

Discussion Paper Series – CRC TR 224

Discussion Paper No. 143
Project C 03

Pricing for the Stars
Dynamic Pricing in the Presence of Rating Systems

André Stenzel*
Christoph Wolf**
Peter Schmidt ***

February 2020
(First version : January 2020)

*University of Mannheim and MaCCI. E-Mail: andre.stenzel[at]uni-mannheim.de
** Bocconi University and IGIER. E-Mail: christoph.wolf[at]unibocconi.it
*** QuantCo. E-Mail: peter.schmidt[at]quantco.com

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Pricing for the Stars*

Dynamic Pricing in the Presence of Rating Systems

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

February 18, 2020

Abstract:

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, they also induce only consumers with a strong taste for the product to purchase, which on average improves reviews. We provide conditions under which the latter effect dominates so that ratings management leads to upward pressure on prices. This upward pressure increases in the sensitivity of the aggregate rating to incoming reviews. As a consequence, recent changes to rating systems may have harmed consumers by increasing long-run price levels.

JEL: D21, D82, L15

Keywords: Rating Systems, Dynamic Pricing, Asymmetric Information

*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 Berlin, Bocconi, Genova, Leicester, Mannheim, OsloMet, EARIE 2019, ESEM 2019, MaCCI Summer Institute 2019, IIOC 2018, EIEF-UNIBO-Igier Workshop 2018, EARIE 2016, MaCCI annual conference 2016, and VfS Jahrestagung 2016. Special thanks go to Johannes Dittrich for making the Steam data available to us. Stenzel gratefully acknowledges funding by the German Research Foundation (DFG) via the CRC TR 224, Project C03.

[†]University of Mannheim and MaCCI. E-Mail: andre.stenzel[at]uni-mannheim.de.

[‡]Bocconi University and IGIER. E-Mail: christoph.wolf[at]unibocconi.it.

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

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, which have 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 23 billion GBP each year, see CMA (2015). The effect on 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 per cent.¹

Aggregate ratings are computed from individual reviews which reflect purchase prices and heterogeneous tastes.² Importantly, these factors determining the review are typically unobservable to future consumers, leaving scope for sellers to use current prices to affect reviews and therefore future quality perceptions and profits. This strategic pricing as a means for online reputation management is the focus of our paper.

We provide a tractable model of dynamic pricing in the presence of rating systems. We analyze how recent changes to the design of rating systems affect pricing incentives and consumer welfare, and focus on the increased sensitivity of aggregate ratings to incoming reviews on platforms such as Amazon and Steam.³ Our model is consistent with stylized facts about the relationship between purchase prices and review scores. Higher prices lead to better ratings when reviews reflect consumers' overall enjoyment of the product, but to worse ratings when consumers primarily judge the value for money. When the rating system becomes more sensitive to recent reviews, firms face a stronger incentive to strategically manage ratings via their pricing decision. This implies that long-run price levels and seller revenues are higher but consumer welfare lower for products such as books where higher prices positively affect reviews. Our results indicate that regulators should consider the implications of the design of rating systems for consumer welfare via the effect on firms' pricing.⁴

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 prior to purchase, and differ in their taste for the good. They form a belief based on current observables consisting of the price and the product's aggregate rating. Consumers that purchase can leave a review. The aggregate rating is a weighted average of the previous rating and the average current review. The weight parametrizes the sensitivity of the rating system to recent reviews.

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).

³Amazon puts a disproportionate emphasis on recent reviews in computing aggregate ratings since 2015, while the video game platform Steam displays a recent average score in addition to the overall rating as of 2016. See wired.com (2019) and Steam (2016) for details.

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

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, i.e., the fraction being equal to zero, corresponds to consumers reporting their gross utility, while maximal price internalization, i.e., the fraction being one, corresponds to net utility reporting. We consider the degree of price internalization as a product-specific feature. Some products are primarily evaluated based on the overall consumption utility they provide, while other—more standardized—products are primarily evaluated based on their value for money.

We capture two empirically documented effects (see, e.g., Cabral and Li (2015) and Zegers (2019)), which, to the best of our knowledge, have not yet been incorporated into a single model. A price increase directly worsens reviews by purchasing consumers. This *direct price effect* is stronger the more reviews internalize the purchase price. However, a higher price also requires a higher expected gross utility of the marginal consumer and thus changes the composition of consumers who purchase. Via this *selection effect*, purchasing consumers on average have a higher 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 only use current observables, consisting of the aggregate rating and price, and their knowledge of the utility and review functions. They form beliefs about the product quality and the set of purchasing consumers to rationalize the observed rating at the *current* price.⁶ Intuitively, consumers treat the model as if it were in a stationary equilibrium. A similar inference rule has been used in the literature (see Crapis et al. (2017) and Besbes and Scarsini (2018)), and is related in spirit to the notion of cursed equilibrium (see Eyster and Rabin (2005)): Consumers do not fully take into account 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 high price internalization, consumers rationalize the same rating at a higher price by inferring the good to be of higher quality; thus, the set of purchasing consumers is less responsive and demand less elastic.

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 therefore balances the effects of prices on flow profits and on the aggregate rating. But how do prices affect future ratings? The direct price effect and the selection effect go in opposite directions. We show that the selection effect dominates whenever the degree of price internalization is below a cutoff, in which case higher prices induce better reviews. For high price internalization, the direct price effect dominates and higher prices induce worse reviews.

This is consistent with the seemingly conflicting empirical evidence on the relationship between prices and reviews. Reviews are likely to feature higher price internalization the more stan-

⁵Note that the selection effect in our setup refers to selection into purchasing, not to selection into reviewing conditional on purchase. We study the latter in an extension.

⁶Full rationality would place high requirements on consumers’ cognitive abilities. They would need to know or form beliefs about how many periods have passed, 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.

standardized the product. In line with this reasoning, higher prices lead to worse reviews for USB sticks (Cabral and Li (2015)), but to better reviews for books (Zegners (2019)). Corroborating evidence is also found in our own empirical analysis of data from the video game platform Steam, see Appendix C. Simple “casual games” feature a negative relationship between prices and reviews, while video games overall display a positive relation.

We show that there is a unique equilibrium, in which 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. We find that rating systems are effective in the sense that consumers learn the quality of the product in the long run despite their misspecified model. When the rating is more sensitive to incoming reviews, a given change in the price level that increases future ratings leads to a greater improvement. The seller’s profits are thus increasing in the sensitivity. How consumers are affected depends on the effect of prices on reviews. A sensitive rating system induces sellers to manage their ratings more. When price internalization is high, they charge lower prices in the long run which benefits consumers. If instead the selection effect dominates, consumer surplus decreases when the rating system becomes more sensitive because sellers raise prices.

To address recent changes to rating systems on Amazon, Steam and other platforms, we model a platform that chooses the sensitivity of its rating systems. The platform maximizes a weighted sum of seller’s profits and consumer surplus.

It is immediate that the platform chooses the highest possible sensitivity whenever both consumers and sellers benefit from high sensitivity, i.e., when the direct price effect dominates. When preferences are misaligned, i.e., when the selection effect dominates, the platform chooses the highest (lowest) possible sensitivity when the weight on consumer surplus is sufficiently low (high).

