Pricing for the Stars*

Dynamic Pricing in the Presence of Rating Systems

André Stenzel† Christoph Wolf‡ Peter Schmidt§

August 17, 2019

Abstract:

We consider dynamic price-setting in the presence of rating systems and asymmetric information about product quality. A long-lived monopolist sells a product of privately known and fixed quality to a sequence of short-lived consumers. In each period, the price charged determines the taste characteristics of purchasing consumers who leave reviews after purchase. Individual reviews are aggregated into ratings observable to future consumers. The price has two effects on future ratings: (i) a direct *price effect* on reviews, and (ii) an indirect *selection effect* by determining the tastes of reviewing consumers. Inference is conducted by looking for the inferred quality and cutoff taste such that purchase decisions are individually rational and the current aggregate rating is matched. We show that rating systems are effective as consumers correctly infer the quality of the product in the long run. However, the design of the rating system affects long-run prices, profits and consumer surplus. If the direct price effect dominates the selection effect, consumers benefit from a rating system which is more sensitive to newly arriving reviews. In contrast, they prefer a more persistent system if the selection effect dominates. Firms are unambiguously better off the more sensitive the rating system.

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, Luis Cabral, Xavier Lambin, Marco Ottaviani, Martin Peitz, Johannes Schneider, Nicolas Schutz, Konrad Stahl and Stefan Weiergräber for helpful comments and suggestions, and appreciate input from seminar audiences at Berlin, Bocconi, Genova, Leicester, IIOC 2018, EARIE 2016, 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. 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@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

Asymmetric information about product quality is pervasive in online markets. As a consequence, platforms such as Amazon.com, eBay or Steam exert extensive effort to design and to improve their rating systems which allow information transmission between consumers.[1] The consumers in turn heavily rely on ratings, which have a significant effect on demand (see e.g. Chevalier and Mayzlin (2006), Cabral and Hortacsu (2010), Luca (2011)). As a natural consequence, firms have strong incentives to manage their rating using the strategic means they have available. In this paper, we provide a theoretical model in which we focus on a specific channel through which a seller affects her product's rating and hence future profits: strategic pricing.[2]

Rating systems allow information to pass from past consumers to future consumer. Importantly, the transmitted information does not solely reflect the inherent quality of the considered product. It is confounded by additional factors such as the taste of the reviewing consumer or the assessment relative to the purchase price (see e.g. De Langhe et al. (2015) and Zegners (2019)). Even if the quality of a product does not change over time, using the information contained in observed ratings is a nontrivial task for consumers, especially as consumers typically observe only aggregate statistics and not all purchase- and review-relevant information about past consumers. This leaves scope for the firm to strategically use its pricing to affect who purchases their product, and the reviews they leave for future consumers. However, the theoretical literature on strategic pricing in the presence of rating systems is surprisingly scarce.[3] At the same time, it is imminently relevant. Given the evolution of rating systems over time (e.g. on Amazon where the aggregate rating displayed now puts more emphasis on more recent reviews), understanding sellers' incentives and resulting prices is key to evaluate whether particular changes to rating systems benefit consumers.

Our theoretical analysis is particularly important given the, at first glance, conflicting empirical assessments of the relationship between prices and ratings. Cabral and Li (2015) document a

[1] Depending on the considered platform, product quality incorporates various considerations such as the actual quality of the considered product (Amazon, Steam) or the reliability of the seller in terms of shipping and handling (eBay). Our focus is on the inherent quality of a product.

[2] In our setup, we abstract from explicit price signalling considerations. There is a substantive literature assessing settings in which prices may serve as an effective means to convey information about product quality in markets with asymmetric information: Products of different quality are sold at different prices which not only reflect differences in production costs but also allow consumers to infer (something about) the underlying quality when making their purchase decision (see e.g. Bagwell and Riordan (1991), Osborne and Shapiro (2014), Wolinsky (1983)). In the context of online markets, price signaling seems to be unable to (fully) resolve the asymmetric information issue – in fact, if prices were perfectly informative about product quality there would be no need for an effective rating system. Indeed, Bhargava and Feng (2015) show that price signalling and reviews are substitutes.

