that regulators should carefully consider the effects changes in the design of rating systems on pricing incentives, as they affect several important market outcomes.⁴

Importantly, we show that different levels of sensitivity qualitatively change the patterns of prices and ratings over time. An increase in the sensitivity can incentivize sellers to engage in cyclical build-up and milking of their rating stock—that is, it can lead to the endogenous emergence of price-rating cycles—which slows consumer learning. Therefore, recent shifts by platforms towards more sensitive rating systems (see, e.g., Steam 2016, wired.com 2019) may have harmed consumers in two ways: through higher prices and through slower learning.

From a managerial perspective, our analysis suggests that platforms should enter markets with comparatively insensitive rating systems weighing all reviews equally, and shift towards disproportionately emphasizing recent reviews as the platforms mature. Sellers, in turn, should tailor their pricing strategies to the characteristics of the rating systems of platforms they are operating on. Cyclical pricing to boost the rating stock which can then be capitalized upon is

¹There is ample additional evidence of a significant demand effect of reviews. For example, Chevalier and Mayzlin (2006) shows such an effect on Amazon.com and Barnesandnobles.com.

²See, for example, De Langhe et al. (2015) or Li and Hitt (2010).

³For the first point, see, e.g. Bondi (2019), Zegners (2019), Carnehl et al. (2022), Luca and Reshef (2021). For the second point, note that only since 2015 has Amazon placed a disproportionate emphasis on recent reviews in computing aggregate ratings, while the video game platform Steam has displayed a recent average score in addition to the overall rating since 2016. See wired.com (2019) and Steam (2016, 2021) for details.

⁴These effects should be considered in addition to other, potentially illegal means of strategically affecting ratings that warrant scrutiny. For example, fake reviews have attracted significant attention from researchers and policymakers. Luca and Zervas (2016) estimates approximately 16% of Yelp reviews to be suspicious. Many competition authorities have dealt with fake reviews, e.g., CMA (2015), EPRS (2017), FTC (2019). Notably, the channels presented in our paper persist even when fake reviews can be eradicated.

beneficial only if rating systems emphasize recent reviews sufficiently—otherwise, gradual price adjustments are optimal.

In our model, a long-lived monopolist sells a good of fixed and privately known quality to short-lived consumers. Consumers value quality but do not observe it before purchase and differ in their taste for the good. They form a belief about the quality based on current observables consisting of the price and the product’s aggregate rating. Consumers who purchase can leave a review. The aggregate rating is a weighted average of the previous rating and the average incoming review. The weight parametrizes the sensitivity of the rating system to recent reviews.

Reviews consist of a unidimensional score equal to the gross utility of consumption less a fraction of the purchase price so that, all else equal, a higher price induces a worse review. The value of this fraction reflects the degree to which reviews internalize the purchase price. Minimal price internalization corresponds to consumers reporting gross utility, while maximal price internalization corresponds to net utility reporting.

The degree of price internalization is a flexible, application-specific parameter that affects the relative strength of two effects determining the impact of prices on reviews. The first effect is the direct *price effect*: a price increase worsens reviews by lowering the value for money. The second effect is the indirect *selection effect*: a higher price affects the composition of purchasing consumers—the purchasing consumers, on average, have a higher expected gross utility—which improves reviews.⁵

We assume that consumers form expectations about the product’s quality using an intuitive non-Bayesian inference rule: they use only current observables, consisting of the aggregate rating and price, and their knowledge of the utility and review functions.⁶ Intuitively, consumers treat the model as if it were in a stationary equilibrium. A similar inference rule has been used in the literature (e.g., Crapis et al. 2017, Besbes and Scarsini 2018) and is related in spirit to the notion of cursed equilibrium (Eyster and Rabin 2005): consumers do not fully take into account all of the information contained in current observables. We show that the inference is uniquely determined and that the resulting demand increases in the rating and decreases in the price. The demand elasticity is determined by the degree of price internalization in the review function. With higher price internalization, consumers rationalize the same rating at a higher price by inferring the good to be of higher quality; the set of purchasing consumers is less responsive and demand less elastic.

Therefore, the degree of price internalization in reviews determines the relative strength of the direct price effect and the selection effect. When reviews heavily internalize the price, the direct price effect is mechanically strong. Simultaneously, demand is unresponsive and the selection effect weak. Thus, lower prices lead to better reviews. The converse reasoning applies when the internalization is low—the selection effect dominates and higher prices lead to better reviews.

Our model provides a novel rationalization of seemingly conflicting empirical findings regarding the impact of price changes on reviews: higher prices have been shown to lead to better reviews for books (Bondi 2019, Zegners 2019), while lower prices have been shown to lead to better reviews for USB sticks (Cabral and Li 2015) and lower-priced listings on Airbnb (Carnehl et al.

⁵The selection effect in our setup refers to selection into purchasing, not selection into reviewing conditional on purchasing. We study the latter in an extension, see Section 5.

⁶Full rationality would place high requirements on consumers’ cognitive abilities. They would need to know or form beliefs about how many periods have passed, the prior beliefs of consumers about product quality, the price path of the seller (which depends on the seller’s cost and discount rate as well as the solution of a dynamic signaling game with ratings), and how reviews are aggregated into ratings. We show that our results carry over to a setting with fully Bayesian consumers, see Section 5.

2022).⁷ Our model also qualitatively matches the observation that reviews move opposite prior quality expectations (Ho et al. 2017, Hui et al. 2021).

As ratings are persistent while prices are not, future demand is affected through current prices via the induced reviews, which influence the future aggregate rating. In each period, the seller balances the current price’s effects on flow profits and the aggregate rating. The latter effect puts downward pressure on prices when the direct price effect dominates and upward pressure when the selection effect dominates.

Our main theoretical contribution is the analytical characterization of the unique equilibrium in our framework. We show that rating systems are effective: consumer beliefs converge linearly to the actual product quality despite their misspecified model. Nevertheless, the long-run price level, profits, and consumer surplus depend on the interplay between the relative strength of the price and selection effects and the sensitivity of the rating system to incoming reviews.

The characterization of the optimal dynamic pricing strategy and resulting market outcomes allows us to establish insights into consumer learning and managerial implications for sellers and platform operators. First, the sensitivity of the rating system has a direct and non-monotonic impact on consumers’ speed of learning.⁸ A marginal increase in the sensitivity of the rating system increases the speed of learning if the initial sensitivity is low but decreases the speed of learning when the initial sensitivity is high. The reduction in the speed of learning occurs because a high sensitivity of the rating system induces sellers to engage in a cyclical build-up and milking of ratings. Consequently, there is an interior sensitivity level that maximizes the speed of learning.

Second, the sensitivity impacts long-run price levels via the sellers’ dynamic pricing incentives. In the steady state, the seller trades off the effect of the current price on flow payoffs and future ratings, where the sensitivity of the rating system determines the weight placed on future ratings. When the rating is more sensitive to incoming reviews, a given change in the price level leads to a stronger rating adjustment. The seller’s profits are thus increasing in the sensitivity and sellers prefer rating systems to be as sensitive as possible. How long-run price levels and consumers are affected depends on the effect of prices on reviews. When price internalization is high, sellers charge lower prices. When price internalization is low, sellers charge higher prices. The sensitivity which maximizes consumer surplus is, therefore, either the highest or the lowest feasible sensitivity depending on the degree of price internalization.

Third, these considerations inform a platform operator’s optimal design of the rating system. The optimal sensitivity depends on the relative weight placed on seller profits, consumer surplus, and consumer learning, as well as the distribution of the degree of price internalization in reviews across the platform’s product portfolio. We expect that relatively young platforms place significant weight on consumer surplus and learning as they aim to attract a sizeable customer base. At the same time, it seems reasonable for a more established platform to place less weight on consumers. At this point, our analysis predicts a shift to more sensitive rating systems, which boost seller profits and thus royalties to the platform. This hypothesis is in line with, for example, Amazon changing its review aggregation in 2015 to emphasize more recent reviews disproportionately.⁹

⁷Many studies have looked at the impact of prices on restaurant reviews on Yelp.com and found heterogeneous effects; see Byers et al. (2012), Li (2016), Luca and Reshef (2021).

⁸In our setup, the speed of learning corresponds to the speed of convergence of the beliefs that consecutive consumer generations hold about the product’s quality; that is, how quickly information flows across consumer generations.

⁹The considerations driving these predictions naturally need to be balanced with additional factors, such as benefits of increased sensitivity when product quality is time-variant (see Kovbasyuk and Spagnolo 2018), and other platform design choices, such as the order of search results steering consumers towards specific products.

Fourth, we derive qualitative implications for sellers who can frequently adjust prices in response to rating changes. Responding to incoming reviews with price changes is profitable for sellers as demand varies with the current rating level.¹⁰ While this also obtains for a myopic seller maximizing flow profits, we illustrate a novel insight for a forward-looking seller. When higher prices lead to worse reviews—i.e., if the price effect dominates—a seller should only moderately raise prices following a rating increase. At the same time, she should strongly cut prices following a rating decrease.¹¹ This asymmetry is due to the ratings management objective. Although the seller wants to exploit the higher demand following a rating improvement, she does not want to raise prices too much to moderate the detrimental effect of current prices on future ratings. More sensitive rating systems amplify this concern. Facing a more sensitive rating system, a seller should respond with even more conservative price increases to rating improvements and more aggressive discounts when bad reviews arrive. Moreover, it is more sensitive rating systems that induce sellers to engage in cyclical pricing to build up and exploit the rating stock, e.g., via frequent sales and discounts. Less sensitive rating systems incentivize sellers to undertake more gradual price adjustments.

Fifth, the above profit-maximizing pricing strategies require sellers to understand the price internalization in reviews for their products. If a seller is uncertain about this parameter, she should implement price experimentation or alternative measures to learn about the causal effect of prices on reviews. Similarly, a platform operator should invest in learning the distribution of price internalization across products on her platform to choose the rating system’s sensitivity optimally.

In a series of extensions, we show that our results are robust to accounting for additional considerations, such as the number of reviews in the updating process, stochastic reviews, general distributions of horizontal preferences, and competition. Perhaps most importantly, our results carry over to settings in which consumers are fully Bayesian and in which consumers differ in their marginal valuation of quality. In all extensions, the key market outcomes, such as prices, profits, consumer surplus, and speed of learning, depend on the interplay between the price internalization in reviews, the seller’s dynamic pricing incentives, and the sensitivity of the rating system to incoming reviews. Therefore, our main qualitative findings remain unchanged.

Related Literature We study the incentives of sellers to manage their ratings through strategic pricing, which relates to the literature on reputation management (see Bar-Isaac and Tadelis 2008, for an overview). In our model, quality is fixed (see, e.g., Cabral and Hortacsu 2010, Board and Meyer-ter Vehn 2013, for settings where quality is affected by sellers’ strategic choices), and information transmission occurs via induced ratings (instead of certifiers as in Marinovic et al. 2018). We abstract from and complement the vast literature on direct price signaling (see, e.g., Wolinsky 1983, Bagwell and Riordan 1991, Osborne and Shapiro 2014), which empirically has been shown to be complementary to information transmission via rating systems (Bhargava and Feng 2015).¹²

Our model flexibly combines the direct price effect (first formalized in a two-period model by Li and Hitt 2010) and a selection effect (introduced separately in a two-period model in Li

¹⁰Such dynamic responses in the pricing are, by assumption, not captured in the fixed-price strategies considered in Li and Hitt (2008), Crapis et al. (2017). We contrast the optimal dynamic strategy with optimal fixed pricing in Section 4.

¹¹Clearly, in the opposite case of a dominant selection effect, i.e., when higher prices lead to better reviews, the price increase in response to a rating increase should be substantial while the price decrease in response to a rating decrease should be moderate only.

¹²The literature explicitly incorporating both price signaling concerns and rating systems is scarce. A notable exception is Martin and Shelegia (2021), who study the interaction of price signaling with ratings in a two-period model. A high-quality seller may have an incentive to mimic a lower-quality seller to improve future reviews, which may induce the latter to engage in loss-leadership to separate itself from the former.

and Hitt 2008). Empirically documented positive correlations between reviews and purchase prices (Byers et al. 2012, Zegners 2019) are not consistent with Li and Hitt (2010) but can be rationalized by our model and Li and Hitt (2008). In Li and Hitt (2008), the selection effect reflects the correlation between consumers’ taste for a product and their harshness in reviewing. We therefore differ by providing an alternative economic mechanism rationalizing empirically documented patterns.¹³

We view our model as plausible for many markets and products because it matches the findings in the empirical literature on the content of online reviews (De Langhe et al. 2015, Luca and Reshef 2021, Jeziorski and Michelidaki 2021): reviews reflect not only quality (as in, e.g., Maglaras et al. 2020) but also additional considerations such as consumer tastes and the purchase price, and are not simply a function of consumer net utility (as posited, e.g., by Acemoglu et al. 2019, Ifrach et al. 2019, who study social learning from rating systems accounting for reviewer selection). We flexibly parametrize the relative strength of the direct price and selection effect to rationalize the empirically documented heterogeneity in reviewing behavior depending on product characteristics.¹⁴

Variations of the price and selection effect also feature in models by Crapis et al. (2017), who characterize conditions under which consumers learn product quality from reviews, and Feng et al. (2019), who focus on the price trajectory in a setting where consumers are long-lived. Our paper differs by explicitly analyzing the interplay between sellers’ strategic pricing incentives and the design of the rating system parameterized by the sensitivity to recent reviews. This interplay is key for our insights as it gives rise to endogenous price-rating cycles and shows the product-specific consequences of recent shifts to more sensitive rating systems. It also differentiates our work from concurrent work by Shin et al. (2022), who approximate a monopolist’s revenue maximization problem by using a fluid model and assess the value of dynamic pricing when consumer reviews reflect either product quality or value for money.

Our work is also related to recent work on the design of rating systems. Jiang and Guo (2015) relates strategic pricing incentives to rating systems’ cardinality and granularity. Kovbasyuk and Spagnolo (2018) shows that a low memory of ratings can be optimal if product or service quality changes over time, as it prevents inefficient exit. Che and Hörner (2018) studies the optimal design of recommender systems to incentivize collaborative learning. Luca (2017) and Dai et al. (2018) discuss the design of rating systems, including issues like reviews being selected and ways to aggregate reviews into rating statistics. Klein et al. (2016) empirically evaluates a change in the design of eBay’s feedback mechanism. Our analysis also relates to the choice between posted prices and auctions in online markets (see Einav et al. 2018). Finally, our work relates to Bonatti and Cisternas (2019) who analyze aggregate scores about purchasing histories that inform short-lived firms about consumers’ evolving willingness to pay. In our paper, the seller holds an informational advantage over boundedly rational and short-lived buyers, as in, e.g., von Thadden (1992).

The remainder of the paper is structured as follows. We set up the model and discuss the effects of prices on reviews in Section 2. The dynamic pricing considerations and the derivation of the equilibrium are given in Section 3. We discuss the implications in Section 4 and extensions in Section 5. Section 6 concludes the paper.

¹³Notably, a positive correlation between prices and review scores in Li and Hitt (2008) requires that consumers with the highest realized surplus leave the worst reviews. In addition, Luca and Reshef (2021) find no evidence of a significant selection of harsher reviews following price increases in the market for restaurants.

¹⁴See, e.g., Byers et al. (2012), Cabral and Li (2015), Li (2016), Bondi (2019), Zegners (2019), Carnehl et al. (2022), Luca and Reshef (2021).

2. Model

We consider a long-lived monopolistic producer of a good with a privately known fixed quality. Consumers are short-lived and exhibit horizontal differentiation in their taste for the good. A review and rating system allows for information transmission across consumer generations. We next describe the model before discussing the main modeling assumptions.

Time Time is discrete, $t \in \{1, 2, 3, \dots, T\}$, $T \leq \infty$.

Seller The seller wishes to sell a good of exogenously given and fixed quality θ , where θ is distributed according to a cumulative distribution function F , $\theta \sim F(\cdot)$ on $\Theta \equiv [\underline{\theta}, \bar{\theta}]$. The realization of θ is private information of the seller. In each period t , the seller decides on the price p_t . The marginal costs of production are independent of quality and normalized to 0. The seller is risk neutral and has discount factor $\delta \in (0, 1)$.

Consumers In each period t , a unit mass of risk-neutral consumers who live for one period arrives. Consumers value quality and differ in their taste for the good that the seller offers. Each consumer i has type $\omega_i \sim U[0, 1]$. The gross utility of a consumer is $u(\theta, \omega_i) = \theta + \omega_i$, so the utility is increasing in quality and taste. A consumer's utility net of the price paid, p , is

$$u_i = \theta + \omega_i - p. \quad (1)$$

When all consumers hold the same beliefs, there is a cutoff consumer $\tilde{\omega}$ such that all consumers with $\omega \geq \tilde{\omega}$ purchase the good and all consumers with $\omega < \tilde{\omega}$ do not.

Reviews and Rating System Information transmission across periods occurs via a review and rating system, which is structured as follows. For tractability, we assume that every consumer who purchases the good leaves a review with the same positive probability. If consumer i leaves a review ψ_i after purchase, the review is

$$\psi_i = \theta + \omega_i - \kappa p, \quad (2)$$

with $\kappa \in [0, 1)$. In line with the empirical evidence, we assume that the review reflects the consumer's gross utility of consuming the product minus a component that depends on the purchase price. The higher the price at which the consumer purchases a product, the worse the review left. We discuss the motivation and rationale for the flexible interior degree of price internalization in reviews in detail in Section 2.1. Note that for $\kappa \rightarrow 1$, each consumer reports her net utility, i.e., her surplus. For $\kappa = 0$, reviews reflect gross utilities instead. The degree of price internalization κ may vary across different products, which is important to keep in mind for two specific reasons. First, when relating the model's predictions to the empirical evidence, different degrees of κ can rationalize the different empirical correlations between prices and reviews documented for different products. Second, our main results crucially depend on the level of κ —implications may thus differ for different products. We adopt an equal weight of quality and taste for expositional purposes only. Introducing weights, potentially differing between utility and review function, does not affect the mechanisms or qualitative results.

The average review in a given period is used to update the aggregate rating.¹⁵ This average review is equal to the review left by the consumer with average taste $\omega^e(\omega_t^*) = E[\omega | \omega \geq \omega_t^*] =$

¹⁵Because only the average review matters for updating the aggregate rating, the specific probability with which an individual consumer leaves a review does not matter due to the law of large numbers. Our results are qualitatively robust to accounting for the number of reviews in the updating process; see Section 5.

$\frac{1+\omega_t^*}{2}$, where ω_t^* is the taste of the marginal consumer who purchases. The rating system is characterized by the mapping from the current aggregate rating $\bar{\psi}_t$ and current average review $\psi_t = \theta + \frac{1+\omega_t^*}{2} - \kappa p_t$ into the next period's aggregate rating $\bar{\psi}_{t+1}$. We consider a specific class of rating systems following the updating rule

$$\bar{\psi}_{t+1} = (1 - \sigma)\bar{\psi}_t + \sigma\psi_t. \quad (3)$$

We parametrize the initial attitude by consumers via an initial rating $\bar{\psi}_1 \geq 0$.¹⁶ This yields

$$\bar{\psi}_t = (1 - \sigma)^{t-1} \bar{\psi}_1 + \sum_{\tau=1}^{t-1} (1 - \sigma)^{t-1-\tau} \sigma \psi_\tau. \quad (4)$$

In (3) and (4), $\sigma \in [\underline{\sigma}, 1]$ parametrizes the sensitivity of the rating system to incoming reviews. The higher σ is, the more responsive the updated rating to incoming reviews and, correspondingly, the lower the weight on older reviews. In the extreme cases, $\sigma = 1$ corresponds to a limited memory rating in which the rating reflects only the last period's average review, while $\sigma = \underline{\sigma} > 0$ denotes the system with the lowest sensitivity.¹⁷ The inclusion of σ as a parameter allows us to assess the recent pushes by online platforms such as Amazon to have more recent reviews matter more for the displayed aggregate rating—in our context, this aspect would be captured by an increased σ .

Timing of the Stage Game The timing of a given period is as follows: the seller observes the current state of the market as characterized by the aggregate rating $\bar{\psi}_t$ and sets the price p_t at which she is willing to sell. Consumers then observe p_t and $\bar{\psi}_t$ and decide whether to purchase the good—a consumer purchases if and only if her expected net utility is weakly positive. If consumers choose to purchase, they realize their net utility, as in (1), and potentially leave a review, as in (2).

Technical Assumptions For technical purposes, we impose the following assumptions. First, we require $\underline{\theta} < -1$ to ensure that there are quality levels such that no consumer would purchase the good irrespective of taste. Second, we require $\bar{\theta} < \infty$ to ensure the boundedness of profits in each period. Finally, we restrict attention to σ such that $\kappa < 1 - \sigma/2$. This sufficient condition ensures the existence of a stationary equilibrium.¹⁸ Moreover, the initial rating has to be sufficiently high, $\bar{\psi}_1 \geq 0$, such that consumers will purchase in the first period at non-negative prices.

Consumer Inference A central requirement is to specify how consumers conduct quality inference given their observables. Recall that consumers live for one period and observe only the current price and rating. Hence, we need to specify how consumers rationalize the current combination, given that they do not observe the path of prices and ratings. Past prices are not observable to consumers on many online sales platforms, such as Amazon and Steam. Moreover, past ratings are not directly observable and can be computed only via time-consuming analysis of individual time-stamped reviews. While, in principle, we could assume that consumers are fully rational and solve the seller's problem from time $t = 0$ onwards, a fully rational consumer would have to solve a dynamic signaling game with a rating system, which is a highly complicated problem. Instead, we assume that consumers try to rationalize the aggregate rating while

¹⁶This parametrization is for expositional purposes only. The results are unchanged if we add an initial period at $t = 0$ in which consumers hold an exogenous belief about quality $\bar{\mu}_0 > -1$ and where selling to these consumers induces a first-period rating $\bar{\psi}_1$. By abstracting from this initial step, we simplify the exposition, as it allows us to treat all periods identically.