In a series of extensions, we show that the tension between direct price and selection effect and the main takeaways of our analysis remain qualitatively unaffected. We analyze stochastic ratings, selection into reviewing, different distributions over idiosyncratic preferences, non-linear price effects as well as competition between sellers.

Given the vast portfolio and turnover of online platforms, our results suggest that the shift towards more sensitive rating systems may have harmed a substantial amount of consumers. Returning to our opening quote, it is crucial to assess the relation between prices and reviews empirically—that is, the relative importance of price and selection effect—to understand whether changes to rating systems benefit consumers. In this respect, our model may be helpful as it lends itself to structural estimation of its primitives due to the closed-form solutions for the value and policy functions.

Related literature. A key ingredient of our model is that reviews reflect not only quality, but also additional considerations such as heterogeneous consumer tastes and purchase prices (see Li and Hitt (2008, 2010) and Bhargava and Feng (2015)). De Langhe et al. (2015) documents a significant difference between aggregate ratings on Amazon.com and objective quality measures from quality scores, even for products where vertical differentiation is likely to be more important

than subjective tastes. Zegners (2019) documents that both purchasing prices and horizontal characteristics matter in determining the review a consumer leaves for a particular ebook. Cabral and Li (2015) finds that the purchasing price is negatively correlated with the resulting review for USB sticks.

We study the incentives of sellers to manage their reviews through strategic pricing. Hence, our paper is closely related to the literature on reputation management, see Bar-Isaac and Tadelis (2008) for an overview. Large parts of the literature have focused on settings in which the strategic choice is whether to invest in quality (Cabral and Hortacsu (2010) and Board and Meyer-ter Vehn (2013)), or whether to suppress or to advertise information (Marinovic et al. (2018)). While we focus on information transmission via induced ratings and the effect of prices on future quality perception, there is a vast literature on direct price signaling (Wolinsky (1983) and Bagwell and Riordan (1991)). Osborne and Shapiro (2014) embed price-signaling considerations in a dynamic context where a monopolist chooses both quality and price. While similar considerations are present in our model, the review and rating system forms the basis for the seller's strategic actions that affect future quality inferences in our setting.

There is a growing body of literature addressing social learning from rating systems in the presence of a selection effect, see, e.g., Ifrach et al. (2019) and Acemoglu et al. (2019) who provide conditions under which Bayesian consumers successfully learn product qualities over time, as well as Bondi (2019). In contrast to these papers, we focus on the pricing decision of a privately informed monopolist. This is closely related to Crapis et al. (2017).⁷ However, our paper differs in three dimensions. First, we allow the seller to change the price dynamically in every period instead of at most two times. Second, we allow for a product-specific degree to which reviews reflect the purchase price. With this, we can capture the various documented correlations between prices and reviews depending on the relative strength of price and selection effect. Third, we explicitly consider the sensitivity of rating systems to recent reviews as a design parameter and assess the implications for long-run consumer welfare.

Bonatti and Cisternas (2018) analyze aggregate scores about consumers' purchasing histories which are informative to short-lived firms about consumers' evolving willingness-to-pay. As such, the price-setter has an informational disadvantage and cannot price strategically to affect future demand. This is in contrast to our paper, where the seller holds an informational advantage over a sequence of boundedly rational buyers as in, e.g., von Thadden (1992).

We also relate to recent work on the design of rating systems. Kovbasyuk and Spagnolo (2018) show that low memory of ratings can be optimal if the quality of a product or service changes over time, as it prevents inefficient exit from the market. Che and Hörner (2017) study the optimal design of recommender systems to incentivize collaborative learning. Luca (2017) and Dai et al. (2018) discuss several issues regarding the design of rating systems including reviews being selected and how to aggregate reviews into rating statistics. Klein et al. (2016) empirically evaluate 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). Auctions may limit the potential of sellers to engage in strategic pricing.

⁷Strategic pricing to exploit the selection effect is also analyzed in Matysková and Šípek (2017) in a two-period game in which consumers exhibit cursed beliefs.

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 in Section 3. We discuss platform incentives in Section 4 and extensions in Section 5. Section 6 concludes.

2. Model

We consider a monopolistic long-lived producer of a good with 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.

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 cdf 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, the seller decides on the price p_t . Marginal costs of production are independent of quality and normalized to 0. The seller is risk-neutral and discounts future profits at a rate $\delta \in (0, 1)$.

Consumers In each period t , there is a new unit mass of risk-neutral consumers who live for one period. Consumers value quality, and differ with respect to their personal taste for the good which 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 that 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 consumer 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 via a review and rating system is structured as follows. If consumer i purchased the good and leaves a review ψ_i according to the review function

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

with $\kappa \in [0, 1)$. In line with the empirical evidence (see, e.g., Li and Hitt (2008, 2010)), 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 lower is the review left. We regard this as realistic behavior of consumers because it relates the enjoyment of a product to its price. The degree of price internalization κ may vary across different products, which is important when relating the model's predictions to the empirical evidence. Note that for $\kappa \rightarrow 1$, each consumer reports her net utility, i.e., her individual surplus. For $\kappa = 0$, reviews reflect gross utilities instead. 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.

For tractability, we assume that every consumer who purchases the good leaves a review with the same probability. The average review in a given period is used to update the aggregate rating. This average review is equal to the review left by the consumer with average taste $\omega^e(\omega_t^*) = E[\omega | \omega \geq \omega_t^*] = \frac{1+\omega_t^*}{2}$, where ω_t^* is the taste of the marginal consumer who purchases. The rating system is characterized by the mapping from current aggregate rating $\bar{\psi}_t$ and current average review $\psi_t = \theta + \frac{1+\omega_t^*}{2} - \kappa p_t$ into next period's aggregate rating $\bar{\psi}_{t+1}$. We consider a specific class of rating systems following the updating rule

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

Given an initial rating $\bar{\psi}_1$, this can be rewritten as

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

In (3) and (4), $\sigma \in [\underline{\sigma}, 1]$ parametrizes the sensitivity of the rating system to incoming reviews. The higher σ , the more responsive is 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 only consists of last period's average review, while $\sigma = \underline{\sigma} > 0$ denotes the system with the lowest sensitivity.⁸ Rating systems with varying degrees of sensitivity σ allow us to assess the recent pushes by online platforms such as Amazon.com to have more recent reviews matter more for the displayed aggregate rating – in our context, this 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 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 or not. If consumers choose to purchase, they realize their net utility as in (1) and leave a review as in (2).

Technical Assumptions For technical purposes, we impose the following assumptions. First, we require $\underline{\theta} < -1$. This ensures that there are quality levels such that no consumer should purchase the good irrespective of taste. Second, we require $\bar{\theta} < \infty$ to ensure boundedness of profits in each period. Finally, we restrict attention to σ such that $\sigma < 1 - \kappa$ whenever $\kappa > \frac{1}{2}$. This ensures concavity of the flow payoff in the state variable in the dynamic programming problem.

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

⁸We impose $\underline{\sigma} > 0$ because the rating would be invariant to incoming reviews for $\sigma = 0$.

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 could only be computed 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 rating systems which is a highly complicated problem. Instead, we assume that consumers try to rationalize the aggregate rating while supposing that past consumers were faced with 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 (see Eyster and Rabin (2005)) in 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 the cutoff type of purchasing consumers ω^* . This is because, despite the use of a heuristic, consumers are cognizant of the fact that reviews are driven by the characteristics of past consumers who purchased the good. Inference about the quality cannot be conducted in isolation from inference about the set of purchasing consumers.