[3] A notable exception is concurrent work by Acemoglu et al. (2019) who consider strategic pricing by a firm without asymmetric information and, hence, symmetric learning. In contrast, we consider a setting in which the seller privately knows her product's quality.

negative relationship, i.e. higher prices leading to lower ratings, in an experiment selling USB sticks on eBay. In contrast to this, Zegners (2019) documents that self-published e-Books feature a positive relationship, whereby books obtain a higher rating when they are sold for positive prices instead of given away for free. Using data from the video game platform Steam, which allows us to match purchase prices to reviews for a given game, we document that the overall relationship is similarly positive. However, the relation is reversed when looking only at "casual games".[4]

How can these relationships be rationalized? The association of higher prices with higher ratings can be motivated by a *selection effect*.[5] If consumers are heterogeneous in their tastes, and if reviews reflect consumers' overall enjoyment of a product (i.e. both inherent quality of the good and taste), a higher price – all else equal – leads to a higher cutoff taste of purchasing consumers and hence a higher review. The reverse relationship in turn can be motivated by a *direct price effect* where the same good purchased by the same consumer at a lower price leads to a higher review. For different goods, a different effect may dominate. Both USB sticks and casual games are products where differentiated tastes by consumers play less of a role. The price effect hence dominates and high reviews are related to low prices. This is reversed for books and non-casual games. The selection effect dominates and high reviews are related to high prices.

Our theoretical setting incorporates both the selection effect and the direct price effect. We set up our model in Section 2. A long-lived monopolist sells a good of fixed quality to short-lived consumers. Consumers value quality but do not directly observe it prior to consumption. Moreover, they exhibit horizontal differentiation and differ in their taste for the product. A consumer contemplating purchase of the good forms her belief about the product's quality using the current price together with the current rating which is a summary statistic of reviews provided by consumers who purchased the good in prior periods. These reviews, crucially, may depend not only on the inherent quality of the good, but also reflect the taste for the good, and how satisfied consumers were relative to the purchase price. As these characteristics matter and are not observable to future consumers, quality inference is confounded and the seller may strategically set prices to affect future inference and thus profits.[6]

In our model, consumers conduct inference by solving a two-dimensional fixed point problem: They look for the combination of product quality and the set of purchasing consumers which at the current price would lead to the current aggregate rating. While this inference implicitly renders consumers boundedly rational, we deem it sensible for multiple reasons. First, it relies on the observables – aggregate rating and current price – typically available to consumers on online platforms. Second, our method of inference is equivalent to consumers acting as if they were in a stationary equilibrium, for which we show that inference is correct.[7] We show that this

[4] Casual games are "typically distinguished by simple rules and by reduced demands on time and learned skill", see https://en.wikipedia.org/wiki/Casual_game. For the empirical analysis, see Appendix C. For additional evidence for the ambiguous relationship between prices and ratings, see Juan et al. (2018).

[5] The term is similarly used in concurrent works by Zegners (2019) and Acemoglu et al. (2019).

[6] This is in line with findings by De Langhe et al. (2015) who show a significant difference between aggregate ratings on Amazon.com and objective quality measures provided by *Consumer Reports* quality scores, even for products where vertical differentiation is likely to be more important than subjective tastes.

[7] Note that for full rationality consumers would need to know their position in time, i.e., how many consumers purchased prior to them, and additionally form beliefs about the entire price and review paths, which implicitly requires solving the firm's optimization problem for any given possible state.

inference is uniquely determined, and derive how the current price affects the inferred quality, cutoff taste, and induced reviews. Notably, whether higher reviews are achieved via high or low prices depends on the relative strength of the *direct price effect*: higher prices increase the cutoff consumer and thereby the average gross utility of reviewing consumers but the direct price effect might overcompensate this *selection effect* if sufficiently strong.