¹⁷We impose $\underline{\sigma} > 0$ because the rating would be invariant to incoming reviews for $\sigma = 0$.

¹⁸In addition, we rule out one non-generic value of κ , $\kappa \neq (2 - 3\sigma)/(2 - 2\sigma)$.

supposing that past consumers faced the identical situation they find themselves in. A similar assumption is used in Crapis et al. (2017). It is in the spirit of the notion of cursed equilibrium (Eyster and Rabin 2005) that consumers do not fully take the information contained in current observables into account.

Assumption 1 (Quality inference by consumers) *Consumers conduct quality inference by imposing that all past consumers faced the same aggregate rating/price combination they currently see. As such, their inference consists of a pair (μ^*, ω^*) of inferred quality μ^* and inferred cutoff taste ω^* such that*

$$\psi(\mu^*, \omega^e(\omega^*), p_t) = \bar{\psi}_t \quad (\text{CONS})$$

$$u(\mu^*, \omega^*) = p_t. \quad (\text{RAT})$$

Note that inference consists not only of forming a belief about the quality of the good, μ^* but also of the cutoff type of purchasing consumers ω^* . This is because, despite using a heuristic, consumers are cognizant that reviews are driven by the characteristics of past consumers who purchased the good. Inference about the quality cannot be conducted in isolation from an inference about the set of purchasing consumers.¹⁹ Given the inference (μ^*, ω^*) , consumer i purchases if and only if her predicted expected utility $u(\mu^*, \omega_i)$ weakly exceeds the price p_t . It follows immediately from (RAT) that the marginal consumer who purchases coincides with the inferred cutoff type of past purchasing consumers ω^* .

While the inference rule is an obvious simplification that mutes effects such as direct price signaling, it allows us to cleanly carve out the tension between the direct price and selection effect. It greatly improves tractability, reducing inference to a two-dimensional fixed point problem. If all past consumers faced the same scenario as current consumers, the inferred quality-cutoff-pair must be such that the aggregate rating is consistent. Contingent upon purchase, the average review left by consumer $\omega^e(\omega^*)$ given that the purchase price was p_t and quality is correctly believed to be μ^* must be consistent with $\bar{\psi}_t$; see (CONS). Moreover, the cutoff type must have been exactly indifferent between purchasing and not purchasing; that is, her gross utility has to be equal to the price; see (RAT). Note that since utility is increasing in taste ω , (RAT) implies that all purchase decisions in the hypothetical scenario were individually rational. An additional advantage of this updating rule is that it is independent of the distribution of qualities F and its support. We assume that consumers believe that the quality is distributed on \mathbb{R} .²⁰

An alternative way of interpreting Assumption 1 is that consumers conduct quality inference by treating the game as if it were in a stationary equilibrium: they deem the good to be of the quality μ^* such that given the induced cutoff type ω^* , the average review is equal to the current aggregate rating $\bar{\psi}_t$. If the game were in a stationary equilibrium, the belief that past consumers faced the same price and rating combination would be correct. This interpretation also highlights that consumers would have no incentive to delay their purchase even if they could do so—given their inference, they predict the situation to be unchanged in future periods.²¹

¹⁹While consumers in Crapis et al. (2017) explicitly form beliefs about product quality only, they implicitly impose that past consumers had the same quality belief as well; thus, the inferred quality uniquely pins down the cutoff taste of purchasing consumers in the past and present. This, in turn, is used in the computation of the likelihood function.

²⁰Note that this allows consumers, in principle, to believe that the quality exceeds the maximal possible quality $\bar{\theta}$. However, we impose this assumption only for tractability reasons. If we were to bound consumers' beliefs at $\bar{\theta}$, we would have to keep track of these belief boundaries in the analysis leading to tedious case distinctions without generating substantial additional insights. At the same time, our main results hold even if we were to introduce such boundaries.

²¹As consumers in our setup use the described heuristic, they are agnostic about what determines the price charged by the seller. It seems reasonable, however, to suppose that they would expect the same price given an unchanged aggregate rating; in a stationary equilibrium, this supposition would again be correct. Moreover,

2.1. Discussion of Modeling Assumptions

We briefly discuss what we consider to be the two central modeling assumptions. We provide additional extensions to highlight the robustness of our analysis and findings in Section 5.

Consumer Inference While our inference assumption implies that consumers are non-Bayesian, we do not consider it to be far from reality. As discussed previously, past prices are not directly observable on online platforms. Individual reviews are often available, but even with historical price data from price-tracking websites (which by itself is cumbersome to obtain), they cannot be directly linked to the price at which the good was purchased. They also rarely mention explicit price points. As reviews are noisy due to horizontal differentiation, the assumption that consumers base their quality inference only on the aggregate rating and the current price seems realistic for a large set of potential consumers. Given that consumer inference is based on these two inputs only, the heuristic used by treating the posted price as part of a stationary equilibrium seems a reasonable approximation. In addition to the substantial requirements full rationality would place on consumers' cognitive abilities, consumers are often uncertain about how many periods have passed and how often the seller changed prices in the past.²² Note that the inference rule, in principle, gives the seller scope to manipulate consumer inference via drastic price adjustments, as these are by assumption undetected by consumers. However, we show that this does not happen under the optimal pricing policy and that the game always converges to a stationary equilibrium in which consumers correctly infer the quality.

While the assumption facilitates analytical tractability, we show the robustness of our main results in Section 5 and Appendix D. In particular, they continue to hold when we allow for a fraction λ of sophisticated consumers. Moreover, the main insights carry over to a similar two-period model with Bayesian consumers and to an alternative continuous-time setup in which consumers update in a Bayesian fashion.

Partial price internalization Assuming a flexible parameter for the degree of price internalization κ is a deliberate modeling choice. Chakraborty et al. (2022) provide evidence from Yelp.com that only 46.7% of consumers mention the price in their review. In line with this observation, Luca and Reshef (2021) find that reviews depend on prices but that they do not reflect net consumption utility perfectly. Both papers suggest that the degree of consumers' price internalization in reviews is interior, i.e., that $\kappa \in (0, 1)$. There are three main rationales which motivate our modeling choice. We discuss them in more detail in Appendix A.

First, during the review process, the gross consumption utility, $\theta + \omega_i$, is more salient for the consumer than the price they paid, p . Consumers typically review a product with a delay after the purchase as they first have to receive and use it. At the time of reviewing, the consumption utility thus features more recently in their memory than the expenditure. Therefore, an interior price internalization can be rationalized by consumers placing a higher weight on the more salient consumption utility.

Second, consumers may differ in the interpretation of a rating system. As Chakraborty et al. (2022) find, a substantial share of consumers does not consider the price as a central component of their review. Other consumers, however, do and mention the price explicitly. Given that individual reviews are aggregated, heterogeneity in whether consumers consider the price as a

consumers have no incentive to strategically review a product negatively today to purchase it at a lower price tomorrow. In our model, each consumer is negligible, and the individual review will not affect the aggregate rating. In addition, platforms take measures against fake reviews, and some platforms only allow for reviews by actual customers.

²²If the game has a stationary equilibrium and consumers are uncertain about the current time period, their naïve best guess is to be in a stationary period, and our imposed inference is correct.

relevant component of reviews can thus generate an interior aggregate degree of price internalization.

Third, Carnehl et al. (2022) and Jeziorski and Michelidaki (2021) provide evidence that different framings of a rating system can lead to different levels of price responsiveness of reviews. Rating subcategories labeled as value for money feature the highest responsiveness, while reviews in other subcategories feature a lower responsiveness. More generally, there is ample heterogeneity in the guidance by platform operators regarding review content (see, e.g., Amazon 2021, Steam 2021). This heterogeneity is closely related to the consumer heterogeneity discussed in the previous paragraph and is likely to affect it.

These rationales imply that there is a wedge between the purchase decision—which is based on a fully internalized price—and the review—which is based only on a partially internalized price. We solve our model with κ as a parameter and derive implications depending on its specific value, which may vary with the product and application considered. The model thus allows for different empirically documented product-specific differences in the degree to which the product price affects online ratings. Carnehl et al. (2022) show heterogeneity in the price-responsiveness of ratings for listings on Airbnb in different market segmentations. In the airline industry, reputation is driven predominantly by overall quality (gross utility) for full-service carriers, while it is value for money (net utility) that is most relevant for budget airlines (Forgas et al. 2010, Rajaguru 2016). Heterogeneity in the price-responsiveness of ratings for different customer groups has also been documented on Yelp.com (Byers et al. 2012, Li 2016). In our model, the value of κ determines how reviews respond to price changes, as we show next.

2.2. Price and Selection Effect

In this part, we solve the consumers' inference and derive the resulting demand function. We formally characterize the price and selection effect and illustrate dynamic pricing incentives.

Explicit Inference We can solve the equation system characterized by (CONS) and (RAT) for the belief of consumers given a pair $(\bar{\psi}, p)$. The unique solution pair $(\tilde{\mu}, \tilde{\omega})$ is

$$\tilde{\mu}(\bar{\psi}, p) = 2\bar{\psi} - 1 - p(1 - 2\kappa) \quad (5)$$

$$\tilde{\omega}(\bar{\psi}, p) = 1 - 2(\bar{\psi} - p(1 - \kappa)). \quad (6)$$

As all consumers form the same beliefs, $\tilde{\omega}$ corresponds to the taste of the marginal consumer indifferent between purchasing and not purchasing given $\bar{\psi}$ and p , which induces demand

$$q(\bar{\psi}, p) = 1 - \tilde{\omega}(\bar{\psi}, p) = 2(\bar{\psi} - p(1 - \kappa)). \quad (7)$$

A higher aggregate rating always increases demand, while a higher price decreases it. Moreover, the inferred quality μ increases in the aggregate rating $\bar{\psi}$. How demand and the inferred quality respond to price changes depends on the price internalization of reviews κ . When κ is high, reviews reflect net consumer surplus. In this case, demand is unresponsive to price changes—when observing a higher price, consumers rationalize the same rating with higher quality. In contrast, when κ is low, reviews reflect gross surplus. Consumers rationalize a rating at a higher price predominantly with a more favorable selection of consumers, reducing the quality belief.

Induced Average Review: Selection Effect vs. Direct Price Effect To understand the dynamic pricing incentives, we first derive the effect of the current price on the induced reviews. The impact of a marginal price change on the induced average review is given by

$$\frac{d\psi}{dp} = \frac{d\omega^e}{d\tilde{\omega}} \frac{d\tilde{\omega}}{dp} - \kappa. \quad (8)$$

Changing the price has two effects. First, it directly affects the average review, as reviews incorporate the purchase price. This is the marginal direct price effect, which equals $-\kappa$. In addition, there is a selection effect. By changing the price, the seller changes the cutoff consumer via the inference and thus the taste of the average consumer, which determines the average review. As established, the degree to which prices are incorporated into reviews determines the responsiveness of demand and thus the strength of the selection effect. Depending on κ , consumers rationalize a given rating at a given price primarily via the product's quality (high κ) or the taste of purchasing consumers (low κ). We obtain for the marginal selection effect that $\frac{d\omega^e}{d\tilde{\omega}} \frac{d\tilde{\omega}}{dp} = 1 - \kappa$; the larger κ is, the less responsive the demand, and thus, the weaker the selection effect.

Applying the demand function derived in (7), we obtain the average review

$$\psi(\theta, \omega^e(\tilde{\omega}), p) = \theta + 1 - \bar{\psi} + p(1 - 2\kappa). \quad (9)$$

A high average review is induced by high prices whenever κ is low ($\kappa < 1/2$), as the selection effect dominates the direct price effect in this case. When κ is high ($\kappa > 1/2$), the reverse is true, and the direct price effect dominates; low prices induce high average reviews. We summarize this observation in the following lemma.

Lemma 1 (Price & Selection Effect) *The direct price effect dominates the selection effect if and only if $\kappa > \frac{1}{2}$. If this is the case, a price increase decreases the induced average review in the current period. For $\kappa < \frac{1}{2}$, the selection effect dominates, and a price increase increases the induced average review.*

Proof. Proof: The proof follows from the preceding discussion. ■

Lemma 1 is an important building block of our analysis. It shows that the relative strength of the direct price and the selection effect in our model depends monotonically on the degree of price internalization, κ . The model thus predicts a positive (negative) correlation between reviews and prices when price internalization is low (high) in a cross-section where reviews can be linked to the prices at which the good was purchased.²³ We, therefore, provide a novel rationalization for the, at first glance, conflicting evidence in the empirical literature.²⁴

Prices have been shown to be positively correlated with reviews for books (Zegners 2019, Bondi 2019) but negatively correlated with reviews for USB sticks (Cabral and Li 2015) and low-priced Airbnb listings (Carnehl et al. 2022). The evidence for restaurants using data from Yelp.com is mixed (Byers et al. 2012, Li 2016, Luca and Reshef 2021). Overall, the empirical evidence suggests that both a positive and a negative correlation between prices and reviews can arise realistically and that the correlation varies across product categories. We expect the price internalization to be high for products that are (i) standardized and targeted at mass markets instead of niche audiences (USB sticks) or (ii) of overall lower quality in markets with vertical segmentation (low-priced Airbnb listings, budget airlines). Conversely, products where taste

²³We address the correlation of price and rating paths over time once we have characterized the solution to the full model in Section 3.

²⁴The framework by Li and Hitt (2008) is similarly consistent with both empirical correlations between purchase prices and reviews. They rationalize a negative correlation by having the consumers most intrinsically inclined towards the product be the harshest in their reviews. While this behavior appears plausible in certain contexts, Luca and Reshef (2021) shows that selection on harshness does not drive reviews on Yelp.com.

plays a substantial role (eBooks) or high-quality products in markets with vertical segmentation (full-service airlines) are more likely to feature a lower degree of price internalization, which implies that high prices are associated with high reviews. More standardized products also tend to be easier to evaluate, which reduces the time between purchase and review. Thus, they feature a higher degree of price internalization when accounting for the salience of the purchase price.

Based on this interpretation, we conduct a limited empirical investigation of the video game market Steam using a matched data-set of 12000 observations where we can link each review to the corresponding purchase price, see Appendix E. We hypothesize that the direct price effect should be more likely to dominate for casual games (simple, lower-quality mass-market products) than for video games as a whole. While the analysis cannot provide conclusive causal evidence due to data limitations, the findings suggest this is the case. While higher prices are associated with higher review scores for video games, this relationship reverses for casual games, where discounts translate into positive reviews.

The tradeoff between the price and the selection effect constitutes the main economic force that drives the seller's strategic pricing incentives. In particular, our results qualitatively differ depending on the dominant effect. From this perspective and given the empirical literature that finds both effects to be plausibly dominant for different products, the flexible product-specific parameter κ is an important modeling tool.

It is furthermore important to note that our model is qualitatively consistent with the empirically well-documented pattern that reviews move opposite to expectations (see, e.g., Ho et al. 2017, Hui et al. 2021, Luca and Reshef 2021). To see this, note that (CONS) requires the aggregate rating to be consistent given the inferred quality and cutoff taste, implying that the difference between the induced average review $\psi(\theta, \omega^e(\tilde{\omega}), p)$ and the current aggregate rating $\bar{\psi}$ is given by

$$\psi(\theta, \omega^e(\tilde{\omega}), p) - \bar{\psi} = \psi(\theta, \omega^e(\tilde{\omega}), p) - \psi(\tilde{\mu}, \omega^e(\tilde{\omega}), p) = \theta - \tilde{\mu}. \quad (10)$$

When considering how the induced average review is affected by a price change, we thus obtain

$$\frac{d\psi}{dp} = \frac{d\psi - \bar{\psi}}{dp} = -\frac{d\tilde{\mu}}{dp}. \quad (11)$$

Whenever consumers have higher expectations about the product quality ($\tilde{\mu} \uparrow$), their reviews are worse ($\psi \downarrow$), and vice versa.²⁵

Illustration in a Two-Period Model To illustrate the impact of the selection and price effect on pricing incentives, consider a two-period version of the model ($T = 2$) with initial rating $\bar{\psi}_1$. Denote the aggregate rating in period 2 by $\bar{\psi}_2$. In period 2, the seller maximizes

$$\max_{p_2} p_2 \cdot q_2 = p_2 \cdot 2(\bar{\psi}_2 - p_2(1 - \kappa)) \quad (12)$$

and thus chooses the myopic monopoly price given by $p_2 = \frac{\bar{\psi}_2}{2(1-\kappa)}$, which yields profits $\pi_2 = \frac{\bar{\psi}_2^2}{2(1-\kappa)}$, highlighting that second-period profits are increasing in the rating at the beginning of that period. Using these continuation profits in period 1, we can write the total profits as

$$\pi_1 = p_1 q_1(\bar{\psi}_1, p_1) + \delta \pi_2(\bar{\psi}_2(p_1)). \quad (13)$$

²⁵Note that this behavior would be reinforced if, e.g., the price internalization level in reviews were explicitly related to consumer expectations so that consumers who are negatively surprised by the product's quality judge it more harshly. While incorporating an endogenous harshness of reviews based on expectations is beyond the scope of this paper, we view it as an interesting avenue to explore in future work.

To understand the seller's pricing problem in period 1, we have to derive the effect of current prices on future profits. Denoting as ψ_1 the average review in the first period, we have $\bar{\psi}_2 = \sigma\psi_1 + (1 - \sigma)\bar{\psi}_1 \stackrel{(9)}{=} \sigma((\theta + 1) - \bar{\psi}_1 + p_1(1 - 2\kappa)) + (1 - \sigma)\bar{\psi}_1$ and

$$\frac{\partial \pi_2(\bar{\psi}_2(p_1))}{\partial p_1} = \frac{\partial \pi_2}{\partial \bar{\psi}_2} \frac{\partial \bar{\psi}_2}{\partial \psi_1} \frac{\partial \psi_1}{\partial p_1} = \sigma \frac{\bar{\psi}_2(p_1)}{(1 - \kappa)} (1 - 2\kappa). \quad (14)$$

Plugging this into the first-order condition yields the optimal period-1 price

$$p_1^* = \frac{\bar{\psi}_1}{2(1 - \kappa)} + \delta \sigma \frac{\bar{\psi}_2(p_1)}{2(1 - \kappa)^2} (1 - 2\kappa). \quad (15)$$

If the seller were to price myopically in the first period, the seller would charge $\frac{\bar{\psi}_1}{2(1 - \kappa)}$. The direction of the distortion is therefore determined by κ . If $\kappa > 1/2$ ($\kappa < 1/2$), the price is lower (higher) than the myopic optimum. As we will see in the next section, these incentives to manage reputation via the pricing decision persist in the infinite horizon model and in conjunction with the sensitivity of the rating system determine the long-run steady state outcome.

3. Optimal Dynamic Pricing

Having established the pricing incentives in a simple two-period version of our model, we move to an infinite horizon to understand both the long-run properties of the game and the transitory dynamics. We show that the game always converges to a stationary equilibrium and that long-run profits and consumer surplus, as well as the speed of consumer learning, are strongly affected by the rating system's sensitivity to new reviews, σ .

The seller solves the problem

$$\max_{(p_t)_{t \geq 0}} \sum_{t=0}^{\infty} \delta^t p_t q(p_t, \bar{\psi}_t) \quad (16)$$

$$\text{s. t. } \bar{\psi}_{t+1} = (1 - \sigma)\bar{\psi}_t + \sigma\psi(p_t, \bar{\psi}_t) \quad (17)$$

$$\bar{\psi}_1 = \bar{\phi} \geq 0. \quad (18)$$

Note that the flow profits are bounded and that, because $\delta \in (0, 1)$, the problem is well defined, and we can write it as a dynamic programming problem (see Stokey et al. 1989, Section 4 and the Appendix). The Bellman equation for this problem is

$$V(\bar{\psi}) = \max_p \left\{ pq(p, \bar{\psi}) + \delta V(\bar{\psi}') \right\} \quad (19)$$

$$\text{s. t. } \bar{\psi}' = (1 - \sigma)\bar{\psi} + \sigma\psi(p, \bar{\psi}). \quad (20)$$

To see the dynamic pricing incentives, consider the derivative of the Bellman equation with respect to the current price.

$$\underbrace{q(p, \bar{\psi}) + p \frac{dq(p, \bar{\psi})}{dp}}_{\text{static monopoly}} + \delta \underbrace{\frac{dV(\bar{\psi}')}{d\bar{\psi}'}}_{\substack{\text{effect of} \\ \text{rating on} \\ \text{future profits}}} \underbrace{\frac{\partial \bar{\psi}'}{\partial \psi} \frac{\partial \psi}{\partial p}}_{\substack{\text{effect of} \\ \text{current price} \\ \text{on future rating}}} = 0. \quad (21)$$

There are two effects. First, flow profits are affected by the increase in the price. Standard static monopoly price effects capture this effect. Second, the price change impacts future profits via

the change in the induced rating. This effect can be decomposed into the (discounted) sensitivity of the future profits to the aggregate future rating ($\frac{dV'}{d\bar{\psi}'}$), the sensitivity of the aggregate rating in the next period to the induced current review ($\frac{\partial \bar{\psi}'}{\partial \psi} = \sigma$), and the effect of the price change on the current review ($\frac{\partial \psi}{\partial p}$).

To solve the problem, we replace the control p using the law of motion for the state $\bar{\psi}$ and treat $\bar{\psi}'$ as the seller's choice variable. The review ψ in any period is linear in p given the current aggregate rating $\bar{\psi}$ and is given by

$$\psi = \theta + 1 - \bar{\psi} + p(1 - 2\kappa), \quad (22)$$

which implies that we can replace the control p with $\bar{\psi}'$, as

$$\bar{\psi}' = (1 - \sigma)\bar{\psi} + \sigma(\theta + 1 - \bar{\psi} + p(1 - 2\kappa)) \iff p = \frac{\bar{\psi}' - \bar{\psi}}{\sigma(1 - 2\kappa)} - \frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa}. \quad (23)$$

The problem of the seller can then be written as

$$V(\bar{\psi}) = \max_{\bar{\psi}'} \left\{ p(\bar{\psi}')q(\bar{\psi}', \bar{\psi}) + \delta V(\bar{\psi}') \right\}. \quad (24)$$

We solve the problem by guessing and verifying that the value function is of the form $V(\bar{\psi}) = c + d\bar{\psi} + e\bar{\psi}^2$, which implies a linear policy function $\bar{\psi}' = a + b\bar{\psi}$. We obtain closed-form solutions for the optimal policy and value function as well as the long-run prices and ratings. This analytic solution allows us to characterize the long-run equilibrium outcomes to which the system converges.