The assumption greatly improves tractability as it reduces inference to a two-dimensional fixed point problem. If all past consumers faced exactly 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} .⁹

An alternative way of interpreting the assumption 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 rating will be exactly 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.

⁹This technical assumption in principle allows consumers to believe that the quality exceeds the maximal possible quality $\bar{\theta}$ and allows us to circumvent specifying boundary solutions in the inference.

While the assumption implies that consumers are non-Bayesian, we do not consider it to be far from reality.¹⁰ As discussed previously, past prices are not directly observable on online platforms. Individual reviews are often available, but they cannot be directly linked to the price at which the good was purchased even with the use of historical price data from price-tracking websites (which by itself is cumbersome to obtain), and rarely mention explicit price points. As they are, moreover, 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 only on these two inputs, the heuristic used by treating the posted price as part of a stationary equilibrium seems a reasonable approximation. Besides 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.¹¹

2.1. Price and Selection Effect

In this part, we solve for the consumers' inference explicitly and derive the demand function. We also obtain a formal characterization of the price and selection effect and illustrate the dynamic pricing incentives.

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

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

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

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 . This induces demand

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

Note that the law of demand is satisfied. A higher aggregate rating always increases demand, while a higher price decreases it. Moreover, the inferred quality μ is increasing in the aggregate rating $\bar{\psi}$. The responsiveness of demand and the inferred quality to price changes both depend on the degree to which reviews reflect the purchase price κ . When κ is high, reviews reflect net consumer surplus. In this case, demand is unresponsive to price changes as consumers rationalize a given rating at a higher price by inferring that the good is of higher quality. In contrast, when κ is low, reviews reflect gross surplus and consumers rationalize a given rating at a higher price via a different selection of consumers – the inferred cutoff taste $\tilde{\omega}$ increases and demand adjusts, while inferred quality decreases.

¹⁰It was in fact inspired through introspection and examination of online purchase decisions of a subset of the authors.

¹¹If 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 would be correct.

To understand this behavior, it is helpful to write (CONS) and (RAT), which determine the inference, as

$$u(\tilde{\mu}, \omega^e(\tilde{\omega})) = \bar{\psi} + \kappa p \quad (\text{CONS}')$$

$$u(\tilde{\mu}, \tilde{\omega}) = p \quad (\text{RAT}')$$

Consistency requires that the average consumer, characterized by taste $\omega^e(\tilde{\omega}) = \frac{1+\tilde{\omega}}{2}$, enjoys a gross utility equal to the observed aggregate rating $\bar{\psi}$ corrected for the degree to which reviews reflect the price. Note that the higher the price p , the higher the required gross utility of the average consumer. Rationality requires that the gross utility of the marginal consumer equals the price. Importantly, the utility of the marginal consumer reacts more strongly to changes in the cutoff taste $\tilde{\omega}$ than the utility of the average consumer, while they show the same reaction to changes in the inferred quality $\tilde{\mu}$. How a price increase at a given rating is rationalized therefore depends on κ .

To illustrate this, consider a price increase. When κ is large and close to 1, this causes the right-hand-sides of both (CONS') and (RAT') to increase by roughly the same amount. The gross utilities of the average and marginal consumer thus also have to increase by the same amount, which due to the differential reaction to changes in the cutoff taste $\tilde{\omega}$ can only be facilitated via changes in the inferred quality $\tilde{\mu}$. The converse is true when κ is small and close to 0. The right-hand-side of (CONS') is unaffected while that of (RAT') increases. The average consumer's utility has to stay constant while the marginal consumer's utility has to increase, which can only be facilitated by a simultaneous increase in the cutoff taste $\tilde{\omega}$ and decrease in the inferred quality $\tilde{\mu}$.

Induced Average Review: Selection Effect vs. Direct Price Effect Towards analyzing the dynamic pricing incentives, it is important to look at the effect of the current price on the induced reviews. As seen previously, the price the seller charges affects the current period's set of purchasing consumers. The induced average review is given by

$$\psi(\theta, \omega^e(\tilde{\omega}), p) = \theta + \omega^e(\tilde{\omega}) - \kappa p, \quad (8)$$

where $\tilde{\omega}$ is determined via the inference and depends on $\bar{\psi}$ and p . Plugging in (6),

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

In (9), $\theta + 1$ is the maximal possible utility, i.e. that of a consumer with taste $\omega = 1$. As is evident, κ determines whether a high average review is induced by high or low prices. This is because κ determines the strength of both the direct price and selection effect. When assessing the impact of a marginal price change on the induced average review, we obtain

$$\frac{d\psi}{dp} = \frac{d\omega^e}{d\tilde{\omega}} \frac{d\tilde{\omega}}{dp} - \kappa. \quad (10)$$

Changing the price has two effects. First, it directly affects the average review as these incorporate the purchase price. This is the marginal direct price effect, which is equal to $-\kappa$. In

addition, there is the selection effect. By changing the price, the seller also changes the cutoff consumer via the inference and hence the taste of the average consumer which determines the average review. As established above, the degree to which prices are incorporated into reviews determines the responsiveness of demand and thus the strength of the selection effect. This is because 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 $\frac{d\omega^e}{d\bar{\omega}} \frac{d\bar{\omega}}{dp} = 1 - \kappa$; the larger κ , the less responsive demand and thus the lower the selection effect.

Overall, this implies that a high average review is induced by high prices whenever κ is low, as the selection effect dominates the direct price effect in this case. When κ is high, the reverse is true and the direct price effect dominates; it is low prices which 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. Follows from the preceding discussion. ■

Lemma 1 predicts a positive (negative) correlation of reviews and prices when the price internalization is high (low). We consider the degree of price internalization to be a product-specific parameter. For standardized goods, such as USB-sticks, the price is more important in the review as these goods are evaluated according to their value for money. In contrast, this is of lesser importance for books and other non-standardized goods. With this reasoning, our model can rationalize the otherwise conflicting evidence in the empirical literature, see Cabral and Li (2015) and Zegers (2019). We perform our own empirical analysis in Appendix C using a unique data set that allows us to match individual reviews to the purchasing price on the video game platform Steam. Similar to books, video games in general are non-standardized and exhibit a substantial idiosyncratic component. In line with this, we find that the overall correlation between prices and reviews for video games is positive. However, the correlation reverses when we restrict the sample to casual games, which are more standardized. To the best of our knowledge, our paper is the first to rationalize both correlations within a single model.