This enables us to assess the pricing incentives of a strategic firm which trades off myopic and future profit considerations in Section 3. We consider a class of rating systems with varying sensitivity to new reviews. We show that irrespective of the rating system, we converge to a stationary equilibrium in which the firm charges the same price, the same set of consumers purchase, and the same average review equal to the aggregate rating is left in every period. In this stationary equilibrium, consumers correctly infer the product quality. However, the *long run rating* and *long run price level* depend on the type of rating system in place, which affects firm profits and consumer surplus.

The firm balances flow payoff and future profit considerations. This amounts to balancing exploitation of the current rating and strategic reputation management via the induced reviews. Crucially, the less sensitive the rating system, the more the firm has to "invest" by deviating from the myopically optimal price to obtain a given future rating. For the seller, a rating system is hence unambiguously better the more sensitive it is to new reviews. For consumers, however, this depends on whether the firm has an incentive to over- or underprice relative to the myopically optimal price, which in turn depends on whether the *direct price effect* or *selection effect* dominates. If reviews do not depend much on the price, i.e. if the direct price effect is relatively weak, the selection effect dominates and a higher price induces a higher rating. Hence, the firm has an incentive to price high, which is exacerbated the more sensitive the aggregate rating is to new reviews. An insensitive rating system is thus beneficial for consumers as it decreases the long-run price level and increases long-run consumer surplus. The converse is true whenever reviews are very sensitive to price, i.e. if the direct price effect is larger than the selection effect. The dynamic considerations provide an incentive for the firm to lower the price, and a rating system emphasizing recent reviews strengthens this effect which benefits consumers.

We discuss several extensions and modifications of our setup in Section 3.3. In particular, we show that our findings are robust to explicitly accounting for the number of reviews when updating ratings, for not all consumers leaving a review, or for more general distributions of taste heterogeneity. They are similarly unaffected when ratings are stochastic. We can also embed our considerations into a competitive setup, where firms additionally have an incentive to price such that the own rating is improved relative to the rival's. This typically increases the incentive to charge high prices. Finally, we show that convergence within our model always occurs linearly, with the rate decreasing in the sensitivity of the rating system. As such, consumers "learn" the product's true quality faster whenever the rating system is more responsive to new reviews.

Our analysis highlights a potential drawback of review systems placing more weight on more recent reviews when product quality is fixed and the selection effect is more important than the direct price effect. For these settings, the negative impact of responsiveness on long-run consumer surplus is novel, and complementary to the literature on beneficial effects of these types of ratings with decay (see e.g. Kovbasyuk and Spagnolo (2018) where decay is beneficial as product quality may change over time). Crucially, these drawbacks are *not* driven by wrong quality inference, but depend on the relative importance of the direct price and selection effect.

Overall, our analysis implies that recent changes to rating systems, which put increased weight on more recent reviews in the displayed aggregate ratings, may have adverse effects for consumers and thus warrant further investigation. We explicitly consider the incentives of a platform which chooses its rating system to maximize a weighted average of long-run consumer surplus and long-run firm profits in 4. We show that the platform typically favors consumers even when there is a wedge between the firm's interests (who always prefer a rating system more responsive to new reviews) and consumers (who only do so when the direct price effect is large). When the direct price effect is large, consumers and firms both benefit from a responsive rating system and the platform accommodates this. When the direct price effect is small, consumers instead prefer a persistent system. Notably, the platform indeed chooses a persistent system even if it cares about the firm twice as much as it does about consumers, it always prefers to set the rating system which maximizes consumer welfare. Only if the weight that the platform puts on consumer surplus is sufficiently low is this decision reversed and an interior sensitivity, or even the firm-optimal most sensitive system possible is chosen.