Proposition 1 *For sellers of type $\theta > -1$ and $\bar{\psi}_1 \geq 0$, there is a unique stationary equilibrium that is characterized by long-run ratings, prices, and beliefs with*

$$\Psi = \frac{(\theta + 1)(2(1 - \delta)(1 - \kappa) + \delta\sigma(3 - 2\kappa))}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)} \quad (25)$$

$$\tilde{p} = \frac{(\theta + 1)((1 - \delta) + 2\delta\sigma)}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)} \quad (26)$$

$$\tilde{\mu} = \theta. \quad (27)$$

Hence, the rating system is effective, and consumers learn the quality.

Proof. Proof: See Appendix B. ■

Proposition 1 shows that the market converges to a unique stationary equilibrium. The rating system effectively alleviates the asymmetric information problem, and consumers learn the quality of the good in the long run. However, despite consumers learning θ in the long run, the long-run prices depend on the details of the rating system; i.e., they depend on σ . In particular, the product-specific degree of price internalization κ determines how they qualitatively depend on this sensitivity. We discuss this dependence in Corollary 1.

Moreover, sellers with qualities that are so low that they could not sell under full information ($\theta < -1$, which implies that even a consumer with $\omega = \bar{\omega} = 1$ enjoys a negative gross utility) eventually leave the market. They always price such that the rating is declining over time until they cannot make a profit.²⁶

²⁶Recall that demand is given by $q(\bar{\psi}, p) \stackrel{(7)}{=} 2(\bar{\psi} - p(1 - \kappa))$. A rating $\bar{\psi} \leq 0$ implies that the seller cannot sell at non-negative prices and, therefore, cannot generate any reviews. Moreover, the average review conditional on quality, rating and the current price is given by $\psi(\theta, \omega^e(\bar{\omega}), p) \stackrel{(9)}{=} \theta + \omega^e - \kappa p$. From $\omega \sim U[0, 1]$, we have $\omega^e \leq 1$, and thus reviews will always be negative for $\theta < -1$. Thus, the rating for a seller with $\theta < -1$ will decrease until $\bar{\psi} \leq 0$, at which point the seller exits.

Using (26), we can assess how long-run prices are affected by the rating system as parametrized by σ . As consumers' quality inferences are correct, the price level directly determines the long-run consumer surplus—a higher price at the same quality is associated with a higher cutoff type due to (CONS) and thus decreases consumer surplus. Moreover, we can assess the effect of σ on the long-run profits

$$\tilde{\pi} = \tilde{p} \cdot q(\Psi, \tilde{p}) = \frac{2((1-\delta) + 2\delta\sigma)(\delta\sigma + (1-\delta)(1-\kappa))}{(4\delta\sigma + (1-\delta)(3-2\kappa))^2}(\theta + 1)^2. \quad (28)$$

Corollary 1 *The comparative statics with respect to the sensitivity of the rating system, σ , are as follows.*

- (a) *Prices are increasing in σ whenever the direct price effect in the reviews is not too large and decreasing otherwise; $\frac{d\tilde{p}}{d\sigma} > 0$ when $\kappa < \frac{1}{2}$, and $\frac{d\tilde{p}}{d\sigma} < 0$ when $\kappa > \frac{1}{2}$.*
- (b) *Consumer surplus is decreasing in σ when the direct price effect in the rating is not too large and increasing otherwise; $\frac{d\tilde{CS}}{d\sigma} < 0$ when $\kappa < \frac{1}{2}$, and $\frac{d\tilde{CS}}{d\sigma} > 0$ when $\kappa > \frac{1}{2}$.*
- (c) *Long-run profits $\tilde{\pi}$ are strictly increasing in σ for $\kappa \neq \frac{1}{2}$.*
- (d) *The long-run rating Ψ is increasing in σ .*

Proof. Proof: See Appendix B. ■

Corollary 1 contains the first main implications of the paper. In the stationary equilibrium, the seller balances flow payoff and future profit considerations, which amounts to balancing the exploitation of the current rating and strategic reputation management via the induced reviews. The less sensitive to new reviews the rating system is, i.e., the lower σ is, the more the seller has to invest by deviating from the myopically optimal price to obtain a given next-period rating. Therefore, a lower sensitivity has an unambiguously negative effect on the seller's profits.

To build intuition, consider a stationary equilibrium with a particular sensitivity σ . After an increase in the sensitivity, the seller could continue charging the previously optimal long-run price, and the ratings would stay constant. However, the seller can now obtain a higher rating by a smaller deviation from the old long-run price. In line with this reasoning, the long-run rating Ψ is increasing in σ .

The price level alone determines the long-run consumer surplus because the quality inferences are correct for consumers. The pricing incentives, in turn, depend on whether the direct *price effect* or *selection effect* dominates, which is determined by the degree κ to which the reviews internalize the purchase price. For low κ , the price has only a small direct effect. The selection effect is more relevant, and the seller has an incentive to price higher than the myopic optimum to induce higher future ratings (see also the two-period case in (15)). This incentive is mitigated by a lower σ , as a particular period has a smaller effect on the rating. The price level hence is increasing in σ , and consumers benefit from having the rating reflect past purchases equally instead of putting more weight on more recent reviews.

The converse is true when κ is large. The direct price effect dominates, future profit considerations incentivize the seller to price below the myopically optimal price, and a high σ benefits consumers as it increases the relevance of future considerations in the seller's optimization problem.

Price and Ratings Paths & Speed of Convergence The closed-form solution for the optimal policy allows us to analytically assess the comovement of prices and ratings over time and the speed of consumer learning. The consumer’s learning speed in our setup corresponds to the convergence speed of the sequence of ratings and beliefs. Recall that the value function of the form $V(\bar{\psi}) = c + d\bar{\psi} + e\bar{\psi}^2$ translates into a law of motion for the aggregate rating $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$ and that we derive analytical expressions for the parameters a and b in Appendix B.

Proposition 2

- (i) *Prices and ratings always comove. The price and ratings paths are monotonic provided the sensitivity of the rating system is not too high, $\sigma < \bar{\sigma} := \max\{\frac{1}{2}, \frac{2-2\kappa}{3-2\kappa}\}$. Otherwise, the price and rating paths feature cycles.*
- (ii) *Ratings and beliefs converge linearly at rate $|b|$. The speed of convergence is increasing in the sensitivity for $\sigma < \bar{\sigma}$, and decreasing in the sensitivity for $\sigma > \bar{\sigma}$.*

Proof. Proof: See Appendix B.3. ■

Proposition 2 is important for several reasons. First, it establishes that prices and ratings always comove. Whenever the aggregate rating increases from one period to the next, the price follows suit, and vice versa. Second, the shapes of the price and rating paths depend on the sensitivity of the rating system and the degree of price internalization. Third, the speed of learning non-trivially depends on the rating system’s sensitivity, which is particularly important for platforms that want to attract consumers. We discuss these observations in more detail below.

A seller should always raise (lower) her price in response to a rating increase (decrease) independent of whether the price or selection effect dominates. To understand this, recall that the price simultaneously determines flow profits and manages future ratings. Strategic ratings management induces the seller to distort the price away from the static monopoly price, with the direction of the distortion depending on the relative strength of the price and selection effect. However, an increased (decreased) aggregate rating shifts the demand curve outward (inward) and thereby increases (decreases) the myopic monopoly price. The corollary establishes that the latter effect always dominates. Even if a dominant selection effect may intuitively suggest to increase prices following a drop in the aggregate rating—to rebuild reputation—this is not the optimal dynamic price policy. The dominant selection effect only attenuates the price decrease.

The shape of the price and rating paths critically depends on the sensitivity of the rating system to incoming reviews, as well as the initial attitude of consumers. If the rating system is relatively insensitive, price and rating paths are monotonic. When the initial attitude towards a product is low relative to the actual quality, $\bar{\psi}_1 < \Psi$, the seller gradually builds up her reputation. The initial price is low and increases as the aggregate rating improves. Analogously, the seller gradually milks her excessive reputation whenever the initial attitude towards a product is high, and prices and ratings decrease over time.

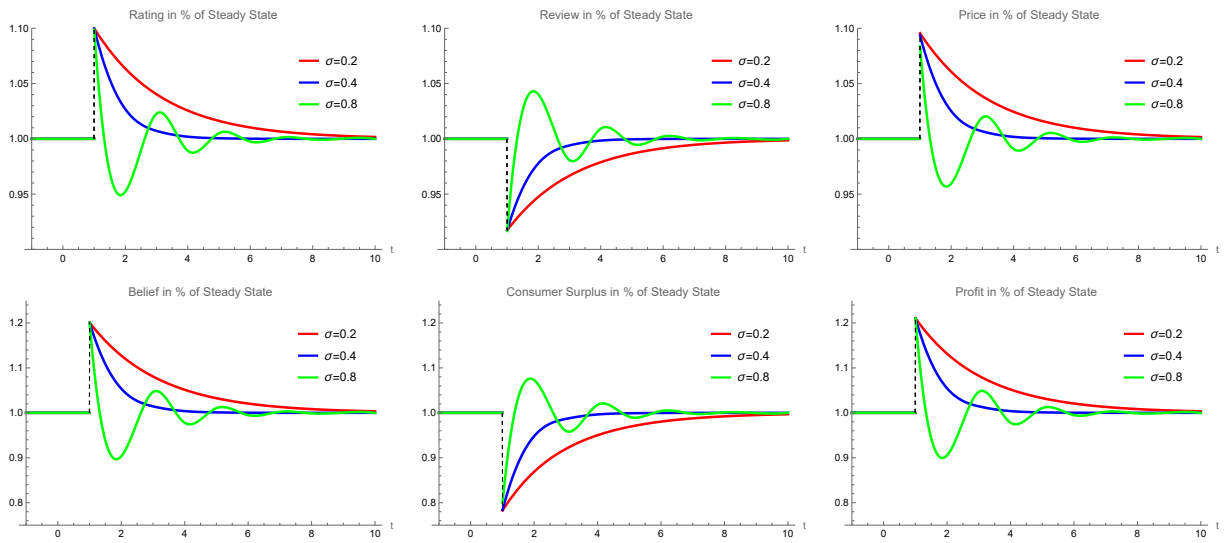
This behavior markedly differs when the rating system is relatively sensitive to incoming reviews: the seller has an incentive to alternate between building up and milking her reputation strategically, and price and rating cycles emerge. Nonetheless, there is convergence in the long run; the amplitude of the cycles shrinks over time. Importantly, cycles can arise for any degree of price internalization κ in reviews.²⁷ As the cutoff sensitivity $\bar{\sigma}$, above which cyclicity arises, is weakly decreasing in κ , it is relatively more likely that cycles arise for products featuring a dominant price effect.

²⁷This follows immediately from observing that the cutoff sensitivity $\bar{\sigma}$ lies strictly below the upper bound given by the technical assumption $\kappa < 1 - \frac{\sigma}{2} \iff \sigma < 2 - 2\kappa$.

Quality beliefs converge linearly regardless of the sensitivity of the rating system. However, the marginal effect of an increase in sensitivity on the speed of convergence depends on whether the paths are monotonic or cyclical. In the former case, a marginal increase in sensitivity increases the speed at which beliefs converge: building up to the “correct” long-run reputation is facilitated by a higher sensitivity whenever the initial rating stock is too low, and milking an excessively high initial rating stock leads to a quicker adjustment. In contrast, an increase in sensitivity has a detrimental effect on the speed of convergence if the seller engages in cyclical build-up and milking of reputation. In this case, a higher sensitivity facilitates the build-up and exploitation, amplifies the cycles, and slows the speed of convergence.

We illustrate these results in Figure 1, which depicts the reaction of the relevant outcomes to a positive shock to the steady state aggregate rating for a non-standardized product for which the selection effect dominates.²⁸ The seller benefits from the initial shock irrespective of the rating system’s sensitivity. She charges a higher price and reaps additional profits at the expense of consumer surplus because consumers mistakenly infer the product to be of higher quality than it actually is. When the rating system is sufficiently insensitive to incoming reviews ($\sigma = 0.2$ and $\sigma = 0.4$), the price and rating paths are monotonic, and the seller gradually milks the excessive rating. The more sensitive the rating system is, the more gradual this milking. However, for sufficiently sensitive rating systems, two changes materialize. First, the seller strategically alternates between building up and milking the aggregate rating. Second, this cyclicity in the induced price and rating paths harms the speed of convergence. Learning is slower than under less sensitive rating systems that do not feature cyclicity.²⁹

Figure 1: Impulse response functions given a 10% shock to the steady-state rating



4. Implications

In this section, we discuss implications from our analysis, both for sellers deciding about their pricing strategies and platform operators designing their rating systems.

²⁸The specific parameters are $\kappa = 0.4$, $\delta = 0.95$ and $\theta = 1$. The figure displays all measures in percentage terms of the long-run steady state outcome; the changing sensitivity naturally also impacts the level of the long-run price, profits, rating, and consumer surplus—given that we consider the case of a non-standardized product with a dominant selection effect, prices are increasing and consumer surplus decreasing in sensitivity.

²⁹The impact of the sensitivity becomes even more prominent when considering the shock affecting the in-period review only, see Figure 2 in Appendix C. In this case, a higher sensitivity naturally leads to a larger impact on the aggregate rating and hence on the initial distortion in prices, beliefs, profits, and consumer surplus.

4.1. Dynamic Pricing and Ratings Management

Online platforms facilitate a frequent adjustment of prices, especially compared to traditional brick-and-mortar stores, where word-of-mouth plays a similar role to reviews on online platforms (see, for example, Gorodnichenko and Talavera 2017). Consequently, prices are a valuable strategic tool not only to adjust to changes in demand—e.g., due to changes in ratings—but also to manage the sellers’ reputation by using current prices to affect future ratings persistently. Based on our analysis, we discuss how sellers can profit from frequent price adjustments accounting for the effect of prices on ratings and highlight the practical differences to alternative pricing strategies.

Our analytical characterization of the optimal dynamic pricing strategy allows us to compare it to two natural benchmarks: one in which the seller dynamically adjusts prices but ignores the effect current prices have on future ratings, and one in which she commits to a constant price over time (as in, e.g., Li and Hitt 2008, Crapis et al. 2017). The detailed derivations are in Appendix B.5.

Naturally, sellers benefit from strategic pricing in our setup: current prices affect future ratings and, thereby, future profits. Accounting for this interaction in the pricing decision cannot make sellers worse off, i.e., myopic dynamic pricing must be worse than strategic dynamic pricing. Notably, the comparison yields practical implications about how sellers should respond to changes in the rating. While sellers should increase prices when ratings go up (see Proposition 2), the magnitude of the price reaction depends on the degree of price internalization in reviews.

Suppose that price internalization is high such that the price effect dominates the selection effect.³⁰ Two considerations determine the seller’s response to an increase in the rating. First, the rating increase shifts demand outward, leading to upward pressure on the price. Second, the ratings management incentive counteracts the upward pressure as higher prices lead to worse reviews. A strategic seller considers the latter and will respond with a more moderate price increase compared to a myopic seller. If the rating system is more sensitive, the weight on the ratings management incentive is more significant and the price increase even smaller. In contrast, if a seller faces a rating decrease, she will respond with a price decrease. In contrast to a myopic seller, the strategic seller will react more aggressively and offer a substantial discount; the more sensitive the rating system, the larger the discount.

These differences illustrate that a seller that wants to use a dynamic, rating-dependent pricing strategy to balance flow-profit maximization with ratings management must know the degree of price internalization for reviews of her product on the platform. If uncertain, she should invest in price experimentation to learn about this parameter. The value of such experimentation is the difference between the value from flow-profit maximization and the value of our optimal dynamic pricing strategy, which we compute in Appendix B.5.

We can also compare the optimal dynamic pricing strategy with a fixed-price strategy which accounts for the effect of prices on reviews. This comparison generates insights in settings where price adjustments are sufficiently expensive (e.g., due to monitoring costs), so their benefits need to be weighed against their cost.

Naturally, total discounted profits are higher if the seller employs a dynamic and strategic pricing strategy than if she commits to a fixed price. With a fixed-price strategy, the seller has to balance the exploitation and management of reputation over time with a single price. Suppose that the initial rating is relatively high and that the price effect dominates. In this case, the seller has an incentive to milk the initial reputation with relatively high prices. As ratings, and thus beliefs, always converge in the long run, the seller cannot exploit the high initial rating too much as excessive long-run prices would lead to low long-run profits. By following a dynamic

³⁰The case with low price internalization and a dominant selection effect is the reverse analog of the case discussed.

and strategic pricing strategy, the seller can finetune prices to the current rating and thereby extract higher profits. Moreover, because consumers are short-lived, there is no cost to the lack of commitment relative to setting a fixed price across periods.³¹

Therefore, the seller benefits from the dynamic pricing strategy by balancing reputation build-up and exploitation incentives over time. If there are periods of excessively high ratings, the seller can choose high prices, and once ratings start declining, adjust the price to the new reputation level. More generally, the exact mode of balancing these considerations depends on the sensitivity of the rating system: sufficiently sensitive rating systems induce sellers to engage in cyclical pricing, e.g., via frequent sales, to boost ratings before capitalizing on the high rating stock (see Proposition 2). When rating systems are less sensitive, sellers more generally prefer gradual price adjustments.

4.2. Platform-Optimal Sensitivity of Rating System

Our analysis also yields implications for platform operators who seek to design an optimal rating system. We focus on the sensitivity of the rating system to incoming reviews, σ , as the design parameter that can be adjusted. We consider the platform's objective to be an increasing function of seller profits (as platforms participate through fees and royalties), consumer surplus (as consumers like to obtain the same product for lower prices) and the speed of consumer learning (as better quality inferences attract a larger customer base).

Our results demonstrate how the sensitivity choice affects these three components. A more sensitive rating system benefits sellers as it simplifies reputation management. At the same time, consumers might suffer from an increased sensitivity because it leads to upward pressure on long-run prices for products with low price internalization in reviews. Moreover, the speed of learning by consumers is maximized at intermediate sensitivities as highly responsive rating systems lead to price and rating cycles. While a derivation of the optimal sensitivity requires parametric assumptions about the platform's preferences and the distribution of product features on the platform, we can nevertheless discuss how changes in the platform's objective affect the optimal sensitivity. Appendix B.6 contains a simplified example in which the platform maximizes a weighted average of consumer surplus and seller profits for a single product or unrelated but homogenous products.³²

For example, as a platform grows and matures, it is likely to eventually place a lower weight on the consumer side because it has already established a loyal customer base (which may be locked in, e.g., due to network effects or switching costs). When the platform correspondingly shifts to focus increasingly on seller profits, it should increase the sensitivity of the rating system. This observation is in line, for example, with recent changes to the rating systems of Amazon (in 2015, see wired.com 2019) and Steam (in 2016, see Steam 2016). Importantly, these shifts are likely to have harmed a substantial amount of consumers by raising prices for products with a dominant selection effect.

Our results can help identify markets suitable for potential entry by platform operators or existing platforms that can improve their performance by adjusting their rating systems. One such example is that of vacation rentals. As the arguably dominant platform, Airbnb features a rating system in which the salient aggregate rating is a simple average of all past reviews; that is,

³¹It is noteworthy that the optimal fixed price can nevertheless lead to higher profits in the long-run steady state. This is because the seller chooses the fixed price to maximize total discounted profits, which leads to long-run profits that depend on the chosen price. For any chosen price, the rating eventually converges, implying that consumer beliefs converge to the actual quality. The seller could thus always commit to the full-information monopoly price and obtain the associated monopoly profits in the long run but distorts the price away from the full-information price if she discounts the future.

³²Accounting for the speed of convergence would greatly complicate the exposition of the forces at play.

more recent reviews receive the same weight as old reviews in a listing’s rating. Given empirical evidence that the price effect overall dominates the selection effect on Airbnb (see Carnehl et al. 2022), our results suggest that an entrant offering a rating system that emphasizes recent reviews may be desirable for both consumers and hosts.³³

Further rating design considerations Another—unmodelled—effect that enters platforms’ considerations is that a rating system sensitive to incoming reviews is less robust to “outlier reviews”. We show in Section 5 that our results are unaffected when allowing for reviews to be noisy; adding noise does not affect the interplay between sensitivity and long-run outcomes within our framework. This finding implies that the only added consideration when reviews are noisy is the trajectory after a large shock. Under a more sensitive rating system, noise is naturally more impactful. However, the trajectory is also affected by the seller’s dynamic incentives, as discussed in Proposition 2, and the variance in consumer beliefs is not necessarily lower with a less sensitive rating system. The reason is that a decrease in the sensitivity may slow consumers’ learning as sellers adjust their pricing strategies.

The platform has to consider its entire product portfolio when considering its rating system. Likely, it will feature both products with a dominant selection effect and products with a dominant price effect. In this case, the platform may use different instruments to mitigate adverse and amplify desirable consequences of a particular rating sensitivity. A natural example of this is the steering of consumers to particular products, e.g., via the order in which search results are displayed. By identifying the product characteristics such that consumers benefit from or are detrimentally affected by the rating system, our results inform the design of the algorithm that governs steering. For example, the platform may steer consumers towards products with a high selection effect when the sensitivity is high and when the platform cares more about sellers than consumers. Conversely, it may steer them away from such products if the platform is more interested in consumer surplus. In general, we view this as an interesting avenue for future research, as steering effects feed back into the choice of the optimal sensitivity of the rating system.