Illustration in a two-period model. To illustrate the effect 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

$$\max_{p_2} p_2 \cdot q_2 = p_2 \cdot 2 (\bar{\psi}_2 - p_2(1 - \kappa)) \quad (11)$$

and thus chooses the myopic monopoly price given by $p_2 = \frac{\bar{\psi}_2}{2(1-\kappa)}$, which yields as profits

$$\pi_2 = \frac{\bar{\psi}_2^2}{2(1-\kappa)}. \quad (12)$$

(12) shows that second-period profits are increasing in the rating at the beginning of that period. Moving to period 1, we can write the profits as

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

with first-order condition

$$2(\bar{\psi}_1 - 2(1 - \kappa)p_1) + \delta \frac{\partial \pi_2(\bar{\psi}_2(p_1))}{\partial p_1} = 0. \quad (14)$$

To understand the seller's pricing problem in period 1, we have to derive the effect of current prices on future profits. Denoting 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

$$\frac{\partial \pi_2(\bar{\psi}_2(p_1))}{\partial p_1} = \frac{\partial \pi_2}{\partial \bar{\psi}_2} \frac{\partial \bar{\psi}_2}{\partial \psi_1} \frac{\partial \psi_1}{\partial p_1} \quad (15)$$

$$= \sigma \frac{\bar{\psi}_2(p_1)}{(1 - \kappa)} (1 - 2\kappa). \quad (16)$$

Plugging this into the first-order condition (14) yields for the optimal period-1 price

$$p_1^* = \frac{\bar{\psi}_1}{2(1 - \kappa)} + \delta \sigma \frac{\bar{\psi}_2(p_1)}{2(1 - \kappa)^2} (1 - 2\kappa). \quad (17)$$

If the seller would price myopically in the first period, it 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 seller prices lower (higher) relative to the myopic optimum.

3. Dynamic Pricing and Long-Run Properties

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

The seller solves the problem

$$\max_{(p_t)_{t \geq 0}} \sum_{t=0}^{\infty} \delta^t p_t q(p_t, \bar{\psi}_t) \quad (18)$$

$$\text{s. t. } \bar{\psi}_t = (1 - \sigma)\bar{\psi}_t + \sigma \psi(p_t, \bar{\psi}_t) \quad (19)$$

$$\bar{\psi}_0 = \bar{\psi}. \quad (20)$$

Note that the flow-profits are bounded and, 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 given by

$$V(\bar{\psi}) = \max_p \left\{ pq(p, \bar{\psi}) + \delta V(\bar{\psi}') \right\} \quad (21)$$

$$\text{s. t. } \bar{\psi}' = (1 - \sigma)\bar{\psi} + \sigma\psi(p, \bar{\psi}). \quad (22)$$

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

$$\underbrace{q(p, \bar{\psi}) + p \frac{dq(p, \bar{\psi})}{dp}}_{\text{static monopoly}} + \delta \underbrace{\frac{dV(\bar{\psi}')}{d\bar{\psi}'}}_{\substack{\text{effect of} \\ \text{rating on} \\ \text{future profits}}} \underbrace{\frac{\partial \bar{\psi}'}{\partial \psi} \frac{\partial \psi}{\partial p}}_{\substack{\text{effect of} \\ \text{current price} \\ \text{on future rating}}}. \quad (23)$$

There are two effects. First, flow profits are affected by the increase in the price. This is captured by standard static monopoly price effects. Second, the price change impacts future profits via the change in the induced rating. This effect in turn can be decomposed into the (discounted) sensitivity of the future profits to the aggregate rating next period ($\frac{dV'}{d\bar{\psi}'}$), 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 the price change has 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 choice of the seller. The review ψ in any period is linear in p given current aggregate rating $\bar{\psi}$ and given by

$$\psi = \theta + 1 - \bar{\psi} + p(1 - 2\kappa) \quad (24)$$

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

$$\bar{\psi}' = (1 - \sigma)\bar{\psi} + \sigma(\theta + 1 - \bar{\psi} + p(1 - 2\kappa)) \quad (25)$$

$$\Leftrightarrow p = \frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa} + \frac{\bar{\psi} - \bar{\psi}'}{\sigma(1 - 2\kappa)}. \quad (26)$$

The problem of the seller can then be written as

$$V(\bar{\psi}) = \max_{\bar{\psi}'} \left\{ p(\bar{\psi}')q(\bar{\psi}', \bar{\psi}) + \delta V(\bar{\psi}') \right\}. \quad (27)$$

We solve the problem by guessing 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 $\bar{\psi}' = a + b\bar{\psi}$, and verifying that this is indeed true. We obtain closed-form solutions for the optimal policy and value function as well as the long-run prices and ratings. This allows us to characterize the stationary equilibrium that we converge to.

Proposition 1 *For sellers of type $\theta > -1$, there is a unique stationary equilibrium that is*

characterized by long-run ratings, prices and beliefs

$$\tilde{\Psi} = \frac{(\theta + 1)(2(1 - \delta)(1 - \kappa) + \delta\sigma(3 - 2\kappa))}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)} \quad (28)$$

$$\tilde{p} = \frac{(\theta + 1)((1 - \delta) + 2\delta\sigma)}{4\delta\sigma + (1 - \delta)(3 - 2\kappa)} \quad (29)$$

$$\tilde{\mu} = \theta. \quad (30)$$

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

Proof. See Appendix A. ■

Proposition 1 shows that there is a unique stationary equilibrium to which the market converges. In particular, the rating system is effective in alleviating 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, long-run prices depend on the details of the rating system, i.e., depend on σ .

Moreover, sellers with qualities that are so low that they should not sell under full information ($\theta < -1$; this implies that even a consumer with $\omega = \bar{\omega} = 1$ enjoys a negative gross utility) will eventually leave the market. It can be shown that these sellers always price such that the rating is declining over time until they cannot make positive profits.

Using (29), we can assess how long-run prices are affected by the rating system as parametrized by σ . As consumers' quality inference is correct, the price level directly determines 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

$$\begin{aligned} \tilde{\pi} &= \tilde{p} \cdot q(\tilde{\Psi}, \tilde{p}) \\ &= \frac{2((1 - \delta) + 2\delta\sigma)(\delta\sigma + (1 - \delta)(1 - \kappa))}{(4\delta\sigma + (1 - \delta)(3 - 2\kappa))^2} (\theta + 1)^2. \end{aligned} \quad (31)$$

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

Proof. See Appendix A. ■

Corollary 1 contains the main implications of the paper. In the stationary equilibrium, the seller balances flow payoff and future profit considerations, which amounts to balancing exploitation of the current rating and strategic reputation management via the induced reviews. The less

sensitive the rating system is to new reviews, i.e. the lower σ , the more the seller would have 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, it 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 $\tilde{\Psi}$ is increasing in σ . This also provides a testable implication of our model. Cross-sectionally, the average rating of products should increase following a change to the rating system which emphasizes recent reviews.

For consumers, the price level determines the long-run consumer surplus because inference is correct. How consumer surplus depends on the sensitivity therefore depends on whether the seller has an incentive to over- or underprice relative to the myopically optimal price. 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 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 high future ratings (see also the two-period case in (17)). This incentive is mitigated by a lower σ as an individual 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. Whenever reviews respond strongly to the purchase price, an emphasis on more recent reviews as implemented by, for example, Amazon and Steam in recent years, is beneficial for consumer surplus.

Speed of Convergence Of natural interest is the speed at which pricing and ratings converge. As discussed, we are able to establish that our value function takes the form $V(\bar{\psi}) = c + d\bar{\psi} + e\bar{\psi}^2$ which translates into a law of motion for the rating of the form $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$. We provide explicit characterizations for b and establish $|b| < 1$ in Appendix A. From this, it follows immediately that ratings converge linearly with rate $|b|$.