1.1. Related Literature

Our work primarily relates to articles assessing strategic firm behavior in the presence of review or reputation concerns, see e.g. Bar-Isaac and Tadelis (2008) for an overview. Board and Meyer-ter Vehn (2013) show that a seller exhibits a reputation build-up and exploitation behavior in a setting with costly investment in quality. Crucially, their setup abstracts from strategic price-setting; firms extract the current willingness-to-pay determined by the quality inference. The price-path is thus implicitly predicted to follow the quality belief and, due to the build-up and exploitation strategy, exhibits a hump-shape. By contrast, we consider a fixed quality and focus purely on the strategic impact of pricing on reputation as measured by the quality inference conducted via the rating system. Marinovic et al. (2018) extend the analysis by Board and Meyer-ter Vehn (2013) by considering a framework in which quality can be certified (and certification is costly). Cabral and Hortacsu (2010), building on an earlier working paper version, Cabral and Hortacsu (2004), consider effort decisions by sellers on eBay. They show a similar mechanism in terms of managing reputation. For certain parametrizations, a low-type seller exerts effort and thus builds up reputation by mimicking the high type only until her type is discovered. The model's prediction about seller exit being preceded by a period of high likelihood of negative reviews, as well as the general negative impact of a negative review on sales, prices, and likelihood to remain in the market, are empirically verified.

Our analysis also draws from and complements the empirical literature on the impact of review and rating systems on online platforms (see e.g. Li and Hitt (2008), Ho et al. (2017)). Ratings have been shown to substantially impact demand in a variety of contexts such as restaurants (on Yelp.com, Anderson and Magruder (2012)), books (in a comparative study of Amazon.com and Barnesandnoble.com, Chevalier and Mayzlin (2006)), and online service marketplaces (Moreno and Terwiesch (2014)). However, ratings have been shown to not simply reflect the inherent quality of a good: De Langhe et al. (2015) show a significant difference between aggregate ratings on Amazon.com and objective quality measures provided by *Consumer Reports* quality scores, even for products where vertical differentiation is likely to be more important than subjective tastes. More generally, Bhargava and Feng (2015) estimate the relative impact of the external information environment incorporating potential reviews, and signaling via price distortions. They show that an effective rating system and price signaling are substitutes in terms of information provision: More external information such as ratings are associated with decreased price distortion.

While we focus on signaling via induced ratings and the indirect effect prices have on future quality inference, there is a vast literature on direct price signaling (see Wolinsky (1983), Bagwell and Riordan (1991)). Osborne and Shapiro (2014) embed price-signaling considerations in a dynamic context where a monopolist chooses both quality and price – consumers thus dynamically learn about the relation between quality and price and the firm strategically affects this inference. A similar consideration is present in our model. However, in our setting the exogenous review and rating system in place forms the basis for the firms' strategic actions impacting future quality inferences from price and ratings.

A central role in our analysis is played by the *selection effect* whereby prices serves as a selection device amongst consumers with heterogenous tastes, thereby influencing reviews and future ratings and profits. This effect is also present in concurrent work by Zegners (2019) and Acemoglu et al. (2019). Zegners (2019) considers the choice by sellers of self-published e-books to give book away for free, or for a positive price and shows that free e-books suffer from lower ratings (while earning higher "sales" and thus more reviews). Acemoglu et al. (2019) study a social learning model in which reviews instead of actions are observed. Horizontally differentiated and fully rational consumers take the selection effect into account and conditions under which the quality is learned are provided. While strategic pricing is addressed in an extension, it considers a seller who chooses prices without knowledge of their own quality and thus also only learns from reviews. This is in contrast to our paper, where the seller holds an informational advantage over a sequence of boundedly rational buyers.[8]

We also relate to recent work on the design of rating systems, which – different to us – feature stochastically changing states of the world. Kovbasyuk and Spagnolo (2018) show that low memory of ratings can be optimal if the quality of a product or service changes over time preventing inefficient exit of the market. Bonatti and Cisternas (2018) study the effect of aggregate scores about consumers' purchasing histories which is informative to short-lived firms about the consumer's evolving willingness-to-pay. Consumers reduce their purchases to signal a lower willingness-to-pay to future firms inducing lower prices.

[8] A similar problem is tackled in e.g. von Thadden (1992).

2. Model

In this section, we present our model. We discuss several extensions and how they affect the results in Section 3.3.

We consider a monopolistic long-lived producer of a good with privately known and 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 periods.