4.3. Testable Implications & Empirical Considerations

Both Corollary 1 and Proposition 2 provide testable implications of our model that allow for an assessment of its validity and can guide policy and managerial recommendations, respectively. Comparing outcomes before and after a change in rating systems towards increased sensitivity, the model predicts that the average rating of products should increase cross-sectionally. In addition, price levels should decrease for standardized products (high κ) and increase for less standardized products (low κ). Unfortunately, we are not aware of datasets that contain the detailed price and review data necessary, allow to link individual reviews to their purchase prices, and cover periods of changes in the design of rating systems along the sensitivity dimension; the Steam data used in our brief empirical analysis in Appendix E only covers the year 2017 which is after Steam started displaying the recent average review score.

We obtain additional predictions from Proposition 2 regarding different price and rating paths depending on consumers’ initial attitude towards a product. When a product is believed to be of too high quality initially, both the rating and price should decline over time. However, when consumers are initially skeptical, the seller of a high-quality product should introduce the product at a low price, which should gradually increase as the rating improves. Finally,

³³The issue of overcoming an incumbent’s advantage due to network effects and switching costs nonetheless remains. In addition to the salient aggregate rating, Airbnb also offers various rating subcategories that differ in their sensitivity to the purchase price; see Carnehl et al. (2022) for details. This may similarly help to explain the lack of movement towards a more sensitive rating system.

Proposition 2 provides guidance on inferring either the sensitivity of the rating system, if the researcher has a good idea about the price internalization κ , or the degree of price internalization if the researcher has a good idea about the rating system’s sensitivity. If the paths feature cycles, the sensitivity must be relatively high given some κ . If the paths do not feature cycles, but the rating system is known to be relatively sensitive, the price internalization in the reviews must be relatively low.

Drawing these inferences is important as it allows to predict, e.g., the performance of potential interventions by policymakers or the optimal entry pricing strategy for a new product.

5. Extensions

While our baseline model allows the derivation of sharp results due to its tractability, it abstracts from several effects at play in reality. In this section, we discuss several extensions and illustrate the robustness of our findings. We are predominantly interested in the interplay of prices, ratings, and rating design and its consequences for the seller and her optimal pricing strategy and platforms designers and the optimal sensitivity of the rating to incoming reviews.

Thus, we mainly illustrate the robustness of the fundamental forces behind Corollary 1. Suppose the seller has a dynamic incentive to lower her price to increase future aggregate ratings and thus profits. An increase in the sensitivity of the rating system amplifies this downward price pressure. Therefore, the long-run price level decreases in the sensitivity. Consequently, consumers benefit in any such stationary equilibrium when they correctly infer the product quality. A more sensitive rating system negatively affects consumers via increased price levels if the dynamic incentive instead yields an upward pressure on prices.

Because the direction of sellers’ strategic price distortions is governed by the relative strength of the price and the selection effect, the subsequent discussions focus on how the modifications to the baseline model impact this tradeoff, that is, whether dynamic considerations induce upward or downward pressure on the price. Throughout, we concentrate on the strategic forces at play and relegate the mathematical derivations to Appendix D.

Number of reviews We can allow the rating updating rule to account for the number of reviewing consumers proxied by the per-period quantity, see Appendix D.1. Because higher quantities always obtain from lower prices, this introduces an asymmetry in the impact on price and rating paths depending on whether the price or selection effect dominates. Accounting for the number of reviews leads to increased long-run prices for products with a dominant selection effect but decreased long-run prices for products with a dominant price effect. Aside from this, the fundamental tradeoff that determines the direction of the seller’s strategic pricing incentives is unaffected.

Non-uniform tastes & review behavior Our results are also robust to allowing consumers’ idiosyncratic tastes not to be uniformly distributed or to have the likelihood of providing a review depending on their taste or the set of purchasing consumers (this allows to rationalize reviewing behavior bimodal in consumer satisfaction as documented by, e.g., Bolton et al. 2004, Dellarocas and Wood 2008, Hui et al. 2021). Both modifications alter the selection effect which reflects these considerations, see Appendix D.2 and Appendix D.3 for details. While it depends on distributional assumptions whether the selection effect is amplified or attenuated, the fundamental tension with the direct price effect persists and determines sellers’ dynamic pricing incentives.

Heterogeneous degree of price internalization in reviews We can also allow the degree of price internalization to differ across consumers and potentially be correlated with their idiosyncratic tastes. This requires consumers and the seller to form an expectation about the reviewing consumers’ average degree of price internalization. We show in Appendix D.4 that this introduces a new source of price pressure: If taste and the degree of price internalization are negatively correlated, a high price selects consumers with not only high taste but also those with a low price internalization which improves review scores. This amplifies the upward pressure on the price from the selection effect. Naturally, downward pressure on prices arises if taste and degree of price internalization are positively correlated, i.e., if high-taste consumers more prominently assess the value for money a good provides.

Sophisticated consumers & noisy reviews Our results are also robust to letting a fraction of consumers be sophisticated in that they can back out the quality of a product by tracking recent prices and observing recent reviews (see Appendix D.5) and to modifying the review functions to exhibit noise (see Appendix D.6). We can also incorporate non-linear price effects. While such adjustments preclude us from obtaining closed-form expressions for the long-run stationary equilibrium, we can still solve the model numerically by value function iteration and verify that consumers nevertheless learn the quality of the product in the long run. For example, when the price is evaluated relative to some reference price, $\psi_t = \theta + \tilde{\omega}_t^e - \kappa(\bar{p} - p_t)^2$, the effect of the price on the resulting review depends on both κ and how deep the discount is.³⁴ The higher κ is, the lower the critical discount level such that higher discounts induce better reviews.

Competition The effects also carry over to a competitive setting, both when consumers’ reviews are based only on the price and quality of the product they purchased and when they use the rival product’s price as a reference price (reviews reflecting a reference price also features, e.g., in Li and Hitt 2010). Sellers’ pricing, in this case, takes additional forces into account. First, the own price affects the selection of purchasing consumers of the rival; in particular, raising the own price shifts consumers to the rival that have a preference lower than average. This gives an incentive to raise prices. However, when the own price, in addition, serves as a reference price for the rival’s, there is an incentive to decrease prices to affect the rating scores of the rival negatively. We show that this reinforces strategic pricing considerations in that the rival’s rating is lowered by a price increase if and only if the own rating benefits from it, see Appendix D.7.

Bayesian consumers In our setup, consumers are non-Bayesian and use a heuristic for quality inference. While we believe that the heuristic is a sensible approximation of real-world consumer behavior, the same tradeoff between price and selection effect materializes when consumers are Bayesian (but ignore direct price signaling considerations) and base their quality inference on noisy ratings, as we discuss in Appendix D.8.³⁵ This is because it is the resolution of the tradeoff between price and selection effect, which shifts the mean of the distribution of the induced aggregate rating.

Heterogeneous tastes for quality Finally, we can let consumers differ in their marginal valuation of quality. While this leaves qualitative findings of our baseline model unaltered, see

³⁴In particular, the effect of the discount on the reviews is given by $\frac{d\psi(p, \bar{\psi})}{d(\bar{p} - p)} = 4\kappa(\bar{p} - p) - 1$. More generally, Stenzel and Wolf (2016) provides conditions under which consumer inference in the vein of the present paper is uniquely determined for a flexible class of potentially nonlinear utility and review functions.

³⁵When consumers are fully Bayesian, some noise component in ratings is necessary as they would otherwise perfectly learn product quality from observing the last period’s average review only.

Appendix D.9, it introduces a dependency of the dominant effect on the current rating.³⁶ This is because milking a high reputation via a high price is more costly with vertical differentiation: A high price selects consumers who are, on average, very sensitive to quality and who heavily penalize the seller when their expectations are not met, thus weakening the selection effect when ratings are high. Overall, as higher quality products enjoy better ratings, the price effect is thus more likely to dominate for these products. Note that this observation is conditional on a particular degree of price internalization in reviews and is thus not at odds with the empirical findings that show that price internalization in reviews is lower for high-quality products in markets with vertical segmentation.

6. Conclusion

We develop a model of dynamic pricing in the presence of rating systems. Our model flexibly captures two key effects of prices on ratings: a *selection effect*—a higher price improves reviews as it induces consumers to purchase who are more positively inclined towards the product—and a direct *price effect*—a higher price directly lowers reviews as consumers evaluate a good partially based on its purchase price.

We solve the infinite-horizon dynamic pricing problem and show that the equilibrium is unique and depends on two important economic parameters: (i) the degree of price internalization in reviews and (ii) the sensitivity of the rating system to incoming reviews. Consumers correctly infer the quality of the product in the long run despite using a misspecified model. The characterization allows us to study properties of the long-run outcomes, such as prices, profits, and consumer surplus, as well as properties of the trajectory to the long run, such as the speed of consumer learning and the price and rating paths. We derive implications for sellers, platforms, and regulators and highlight that they depend on the relative strength of the price and the selection effect.

We show that sellers benefit from engaging in dynamic ratings-based pricing and that reputation management concerns asymmetrically impact the price adjustments following shifts in demand due to rating changes. Implementing sophisticated dynamic pricing strategies requires knowledge about the product-specific degree of price internalization in customer reviews and may encourage price experimentation.

Platforms, in turn, need to decide on the sensitivity of aggregate ratings to incoming reviews when designing their rating system. Recent shifts towards more sensitive rating systems by established platforms align with our findings: sellers benefit from a higher sensitivity as it facilitates reputation management via strategic pricing. At the same time, increases in sensitivity have ambiguous effects on consumer surplus and the speed of consumer learning. Our analysis can help identify platforms and markets that may benefit from altering their rating system or are vulnerable to entrants who compete by virtue of their rating system. We view a more detailed formal analysis both of individual platforms' incentives as well as a multi-platform setting as an exciting and promising avenue for future research.

References

- ACEMOGLU, D., A. MAKHDOUNI, A. MALEKIAN, AND A. OZDAGLAR (2019): “Learning From Reviews: The Selection Effect and the Speed of Learning,” *working paper*.
- AMAZON (2021): “Amazon Community Guidelines,” *retrieved from* <https://www.amazon.com/gp/help/customer/display.html?nodeId=GLHXEX85MENUE4XF>.

³⁶In a model of horizontal differentiation with an analogous review function, it is straightforward to verify that this is not the case and that the comparison of the effects is *independent* of the rating.

- BAGNOLI, M. AND T. BERGSTROM (2005): “Log-concave probability and its applications,” *Economic theory*, 26, 445–469.
- BAGWELL, K. AND M. H. RIORDAN (1991): “High and declining prices signal product quality,” *The American Economic Review*, 224–239.
- BAR-ISAAC, H. AND S. TADELIS (2008): “Seller reputation,” *Foundations and Trends in Microeconomics*, 4, 273–351.
- BESBES, O. AND M. SCARSINI (2018): “On information distortions in online ratings,” *Operations Research*, 66, 597–610.
- BHARGAVA, H. K. AND J. FENG (2015): “Does better information lead to lower prices? Price and Advertising Signaling under External Information about Product Quality,” *working paper*.
- BOARD, S. AND M. MEYER-TER VEHN (2013): “Reputation for quality,” *Econometrica*, 81, 2381–2462.
- BOLTON, G. E., E. KATOK, AND A. OCKENFELS (2004): “How effective are electronic reputation mechanisms? An experimental investigation,” *Management Science*, 50, 1587–1602.
- BONATTI, A. AND G. CISTERNAS (2019): “Consumer Scores and Price Discrimination,” *The Review of Economic Studies*, rdz046.
- BONDI, T. (2019): “Alone, Together: Product Discovery Through Consumer Ratings,” *SSRN working paper*.
- BORDALO, P., N. GENNAIOLI, AND A. SHLEIFER (2020): “Memory, attention, and choice,” *The Quarterly journal of economics*, 135, 1399–1442.
- (2022): “Salience,” *Annual Review of Economics*, 14, 521–544.
- BYERS, J. W., M. MITZENMACHER, AND G. ZERVAS (2012): “The groupon effect on yelp ratings: a root cause analysis,” in *Proceedings of the 13th ACM conference on electronic commerce*, 248–265.
- CABRAL, L. AND A. HORTACSU (2010): “The dynamics of seller reputation: Evidence from ebay,” *The Journal of Industrial Economics*, 58, 54–78.
- CABRAL, L. AND L. LI (2015): “A dollar for your thoughts: Feedback-conditional rebates on eBay,” *Management Science*, 61, 2052–2063.
- CARNEHL, C., M. SCHAEFER, A. STENZEL, AND K. TRAN (2022): “Value for Money and Selection: How Pricing Affects Airbnb Ratings,” *IGIER working paper*.
- CHAKRABORTY, I., M. KIM, AND K. SUDHIR (2022): “Attribute sentiment scoring with online text reviews: Accounting for language structure and missing attributes,” *Journal of Marketing Research*, 59, 600–622.
- CHE, Y.-K. AND J. HÖRNER (2018): “Recommender systems as mechanisms for social learning,” *The Quarterly Journal of Economics*, 133, 871–925.
- CHETTY, R., A. LOONEY, AND K. KROFT (2009): “Salience and taxation: Theory and evidence,” *American economic review*, 99, 1145–1177.
- CHEVALIER, J. A. AND D. MAYZLIN (2006): “The effect of word of mouth on sales: Online book reviews,” *Journal of Marketing Research*, 43, 345–354.
- CMA (2015): “Online Reviews and Endorsements,” *Competition & Markets Authority*.
- CRAPIS, D., B. IFRACH, C. MAGLARAS, AND M. SCARSINI (2017): “Monopoly pricing in the presence of social learning,” *Management Science*, 63, 3586–3608.
- DAI, W. D., G. JIN, J. LEE, AND M. LUCA (2018): “Aggregation of consumer ratings: an application to Yelp.com,” *Quantitative Marketing and Economics*, 16, 289–339.
- DE LANGHE, B., P. M. FERNBACH, AND D. R. LICHTENSTEIN (2015): “Navigating by the stars: Investigating the actual and perceived validity of online user ratings,” *Journal of Consumer Research*, 42, 817–833.
- DELLAROCAS, C. AND C. A. WOOD (2008): “The sound of silence in online feedback: Estimating trading risks in the presence of reporting bias,” *Management Science*, 54, 460–476.
- EINAV, L., C. FARRONATO, J. LEVIN, AND N. SUNDARESAN (2018): “Auctions versus Posted Prices in Online Markets,” *Journal of Political Economy*, 126, 178–215.
- EPRS (2017): “Online consumer reviews: The case of misleading or fake reviews,” *European Parliament Research Services*.
- EYSTER, E. AND M. RABIN (2005): “Cursed equilibrium,” *Econometrica*, 73, 1623–1672.
- FENG, J., X. LI, AND X. ZHANG (2019): “Online product reviews-triggered dynamic pricing: theory and evidence,” *Information Systems Research*, 30, 1107–1123.

- FINKELSTEIN, A. (2009): “E-ztax: Tax salience and tax rates,” *The Quarterly Journal of Economics*, 124, 969–1010.
- FORGAS, S., M. A. MOLINER, J. SÁNCHEZ, AND R. PALAU (2010): “Antecedents of airline passenger loyalty: Low-cost versus traditional airlines,” *Journal of Air Transport Management*, 16, 229–233.
- FTC (2019): “Cosmetic company wrote fake reviews,” *retrieved from <https://www.consumer.ftc.gov/blog/2019/10/cosmetic-company-wrote-fake-reviews/>*.
- GOEREE, J. K. AND T. OFFERMAN (2003): “Competitive bidding in auctions with private and common values,” *The Economic Journal*, 113, 598–613.
- GORODNICHENKO, Y. AND O. TALAVERA (2017): “Price setting in online markets: Basic facts, international comparisons, and cross-border integration,” *American Economic Review*, 107, 249–82.
- HO, Y.-C., J. WU, AND Y. TAN (2017): “Disconfirmation effect on online rating behavior: A structural model,” *Information Systems Research*, 28, 626–642.
- HUI, X., T. J. KLEIN, K. STAHL, ET AL. (2021): “When and Why Do Buyers Rate in Online Markets?” Tech. rep., University of Bonn and University of Mannheim, Germany.
- IFRACH, B., C. MAGLARAS, M. SCARSINI, AND A. ZSELEVA (2019): “Bayesian social learning from consumer reviews,” *Operations Research*, 67, 1209–1221.
- JEZIORSKI, P. AND K. MICHELIDAKI (2021): “Impact of pricing on online reviews: A field experiment,” *Available at SSRN 4087034*.
- JIANG, Y. AND H. GUO (2015): “Design of consumer review systems and product pricing,” *Information Systems Research*, 26, 714–730.
- KLEIN, T. J., C. LAMBERTZ, AND K. O. STAHL (2016): “Market transparency, adverse selection, and moral hazard,” *Journal of Political Economy*, 124, 1677–1713.
- KOVBASYUK, S. AND G. SPAGNOLO (2018): “Memory and markets,” *SSRN working paper*.
- LEKKERKERKER, C. G. (1953): “A property of logarithmic concave functions,” *Indagationes Mathematicae*, 15, 505–521.
- LI, X. (2016): “Could deal promotion improve merchants’ online reputations? The moderating role of prior reviews,” *Journal of Management Information Systems*, 33, 171–201.
- LI, X. AND L. M. HITT (2008): “Self-selection and information role of online product reviews,” *Information Systems Research*, 19, 456–474.
- (2010): “Price effects in online product reviews: An analytical model and empirical analysis,” *MIS Quarterly*, 809–831.
- LUCA, M. (2016): “Reviews, reputation, and revenue: The case of Yelp. com,” *Harvard Business School NOM Unit Working Paper*.
- (2017): “Designing Online Marketplaces: Trust and Reputation Mechanisms,” *Innovation Policy and the Economy*, 17, 77–93.
- LUCA, M. AND O. RESHEF (2021): “The Effect of Price on Firm Reputation,” *Management Science*.
- LUCA, M. AND G. ZERVAS (2016): “Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud,” *Management Science*, 62, 3412–3427.
- MAGLARAS, C., M. SCARSINI, AND S. VACCARI (2020): “Social Learning from Online Reviews with Product Choice,” *Working paper*.
- MARINOVIC, I., A. SKRZYPACZ, AND F. VARAS (2018): “Dynamic certification and reputation for quality,” *American Economic Journal: Microeconomics*, 10, 58–82.
- MARTIN, S. AND S. SHELEGIA (2021): “Underpromise and Overdeliver?-Online Product Reviews and Firm Pricing,” *International Journal of Industrial Organization*, 102775.
- MERKLE, M. (1998): “Convolutions of logarithmically concave functions,” *Publikacije Elektrotehničkog fakulteta. Serija Matematika*, 113–117.
- OSBORNE, M. AND A. H. SHAPIRO (2014): “A dynamic model of price signaling, consumer learning, and price adjustment,” *working paper*.
- RAJAGURU, R. (2016): “Role of value for money and service quality on behavioural intention: A study of full service and low cost airlines,” *Journal of Air Transport Management*, 53, 114–122.
- SHIN, D., S. VACCARI, AND A. ZEEVI (2022): “Dynamic pricing with online reviews,” *Management Science*.
- STEAM (2016): “Customer Review System Updated To Show Recent Reviews And Summary,” *retrieved from <https://store.steampowered.com/news/21695/>*.

- (2021): “User Reviews (Steamworks Documentation),” *retrieved from <https://partner.steamgames.com/doc/store/reviews>*.
- STENZEL, A. AND C. WOLF (2016): “Consumer Rating Dynamics,” *Working Paper*.
- STOKEY, N. L., R. LUCAS, AND E. PRESCOTT (1989): “Recursive methods in dynamic economics,” *Cambridge, MA: Harvard University*.
- VON THADDEN, E.-L. (1992): “Optimal pricing against a simple learning rule,” *Games and Economic Behavior*, 4, 627–649.
- WIRED.COM (2019): “What Do Amazon’s Star Ratings Really Mean?” *retrieved from <https://www.wired.com/story/amazon-stars-ratings-calculated/>*.
- WOLINSKY, A. (1983): “Prices as signals of product quality,” *The Review of Economic Studies*, 647–658.
- ZEGNERS, D. (2019): “Building an Online Reputation with Free Content: Evidence from the E-book Market,” *SSRN working paper*.
- ZHU, F. AND X. ZHANG (2010): “Impact of online consumer reviews on sales: The moderating role of product and consumer characteristics,” *Journal of Marketing*, 74, 133–148.

A. Details on Price Internalization κ

Here, we provide further details on our underlying motivation for the reduced-form parameter κ measuring the degree to which consumers internalize the price in their review. We provide three simple microfoundations that can give rise to an interior level of price internalization in the review: (i) different salience of the consumption utility and the price component in consumers’ assessment of the product, (ii) consumer heterogeneity in product assessment, and (iii) the framing of review system by the platform.

Salience in Review Formation While consumers pay the price for the good p at the time of purchase, they generally review the product only after they have used it for some time. Thus, the consumer experienced the price some time $\Delta > 0$ before the review decision.³⁷ In contrast, the last experience of the good itself occurred more recently. We therefore expect the gross consumption utility of the consumer to be more salient in the review formation.

To fix ideas, suppose consumers attach a higher weight to more salient product characteristics; such an interpretation is reminiscent of salience models, see Bordalo et al. (2022) for an overview, and relates to empirical evidence that non-salience of price components distorts consumer choices (see, e.g., Chetty et al. 2009, Finkelstein 2009).³⁸ Normalizing the weight on the gross utility to be equal to one, denote the consumers’ salience discount on the price component by $\beta^\Delta \in (0, 1)$, where $\beta \in (0, 1)$ parametrizes the rate at which past experiences are discounted. Then, consumers’ salience-weighted net consumption utility is

$$\psi_i = \theta + \omega_i - \beta^\Delta p. \quad (29)$$

Interpreting our price-internalization parameter κ as measuring the salience of the price component, we obtain $\kappa = \beta^\Delta \in (0, 1)$. Note that this then maps into an average per-period rating of

$$\psi_t = \theta + \mathbb{E}[\omega | \omega \geq \omega_t^*] - \beta^\Delta p, \quad (30)$$

which is used to update the aggregate rating observable to consumers.

³⁷We assume Δ to be identical across consumers for ease of exposition. Accounting for heterogeneous Δ_i is straightforward; the observed aggregate rating would reflect the *average* salience discount $\mathbb{E}[\beta^{\Delta_i} | \omega_i \geq \omega_t^*]$ in (30).