Proposition 2 *Ratings converge linearly to Ψ at rate $|b|$. For $\sigma < \frac{1}{2}$, $b > 0$ and $\frac{\partial b}{\partial \sigma} < 0$.*

Proof. See Appendix A. ■

When assessing the comparative statics, we restrict attention to comparatively low sensitivity levels as this ensures that sellers do not have an incentive to strategically induce oscillating ratings, which would materialize for $b < 0$. The main takeaway is that for these parametrizations a lower sensitivity increases the speed at which ratings and hence prices and inference converge.

Similarly, the lower the sensitivity the slower everything reverts back to the stationary long-run outcome following a (zero-probability) deviation from it. This can be seen in the impulse response functions of a 10%-shock to the long-run rating in Figure 1.

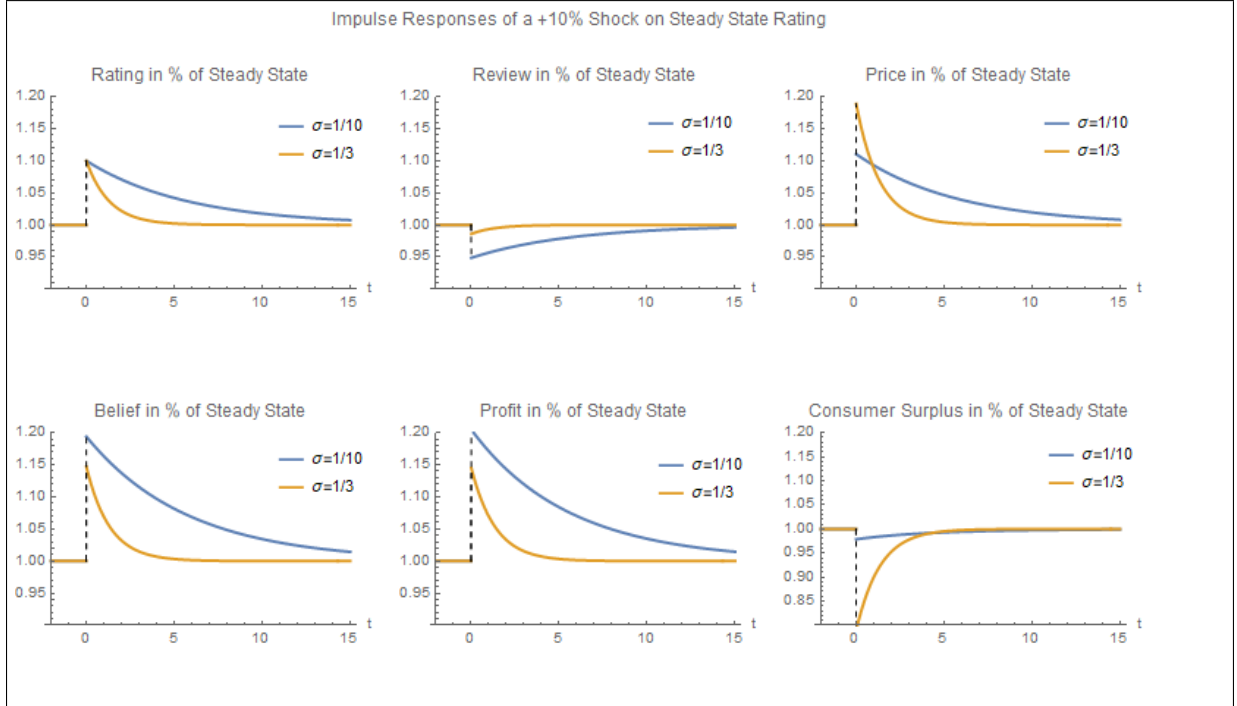


Figure 1: Impulse Response Functions given a 10% shock to the steady state rating

4. Platform Incentives

Given the closed-form expressions for the long-run consumer surplus and seller profits, we are able to speak to the incentives of a platform who maximizes a weighted sum of the two. Specifically, we consider a platform who chooses the sensitivity of the rating system to new reviews, i.e., sets σ , to maximize

$$\pi^P = w_c \cdot \tilde{CS} + (1 - w_c) \cdot \tilde{\pi}, \quad (32)$$

where \tilde{CS} 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 think of this as the objective of a social planner who maximizes total surplus ($w_c = \frac{1}{2}$), a regulator who focuses on consumer surplus ($w_c \rightarrow 1$), or a platform operator who maximizes seller profits ($w_c \rightarrow 0$), for example because it receives a commission.

It is straightforward from Corollary 1 that the highest possible σ maximizes π^P whenever the direct price effect is large, i.e. $\kappa > \frac{1}{2}$, 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 downward pressure on prices so as to manage the rating. However, when the direct price effect is small, the interests diverge. The

seller prefers a responsive rating system, i.e. a high σ , while consumers are better off whenever σ is low.

Proposition 3 (Platform Incentives) *A platform maximizing $\pi^P(\sigma)$ as given in (32) chooses the highest sensitivity σ if*

(i) *the direct price effect is strong ($\kappa > 1/2$), or*

(ii) *the direct price effect is weak ($\kappa < 1/2$) and the weight on the consumer is sufficiently low ($w_c < \frac{3-\delta-4(1-\delta)\kappa}{(1-\delta)(1-2\kappa)}$).*

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

$$\tilde{\sigma} = \frac{(1-\delta)(1-2\kappa-(3-4\kappa)w_c)}{2w_c\delta}. \quad (33)$$

Proof. See Appendix A. ■

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 its maximization. Importantly, this misalignment can only materialize whenever the direct price effect in the review parametrized by κ is not too large, $\kappa < \frac{1}{2}$.

To choose the seller-optimal sensitivity, i.e. to let the rating system be maximally sensitive to incoming reviews despite consumers preferring the opposite, the weight on the consumers 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 $w_c > \frac{1}{3}$. This is 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 which have a dominant selection effect, it has to be that the platform puts more than twice as much weight on consumer surplus than on profits.

5. Extensions

In this subsection, we discuss several extensions of our baseline model and, in particular, show that the mechanisms driving our main results are present and remain the driving forces in these settings.

Number of reviews. In practice, the number of reviewing consumers, both at a given point in and over time, is relevant for the computation and evolution of the aggregate rating. If more consumers review the product, the weighted average will move more strongly than if few do. The previous model abstracted from this consideration for tractability reasons. It can, however, be incorporated into our setting by utilizing a modified updating rule

$$\bar{\psi}_{t+1} = (1 - q_t\sigma)\bar{\psi}_t + q_t\sigma\psi_t, \quad (34)$$

where q_t is the number of consumers in period t . In this sense, the rating reacts more strongly when more consumers purchase, while we still allow the general sensitivity of the rating system to vary. We leave consumers' inference unchanged; they continue to treat the observed price and rating as quasi-stationary. Unfortunately, we cannot apply our results directly because we obtained them through a guess and verify procedure with a linear policy function and, with the present formulation, the objective is more complicated and includes a quadratic term on the control which precludes us from finding closed form solutions.¹² 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

$$\begin{aligned} \frac{dV(\bar{\psi}_t)}{dp_t} &= q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t + \delta \frac{dV_{t+1}}{dp_t} \\ &= \underbrace{q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t}_{\text{static monopoly pricing}} + \delta \underbrace{\frac{dV_{t+1}}{d\bar{\psi}_{t+1}}}_{\text{effect of reviews on CV}} \left(\underbrace{\sigma q_t(p_t) \frac{d\psi_t(p_t)}{dp_t}}_{\substack{\text{better reviews} \\ \rightarrow \\ \text{higher ratings}}} - \underbrace{\sigma \frac{dq_t(p_t)}{dp_t} (\bar{\psi}_t - \psi_t(p_t))}_{\text{number of reviews effect}} \right). \end{aligned} \quad (35)$$

Contrasting this with the first-order condition in the original model, the only newly appearing 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. The main takeaway is that while an additional effect materializes, the main mechanisms driving our results in the baseline model continue to be present and relevant for the direction of price effects on reviews.