Time Time is discrete, $t \in \{0, 1, 2, 3, \ldots, T\}$, $T \leq \infty$. The economy consists of a single long-lived seller and a unit mass of short-lived consumers in any given period.

Seller The seller ('She') wishes to sell a good of exogenously given quality θ, where $\theta \sim F(\cdot)$ on $\Theta \equiv [\underline{\theta}, \bar{\theta}]$ is distributed according to a cdf F. The realization of θ is private information to 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 unit mass of risk-neutral consumers. Consumers are short-lived and only present for one period. Consumers value quality, and are horizontally differentiated with respect to their personal liking 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 given by $u(\theta, \omega_i) = \alpha\theta + \beta\omega_i$, where $\alpha > 0, \beta > 0$, so that consumers value quality and taste. A consumer's net utility u, derived as the gross consumption utility $u(\cdot)$ net of price, is characterized by

$$u = \alpha\theta + \beta\omega_i - p. \tag{1}$$

When all consumers hold the same beliefs, which we verify with our inference below, there will be 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 While consumers are short-lived, we allow for information transmission across periods via a review and rating system. The review ψ left by a consumer i conditional on purchase is given by

$$\psi_i = \alpha\theta + \beta\omega_i - \kappa p. \tag{2}$$

with $\kappa \in [0, 1)$. Hence, we assume that the consumer reports her gross utility of consuming the product minus a component that depends on the price. The underlying idea is that 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.

For tractability, we assume that every consumer who purchases the good leaves a review, and that only the average review in a given period is used to update the aggregate rating. This average review is equal to the one left by the consumer with average taste $\omega^e(\omega_t^*) = E[\omega|\omega > \omega_t^*] = \frac{1+\omega_t^*}{2}$, where ω_t^* is the taste of the marginal consumer who purchases. The rating system is hence characterized by the mapping from current aggregate rating $\bar{\psi}_t$ and current average review $\psi_t = \alpha\theta + \beta\frac{1+\tilde{\omega}_t}{2} - \kappa p_t$ into next period's aggregate rating $\bar{\psi}_{t+1}$. We denote this mapping by ρ_t, allowing it in principle to depend on the time period t: $\bar{\psi}_{t+1} = \rho^t\left(\bar{\psi}_t, \psi_t\right)$.

In this paper, we consider a specific class of rating systems which are time-independent, i.e. characterized by a time-invariant mapping ρ. In particular, we consider the following updating rule

$$\rho(\bar{\psi}_t, \psi_t) = \bar{\psi}_{t+1} = \frac{m-1}{m}\bar{\psi}_t + \frac{1}{m}\psi_t. \tag{3}$$

Given an initial rating $\bar{\psi}_0$, this can be rewritten as

$$\bar{\psi}_t = \left(\frac{m-1}{m}\right)^t \bar{\psi}_0 + \sum_{\tau=0}^{t-1}\left(\frac{m-1}{m}\right)^{t-1-\tau}\frac{\psi_\tau}{m}. \tag{4}$$

In (3) and (4), m parametrizes the sensitivity of the rating system to incoming reviews. The higher m, the less responsive is the updated rating to the current period reviews, and conversely the higher is the weight on older reviews. The opposite occurs for low m; the case $m = 1$ corresponds to a limited memory rating in which the rating only consists of last period's average review. In terms of interpretation, rating systems with varying degrees of sensitivity as given by m allows us to assess the recent pushes by online platform such as Amazon.com to have more recent reviews matter more for the displayed aggregate rating – in our context, this would be captured by a reduced m.

Technical Assumptions For technical purposes, we impose the following assumptions. First, we require $\underline{\theta} < -\frac{\beta}{\alpha}$. 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 m such that $m > \frac{1}{1-\kappa}$ whenever $\kappa > \frac{1}{2}$. This ensures concavity of the flow payoff in the state variable in the dynamic programming approach.

Timing of the stage game The timing of a given period is as follows: The firm observes the current state of the market characterized by the aggregate rating $\bar{\psi}_t$ and sets the price p_t at which it 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).