³⁸The setup described in the following is most closely related to Bordalo et al. (2020) which does not impose a normalization of the sum of the salience-distorted weights. Note that we focus on the “prominence” of different attributes (see Bordalo et al. 2022, Section 3.3) and not on the “surprise”, i.e., evaluation of the individual components relative to reference points.

Consumer Heterogeneity in Reviewing An alternative and complementary interpretation of the interior price internalization in reviews derives from consumer heterogeneity in their product assessment. Chakraborty et al. (2022) provides evidence from Yelp.com that only 46.7% percent of consumers mention the price in their reviews. Thus, consumers appear heterogeneous in whether they consider the price to be a relevant component to consider in the review. Under this interpretation, we can interpret the aggregated rating statistic as consisting of a share of reviews $\nu \in (0, 1)$ featuring the price and a share of reviews $1 - \nu$ not featuring the price. If the price is featured in the review, the consumer reports her net utility $\theta + \omega_i - p$. If the price is not featured in the review, the consumer reports her gross utility $\theta + \omega_i$.³⁹ Assuming that whether a consumer considers the price in the review is uncorrelated with the idiosyncratic taste component ω_i , we recover our model with an average per-period review used to update the aggregate rating observable to consumers.⁴⁰

$$\psi_t = \nu \cdot (\theta + \mathbb{E}[\omega | \omega \geq \omega_t^*] - p) + (1 - \nu) \cdot (\theta + \mathbb{E}[\omega | \omega \geq \omega_t^*]) = \theta + \mathbb{E}[\omega | \omega \geq \omega_t^*] - \nu p. \quad (31)$$

Interpreting our price-internalization parameter κ as measuring the share of consumers featuring the price in their review, we obtain $\kappa = \nu \in (0, 1)$.

The heterogeneity can arise for different reasons. For example, a wealthy consumer might assess a product less on the basis of its value for money but rather by its consumption utility, while a less wealthy consumer might consider the price to be a very important characteristic when assessing a product. Alternatively, consumers might interpret a rating system as assessing solely a product's consumption value, while others might interpret it as assessing the product's value for money. The latter point brings us to the final interpretation.

Framing of Rating System Beyond an intrinsic consumer heterogeneity, there are other reasons why consumer reviews may feature the price more or less prominently. In particular, the framing of the rating system by the platform affects how consumers review a product. Jeziorski and Michelidaki (2021) and Carnehl et al. (2022) provide evidence that reviews across different rating subcategories feature different responsiveness to prices; in particular, value-for-money rating categories, which explicitly ask consumers to assess the overall consumer experience relative to the price, feature the highest price responsiveness. By explicitly guiding consumers in what is to be assessed, a platform can affect reviewing behavior. The platform, abstractly, can suggest consumers to review gross utility or net utility. However, the platform can also use more ambiguous wording in the description of its rating system (see, for example, Amazon 2021) leading to a potential heterogeneity of interpretations as suggested in the previous paragraph.

Interpreting our price-internalization parameter κ as measuring the degree to which the platform highlights the desired prominence of the price component in consumer ratings microfound the framing interpretation presented here.

B. Proofs

B.1. Derivation and Proof of Proposition 1

We proceed by guessing and verifying the value function which is unique (Stokey et al. (1989), Theorem 4.3). The theorem applies because our setup satisfies Assumption 4.3 and 4.4 therein, that is, the state space is a convex subset of \mathbb{R} , the correspondence mapping into future states is non-empty, compact-valued and continuous. Moreover, flow profits are bounded and we have

³⁹Note that these are also the polar cases considered in concurrent work by Shin et al. (2022).

⁴⁰We discuss heterogeneity in the degree of price internalization, allowing for correlation with the idiosyncratic price component, further in Appendix D.4.

discounting. Taking this as given, we guess that the value function is of the form $V(\bar{\psi}) = c + d\bar{\psi} + e\bar{\psi}^2$. As discussed, we replace the price as a control by the rating tomorrow such that

$$p = \frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa} + \frac{\bar{\psi} - \bar{\psi}'}{\sigma(1 - 2\kappa)} \quad (32)$$

$$q = \frac{2(\kappa - 1)(\sigma(\theta + 1) - \bar{\psi}') + 2\bar{\psi}(\kappa + \sigma - 1)}{(2\kappa - 1)\sigma} \quad (33)$$

and the Bellman equation becomes

$$V(\bar{\psi}) = \max_{\bar{\psi}'} \left(\frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa} + \frac{\bar{\psi} - \bar{\psi}'}{\sigma(1 - 2\kappa)} \right) \frac{2(\kappa - 1)((\theta + 1)\sigma - \bar{\psi}') + 2\bar{\psi}(\kappa + \sigma - 1)}{(2\kappa - 1)\sigma} + \delta V(\bar{\psi}'). \quad (34)$$

Differentiating the guessed value function and shifting it one period forward yields

$$V'(\bar{\psi}) = d + 2e\bar{\psi}. \quad (35)$$

Plugging this into the differentiated Bellman equation and solving for $\bar{\psi}'$ delivers

$$\bar{\psi}' = \frac{\sigma(4(1 - \kappa)(\theta + 1) + d\delta(1 - 2\kappa)^2\sigma)}{4(1 - \kappa) - 2\delta e(1 - 2\kappa)^2\sigma^2} - \frac{4(1 - \kappa) - 2(2\kappa(1 - \sigma) + 3\sigma - 2)}{2\delta e(1 - 2\kappa)^2\sigma^2} \bar{\psi} \quad (36)$$

for the law of motion of the rating. Using this law of motion in the Bellman equation and applying the guess on both sides yields an equation system for the undetermined coefficients (c, d, e) that has to be solved. A Mathematica file calculating the expressions can be obtained from the authors' websites.

The solutions for c , d , and e are complicated expressions and omitted here for brevity. More instructive is the induced law of motion given by

$$\begin{aligned} \bar{\psi}' = & \frac{\sigma(\theta + 1)(\delta(3 - 2\kappa)\sigma + 2(1 - \delta)(1 - \kappa))}{2\delta\sigma^2 + (1 - \kappa)(1 - \delta) + \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}} \\ & + \frac{1 - \kappa + \delta(2\sigma - 1)(\kappa + \sigma - 1) - \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}}{\delta(2(1 - \kappa) - (3 - 2\kappa)\sigma)} \bar{\psi}. \end{aligned} \quad (37)$$

Denote

$$a \equiv \frac{\sigma(\theta + 1)(\delta(3 - 2\kappa)\sigma + 2(1 - \delta)(1 - \kappa))}{2\delta\sigma^2 + (1 - \kappa)(1 - \delta) + \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}} \quad (38)$$

$$b \equiv \frac{1 - \kappa + \delta(2\sigma - 1)(\kappa + \sigma - 1) - \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}}{\delta(2(1 - \kappa) - (3 - 2\kappa)\sigma)}. \quad (39)$$

so that we can write $\bar{\psi}' = a + b\bar{\psi}$. Given an initial rating $\bar{\psi}_1$, we can hence write

$$\bar{\psi}_\tau = \left(a \cdot \sum_{i=0}^{\tau-2} b^i \right) + b^{\tau-1} \bar{\psi}_1 \quad (40)$$

and thus, using $|b| < 1$ which follows from the maintained technical assumption $\kappa < 1 - \frac{\sigma}{2}$,

$$\lim_{\tau \rightarrow \infty} \bar{\psi}_\tau = \frac{a}{1 - b} = \frac{(\theta + 1)(\delta\sigma(3 - 2\kappa) + 2(1 - \delta)(1 - \kappa))}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)} \equiv \Psi. \quad (41)$$

At this long-run rating, we can use (23) and obtain the long-run price

$$\tilde{p} = p(\Psi, \Psi) = \frac{((1 - \delta) + 2\delta\sigma)(\theta + 1)}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)}. \quad (42)$$

Rewriting Ψ and $p(\Psi, \Psi)$ yields the expressions in Proposition 1. Moreover, it immediately follows from (5) that $\mu(\Psi, \tilde{p}) = \theta$.

Uniqueness follows from the quadratic value function and the fact that it is attained by only two (linear) pricing policies one of which diverges and yields infinite or negative prices. Hence, there is only one feasible optimal policy that solves the seller's problem.

B.2. Proof of Corollary 1

Differentiating (26) gives

$$\frac{\partial \tilde{p}}{\partial \sigma} = \frac{\overbrace{2(1-\delta)\delta(\theta+1)}^{>0}}{((1-\delta)(3-2\kappa) + 4\delta\sigma)^2} \cdot (1-2\kappa), \quad (43)$$

so that the sign depends on the sign of $1-2\kappa$. This gives (a) and via the relation to CS (b). For (c), we differentiate (28) and get

$$\frac{\partial \tilde{\pi}}{\partial \sigma} = \frac{\overbrace{2(1-\delta)^2\delta(\theta+1)^2}^{>0}}{((1-\delta)(3-2\kappa) + 4\delta\sigma)^3} \cdot (1-2\kappa)^2, \quad (44)$$

which is unambiguously weakly (strictly for $\kappa \neq \frac{1}{2}$) positive. The same is true for (d), where we obtain

$$\frac{\partial \Psi}{\partial \sigma} = \frac{\overbrace{(1-\delta)\delta(\theta+1)}^{>0}}{((1-\delta)(3-2\kappa) + 4\delta\sigma)^2} \cdot (1-2\kappa)^2. \quad (45)$$

B.3. Proof of Proposition 2

Recall that $\bar{\psi}_{t+1} - \bar{\psi}_t = a - (1-b)\bar{\psi}_t$ and hence that ratings are increasing whenever $\frac{a}{1-b} = \Psi > \bar{\psi}_t$. Using (23), the policy function for ratings implies for the evolution of prices that

$$p_{t+1} - p_t = \frac{1}{1-2\kappa} \left(2(\bar{\psi}_{t+1} - \bar{\psi}_t) + \frac{\bar{\psi}_{t+2} - \bar{\psi}_{t+1}}{\sigma} - \frac{\bar{\psi}_{t+1} - \bar{\psi}_t}{\sigma} \right) \quad (46)$$

With $\bar{\psi}_{t+2} - \bar{\psi}_{t+1} - (\bar{\psi}_{t+1} - \bar{\psi}_t) = -(1-b)(\bar{\psi}_{t+1} - \bar{\psi}_t)$, this reduces to

$$p_{t+1} - p_t = \frac{1}{1-2\kappa} \cdot (\bar{\psi}_{t+1} - \bar{\psi}_t) \cdot \left(2 - \frac{1-b}{\sigma} \right) = \frac{2\sigma - (1-b)}{(1-2\kappa)\sigma} (\bar{\psi}_{t+1} - \bar{\psi}_t). \quad (47)$$

Plugging in for b and simplifying allows us to show that $\frac{2\sigma - (1-b)}{(1-2\kappa)\sigma} > 0$ so that prices and ratings always comove. Finally, note that for $b > 0$ we have that

$$\bar{\psi}_{t+1} = a + b\bar{\psi}_t > \frac{a}{1-b} = \Psi \quad (48)$$

$$\iff b\bar{\psi}_t > \frac{ab}{1-b} \quad (49)$$

$$\iff \bar{\psi}_t > \frac{a}{1-b}, \quad (50)$$

so that movement is monotonic over time provided $b > 0$. In contrast, we have $\bar{\psi}_{t+1} < \Psi \iff \bar{\psi}_t > \Psi$ for $b < 0$. Denoting by $\bar{\sigma}$ the bound on σ so that $b < 0 \iff \sigma > \bar{\sigma}$, the Corollary follows. To derive $\bar{\sigma}$, recall that

$$b = \frac{1 - \kappa + \delta(2\sigma - 1)(\kappa + \sigma - 1) - \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}}{\delta(2(1 - \kappa) - (3 - 2\kappa)\sigma)} \quad (51)$$

Our maintained assumption $\kappa < 1 - \frac{\sigma}{2}$ ensures that $b > 0$ for $\sigma < \frac{1}{2}$. For $\sigma > \frac{1}{2}$, observe that the denominator of b is positive if and only if $\kappa < \frac{2-3\sigma}{2-2\sigma}$, where the right hand side is strictly negative for $\sigma > \frac{2}{3}$. The numerator of b in turn is positive iff $\sigma > \frac{2}{3}$. This allows us to conclude that $b < 0$ if (i) $\sigma > \frac{2}{3}$, or if (ii) $\sigma \in (\frac{1}{2}, \frac{2}{3})$ and $\kappa > \frac{2-3\sigma}{2-2\sigma} \iff \sigma > \frac{2-2\kappa}{3-2\kappa}$. We can thus define $\bar{\sigma} \equiv \max\{\frac{1}{2}, \frac{2-2\kappa}{3-2\kappa}\}$. Collecting results yields Proposition 2.

B.4. Speed of Convergence

Linear convergence of a sequence $\{y_t\}$, which converges to y , at rate ν requires that $\nu = \lim_{t \rightarrow \infty} \frac{y_t - y}{y_{t-1} - y}$. Given that $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$, we have $\Psi = \frac{a}{1-b}$ and thus

$$\frac{|\bar{\psi}_t - \bar{\psi}|}{|\bar{\psi}_{t-1} - \bar{\psi}|} = \frac{|a + b\bar{\psi}_{t-1} - \frac{a}{1-b}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (52)$$

$$= \frac{|a \cdot \left(1 - \frac{1}{1-b}\right) + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (53)$$

$$= \frac{|-\frac{ab}{1-b} + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (54)$$

$$= \frac{|b \cdot \left(\bar{\psi}_{t-1} - \frac{a}{1-b}\right)|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (55)$$

$$= |b|. \quad (56)$$

Similarly, we can use (5), (23) and $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$ to obtain for the belief held by consumers in period t , μ_t ,

$$\mu_t = 2\bar{\psi}_t - 1 - (1 - 2\kappa) \left[\frac{\bar{\psi}_{t+1} - \bar{\psi}_t}{\sigma(1 - 2\kappa)} - \frac{\theta + 1 - 2\bar{\psi}_t}{1 - 2\kappa} \right] = \theta - \frac{\bar{\psi}_{t+1} - \bar{\psi}_t}{\sigma} = \theta - \frac{a - (1 - b)\bar{\psi}_t}{\sigma}, \quad (57)$$

so that

$$\frac{|\mu_t - \theta|}{|\mu_{t-1} - \theta|} = \frac{\left| \frac{a - (1-b)(a + b\bar{\psi}_{t-1})}{\sigma} \right|}{\left| \frac{a - (1-b)\bar{\psi}_{t-1}}{\sigma} \right|} \quad (58)$$

$$= \frac{|a - (1-b)(a + b\bar{\psi}_{t-1})|}{|a - (1-b)\bar{\psi}_{t-1}|} \quad (59)$$

$$= \frac{|b(a - (1-b)\bar{\psi}_{t-1})|}{|a - (1-b)\bar{\psi}_{t-1}|} \quad (60)$$

$$= |b|. \quad (61)$$

For $\sigma \neq \frac{2-2\kappa}{3-2\kappa}$, we obtain $\frac{\partial b}{\partial \sigma} < 0$. The detailed calculations are extensive and omitted here for brevity; they are verified in the supplementary Mathematica file. Whenever $b < 0$, a marginal increase in sensitivity thus harms the speed of convergence as $|b|$ increases.

B.5. Comparison to Alternative Pricing Strategies

Limit Profits and Consumer Surplus under Myopic Pricing To compare the long-run profits and consumer surplus, we first need to characterize the long-run steady state under myopic pricing and fixed pricing. A seller myopically maximizing flow profits sets the price $p_t = \frac{\bar{\psi}_t}{2(1-\kappa)}$ in each period, which determines the inference as a function of the rating stock, and the induced per-period average review. Specifically, we obtain

$$\begin{aligned}\tilde{\mu}_t^m &= \frac{3-2\kappa}{2-2\kappa}\bar{\psi} - 1 \\ \tilde{\omega}_t^m &= 1 - \bar{\psi} \\ \psi_t^m &= 1 + \theta - \frac{\bar{\psi}}{2(1-\kappa)}\end{aligned}\tag{62}$$

The aggregate rating hence evolves according to

$$\bar{\psi}_{t+1} = (1-\sigma)\bar{\psi}_t + \sigma\psi_t = \sigma(1+\theta) + \left(1-\sigma - \frac{1}{2(1-\kappa)}\sigma\right)\bar{\psi}_t\tag{63}$$

and converges to

$$\bar{\psi}_t = \bar{\psi}_{t+1} = \Psi^m = \frac{2(1+\theta)(1-\kappa)}{3-2\kappa}.\tag{64}$$

The seller's profits in this steady state are given by $\tilde{\pi}^m = \frac{2(1+\theta)^2(1-\kappa)}{(3-2\kappa)^2}$, and consumer surplus by $\tilde{CS}^m = 2\frac{(1+\theta)^2(1-\kappa)^2}{(3-2\kappa)^2}$. Observe that these are independent of the sensitivity σ .

For the profits and consumer surplus in the steady state under fixed pricing, we first need to compute the optimal fixed price, for which we need the total discounted profits.

Total Discounted Profits and Consumer Surplus To derive the total discounted profits and consumer surplus, we make use of the analytical expressions for the law of motion of the aggregate rating, from which we can compute the flow profits and consumer surplus in each of the three pricing regimes.

Given any of the pricing strategies, the law of motion of the rating has the form $\bar{\psi}_{t+1} = \tilde{a} + \tilde{b}\bar{\psi}_t$, where \tilde{a} and \tilde{b} take different values depending on the considered strategy. Given this law of motion, we can write the rating in any period t as

$$\bar{\psi}_t = \frac{1 - \tilde{b}^{t-1}}{1 - \tilde{b}}\tilde{a} + \tilde{b}^{t-1}\bar{\psi}_1.\tag{65}$$

Based on this formulation, the price and quantity in any period t can be obtained from (32) and (33), which yields flow profits in each period π_t of

$$\begin{aligned}\pi_t &= \frac{2 \left[\tilde{b}(\tilde{a} - (1-\tilde{b})(1+\theta)(1-\kappa))\sigma + \tilde{b}^t((1-\tilde{b})(1-\kappa) - \sigma)(\tilde{a} - (1-\tilde{b})\bar{\psi}_1) \right]}{(1-\tilde{b})^2\tilde{b}^2(1-2\kappa)^2\sigma^2} \\ &\quad \times \left[\tilde{b}((1-\tilde{b})(1+\theta) - 2\tilde{a})\sigma + \tilde{b}^t(1-\tilde{b}-2\sigma)((1-\tilde{b})\bar{\psi}_1 - \tilde{a}) \right]\end{aligned}\tag{66}$$

and consumer surplus CS_t of

$$\begin{aligned}CS_t &= \frac{2 \left[\tilde{b} \left((1+\theta) - \tilde{a} - \tilde{b}(1+\theta)(1-\kappa) - (1+\theta)\kappa \right) \sigma - \tilde{b}^t \left((1-\tilde{b})(1-\kappa) - \sigma \right) (\tilde{a} - (1-\tilde{b})\bar{\psi}_1) \right]}{(1-\tilde{b})^2\tilde{b}^2(1-2\kappa)^2\sigma^2} \\ &\quad \times \left[\tilde{b}((1+\theta) - \tilde{a} - \tilde{b}(1+\theta)(1-\kappa) - (1+\theta)\kappa)\sigma + \tilde{b}^t \left((1-\tilde{b})\kappa - \sigma \right) ((1-\tilde{b})\bar{\psi}_1 - \tilde{a}) \right].\end{aligned}\tag{67}$$

We can use these closed-form expressions to obtain analytical expressions for the total discounted profits $\Pi = \sum_{t=1}^{\infty} \delta^{t-1} \pi_t$ and consumer surplus $TCS = \sum_{t=1}^{\infty} \delta^{t-1} CS_t$ at $t = 1$, respectively. The resulting expressions are complicated and omitted here for brevity. Based on the expressions, we simply need to obtain \tilde{a} and \tilde{b} for each of the respective pricing strategies.

In the case of strategic pricing, these are given by (38) and (39). For myopic pricing, we have $p_t = \frac{\bar{\psi}_t}{2(1-\kappa)}$ which induces $\bar{\psi}_{t+1} = (1 + \theta)\sigma + (1 - \frac{(3-2\kappa)\sigma}{2(1-\kappa)})\bar{\psi}_t$, i.e. $\tilde{a} = a_m = (1 + \theta)\sigma$ and $\tilde{b} = b_m = (1 - \frac{(3-2\kappa)\sigma}{2(1-\kappa)})$.

Finally, consider that the seller is restricted to charge a fixed price, denoted \hat{p} . This yields a law of motion for the rating of $\bar{\psi}_{t+1} = \sigma(1 + \theta - (1 - 2\hat{p})\kappa) + (1 - 2\sigma)\bar{\psi}_t$. From this, we can obtain the total discounted profits Π_f as a function of the initial rating $\bar{\psi}_1$ and the price \hat{p} as

$$\Pi_f(\hat{p}, \bar{\psi}_1) = 2 \frac{\hat{p}(\delta\sigma(1 + \theta) + (1 - \delta)\bar{\psi}_1 - \hat{p}(1 - \delta)(1 - \kappa) - \sigma\delta\hat{p})}{(1 - \delta)(1 - (1 - 2\sigma)\delta)}. \quad (68)$$

Maximizing this yields the optimal fixed price maximizing total discounted profits, p_f

$$p_f = \frac{\delta((1 + \theta)\sigma - \bar{\psi}_1) + \bar{\psi}_1}{2(1 - \kappa) - 2\delta(1 - \kappa - \sigma)}, \quad (69)$$

which we can plug into the law of motion to compute the total discounted profits Π_f and consumer surplus TCS_f given the optimal fixed price. Note that this price has the property that it is increasing in the sensitivity if and only if the initial attitude towards the product is not too high, and that the same applies with respect to the seller's patience.

$$\frac{\partial p_f}{\partial \sigma} = \frac{\delta(1 - \delta)((1 + \theta)(1 - \kappa) - \bar{\psi}_1)}{2(1 - \kappa - \delta(1 - \kappa - \sigma))^2} > 0 \iff \bar{\psi}_1 < (1 + \theta)(1 - \kappa). \quad (70)$$

$$\frac{\partial p_f}{\partial \delta} = \frac{\sigma((1 + \theta)(1 - \kappa) - \bar{\psi}_1)}{2(1 - \kappa - \delta(1 - \kappa - \sigma))^2} > 0 \iff \bar{\psi}_1 < (1 + \theta)(1 - \kappa). \quad (71)$$

Limit Profits and Consumer Surplus under Fixed Pricing For any fixed price \tilde{p} , the limit rating under this fixed price can easily be computed from the law of motion as $\frac{1}{2}(1 + \theta + (1 - 2\kappa)\tilde{p})$. From this, it is immediately apparent that the seller can always induce monopoly profits $\frac{(1+\theta)^2}{4}$ in the limit by committing to a price $\frac{1+\theta}{2}$ —this would exceed the long-run profits under strategic pricing. Note that this is indeed the limiting outcome for optimal dynamic pricing when the seller becomes infinitely patient; $\lim_{\delta \rightarrow 1} \tilde{p} = \frac{1+\theta}{2}$ with the corresponding limit rating $\Psi \xrightarrow{\delta \rightarrow 1} \frac{(1+\theta)(3-2\kappa)}{4}$.