Distribution over reviewing agents. In the main part of the paper we have assumed that every purchasing consumer reviews the product with the same probability. Similar to the number of reviews not mattering, we made this assumption for expositional and tractability purposes. There is ample empirical evidence that reviews are not uniformly distributed over consuming agents' satisfaction but rather bimodal on the extremes, see e.g. Bolton et al. (2004) and Dellarocas and Wood (2008). To incorporate this, suppose that the probability of reviewing is given by $f_\psi(\omega; \tilde{\omega})$, i.e., the probability to review depends on the consumer's idiosyncratic taste and the cutoff consumer.¹³ We assume that f is continuously differentiable in both its arguments and strictly positive on its support. Moreover, if the number of purchasing consumers decreases ($\tilde{\omega}$ increases), the average reviewing consumer, $w^e(\tilde{\omega}) \equiv \int_{\tilde{\omega}}^1 w \frac{f_\psi(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(w, \tilde{\omega})} d\omega$, increases, but not too much, $\frac{dw^e(\tilde{\omega})}{d\tilde{\omega}} \in (0, 1)$. Under these assumptions, the consumers' inference and demand is given

¹²Both the quantity and the review are linear in the price and multiplied with each other.

¹³The reason we include the cutoff consumer is that empirically the agents that are most likely to review are those with the most extreme utilities conditional on purchase.

by the solution to the equation system

$$\alpha\mu + \beta\tilde{\omega} = p \quad (36)$$

$$\alpha\mu + \beta \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega; \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega - \kappa p = \bar{\psi}. \quad (37)$$

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

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = - \begin{pmatrix} \frac{\partial u}{\partial \theta} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} & \frac{\partial \psi}{\partial \tilde{\omega}} \end{pmatrix}^{-1} \begin{pmatrix} -1 \\ \frac{\partial \psi}{\partial p} \end{pmatrix}. \quad (38)$$

Inverting the matrix and plugging in the partial derivatives yields

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = \begin{pmatrix} \kappa - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega \\ \frac{1}{\alpha} \frac{1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega}{1 - \kappa} \\ \frac{1}{\beta} \frac{1 - \kappa}{1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega} \end{pmatrix}, \quad (39)$$

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, 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

$$\frac{dV(\bar{\psi}_t)}{dp_t} = q_t(p_t) + \frac{dq_t(p_t)}{dp_t} p_t + \delta \frac{dV_{t+1}}{dp_t} \quad (40)$$

$$= 1 - \tilde{\omega}(p_t) - \frac{1}{\beta} \frac{1 - \kappa}{1 - \frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega} p_t \quad (41)$$

$$+ \delta \frac{dV_{t+1}}{d\bar{\psi}_{t+1}} \sigma \left(\frac{d}{d\tilde{\omega}} \int_{\tilde{\omega}}^1 \omega \frac{f_{\psi}(\omega, \tilde{\omega})}{\int_{\tilde{\omega}}^1 f_{\psi}(\omega, \tilde{\omega}) d\omega} d\omega - \kappa \right) \quad (42)$$

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, $\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} = \frac{1}{2}$, so that this is consistent with the original findings.

Distribution of horizontal preferences. In a similar vein to the distribution of reviewing consumers, we can incorporate different distributions of horizontal preferences. The intuition for the pricing incentives remain unchanged if we instead assume that ω is 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

$$\alpha\mu + \beta\tilde{\omega} = p \quad (43)$$

$$\alpha\mu + \beta \int_{\tilde{\omega}}^1 \omega \frac{g(\omega; \tilde{\omega})}{\int_{\tilde{\omega}}^1 g(\omega, \tilde{\omega})} d\omega - \kappa p = \bar{\psi}. \quad (44)$$

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

$$\frac{dV(\bar{\psi}_t)}{dp_t} = 1 - \tilde{\omega}(p_t) - \frac{1}{\beta} \frac{1 - \kappa}{1 - \frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}}} p_t + \delta \frac{dV_{t+1}}{d\bar{\psi}_{t+1}} \sigma \left(\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} - \kappa \right) \quad (45)$$

and it is immediate that 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.

Stochastic ratings. Ratings arguably have a stochastic component in reality, which is not correlated with the observables. In particular, it is not necessarily the case that reviews in a given period represent the average consumer. To address this issue, we study the case of mean zero noise in the reviews in each period such that the induced review which enters the rating in a given period t is given by

$$\psi_t = \alpha\theta + \beta\tilde{\omega}_t^e - \kappa p_t + \varepsilon_t \quad (46)$$

where $\varepsilon_t \in \{-\epsilon, \epsilon\}$ and $\Pr(\epsilon) = \frac{1}{2}$. We can show analytically and numerically that the results remain qualitatively unchanged. Convergence to a stationary equilibrium still obtains (albeit incorporating the reaction to the individual shocks) and comparative statics with respect to m are as in the case without noise. The formal derivations and a Matlab program with the value function iteration for a numerical solution can be obtained from the authors upon request.

Non-linear price effect. We can also address the case in which a price discount ($\bar{p} - p_t$) has a non-linear effect on the review such that the review function is given by

$$\psi_t = \alpha\theta + \beta\tilde{\omega}_t^e - \kappa(\bar{p} - p_t)^2. \quad (47)$$

Unfortunately, this precludes us from obtaining closed form expressions for the long-run stationary equilibrium. However, we solve the model numerically by value function iteration and verify that consumers nevertheless learn the quality of the product in the long run. Matlab-Files are available upon request. The effect of discounts on reviews is then given by

$$\frac{d\psi(p, \bar{\psi})}{d(\bar{p} - p)} = 4\kappa(\bar{p} - p) - 1 \quad (48)$$

and the effect depends on how deep the discount is. The higher κ , the lower is the critical discount level such that higher discounts induce better reviews, resembling similar comparative statics in the effect of prices on reviews from the main part of the paper.

More generally, a previous version of this article provides conditions under which consumer inference in the vein of the present paper is uniquely determined for a flexible class of potentially

non-linear utility and review functions, see Stenzel and Wolf (2016).

Competition The main effects also carry over to a setting in which multiple firms compete. To illustrate this, consider the following stylized setup. 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 . Utilities given this taste and reviews are as in our baseline model.

Each firm i starts with an initial rating $\bar{\psi}_i^0$. For simplicity, let the firms compete in two consecutive periods without discounting, $t \in \{0, 1\}$. In each period, firms simultaneously set prices. Consumers conduct inference similarly to the monopolistic setup: They treat the game as quasi-stationary 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.¹⁴

Formally, inference is determined by looking for the triple μ_1^t, μ_2^t, x_c^t which 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 B.