Consumer Inference A central requirement is to specify how consumers conduct quality inference given their observables. Recall that consumers only observe the current price and rating.

Hence, we need to specify how consumers rationalize the current combination given that they do not observe the past path of prices and ratings. That past prices and ratings are not observable is true for many online sales platforms such as Amazon and Steam. While, in principle, we could assume that consumers are fully rational and solve the firms' problem from time $t = 0$ onwards, we assume that consumers apply a heuristic.

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*

$$\psi\left(\mu^*, \omega^e(\omega^*), p_t\right) = \bar{\psi}_t \tag{CONS}$$

$$u\left(\mu^*, \omega^*, p\right) = p_t. \tag{RAT}$$

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 purchasing past consumers and in particular their taste for the product: 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 purchase occurred at price 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 weakly 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 $\bar{\psi}_t$.

While the assumption is non-standard, we do not consider it to be too far removed 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

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

[10] In fact, it was inspired by introspection of a subset of the authors while making a purchase decision online.

due to horizontal differentiation, the assumption that consumers base their quality inference only on the aggregate rating and 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. We verify subsequently that the game indeed always converges towards a stationary equilibrium in which quality inference is correct. Finally, consumers are often unclear about how many periods have passed and how often the firm changed prices in the past.[11]

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

$$\mu(\overline{\psi}, p) = \frac{2\overline{\psi} - \beta - p(1-2\kappa)}{\alpha} \tag{5}$$

$$\tilde{\omega}(\overline{\psi}, p) = 1 - \frac{2}{\beta}\left(\overline{\psi} - p(1-\kappa)\right) \tag{6}$$

This induces a quantity sold at price p given aggregate rating $\overline{\psi}$ of

$$q(\overline{\psi}, p) = \frac{2}{\beta}\left(\overline{\psi} - p(1-\kappa)\right). \tag{7}$$

The comparative statics of the inference and quantity are natural. The belief about quality increases in the observed rating; hence, better reviews induce higher beliefs of consumers. Observing a low price reduces the belief about the quality when κ is sufficiently low ($\kappa < 1/2$) as a low price in this case induces consumers to believe that the cutoff consumer in the past was low. Therefore, a particular rating is interpreted more positively at low prices. If, however, the effect of the price on the rating is high, i.e. $\kappa > 1/2$, higher prices induce higher beliefs of the consumer because the direct price effect of the rating is stronger than the selection effect. The quantity always decreases in the price given a rating and increases in the rating for any price – this exemplifies the incentive for the firm to positively influence future ratings as it allows to increase its future profits.

3. Dynamic Pricing and Long-Run Properties

To illustrate dynamic pricing incentives, consider a given period t. The firm aims to maximize the sum of flow and discounted future profits. Future profits in turn depend on two state variables: the period, t, and the aggregate rating, which the firm can influence strategically. Denote the value of future profits as $V_{t+1}(\bar{\psi}_{t+1})$ and we obtain for the firm's maximization problem:

[11] If the game has a stationary equilibrium and consumers are uncertain about their position in the timing, their naïve best guess is to be in the stationary period and our imposed inference would be correct.

$$V_t(\bar{\psi}_t) = \max_{p_t} p_t \cdot q_t(p_t) + \delta V_{t+1}(\bar{\psi}_{t+1}), \tag{8}$$

$$\text{s. t. } \bar{\psi}_{t+1} = \rho_t(\bar{\psi}_t, \psi_t) \tag{9}$$

$$\psi_t = \alpha\theta + \beta\omega_t^e - \kappa p_t. \tag{10}$$

The derivative with respect to price is given by

$$\underbrace{q(p_t, \overline{\psi}_t) + p_t q'(p_t, \overline{\psi}_t)}_{\text{static monopoly}} + \delta \underbrace{V'_{t+1}(\bar{\psi}_{t+1})}_{\text{effect of rating on future profits}} \cdot \underbrace{\frac{\partial \rho_t}{\partial \psi_t} \cdot \frac{\partial \psi_t}{\partial p_t}}_{\text{effect of current price on future rating}}. \tag{11}$$

We thus have two effects: First, flow profits are affected by the increase in the price. This is captured by the standard 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 ($V'_{t+1}(\bar{\psi}_{t+1})$), the sensitivity of the aggregate rating in the next period to the induced current review ($\frac{\partial \rho_t}{\partial \psi_t}$), and the effect the price change has on the current review ($\frac{\partial \psi_t}{\partial p}$).