Plugging in the optimal price p_f , we obtain as the steady state rating Ψ_f and steady state profits \tilde{p}_f and consumer surplus $\tilde{C}S_f$

$$\Psi_f = \frac{(1 + \theta)(2(1 - \delta)(1 - \kappa) + \delta(3 - 2\kappa)\sigma) + (1 - \delta)(1 - 2\kappa)\bar{\psi}_1}{4(1 - \kappa) - 4\delta(1 - \kappa - \sigma)} \quad (72)$$

$$\tilde{p}_f = \frac{(\delta(1 + \theta)\sigma + (1 - \delta)\bar{\psi}_1)((1 + \theta)(2(1 - \delta)(1 - \kappa) + \delta\sigma) - (1 - \delta)\bar{\psi}_1)}{4(1 - \kappa - \delta(1 - \kappa - \sigma))^2} \quad (73)$$

$$\tilde{C}S_f = \frac{((1 + \theta)(2(1 - \delta)(1 - \kappa) - \delta\sigma) - (1 - \delta)\bar{\psi}_1)^2}{8(1 - \kappa - \delta(1 - \kappa - \sigma))^2} \quad (74)$$

Note that \tilde{p}_f is increasing in σ —while the seller charges a fixed price, a higher responsiveness to incoming reviews nonetheless gives scope for that fixed price to strategically induce higher ratings.

Comparison of Limit Profits and Consumer Surplus Restrict attention to $\kappa \neq \frac{1}{2}, \sigma \neq \frac{1}{2}$ and $\sigma \neq \frac{2-2\kappa}{3-2\kappa}$. Given the characterizations for long-run profits $\tilde{\pi}, \tilde{\pi}_m, \tilde{\pi}_f$, we immediately obtain that the value of strategic pricing relative to myopic pricing $\Delta\tilde{\pi}_m$ is given by

$$\Delta\tilde{\pi}_m = \tilde{\pi} - \tilde{\pi}^m = \frac{2\delta(1+\theta)^2(1-2\kappa)^2\sigma((1-\delta)(3-2\kappa)+2\delta\sigma)}{(3-2\kappa)^2((1-\delta)(3-2\kappa)+4\delta\sigma)^2}, \quad (75)$$

which is strictly positive for all $\kappa \neq \frac{1}{2}$. Moreover, for $\kappa \neq \frac{1}{2}$ it holds that $\frac{\partial\Delta\tilde{\pi}_m}{\partial\delta} > 0$, $\frac{\partial\Delta\tilde{\pi}_m}{\partial\theta} > 0$, $\frac{\partial\Delta\tilde{\pi}_m}{\partial\sigma} > 0$, as well as $\frac{\partial\Delta\tilde{\pi}_m}{\partial\kappa} < 0$ for $\kappa < \frac{1}{2}$ and $\frac{\partial\Delta\tilde{\pi}_m}{\partial\kappa} > 0$ for $\kappa > \frac{1}{2}$. In terms of consumer surplus, we obtain that consumer surplus is higher under myopic pricing in the long run if and only if the selection effect dominates, which incentivizes the strategic seller to charge high prices to maintain high ratings, and vice versa for a dominant price effect which incentivizes the strategic seller to lower prices. Specifically, we have

$$\tilde{C}S > \tilde{C}S_m \iff \kappa < \frac{1}{2}. \quad (76)$$

For the comparison between fixed pricing and strategic pricing, the dependence of the optimal fixed price on the initial rating $\bar{\psi}_1$ gives scope for different orderings between long-run profits and long-run consumer surplus under the two regimes, respectively. To understand this, note that the optimal fixed price p_f is increasing in the initial rating $\bar{\psi}_1$; as consumers correctly infer the quality level θ in the long-run steady state under any considered pricing regime, a higher price directly corresponds to a lower consumer surplus. Equating the optimal fixed price p_f with the long-run steady state price under strategic pricing, \tilde{p} , thus yields the cutoff rating such that the long-run consumer surplus under strategic pricing is higher than under optimal fixed pricing. Unsurprisingly, this cutoff is equal to the long-run steady state rating under strategic pricing, Ψ , as the optimal fixed price coincides with the steady state price under dynamic pricing in this case.

Regarding the comparison of profits, it is helpful to observe that $\tilde{\pi}_f$ is quadratic and concave in the initial rating $\bar{\psi}_1$, with two roots equal to Ψ and $\frac{(1+\theta)(4(1-\delta)(1-\kappa)^2+\delta(5-6\kappa)\sigma)}{(1-\delta)(3-2\kappa)+4\delta\sigma} \equiv \hat{\Psi}$ with $\hat{\Psi} < \Psi \iff \kappa > \frac{1}{2}$ and $\hat{\Psi} > \Psi \iff \kappa < \frac{1}{2}$. As $\tilde{\pi}_f > \tilde{\pi}$ whenever $\bar{\psi}_1$ is between those two roots, we can conclude that fixed pricing leads to higher steady state profits for a range of initial ratings just above the long-run steady state rating under strategic pricing whenever the selection effect dominates, $\kappa < \frac{1}{2}$, and just below whenever the direct price effect dominates, $\kappa > \frac{1}{2}$.

More generally, we obtain

$$\begin{aligned} \Delta\tilde{\pi}_f = \tilde{\pi} - \tilde{\pi}_f = & \frac{2(1+\theta)^2(1-\kappa-\delta(1-\kappa-\sigma))(1-\delta(1-2\sigma))}{((1-\delta)(3-2\kappa)+4\delta\sigma)^2} \\ & - \frac{(\delta((1+\theta)\sigma - \bar{\psi}_1) + \bar{\psi}_1)((1+\theta)(2(1-\theta)(1-\kappa) + \delta\sigma) - (1-\delta)\bar{\psi}_1)}{4(1-\kappa-\delta(1-\kappa-\sigma))^2}. \end{aligned} \quad (77)$$

Differentiating this with respect to σ allows us to establish the existence of cutoffs c_1, c_2 such that a marginal increase in sensitivity of the rating system ($\sigma \uparrow$) is more beneficial for long-run steady state profits under strategic pricing than under optimal fixed pricing ($\Delta\tilde{\pi}_f \uparrow$) if and only if $c_1 < \bar{\psi}_1 < c_2$. While both long-run profits are increasing in the sensitivity, strategic pricing benefits more from this whenever the initial attitude is intermediate, while it is optimal fixed on which sensitivity has a higher impact whenever the initial attitude is extremely high or low, respectively. In a similar vein, there are cutoffs c_3, c_4 such that $\Delta\tilde{\pi}_f$ increases in κ if and only if the initial attitude is sufficiently low, $\bar{\psi}_1 < c_3$, or sufficiently high, $\bar{\psi}_1 > c_4$.

Comparison of Total Discounted Profits Restrict attention to $\kappa \neq \frac{1}{2}, \sigma \neq \frac{1}{2}$ and $\sigma \neq \frac{2-2\kappa}{3-2\kappa}$. It can be established analytically that the seller's profits under optimal strategic pricing, Π , weakly exceed those under optimal fixed pricing, Π_f . Moreover, this inequality holds strictly unless the initial rating $\bar{\psi}_1$ coincides with the long-run rating under strategic pricing, Ψ . Given the complexity of the expressions, we again refer to the provided Mathematica file for verification. It is immediate that generically $\Pi > \Pi_m$ as under optimal dynamic pricing the firm can always mimic the myopically optimal strategy but generically prices differently.

Comparison of Total Consumer Surplus We are unable to analytically compare the total consumer surplus under the three pricing regimes. However, any pairwise ordering between TCS , TCS_m and TCS_f is possible. To exemplify this, it suffices to provide specific parameterizations for each ordering.

- TCS vs TCS_f
 - $TCS > TCS_f$ for $\delta = 0.9, \kappa = 0.2, \sigma = 0.3, \theta = 1, \bar{\psi}_1 = 0.9$
 - $TCS < TCS_f$ for $\delta = 0.9, \kappa = 0.2, \sigma = 0.3, \theta = 1, \bar{\psi}_1 = 1.1$
- TCS vs TCS_m
 - $TCS > TCS_m$ for $\delta = 0.7, \kappa = 0.8, \sigma = 0.15, \theta = 1, \bar{\psi}_1 = 0.5$
 - $TCS < TCS_m$ for $\delta = 0.9, \kappa = 0.2, \sigma = 0.3, \theta = 1, \bar{\psi}_1 = 1$
- TCS_m vs TCS_f
 - $TCS_m > TCS_f$ for $\delta = 0.9, \kappa = 0.2, \sigma = 0.3, \theta = 1, \bar{\psi}_1 = 1$
 - $TCS_m < TCS_f$ for $\delta = 0.7, \kappa = 0.8, \sigma = 0.15, \theta = 1, \bar{\psi}_1 = 0.5$

B.6. Example: Platform Incentives

To analyze the incentives of platform designers, we consider a platform that maximizes a weighted sum of long-run profits and long-run consumer surplus. The focus on long-run outcomes is deliberate, as total discounted profits and surplus are naturally affected by the initial attitude towards a given product—this dependence vanishes in the steady state. Moreover, a platform operates with many sellers at different stages in their life cycle. Thus, at any given moment, a significant number of mature products are offered on the platform. Throughout the analysis, we focus on sellers who engage in strategic pricing. We believe that this is the natural assumption, as it is easy to adjust prices online and a seller who enters the platform with a new product does always prefer to price dynamically rather than to commit to a single price.⁴¹

Specifically, we consider a platform that chooses the sensitivity of the rating system to new reviews, i.e., it sets σ , to maximize

$$\pi^P = w_c \cdot \tilde{C}\tilde{S} + (1 - w_c) \cdot \tilde{\pi}, \quad (78)$$

where $\tilde{C}\tilde{S}$ and $\tilde{\pi}$ are long-run consumer surplus and profits and $w_c \in [0, 1]$ is the weight the platform attaches to consumer welfare. Depending on w_c , we can interpret π^P as the objective function of a social planner who maximizes the total surplus ($w_c = \frac{1}{2}$), a regulator who focuses

⁴¹As sellers using optimal fixed pricing also strictly benefit from increased sensitivity in their long-run steady state profits, while long-run profits under myopic pricing are independent of σ , the results would be qualitatively unaffected if the platform features a mix of sellers using different degrees of sophistication in their pricing regimes as long as some sellers price strategically. In the previous subsection, we discussed considerations when facing a change in the frequency of price adjustments by sellers.

on consumer surplus ($w_c \rightarrow 1$), or a platform operator who receives a commission and maximizes seller profits ($w_c \rightarrow 0$). Due to the network effects inherent to multi-sided platforms, a platform operator is unlikely to only care about one side of the market—for example, if consumer surplus were too low, consumers would be likely to leave the platform, in which case the platform's revenue would shrink. Given spillovers from consumer surplus on the future demand on other products, it is reasonable that platforms place an interior weight $w_c \in (0, 1)$ on consumer surplus.⁴²

From Corollary 1, it follows that the highest possible σ maximizes π^P whenever the direct price effect is large, i.e., when $\kappa > \frac{1}{2}$, as both consumers and the seller benefit from a high σ . The seller always prefers the rating system most responsive to recent reviews, while consumers in this case want a sensitive rating system, as this leads to a downward pressure on prices to manage the ratings. However, when the direct price effect is small, interests diverge. The seller prefers a responsive rating system (high σ), while consumers are better off whenever σ is low.

Proposition 3 (Platform Incentives) *A platform maximizing $\pi^P(\sigma)$ chooses the highest sensitivity σ if*

- (i) *the direct price effect is strong ($\kappa > 1/2$) or if*
- (ii) *the direct price effect is weak ($\kappa < 1/2$) and the weight on consumers is sufficiently low*

$$w_c < \frac{3 - \delta - 4(1 - \delta)\kappa}{(1 - \delta)(1 - 2\kappa)}.$$

If neither of the two is satisfied, the platform chooses a sensitivity of $\max\{\underline{\sigma}, \tilde{\sigma}\}$ with

$$\tilde{\sigma} := \frac{(1 - \delta)(1 - 2\kappa - (3 - 4\kappa)w_c)}{2w_c\delta}. \quad (79)$$

Proof. Proof: See below. ■

Proposition 3 is intuitive in that whenever the incentives vis-a-vis maximizing seller profits and consumer surplus are misaligned, the platform either balances the two aspects by choosing an interior σ or fully follows one of the two sides provided that it puts sufficient weight on them in the maximization. Importantly, this misalignment can materialize only when the selection effect dominates the direct price effect ($\kappa < 1/2$).

To choose the seller-optimal sensitivity, i.e., to let the rating system be maximally sensitive despite consumers preferring the opposite, the weight on consumer surplus, w_c , in π^P needs to be sufficiently low. To illustrate this, note that even for $\kappa < \frac{1}{2}$, the consumer-optimal lowest possible sensitivity is chosen already for $w_c > \frac{1}{3}$ because $w_c > \frac{1}{3} > \frac{1-2\kappa}{3-4\kappa}$ implies $\tilde{\sigma} < 0$. Hence, to justify a high-sensitivity rating for platforms that primarily sell products that have a dominant selection effect, the platform must put more than twice as much weight on profits than on consumer surplus.

Proof of Proposition 3 In line with the previous argument for optimality of a high σ whenever $\kappa > \frac{1}{2}$, we restrict attention to $\kappa < \frac{1}{2}$. Using (26) to obtain $\tilde{C}S$ and plugging this together with (28) into (78), we obtain

$$\pi^P = 2(\theta + 1)^2 \cdot \frac{(\delta\sigma + (1 - \delta)(1 - \kappa))((2 - w_c)\delta\sigma + (1 - \delta)(1 - \kappa w_c))}{(4\delta\sigma + (1 - \delta)(3 - 2\kappa))^2} \quad (80)$$

⁴²In particular, to attract customers to purchase on the platform rather than offline, the platform should place a sufficient weight on consumer surplus.

Differentiating this with respect to σ and rearranging, we obtain that

$$\frac{\partial \pi^P}{\partial \sigma} > 0 \iff 2w_c\delta\sigma - (1-\delta)[(1-2\kappa) - (3-4\kappa)w_c] < 0 \quad (81)$$

and analogously

$$\frac{\partial \pi^P}{\partial \sigma} < 0 \iff 2w_c\delta\sigma - (1-\delta)[(1-2\kappa) - (3-4\kappa)w_c] > 0. \quad (82)$$


Define $\tilde{\sigma} = \frac{(1-\delta)[1-2\kappa-(3-4\kappa)w_c]}{2w_c\delta}$ and it follows that the strict maximizer of π^P is obtained at $\tilde{\sigma}$. Noting that $\tilde{\sigma}$ is decreasing in w_c and

$$\tilde{\sigma}|_{w_c = \frac{(1-\delta)(1-2\kappa)}{3-\delta-4(1-\delta)\kappa}} = 1,$$

the proposition immediately follows.

C. Additional Figures

Figure 2: Impulse response functions given a 20% shock to the steady-state review



D. Extensions

D.1. Number of Reviews

In reality, the more consumers review a given product, the more the rating will adjust. This effect can be incorporated into our setting via a modified updating rule

$$\bar{\psi}_{t+1} = (1 - q_t\sigma)\bar{\psi}_t + q_t\sigma\psi_t, \quad (83)$$

where q_t is the number of consumers in period t . Unfortunately, we cannot apply our results directly because we obtained them through a guess and verify procedure with a linear policy function; with the present formulation, however, the objective is more complicated and includes a quadratic term on the control that precludes us from finding closed-form solutions: both the

quantity and the review are linear in the price and multiplied with each other. However, we can assess the effect of making ratings dependent on the number of reviews by studying the seller's first-order condition for prices as

$$\begin{aligned} \frac{dV(\bar{\psi}_t)}{dp_t} &= q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t + \delta \frac{dV_{t+1}}{dp_t} \\ &= \underbrace{q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t}_{\text{static monopoly pricing}} + \delta \underbrace{\frac{dV_{t+1}}{d\bar{\psi}_{t+1}}}_{\text{effect of reviews on CV}} \left(\underbrace{\sigma q_t(p_t) \frac{d\psi_t(p_t)}{dp_t}}_{\substack{\text{better reviews} \\ \rightarrow \\ \text{higher ratings}}} - \underbrace{\sigma \frac{dq_t(p_t)}{dp_t} (\bar{\psi}_t - \psi_t(p_t))}_{\text{number of reviews effect}} \right). \end{aligned} \quad (84)$$

Contrasting this with the first-order condition in the original model, the only new term is the last one. If the induced review is above (below) the current rating, the seller has an incentive to increase (reduce) the current quantity, that is, reduce (increase) the price relative to the case in which the number of reviews does not enter the updating rule. This is because a price increase always decreases the number of purchasing consumers and can hence be used to amplify (attenuate) the effect of inducing a high (low) average review.

D.2. Distribution over Reviewing Agents

In our baseline model, horizontal preferences are distributed uniformly, and consumers rate with an identical probability. Both assumptions are made for expositional and tractability purposes. Nonetheless, it is important to discuss the impact of departures from these assumptions, in particular in light of ample empirical evidence that reviews tend to be bimodal on the extremes of consumer satisfaction; see, e.g., Bolton et al. (2004), Dellarocas and Wood (2008) and Hui et al. (2021).

Within our model, we can incorporate this flexibly by letting the probability of reviewing be given by a function $f_\psi(\omega; \tilde{\omega})$. Let f be continuously differentiable in both its arguments and strictly positive on its support. Moreover, assume that if the number of purchasing consumers decreases ($\tilde{\omega}$ increases), the average reviewing consumer, $w^e(\tilde{\omega}) \equiv \int_{\tilde{\omega}}^1 w \frac{f_\psi(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega, \tilde{\omega})} d\omega$, increases, but not too much: $\frac{dw^e(\tilde{\omega})}{d\tilde{\omega}} \in (0, 1)$. Under these assumptions, consumers' inference and demand are given by the solution to the equation system as

$$\mu + \tilde{\omega} = p \quad (85)$$

$$\mu + \int_{\tilde{\omega}}^1 \omega \frac{f_\psi(\omega; \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega, \tilde{\omega})} d\omega - \kappa p = \bar{\psi}. \quad (86)$$

The implicit function theorem yields the effect of price changes on the solution pair $(\mu, \tilde{\omega})$

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = - \begin{pmatrix} \frac{\partial u}{\partial \theta} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} & \frac{\partial \psi}{\partial \omega} \end{pmatrix}^{-1} \begin{pmatrix} -1 \\ \frac{\partial \psi}{\partial p} \end{pmatrix}. \quad (87)$$

Inverting the matrix and plugging in the partial derivatives yields

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = \begin{pmatrix} \kappa - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_\psi(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega, \tilde{\omega})} d\omega \\ 1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_\psi(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega, \tilde{\omega})} d\omega \\ 1 - \kappa \\ 1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_\psi(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega, \tilde{\omega})} d\omega \end{pmatrix}, \quad (88)$$

where we assume that consumers are aware that the reviews are not given by the average consumer but by a selected sample of consumers.

The seller's pricing decision is affected twofold: first, consumers' inference is different and, therefore, the demand, $(1 - \tilde{\omega})$, reacts differently to price, and second, the pricing has an effect on the selection into reviewing. These two components can be seen in the seller's first-order condition

$$\frac{dV(\bar{\psi}_t)}{dp_t} = q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t + \delta \frac{dV_{t+1}}{dp_t} \quad (89)$$

$$= 1 - \tilde{\omega}(p_t) - \frac{1 - \kappa}{1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega}} p_t \quad (90)$$

$$+ \delta \frac{dV_{t+1}}{d\bar{\psi}_{t+1}} \sigma \left(\frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega - \kappa \right) \quad (91)$$

Note that the main term is again the change in the average reviewing consumer. The relevant consideration for the seller therefore derives from the same forces as in the main part of the model: the selection effect given by the change in the average reviewing consumer and the direct price effect given by κ . Whenever $\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} > \kappa$, higher prices induce better reviews, and the seller has an incentive to price higher than under myopia. Note that given uniformly distributed tastes in the baseline model, $\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} = \frac{1}{2}$, this is consistent with the original findings.