Flow profits in each period naturally depend positively on a firm's own rating at the start of the period, $\bar{\psi}_i^t$, and negatively on the other firm's rating $\bar{\psi}_{-i}^t$. In the first period, this gives an incentive to price such that the own rating at the start of the next period is high, while the opponent's rating is low.

Consider the maximization problem of a firm in the first period. By setting its price, it affects inference and hence its quantity and flow profits. In addition, it affects the induced review for its own product via the price and the selection effect, as in our baseline. However, there is the additional effect that the own price affects the selection of purchasing consumers of the rival firm which exerts influence on that firm's reviews and hence future rating. We show in the Appendix that a high price charged by firm i decreases the rival firm's average review in that period ψ_{-i} . This is intuitive – a higher price, all else equal, pushes the cutoff consumer further away and closer to the rival firm as the indifferent consumer requires a higher taste for the own product to purchase it. This in turn leads to a lower induced review for the rival via the selection effect.

To illustrate this more formally, each firm i in period 0 solves

$$V_i = \max_{p_i^0} p_i^0 \cdot q_i^0 + \pi_i^1(\bar{\psi}_i^1, \bar{\psi}_{-i}^1). \quad (49)$$

¹⁴We implicitly assume full market coverage which facilitates the exposition of the effects.

as booking.com have recently shifted towards more disaggregated ratings allowing a separate assessment of the multiple dimensions affecting a consumer's satisfaction. This removes or at least attenuates the channels present in our model, and, as tastes regarding hotels are likely to be idiosyncratic and thus feature a strong selection effect, removes upward pressure on prices. Moving to disaggregated ratings can thus be interpreted as a means of competing for consumers.

Appendix

A. Proofs

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 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

$$p = \frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa} + \frac{\bar{\psi} - \bar{\psi}'}{\sigma(1 - 2\kappa)} \quad (51)$$

$$q = \frac{2(\kappa - 1)(\sigma(\theta + 1) - \bar{\psi}') + 2\bar{\psi}(\kappa + \sigma - 1)}{(2\kappa - 1)\sigma} \quad (52)$$

and the Bellman equation becomes

$$V(\bar{\psi}) = \max_{\bar{\psi}'} \left(\frac{\theta + 1 - 2\bar{\psi}}{1 - 2\kappa} + \frac{\bar{\psi} - \bar{\psi}'}{\sigma(1 - 2\kappa)} \right) \frac{2(\kappa - 1)((\theta + 1)\sigma - \bar{\psi}') + 2\bar{\psi}(\kappa + \sigma - 1)}{(2\kappa - 1)\sigma} + \delta V(\bar{\psi}'). \quad (53)$$

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

$$V'(\bar{\psi}) = d + 2e\bar{\psi}. \quad (54)$$

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

$$\bar{\psi}' = \frac{\sigma(4(1 - \kappa)(\theta + 1) + d\delta(1 - 2\kappa)^2\sigma)}{4(1 - \kappa) - 2\delta e(1 - 2\kappa)^2\sigma^2} - \frac{4(1 - \kappa) - 2(2\kappa(1 - \sigma) + 3\sigma - 2)}{2\delta e(1 - 2\kappa)^2\sigma^2} \bar{\psi} \quad (55)$$

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

$$\begin{aligned} \bar{\psi}' &= \frac{\sigma(\theta+1)(\delta(3-2\kappa)\sigma+2(1-\delta)(1-\kappa))}{2\delta\sigma^2+(1-\kappa)(1-\delta)+\sqrt{(\delta(1-2\sigma)^2-1)(\delta(\kappa+\sigma-1)^2-(\kappa-1)^2)}} \\ &+ \frac{1-\kappa+\delta(2\sigma-1)(\kappa+\sigma-1)-\sqrt{(\delta(1-2\sigma)^2-1)(\delta(\kappa+\sigma-1)^2-(\kappa-1)^2)}}{\delta(2(1-\kappa)-(3-2\kappa)\sigma)}. \end{aligned} \quad (56)$$

Denote

$$a \equiv \frac{\sigma(\theta+1)(\delta(3-2\kappa)\sigma+2(1-\delta)(1-\kappa))}{2\delta\sigma^2+(1-\kappa)(1-\delta)+\sqrt{(\delta(1-2\sigma)^2-1)(\delta(\kappa+\sigma-1)^2-(\kappa-1)^2)}} \quad (57)$$

$$b \equiv \frac{1-\kappa+\delta(2\sigma-1)(\kappa+\sigma-1)-\sqrt{(\delta(1-2\sigma)^2-1)(\delta(\kappa+\sigma-1)^2-(\kappa-1)^2)}}{\delta(2(1-\kappa)-(3-2\kappa)\sigma)}. \quad (58)$$

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

$$\bar{\psi}_\tau = \left(a \cdot \sum_{i=0}^{\tau-2} b^i \right) + b^{\tau-1} \bar{\psi}_1 \quad (59)$$

and thus, using $|b| < 1$,¹⁶

$$\lim_{\tau \rightarrow \infty} \bar{\psi}_\tau = \frac{a}{1-b} = \frac{(\theta+1)(\delta\sigma(3-2\kappa)+2(1-\delta)(1-\kappa))}{4\delta\sigma+(1-\delta)(3-2\kappa)} \equiv \Psi. \quad (60)$$

At this long-run rating, we can use (26) and obtain the long-run price

$$\tilde{p} = p(\Psi, \Psi) = \frac{((1-\delta)+2\delta\sigma)(\theta+1)}{4\delta\sigma+(1-\delta)(3-2\kappa)}. \quad (61)$$

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.

Proof of Corollary 1 Differentiating (29) gives

$$\frac{\partial \tilde{p}}{\partial \sigma} = \overbrace{\frac{2(1-\delta)\delta(\theta+1)}{((1-\delta)(3-2\kappa)+4\delta\sigma)^2}}^{>0} \cdot (1-2\kappa), \quad (62)$$

¹⁶This follows from $\kappa < \max\{\frac{1}{2}, 1-\sigma\}$ which holds given that $\sigma < 1-\kappa$ whenever $\kappa > \frac{1}{2}$.

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 (31) and get

$$\frac{\partial \tilde{\pi}}{\partial \sigma} = \frac{\overbrace{2(1-\delta)^2\delta(\theta+1)^2}^{>0}}{((1-\delta)(3-2\kappa) + 4\delta\sigma)^3} \cdot (1-2\kappa)^2, \quad (63)$$

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

$$\frac{\partial \tilde{\Psi}}{\partial \sigma} = \frac{\overbrace{(1-\delta)\delta(\theta+1)}^{>0}}{((1-\delta)(3-2\kappa) + 4\delta\sigma)^2} \cdot (1-2\kappa)^2. \quad (64)$$