D.3. Horizontal Preferences

Let ω be distributed on $[\underline{\omega}, \bar{\omega}]$ according to some density $g_{\omega}(\omega)$ with distribution $G(\omega)$. The only change is that consumers have to take into account that previous purchasing consumers are drawn from the distribution G . The expected purchasing consumer is in this case given by $\omega^e(\tilde{\omega}) \equiv \int_{\tilde{\omega}}^{\bar{\omega}} \omega \frac{g(\omega)}{\int_{\tilde{\omega}}^{\bar{\omega}} g(\omega) d\omega} d\omega$. Then, consumer inference is given by the solution to the equation system

$$\mu + \tilde{\omega} = p \quad (92)$$

$$\mu + \int_{\tilde{\omega}}^1 \omega \frac{g(\omega; \tilde{\omega})}{\int_{\tilde{\omega}}^1 g(\omega, \tilde{\omega}) d\omega} d\omega - \kappa p = \bar{\psi}. \quad (93)$$

Hence, the first-order condition for pricing is given by

$$\frac{dV(\bar{\psi}_t)}{dp_t} = 1 - \tilde{\omega}(p_t) - \frac{1 - \kappa}{1 - \frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}}} p_t + \delta \frac{dV_{t+1}}{d\bar{\psi}_{t+1}} \sigma \left(\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} - \kappa \right), \quad (94)$$

and we immediately obtain that the pricing incentives depend on the relative strength of selection $(\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}})$ and direct price effect κ , as before.

D.4. Heterogeneous κ

Let κ_i be a consumer's individual degree to which her review reflects the purchase price, and let this be independent of quality and price but potentially correlated with ω_i . If they are negatively correlated, consumers with a higher idiosyncratic preference place less weight on the price in their review than consumers with a lower idiosyncratic preference.⁴³ Consumers and the

⁴³While this seems like the natural specification, a positive correlation may also be relevant in some contexts in which consumers with a high idiosyncratic preference are more expert consumers and more confidently assess the value for money a product provides. In this case, the effects are the mirror image of those discussed in the text.

seller will form an expectation about the reviewing consumers' κ based on the set of purchasing consumers, $\mathbb{E}[\kappa|\tilde{\omega}]$. This naturally affects the inference of consumers, and the equation system to solve becomes

$$\begin{aligned}\mu + \omega^e(\tilde{\omega}) - \mathbb{E}[\kappa|\tilde{\omega}]p &= \bar{\psi} & (\text{CONS''}) \\ \mu + \tilde{\omega} &= p. & (\text{RAT''})\end{aligned}$$

Applying the implicit function theorem yields the effect of price changes on the solution pair $(\mu, \tilde{\omega})$

$$\left(\frac{d\mu}{d\tilde{\omega}}\right) = \left(\frac{-\frac{1-2\mathbb{E}[\kappa|\tilde{\omega}]-2\frac{d}{d\tilde{\omega}}\mathbb{E}[\kappa|\tilde{\omega}]}{1+2\frac{d}{d\tilde{\omega}}\mathbb{E}[\kappa|\tilde{\omega}]}}{\frac{2(1+\mathbb{E}[\kappa|\tilde{\omega}])}{1+2\frac{d}{d\tilde{\omega}}\mathbb{E}[\kappa|\tilde{\omega}]}}\right), \quad (95)$$

which in contrast to the baseline case reflects that a price change now has a potential impact on the average price internalization parameter. For $\frac{d}{d\tilde{\omega}}\mathbb{E}[\kappa|\tilde{\omega}] = 0$ and $\mathbb{E}[\kappa|\tilde{\omega}] = \kappa$, i.e., absent correlation between the degree of price internalization and idiosyncratic taste, this reduces to the baseline comparative statics. Using (95), we can derive the seller's first order condition

$$\begin{aligned}1 - \tilde{\omega}(p_t) + p_t\left(-\frac{d\tilde{\omega}(p_t)}{dp_t}\right) + \delta\frac{dV_{t+1}}{d\bar{\psi}_{t+1}}\sigma\frac{d\psi_t}{dp_t} \\ = 1 - \tilde{\omega}(p_t) + p_t\left(-\frac{d\tilde{\omega}(p_t)}{dp_t}\right) + \delta\frac{dV_{t+1}}{d\bar{\psi}_{t+1}}\sigma\left(\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}}\frac{d\tilde{\omega}}{dp_t} - \frac{d\tilde{\omega}}{dp_t}\frac{d}{d\tilde{\omega}}\mathbb{E}[\kappa|\tilde{\omega}]p_t - \mathbb{E}[\kappa|\tilde{\omega}]\right).\end{aligned} \quad (96)$$

Inspecting (96) shows that the main forces from our analysis are still present in that the selection effect and direct price effect determine how current prices affect future reviews. However, the direct price effect is altered and consists of two parts. First, we have $-\mathbb{E}[\kappa|\tilde{\omega}]$, which as before implies that higher prices are factored into the reviews negatively with the expected weight. However, there is an additional effect that derives from the potential correlation of individual κ and ω . The seller takes into account that a price change affects the average price internalization in the review. Hence, if κ and ω are positively correlated, a higher price will induce a higher average κ amplifying the direct price effect, while a negative correlation attenuates it. Importantly, if κ_i and ω_i are uncorrelated, all results remain unchanged.

D.5. Sophisticated Consumers

Consider the following variation to the setup. The fraction of consumers that uses the heuristic is given by $(1 - \lambda)$, where $\lambda \in [0, 1]$ is the fraction of sophisticated consumers who “know” the true quality θ . For simplicity, let tastes be uncorrelated with the sophistication. Given a price p , sophisticated consumers hence purchase iff $\theta + \omega_i \geq p \iff \omega_i \geq p - \theta \equiv \tilde{\omega}_s(p)$. With this cutoff taste of purchasing consumers, we can compute the demand, \hat{q} , as a function of price and the aggregate rating (which remains relevant for consumers using the heuristic), as well as the induced average review $\hat{\psi}$.

$$\begin{aligned}\hat{q}(\bar{\psi}, p) &= (1 - \lambda) \cdot q(\bar{\psi}, p) + \lambda q(p) \\ &= 2(1 - \lambda) (\bar{\psi} - p(1 - \kappa)) + \lambda(1 + \theta - p)\end{aligned} \quad (97)$$

$$\begin{aligned}\hat{\psi}(\bar{\psi}, p) &= (1 - \lambda) \cdot \psi(\theta, \omega^e(\tilde{\omega}), p) + \lambda\psi(\theta, \omega^e(\tilde{\omega}_s), p) \\ &= (1 - \lambda) (\theta + 1 - \bar{\psi} + p(1 - 2\kappa)) + \lambda\frac{1}{2} (\theta + 1 + (1 - 2\kappa)p)\end{aligned} \quad (98)$$

This allows us to set up the dynamic programming problem as before. There remains a one-to-one mapping from target ratings to prices, so that we can replace the control p with $\bar{\psi}'$. Using the

same guess and verify approach as in the baseline model, we establish that the value function is quadratic, which implies a linear policy function. We obtain closed form solutions for the optimal policy and value function, and characterize the stationary equilibrium. For brevity, we state the results below without a formal proof; the approach is identical to solving the baseline model and a Mathematica file verifying the results is available from the authors' websites.

Proposition 4 *For sellers of type $\theta > -1$ and $\bar{\psi}_1 \geq 0$, there is a unique stationary equilibrium that is characterized by long-run ratings, prices and beliefs by consumers using the heuristic as*

$$\hat{\Psi} = \frac{(\theta + 1)((1 - \delta)(4 - \lambda - 2(2 - \lambda)\kappa) + \delta(3 - 2\kappa)(2 - \lambda)\sigma)}{2(1 - \delta)(3 - \lambda - 2(1 - \lambda)\kappa) + 4(2 - \lambda)\delta\sigma} \quad (99)$$

$$\hat{p} = \frac{(\theta + 1)(1 - \delta + (2 - \lambda)\delta\sigma)}{(1 - \delta)(3 - \lambda - 2\kappa(1 - \lambda)) + 2(2 - \lambda)\delta\sigma} \quad (100)$$

$$\hat{\mu} = \theta. \quad (101)$$

The rating system is effective, and all consumers learn the quality.

We can show that the comparative statics with respect to the sensitivity of the rating system are as in the baseline model. The seller always benefits from increased sensitivity, while consumers benefit if and only if long-run prices decrease in the sensitivity, which occurs when $\kappa > \frac{1}{2}$ so that the direct price effect dominates. Moreover, the long-run price and consumer surplus converge to the full information benchmark for $\lambda \rightarrow 1$, while they converge to the baseline outcomes as $\lambda \rightarrow 0$. Finally, we can assess the comparative statics with respect to the fraction of sophisticated consumers, where we focus on the price and rating level.

Corollary 2 *The comparative statics with respect to the fraction of sophisticated consumers λ are as follows.*

- (a) *The long-run rating $\hat{\Psi}$ is decreasing in λ .*
- (b) *The long-run price \hat{p} is decreasing in λ if and only if the selection effect dominates, i.e. for $\kappa < \frac{1}{2}$. For $\kappa > \frac{1}{2}$, an increase in the fraction of sophisticated consumers increases the price level.*

Corollary 2 shows that increased sophistication of consumers lowers the rating level because the incentive for strategic ratings management is mitigated—the rating does not affect sophisticated consumers' purchase decisions. The effect of an increase in λ on the price level depends on the relative strength of price and selection effect. When the selection effect dominates, strategic ratings management exerts an upward pricing pressure so that an increase in λ benefits consumers by alleviating this pressure. In contrast, there is downward pressure on prices when the direct price effect dominates, and a lessening of strategic ratings management due to an increase in λ increases price levels.

D.6. Stochastic Ratings

In reality, ratings arguably have a stochastic component that is not correlated with observables. In particular, it is not necessarily the case that reviews in a given period accurately reflect the average consumer's experience.

To address this issue, we study the case of noisy reviews. Consider our benchmark model, but suppose that the ratings contain some aggregate per-period noise ε_t with $\varepsilon_t \sim F_\varepsilon$ which has a mean of zero and is iid across time with variance σ_ε^2 . The per-period review is

$$\psi_t = \theta + \frac{1 + \omega_t^*}{2} + \varepsilon_t. \quad (102)$$

while the rating aggregation remains unchanged

$$\bar{\psi}_{t+1} = (1 - \sigma)\bar{\psi}_t + \sigma\psi_t. \quad (103)$$

Consumers still apply the same inference rule rationalizing rating-price-quality combinations via the consistency and rationality assumptions. The expected review given price p_t remains unchanged because $\mathbb{E}[\varepsilon_t] = 0$.

However, we relabel our control variable to the expected rating, $\tilde{\psi}_{t+1}$ as the noise cannot be controlled by the seller.

$$\bar{\psi}_{t+1} = \tilde{\psi}_{t+1} + \sigma\varepsilon_t. \quad (104)$$

We impose the assumption that sellers with $\theta > -1$ who face a series of negative shocks remain active and potentially sell at negative prices as they almost surely will make positive profits in the long run by reestablishing their reputation. This facilitates the mathematical analysis and allows for a cleaner comparison to the non-stochastic baseline model. We obtain for the Bellman equation

$$V(\bar{\psi}) = \max_p p \cdot q(p, \bar{\psi}) + \delta \mathbb{E}[V(\bar{\psi}'(p))] \quad (105)$$

$$s.t. \bar{\psi}' = \underbrace{(1 - \sigma)\bar{\psi} + \sigma\psi(p)}_{\tilde{\psi}} + \sigma\varepsilon. \quad (106)$$

Solving for the price that induces the expected rating $\tilde{\psi}$, we obtain

$$p(\tilde{\psi}) = \frac{\sigma(1 + \theta) + \bar{\psi} - \tilde{\psi} - 2\sigma\bar{\psi}}{\sigma(1 - 2\kappa)} \quad (107)$$

which results in the following Bellman equation

$$V(\bar{\psi}) = \frac{2(\bar{\psi} - \tilde{\psi} + \sigma(1 + \theta - 2\bar{\psi}))}{\sigma^2(1 - 2\kappa)^2} ((1 - \kappa)(\tilde{\psi} - \bar{\psi} - \sigma(1 + \theta)) + \sigma\bar{\psi}) + \delta \mathbb{E}[V(\tilde{\psi} + \sigma\varepsilon)]. \quad (108)$$

Guess and verify. We again follow a guess-and-verify approach with a quadratic value function

$$\mathbb{E}[V(\bar{\psi})] = c + d\mathbb{E}[\bar{\psi}] + e\mathbb{E}[\bar{\psi}^2] \quad (109)$$

$$\mathbb{E}[V_{\tilde{\psi}}(\bar{\psi})] = d + 2e\mathbb{E}[\bar{\psi}] \quad (110)$$

where we need to keep in mind that the continuation value is an expected value

$$\mathbb{E}[V(\tilde{\psi} + \varepsilon)] = \mathbb{E}[c + d(\tilde{\psi} + \varepsilon) + e(\tilde{\psi} + \varepsilon)^2] \quad (111)$$

$$= c + d\tilde{\psi} + e(\tilde{\psi}^2 + \mathbb{E}[\varepsilon^2]) \quad (112)$$

$$= c + d\tilde{\psi} + e(\tilde{\psi}^2 + \sigma_\varepsilon^2). \quad (113)$$

We next take the first-order condition of the value function with respect to $\tilde{\psi}$ to obtain

$$0 = \frac{d}{d\tilde{\psi}} p(\tilde{\psi}) q(p(\tilde{\psi}), \bar{\psi}) + \delta(d + 2e\tilde{\psi}) \quad (114)$$

and then solve for the policy function $\tilde{\psi}(\bar{\psi})$, which is

$$\tilde{\psi}(\bar{\psi}) = \frac{\sigma(4(1-\kappa)(\theta+1) + d\delta(1-2\kappa)^2\sigma)}{4(1-\kappa) - 2\delta e(1-2\kappa)^2\sigma^2} - \frac{4(1-\kappa) - 2(2\kappa(1-\sigma) + 3\sigma - 2)}{2\delta e(1-2\kappa)^2\sigma^2} \bar{\psi}. \quad (115)$$

Plugging the resulting policy function into the value function yields

$$\begin{aligned} V(\bar{\psi}) = & \delta \frac{4c\delta e(1-2\kappa)^2\sigma^2 + 8c(\kappa-1) + \sigma(-d^2\delta(1-2\kappa)^2\sigma + 8d(\theta+1)(\kappa-1) + 8e(\theta+1)^2(\kappa-1)\sigma)}{4\delta e(1-2\kappa)^2\sigma^2 + 8(\kappa-1)} \\ & + \delta \frac{4e\sigma_\varepsilon^2(\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1))}{4\delta e(1-2\kappa)^2\sigma^2 + 8(\kappa-1)} - \frac{\delta(2\kappa(\sigma-1) - 3\sigma + 2)(d + 2e(\theta+1)\sigma)}{\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1)} \bar{\psi} \\ & + \frac{-2\delta e(2\sigma-1)(\kappa + \sigma - 1) - 1}{\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1)} \bar{\psi}^2. \end{aligned} \quad (116)$$

Equating coefficients, we solve the resulting equation system

$$\begin{aligned} c = & \delta \frac{4c\delta e(1-2\kappa)^2\sigma^2 + 8c(\kappa-1) + \sigma(-d^2\delta(1-2\kappa)^2\sigma + 8d(\theta+1)(\kappa-1) + 8e(\theta+1)^2(\kappa-1)\sigma)}{4\delta e(1-2\kappa)^2\sigma^2 + 8(\kappa-1)} \\ & + \delta \frac{4e\sigma_\varepsilon^2(\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1))}{4\delta e(1-2\kappa)^2\sigma^2 + 8(\kappa-1)} \\ d = & - \frac{\delta(2\kappa(\sigma-1) - 3\sigma + 2)(d + 2e(\theta+1)\sigma)}{\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1)} \\ e = & \frac{-2\delta e(2\sigma-1)(\kappa + \sigma - 1) - 1}{\delta e(1-2\kappa)^2\sigma^2 + 2(\kappa-1)} \end{aligned} \quad (117)$$

which allows us to obtain analytical characterizations of c , d , and e which we omit for brevity. It is straightforward to verify that the induced a and b coincide with those derived in the absence of noise which implies the same price conditional on the current rating stock.

D.7. Competition

Let two firms, $i \in \{1, 2\}$ be located at the end of a Hotelling line of length 1. Consumers are uniformly located on the Hotelling line, and a consumer at location $x \in [0, 1]$ has taste for firm 1 of $(1-x)$ (the distance from the firm's location) and taste for firm 2 of x . The utilities given this taste and reviews are as in our baseline model, i.e., we have

$$u_1(\theta_1, x) = \theta_1 + (1-x) \quad (118)$$

$$u_2(\theta_2, x) = \theta_2 + x \quad (119)$$

$$\psi_1(\theta_1, x, p_1) = \theta_1 + (1-x) - \kappa p_1 \quad (120)$$

$$\psi_1(\theta_2, x, p_2) = \theta_2 + x - \kappa p_2 \quad (121)$$

$$\psi_2(\theta_1, x, p_1) = \theta_1 + x - \kappa p_1 \quad (122)$$

$$\psi_2(\theta_2, x, p_2) = \theta_2 + (1-x) - \kappa p_2. \quad (123)$$

For simplicity, let the firms compete in two consecutive periods without discounting, $t \in \{0, 1\}$. Each firm i starts with an initial rating $\bar{\psi}_i^0$. In each period, the firms simultaneously set prices. We assume full market coverage and let consumers conduct inference similar to the monopolistic setup; they treat the game as quasistationary and look for inferred qualities μ_1^t, μ_2^t and an inferred

cutoff consumer x_c^t such that all consumers up to x_c^t prefer to purchase from firm 1, consumers above x_c^t prefer to purchase from firm 2, and given this preference, the aggregate ratings are matched.

Formally, inference is determined by looking for the triple μ_1^t, μ_2^t, x_c^t that solves

$$u_1(\mu_1^t, x_c^t) - p_1^t = u_2(\mu_2^t, x_c^t) - p_2^t \quad (\text{IND})$$

$$\psi_1(\mu_1, \frac{x_c^t}{2}, p_1^t) = \bar{\psi}_1^t \quad (\text{CONS}_1)$$

$$\psi_2(\mu_2, \frac{1+x_c^t}{2}, p_2^t) = \bar{\psi}_2^t, \quad (\text{CONS}_2)$$

where $\frac{x_c^t}{2}$ and $\frac{1+x_c^t}{2}$ are the average consumers purchasing from firms 1 and 2, respectively. The derivation of profits and induced reviews given the firms' pricing decisions in a given period is relegated to Appendix D.7.

Omitting time superscripts for ease of exposition, we can solve the equation system characterized by (IND), (CONS₁) and (CONS₂) and obtain the inference (μ_1, μ_2, x_c) given $(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)$ as

$$\mu_1 = \frac{(6\kappa - 2)p_1 + 2(1 - \kappa)p_2 + 6\bar{\psi}_1 - 2\bar{\psi}_2 - 3}{4} \quad (124)$$

$$\mu_2 = \frac{(6\kappa - 2)p_2 + 2(1 - \kappa)p_1 + 6\bar{\psi}_2 - 2\bar{\psi}_1 - 3}{4} \quad (125)$$

$$x_c = \frac{1}{2} + (1 - \kappa)(p_2 - p_1) + (\bar{\psi}_1 - \bar{\psi}_2). \quad (126)$$

which induces quantities

$$q_1(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)) = x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2), \quad q_2(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)) = 1 - x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2) \quad (127)$$

and reviews

$$\psi_1 = \frac{1}{4}(4\theta_1 + 3 + (2 - 6\kappa)p_1 - 2(1 - \kappa)p_2 - 2\bar{\psi}_1 + 2\bar{\psi}_2) \quad (128)$$

$$\psi_2 = \frac{1}{4}(4\theta_2 + 3 - 2(1 - \kappa)p_1 + (2 - 6\kappa)p_2 + 2\bar{\psi}_1 - 2\bar{\psi}_2). \quad (129)$$

Note that for both firms, $\frac{\partial \psi_{-i}}{\partial p_i} < 0$, i.e. that a higher price decreases the other firm's review and hence induced rating for the next period. This provides firms with an additional incentive to charge higher prices.

As full market coverage is assumed, firms solve

$$\max_{p_i^1} p_i^1 \cdot q_i(x_c(p_1^1, p_2^1, \bar{\psi}_1^1, \bar{\psi}_2^1)) \quad (130)$$

in the final period. Solving the system obtained from the two firms' first order conditions yields

$$p_1^1 = \frac{3 + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1}{6(1 - \kappa)}, \quad p_2^1 = \frac{3 - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1}{6(1 - \kappa)} \quad (131)$$

which induces profits

$$\pi_1^1 = \frac{(3 + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1)^2}{36(1 - \kappa)}, \quad \pi_2^1 = \frac{(3 - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1)^2}{36(1 - \kappa)} \quad (132)$$

Final-period profits behave as expected. They are increasing in a firm's own rating and decreasing in the other firm's rating. This gives an incentive to in the first period price such that a firm's own rating increases – which is present also in our monopolistic baseline setting – and such that the opposing firm's rating decreases, which is a novel effect in the competitive setting.

Inference and flow profits in the initial period $t = 0$ are obtained from the same equations. Firms hence solve

$$\max_{p_i^0} V_i \equiv p_i^0 \cdot q_i^0 + \pi_i^1(\bar{\psi}_i^1, \bar{\psi}_{-i}^1). \quad (133)$$

Differentiating the objective with respect to p_i allows us to decompose the impact of the firm's own price.

$$\frac{dV_i}{dp_i^0} = \underbrace{\frac{dq_i^0}{dp_i^0} + p_i^0 \frac{dq_i^0}{dp_i^0}}_{\text{flow profit effect}} + \underbrace{\frac{d\pi_i^1}{d\bar{\psi}_i^1} \cdot \frac{d\bar{\psi}_i^1}{\psi_i^0} \frac{d\psi_i^0}{dp_i^0}}_{\substack{\text{dynamic effect via } \textit{own} \text{ rating} \\ >0 \quad >0 \quad \geq 0}} + \underbrace{\frac{d\pi_i^1}{d\bar{\psi}_{-i}^1} \cdot \frac{d\bar{\psi}_{-i}^1}{\psi_{-i}^0} \frac{d\psi_{-i}^0}{dp_i^0}}_{\substack{\text{dynamic effect via } \textit{rival's} \text{ rating} \\ <0 \quad >0 \quad <0}} \quad (134)$$

Importantly, we have the same incentives as in the baseline setup, i.e., the flow profit effect and dynamic effect via the impact of current period pricing on future profits via the firm's own rating, but in addition an incentive to increase prices due to the dynamic effect via the induced reviews and rating for the opponent.