Proof of Proposition 2 Linear convergence of a sequence $\{y_t\}$, which converges to y , at rate μ requires that $\mu = \lim_{t \rightarrow \infty} \frac{y_t - y}{y_{t-1} - y}$. Given that $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$, we have $\Psi = \frac{a}{1-b}$ and thus

$$\frac{|\bar{\psi}_t - \bar{\psi}|}{|\bar{\psi}_{t-1} - \bar{\psi}|} = \frac{|a + b\bar{\psi}_{t-1} - \frac{a}{1-b}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (65)$$

$$= \frac{|a \cdot \left(1 - \frac{1}{1-b}\right) + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (66)$$

$$= \frac{|-\frac{ab}{1-b} + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (67)$$

$$= \frac{|b \cdot \left(\bar{\psi}_{t-1} - \frac{a}{1-b}\right)|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} \quad (68)$$

$$= |b|. \quad (69)$$

We explicitly derived for b that

$$b = \frac{1 - \kappa + \delta(2\sigma - 1)(\kappa + \sigma - 1) - \sqrt{(\delta(1 - 2\sigma)^2 - 1)(\delta(\kappa + \sigma - 1)^2 - (\kappa - 1)^2)}}{\delta(2(1 - \kappa) - (3 - 2\kappa)\sigma)} \quad (70)$$

Recall that $\kappa < \max\{\frac{1}{2}, 1 - \sigma\}$ was assumed, which ensures that $b > 0$ for $\sigma < \frac{1}{2}$. Under the same restrictions, it can be established that $\frac{\partial b}{\partial \sigma} < 0$. The detailed calculations are extensive and omitted here for brevity. They are verified in the supplementary Mathematica file.

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 (29) to obtain $\tilde{C}\tilde{S}$ and plugging this together with (31) into (32), we obtain

$$\pi^P = \frac{2(\theta + 1)^2}{\beta} \cdot \frac{(\delta\sigma + (1 - \delta)(1 - \kappa))((2 - w_c)\delta\sigma + (1 - \delta)(1 - \kappa w_c))}{(4\delta\sigma + (1 - \delta)(3 - 2\kappa))^2} \quad (71)$$

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

$$\frac{\partial \pi^P}{\partial \sigma} > 0 \iff 2w_c \delta \sigma - (1 - \delta) [(1 - 2\kappa) - (3 - 4\kappa)w_c] < 0 \quad (72)$$

and analogously

$$\frac{\partial \pi^P}{\partial \sigma} < 0 \iff 2w_c \delta \sigma - (1 - \delta) [(1 - 2\kappa) - (3 - 4\kappa)w_c] > 0. \quad (73)$$

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

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

the proposition immediately follows.

B. Calculations for Competition Setup

Omitting time superscripts to simplify the exposition, we have in any given period that

$$u_1(\theta_1, x) = \theta_1 + (1 - x) \quad (74)$$

$$u_2(\theta_2, x) = \theta_2 + x \quad (75)$$

$$\psi_1(\theta_1, x, p_1) = \theta_1 + (1 - x) - \kappa p_1 \quad (76)$$

$$\psi_2(\theta_2, x, p_2) = \theta_2 + x - \kappa p_2. \quad (77)$$

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

$$\mu_1 = \frac{(6\kappa - 2)p_1 + 2(1 - \kappa)p_2 + 6\bar{\psi}_1 - 2\bar{\psi}_2 - 3}{41} \quad (78)$$

$$\mu_2 = \frac{(6\kappa - 2)p_2 + 2(1 - \kappa)p_1 + 6\bar{\psi}_2 - 2\bar{\psi}_1 - 3}{41} \quad (79)$$

$$x_c = \frac{1}{2} + (1 - \kappa)(p_2 - p_1) + (\bar{\psi}_1 - \bar{\psi}_2). \quad (80)$$

which induces quantities

$$q_1(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)) = x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2), \quad q_2(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)) = 1 - x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2) \quad (81)$$

and reviews

$$\psi_1 = \frac{1}{4} (4\theta_1 + 3 + (2 - 6\kappa)p_1 - 2(1 - \kappa)p_2 - 2\bar{\psi}_1 + 2\bar{\psi}_2) \quad (82)$$

$$\psi_2 = \frac{1}{4} (4\theta_2 + 3 - 2(1 - \kappa)p_1 + (2 - 6\kappa)p_2 + 2\bar{\psi}_1 - 2\bar{\psi}_2). \quad (83)$$

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

$$\max_{p_i^1} p_i^1 \cdot q_i(x_c(p_1^1, p_2^1, \bar{\psi}_1^1, \bar{\psi}_2^1)) \quad (84)$$

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

$$p_1^1 = \frac{3 + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1}{6(1 - \kappa)}, p_2^1 = \frac{3 - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1}{6(1 - \kappa)} \quad (85)$$

which induces profits

$$\pi_1^1 = \frac{(3 + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1)^2}{36(1 - \kappa)}, \pi_2^1 = \frac{(3 - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1)^2}{36(1 - \kappa)} \quad (86)$$

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

$$\max_{p_i^0} p_i^0 \cdot q_i^0 + \pi_i^1(\bar{\psi}_i^1, \bar{\psi}_{-i}^1), \quad (87)$$

which gives rise to the FOC as discussed in the main part.

C. 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%.¹⁷ Annual revenue of the platform in 2017 is estimated at \$4.3 billion.¹⁸

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

¹⁷<https://expandedramblings.com/index.php/steam-statistics/>

¹⁸<https://www.gamesindustry.biz/articles/2018-03-23-valves-generates-record-breaking-usd4-3bn-from-sales-revenue-in-2017>

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 Table 1.

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:

$$y_{ig} = \lambda_g + \beta \cdot X_{ig} + \delta \cdot P_{ig} + \epsilon_{ig} \quad (88)$$

In (88), 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.

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 (88) 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 (88). 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 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

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

	(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

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

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”.

References

- ACEMOGLU, D., A. MAKHDOUMI, A. MALEKIAN, AND A. OZDAGLAR (2019): “Learning From Reviews: The Selection Effect and the Speed of Learning,” *working paper*.
- 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 (2018): “Consumer scores and price discrimination,” *working paper*.
- BONDI, T. (2019): “Alone, Together: Product Discovery Through Consumer Ratings,” *SSRN working paper*.
- 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.
- CHE, Y.-K. AND J. HÖRNER (2017): “Recommender systems as mechanisms for social learning,” *The Quarterly Journal of Economics*, 133, 871–925.
- 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.
- FTC (2019): “Cosmetic company wrote fake reviews,” *retrieved from <https://www.consumer.ftc.gov/blog/2019/10/cosmetic-company-wrote-fake-reviews/>*.
- IFRACH, B., C. MAGLARAS, M. SCARSINI, AND A. ZSELEVA (2019): “Bayesian social learning from consumer reviews,” *Operations Research*, 67, 1209–1221.
- 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*.
- 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,” *Com (March 15, 2016). 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 G. ZERVAS (2016): “Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud,” *Management Science*, 62, 3412–3427.
- MARINOVIC, I., A. SKRZYPACZ, AND F. VARAS (2018): “Dynamic certification and reputation for quality,” *American Economic Journal: Microeconomics*, 10, 58–82.
- MATYSKOVÁ, L. AND J. ŠÍPEK (2017): “Manipulation of Cursed Beliefs in Online Reviews,” *CERGE-EI Working Paper Series*.
- OSBORNE, M. AND A. H. SHAPIRO (2014): “A dynamic model of price signaling, consumer learning, and price adjustment,” *Tech. rep.*
- STEAM (2016): “Customer Review System Updated To Show Recent Reviews And Summary,” *retrieved from <https://store.steampowered.com/news/21695/>*.
- 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*.