Reference Price We consider a competitive setup in which consumers evaluate products relative to a reference price.⁴⁴ This reference price can emerge in various ways, and we consider two specific alternatives. First, suppose that the reference price is exogenously given and identical for the two products. In this case, the review functions are altered to reflect the reference price denoted \hat{p} , e.g., $\psi_1(\theta_1, x, p_1) = \theta_1 + (1 - x) - \kappa(p_1 - \hat{p})$. Due to the identical reference price and linearity embedded in the Hotelling setup, this variation does not affect the first order conditions and hence the strategic pricing incentives of the two firms.

More interesting is the setup in which the rival firm's price serves as an endogenous reference price. The utility and review functions in this setup are given by

$$u_1(\theta_1, x) = \theta_1 + (1 - x) \quad (135)$$

$$u_2(\theta_2, x) = \theta_2 + x \quad (136)$$

$$\psi_1(\theta_1, x, p_1) = \theta_1 + (1 - x) - \kappa(p_1 - p_2) \quad (137)$$

$$\psi_1(\theta_2, x, p_2) = \theta_2 + x - \kappa(p_2 - p_1) \quad (138)$$

$$\psi_2(\theta_1, x, p_1) = \theta_1 + x - \kappa(p_1 - p_2) \quad (139)$$

$$\psi_2(\theta_2, x, p_2) = \theta_2 + (1 - x) - \kappa(p_2 - p_1), \quad (140)$$

which yields an explicit inference

$$x_c = \frac{1}{2} - (1 - 2\kappa)p_1 + (1 - 2\kappa)p_2 + \bar{\psi}_1 - \bar{\psi}_2 \quad (141)$$

$$\mu_1 = -\frac{1}{4}(3 + (2 - 8\kappa)p_1 - (2 - 8\kappa)p_2 + -6\bar{\psi}_1 + 2\bar{\psi}_2) \quad (142)$$

$$\mu_2 = -\frac{1}{4}(3 + (2 - 8\kappa)p_2 - (2 - 8\kappa)p_1 + -6\bar{\psi}_2 + 2\bar{\psi}_1) \quad (143)$$

and thus induced reviews

$$\psi_1 = \frac{1}{4}(3 + (2 - 8\kappa)p_1 - (2 - 8\kappa)p_2 + 4\theta_1 - 2(\bar{\psi}_1 - \bar{\psi}_2)) \quad (144)$$

$$\psi_2 = \frac{1}{4}(3 + (2 - 8\kappa)p_2 - (2 - 8\kappa)p_1 + 4\theta_1 - 2(\bar{\psi}_2 - \bar{\psi}_1)). \quad (145)$$

⁴⁴Similar to the monopoly version with an exogenous reference price in Li and Hitt (2010).

Due to the additional interaction, we need to restrict attention to $\kappa < \frac{1}{2}$ for profit functions to be well-behaved—however, as it turns out this does not rule out a dominant price effect. With this restriction, second-period equilibrium profits are given by

$$\pi_1 = \frac{(3 + 2\bar{\psi}_1 - 2\bar{\psi}_2)^2}{36 - 72\kappa}, \pi_2 = \frac{(3 - 2\bar{\psi}_1 + 2\bar{\psi}_2)^2}{36 - 72\kappa}, \quad (146)$$

which are increasing in the own firm's rating, and decreasing in the rival's rating.

Plugging this into firm's first-period maximization problems yields the first-order conditions and allows us to determine the effects of a price increase on the induced own and rival review (and thus aggregate rating affecting second-period profits).

$$\frac{d\psi_1}{dp_1} = \frac{1}{2} - 2\kappa, \frac{d\psi_2}{dp_1} = -\frac{1}{2} + 2\kappa, \frac{d\psi_2}{dp_2} = \frac{1}{2} - 2\kappa, \frac{d\psi_1}{dp_2} = -\frac{1}{2} + 2\kappa. \quad (147)$$

The main takeaway is that it is not always a price *increase* which lowers the rival's rating. Instead, via the inference, it may in fact be a price decrease which achieves this because the price serves as a reference point by which consumers evaluate their rival's product. Notably, this implies that strategic pricing incentives are reinforced. Whenever the own rating benefits from a price increase ($\frac{d\psi_1}{dp_1} > 0 \iff \kappa < \frac{1}{4}$), this also lowers the rival's rating ($\frac{d\psi_2}{dp_1} < 0 \iff \kappa < \frac{1}{4}$) and vice versa.

D.8. Bayesian Consumers

Consider a two-period model with a unit mass of short-lived consumers in each period. The product quality θ is privately observed by the firm and distributed according to a distribution $F(\theta)$ which has continuously differentiable log-concave density f . The expectation $\mu_0 := \mathbb{E}_f[\theta]$ constitutes consumers' initial belief about quality. Consumers are horizontally differentiated with taste component ω_i , which is distributed according to the distribution $G(\omega_i)$. Consumer i purchases the good whenever their expected net utility is weakly positive; that is, whenever their current belief μ_t plus their horizontal taste ω_i exceeds the price p_t : $\mu_t + \omega_i - p_t \geq 0$ and a marginal consumer $\tilde{\omega} = p_t - \mu_t$ exists who is just indifferent between purchasing and not purchasing. This individual rationality consideration determines the demand function of the firm given current belief μ_t : $q(\mu_t, p_t) = 1 - G(\tilde{\omega}(p_t - \mu_t))$.

As in the baseline model, let all purchasing consumers rate the product with equal probability and let the average rating be given by $\psi = \theta + \omega^e(p) - \kappa p + \varepsilon$, where $\omega^e = \int_{\tilde{\omega}(p)}^{\infty} \omega g(\omega) d\omega / (1 - G(\tilde{\omega}(p)))$ is the expected taste of purchasing consumers and ε is mean-zero noise drawn independently of θ and p from the distribution $H(\varepsilon)$ with continuously differentiable log-concave density h . For simplicity, we assume that F , G , and H have full support on the real line.

Firms discount the future with discount factor $\delta \in [0, 1]$ and choose their price in the first period taking its own quality and the effect of the price on the rating into account. Restricting attention to pure strategies, prices are a deterministic function of product quality, $p(\theta)$. While first-period consumers base their purchase decision on their initial belief only, second-period consumers use the first-period rating to update their belief about the product's quality via Bayes' rule (i.e., we still abstract from direct price signaling considerations). Note that if there were no noise in the rating, then second-period consumers would be perfectly informed about the firm's quality due to the one-to-one mapping between prices and ratings.⁴⁵ When there is noise in the rating, consumers' second-period belief is non-degenerate.

⁴⁵This statement conditions on an optimal pricing function for each quality type θ . As long as prices are non-stochastic, the rating will generically be strictly monotonic in the quality.

In addition, note that second-period profits depend on the posterior expected quality $\mu(\psi) := \mathbb{E}[\theta|\psi]$ only because of consumers' risk-neutrality. We therefore derive sufficient conditions for the posterior belief to be increasing in the rating. Denote by $\phi(\theta) := \theta + \omega^e(p(\theta)) - \kappa p(\theta)$ the part of the rating which depends product quality.⁴⁶ As θ is drawn from a distribution with log-concave density, it follows that $\phi(\theta)$ has a log-concave density if $\phi(\cdot)$ is an increasing and weakly concave function (see, for example, Bagnoli and Bergstrom 2005). Moreover, given that ε is drawn from a log-concave density, both $\phi(\theta)$ and ε have log-concave densities implying that ψ has a log-concave density as log-concavity is preserved by convolution (see, for example, Merkle 1998, Lekkerkerker 1953). Lemma 1 in Goeree and Offerman (2003) implies that the conditional expectation $\mu(\psi)$ is increasing in ψ as ϕ is an increasing function.

Second-period profits are $\pi(\psi) := \max_p q(\mu(\psi), p)p$ and the first-order condition is sufficient to find the optimal price as the profit function is quasi-concave due to the log-concavity of the horizontal taste density $g(\omega)$.

Under these assumptions, the first-period price is distorted away from the static monopoly price: expected second-period profits are increasing in ψ and a price distortion in the first period which causes a rating increase leads to a first-order stochastic dominance shift in the distribution of ψ .

The direction of the distortion depends on the total derivative of the revenue function with respect to the price; that is, on the sign of

$$\frac{g(\tilde{\omega}(p))}{1 - G(\tilde{\omega}(p))} (\omega^e(p) - \tilde{\omega}(p)) - \kappa, \quad (148)$$

which is the usual comparison of (positive) selection effect and (negative) price effect.

Hence, in equilibrium, the marginal benefit from a rating-increasing price distortion is offset by the marginal cost of the distortion in terms of first-period profit losses.⁴⁷

Note that if the initial belief were parametrized by an initial rating and if the second-period rating were a weighted average of the initial and the first-period rating by consumers with sensitivity σ , then, by the implicit function theorem, the price distortion in the first period would be increasing in the sensitivity σ analogous to our baseline model whenever the first-period value is concave in the first-period price.

D.9. Vertical Differentiation

Model We adjust the utility function and the review function to account for vertical differentiation of customers

$$u(\theta, \omega, p) = \theta\omega - p \quad (149)$$

$$\psi(\theta, \omega, p) = \frac{\theta \frac{1+\omega}{2}}{1 + \kappa p}, \quad (150)$$

where consumers' types ω are uniformly distributed on $[0, 1]$. The adjusted review function can be interpreted as follows: if a consumer were to receive the good for free, she would report her gross utility. However, as the price rises she accounts for this with a penalty to the review.

⁴⁶This specification contains that rational Bayesian consumers infer the price function correctly in equilibrium.

⁴⁷A full characterization would go beyond the scope of the extension and require additional conditions to ensure well-behavedness of the first-period value.

Consumers' inference obtains from solving the adjusted consistency and rationality conditions and yields

$$\mu(p, \Psi) = 2\Psi + p(2\kappa\Psi - 1) \quad (151)$$

$$\omega(p, \Psi) = \frac{p}{2\Psi + p(2\kappa\Psi - 1)}. \quad (152)$$

The belief is increasing in the price whenever κ is large enough given Ψ . This inference induces demand and average review

$$q(p, \Psi) = 1 - \frac{p}{2\Psi + p(2\kappa\Psi - 1)} \quad (153)$$

$$\psi(p, \Psi) = \frac{\theta\Psi}{2\Psi + p(2\kappa\Psi - 1)}. \quad (154)$$

We see that demand is always decreasing in the price. Moreover, $\kappa\Psi < 1$ ensures that the price is positive and that demand is well-behaved between zero and one. Nonetheless, the induced review may be increasing or decreasing in p . In particular, it is decreasing in p whenever the weight on the price in the review κ is large enough, i.e., whenever $2\kappa\Psi > 1$.

To dissect the two effects, consider the change of a review with respect to the price

$$\frac{d\psi}{dp} = \psi \left(\frac{\frac{d}{dp}\omega^e(p)}{\omega^e(p)} - \frac{\kappa}{1 + \kappa p} \right). \quad (155)$$

The relevant consideration comes from comparing the relative strengths of the selection effect, $\frac{\frac{d}{dp}\omega^e(p)}{\omega^e(p)}$, to the direct price effect $\frac{\kappa}{1 + \kappa p}$. In particular, the sign of $\frac{d\psi}{dp}$ is determined by the following condition:

$$\frac{1}{(2\Psi - p(1 - 2\kappa\Psi))} > / < \kappa. \quad (156)$$

In particular, the selection is dominant, i.e. higher prices positively affect reviews, whenever

$$1 > \kappa(2\Psi(1 + \kappa p) - p) \quad (157)$$

$$\iff 1 + \kappa p > \kappa 2\Psi(1 + \kappa p) \quad (158)$$

$$\iff \kappa < \frac{1}{2\Psi} \quad (159)$$

The reason that the rating appears in the relative strength of the selection and the direct price effect derives from precisely the vertical differentiation structure. As quality interacts multiplicatively with the idiosyncratic type, a higher rating, which translates into a higher belief, *rotates* the demand curve (in contrast to a shift in the demand curve in the model with horizontal differentiation). This can be seen by writing the demand curve as a function of the price and the belief which yields $q(p, \mu) = 1 - \frac{p}{\mu}$ in a model of vertical differentiation, as well as by rewriting (156) as $\mu(\Psi, p)\kappa < / > 1$.⁴⁸ Thus, an increase in the belief, which follows any ceteris paribus increase in the rating, will weaken the selection effect.

Two-period model. While we do not derive an analytical solution to the two-period model due to its intractability, we can illustrate the incentives at play within it nevertheless. In the final, second, period, the seller chooses the myopically optimal price

$$p_2(\Psi_2) = \frac{2\Psi_2}{2(1 - \kappa\Psi_2) + \sqrt{2(1 - 2\kappa\Psi_2)}} \quad (160)$$

⁴⁸This is in contrast to $q(p, \mu) = 1 + \mu - p$ in a model of horizontal differentiation where a change in the belief leads to a parallel shift in the demand curve.

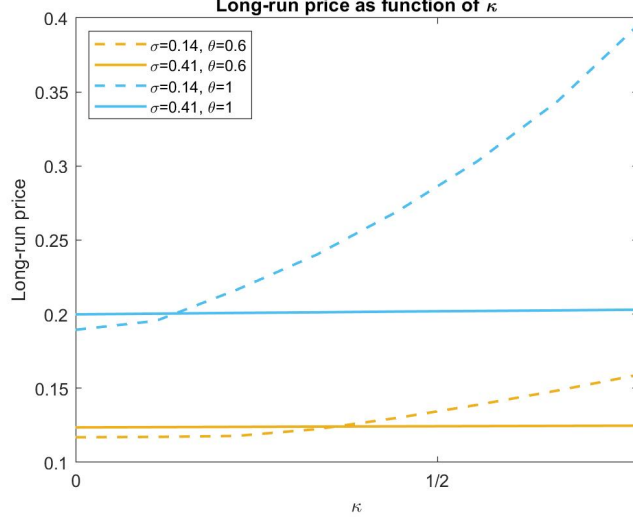


Figure 3: Long-run price as function of κ . Discount rate: $\delta = 0.7$

resulting in a second-period profit of

$$\pi_2(\Psi_2) = \frac{2\Psi_2 \left(3 - 2\kappa\Psi_2 - \sqrt{8(1 - \kappa\Psi_2)} \right)}{(1 - 2\kappa\Psi_2)^2} \quad (161)$$

which is increasing in Ψ_2 given our assumptions. Hence, the seller has an incentive to distort its first-period price to increase her period-two profits via a higher period-two rating.

It follows from the review function induced by a price p , which we computed above, that an increase in the period-one price increases (decreases) reviews if $\kappa\Psi_1 < (>)1/2$. Hence, whenever the weight on the price in the review is sufficiently low, the seller has an incentive to charge higher prices than she would myopically. Moreover, the more responsive the rating system is to incoming reviews the higher the incentive to distort the first-period price and whenever the selection effect dominates, i.e., whenever $\kappa\Psi_1 < 1/2$, consumers dislike higher sensitivity as it leads to higher prices.

However, with longer time horizons, the rating is an endogenous object and, thus, whether the selection or price effect dominates depends on the rating levels that arise over time. For this reason, we solved the problem numerically using standard value function iteration. A numerical illustration is depicted in Figure 3.

Figure 3 highlights the robustness of our findings to a model of vertical differentiation. In addition, it illustrates the dependence of the dominating effect on the interaction of the reviews and the price on the long-run rating, which in turn depends on the quality of the product. Higher quality products will have higher long-run ratings, and thus, the selection effect is less likely to dominate. A higher quality θ and associated higher rating level implies that the direct price effect dominates, and long-run prices are lower with more sensitive rating systems (solid lines) for lower levels of price internalization κ .

E. Empirical Analysis

Steam is an online platform on which video game developers advertise their games and make them available for players to purchase and download. As of 2018, around 20 000 games are offered to 67 million monthly active users, giving *Steam* an estimated market share of the PC

video game market of 50-70%.⁴⁹ Annual revenue of the platform in 2017 is estimated at \$4.3 billion.⁵⁰ Prior studies have shown that reviews matter substantially in the market for video games, see, e.g., Zhu and Zhang (2010).

After purchasing a game on the platform (and only then), players can leave a binary rating (either ‘Recommended’ or ‘Not Recommended’) as well as a written review text. Both are visible to potential buyers on the *Steam* page of the game. After purchasing a game, it is part of a player’s ‘library’ and can be launched through the platform. Some players choose not to make their game libraries private, so that it is publicly visible which games they own.

This empirical section uses a unique dataset, which matches individual players’ purchases of video games to the ratings they left them on the platform, as well as to player characteristics. The dataset was created by crawling through the libraries of around 50 000 players every day from February to August 2017 and registering changes in the libraries as game purchases. Purchase prices were obtained by crawling through all game sites on *Steam* on a daily basis. Finally, using the players’ unique platform identification number, the ratings left by a subset of the purchasers can be matched to their purchase dates and prices, as well as some player-specific variables. The resulting dataset consists of around 12 000 rating-purchase price matches. Observed variables include the full purchase price and discount (if any), whether the rating was positive, how long the purchasing player played the game before writing the review, how many other games the player owns and other player-specific variables. Summary statistics for all observed variables are in Appendix E.

In order to sign and quantify the association between price changes and the probability of receiving a positive rating we use the following regression framework:

$$y_{ig} = \lambda_g + \beta \cdot X_{ig} + \delta \cdot P_{ig} + \epsilon_{ig} \quad (162)$$

In (162), the outcome variable y is the binary rating player i gave game g . λ_g denotes game fixed effects, X_{ig} is a vector of reviewer-game specific control variables, P_{ig} is the price as a fraction of the full price at which i purchased g and ϵ_{ig} denotes the error term. We estimate a linear probability model in order to accommodate the high-dimensional game-fixed effects, λ_g . It is important to account for unobservable game quality using the fixed effects so that we can identify the price effect using variation in the discount rate within games over time. Including these in a logistic regression specification leads to convergence problems with standard software packages.

The rating not only depends on the price, but also on the quality of the game, as well as characteristics of the reviewer. In order to control for the quality of the game, Equation (162) includes game fixed effects. Using game fixed effects requires us to limit the dataset to games for which we observe at least two purchases, leaving 3 746 observations. Observable characteristics of the reviewer-game match, such as for how long she played the game before writing the review, how helpful her review was to other potential buyers and how old the game was at the time of purchase, are included as control variables.

Table 2 shows estimation results for regression (162). Without controlling for reviewer-game specific variables (column (1)), the coefficient on price is positive but insignificant. Including control variables leads to an increase in the coefficient, which is now significantly different from zero at the 90% confidence level. The coefficient of 0.074 implies that discounting the price of a game by 50% is associated with a 3.7% lower probability of receiving a positive review. This result is consistent with the selection effect being the pre-dominant force in the overall sample. Next we split the sample according to whether a game belongs to the “casual” genre or not and re-run the regressions. Casual games are typically straightforward in terms of gameplay and

⁴⁹<https://expandedramblings.com/index.php/steam-statistics/>

⁵⁰<https://www.gamesindustry.biz/articles/2018-03-23-valves-generates-record-breaking-usd4-3bn-from-sales-revenue-in-2017>

fairly interchangeable. They appeal to a narrow range of relatively unsophisticated players, who are less willing to spend time and money on video games. The results of the regression using only observations from purchases of casual games (column (3)) indicate that the direct price effect dominates in this subsample. A higher price is associated with a statistically highly significant reduction in the propensity of receiving a positive review. A discount of 50% translates to a more than 20% increase in the probability of receiving a positive review. The opposite is true for the non-casual games (column (4)). As in the overall sample, higher prices are associated with better reviews for these games, indicating the importance of the selection effect for non-casual games.

Tables

Table 1: Summary Statistics. An observation corresponds to a rating-purchase combination. $N = 3\,738$.

	Mean	SD	p25	p50	p75
Initial price (in \$)	27.3	17.1	15	20	40
Fraction of full price actually paid	0.79	0.28	0.60	1	1
Recommended	0.81	0.40	1	1	1
Age of game at purchase time (in days)	362.6	671.0	3	64	440
Playtime at review (in minutes)	1 138.8	3 754.6	111	414	1134
Number of reviews written	37.4	95.2	11	19	35
Number of owned games	315.9	451.6	88	180	361
Number of ratings for review	12.1	36.4	2	4	9
Fraction Helpful	0.73	0.27	0.5	0.75	1
Length of Review (in Words)	742.6	1 095.9	107	343	906

Table 2: Recommendation Propensity.

The outcome variable is equal to one for a positive review and 0 for a negative review. price refers to the fraction of the undiscounted price at which the game was purchased. game_age is the number of days between purchase and release of the game. review_playtime refers to the number of hours that reviewer had played before writing the review. num_reviews and num_owned_games refer to the number of previously written reviews and the number of games owned by the reviewer. rated and frac_helpful refer to the number of ratings the review received and the fraction that found the review helpful. length is the length of the review in words. For the regressions in columns (3) and (4) the sample was split depending on whether the purchased game belongs to the genre "Casual".

	(1)	(2)	(3)	(4)
	recommend	recommend	recommend	recommend
price	0.041 (0.043)	0.074* (0.042)	0.115*** (0.044)	-0.418*** (0.134)
game_age		-0.001** (0.000)	-0.001** (0.000)	-0.001 (0.001)
review_playtime		0.000** (0.000)	0.000** (0.000)	-0.000 (0.000)
num_reviews		0.000 (0.000)	0.000 (0.000)	-0.000 (0.000)
num_owned_games		-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
rated		-0.001*** (0.000)	-0.001*** (0.000)	-0.002*** (0.001)
frac_helpful		0.596*** (0.024)	0.604*** (0.025)	0.502*** (0.091)
length		-0.000*** (0.000)	-0.000*** (0.000)	0.000 (0.000)
Sample	Full	Full	Non-Casual	Casual
Observations	3738	3738	3413	325
R^2	0.270	0.404	0.403	0.455

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$