3.1. Dynamic Pricing - Finite Horizon

If time is finite, $T < \infty$, the incentive to price strategically to affect future profits depends on the time. In the last period, the firms will always choose the myopic monopoly price given the rating. In all other periods, however, the firm has an incentive to price away from this myopically optimal monopoly price – how strong this incentive is depends on the sensitivity of future ratings vis-a-vis the induced current period reviews, as well as on the remaining periods the firm has to exploit high ratings. This incentive can take two forms: It may either induce the firm to increase the price (if $\frac{\partial \psi_t}{\partial p} > 0$) or to decrease the price ($\frac{\partial \psi_t}{\partial p} < 0$), which in turn depends on the sensitivity of the review function to the price, κ. To illustrate this, we consider the case where $T = 2$.

Illustration in a two-period model. Consider a two-period version of the model ($T = 2$) with initial rating $\overline{\psi}_1$. Denote the rating in period 2 by $\psi_2 = \frac{m-1}{m}\overline{\psi}_1 + \frac{1}{m}\psi_1$, where m characterizes the responsiveness of the rating. In period 2, the firm maximizes

$$\max_{p_2} p_2 \cdot q_2 = p_2 \cdot \frac{2}{\beta}\left(\overline{\psi_2} - p_2(1-\kappa)\right) \tag{12}$$

and therefore chooses the myopic monopoly price given by

$$p_2 = \frac{\overline{\psi}_2}{2(1-\kappa)} \tag{13}$$

which yields as profits

$$\pi_2 = \frac{\overline{\psi}_2^2}{2(1-\kappa)\beta}. \tag{14}$$

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

$$\pi_1 = p_1 q_1(\overline{\psi}_1, p_1) + \delta\pi_2(\overline{\psi}_2(p_1)) \tag{15}$$

with first-order condition

$$\frac{2}{\beta}\left(\overline{\psi}_1 - 2(1-\kappa)p_1\right) + \delta\frac{\partial\pi_2(\overline{\psi}_2(p_1))}{\partial p_1}. \tag{16}$$

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

$$\frac{\partial\pi_2(\overline{\psi}_2(p_1))}{\partial p_1} = \frac{\partial\pi_2}{\partial\overline{\psi}_2}\frac{\partial\overline{\psi}_2}{\partial\psi_1}\frac{\partial\psi_1}{\partial p_1} \tag{17}$$

$$= \frac{\overline{\psi}_2(p_1)}{(1-\kappa)\beta}\frac{1}{m}(1-2\kappa). \tag{18}$$

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

$$p_1 = \underbrace{\frac{\overline{\psi}_1}{2(1-\kappa)}}_{\text{static monopoly price}} + \underbrace{\frac{\overline{\psi}_2(p_1)}{2(1-\kappa)}\frac{\delta}{2m}(1-2\kappa)}_{\text{dynamic rating effect}}. \tag{19}$$

It is easy to see that the firm will change its price away from the static monopoly price if it takes the effect of current prices on future ratings into account. In particular, as profits are increasing in the rating, the firm will deviate from the static monopoly price in a way that increases future ratings. If $\kappa > \frac{1}{2}$, the price effect is sufficiently strong and dominates the selection effect in the review. In this case, the firm will choose a lower price in the first period as this leads to a higher rating entering the second period. By contrast, for $\kappa < \frac{1}{2}$ the selection effect dominates the price effect. The firm will choose a higher price in the first period than the static monopoly price. Clearly, the effects are stronger if the firm discounts less (high δ) and the rating is more responsive to current reviews (low m). An important takeaway from this illustrative analysis is that the direction of the price distortion relative to the myopically optimal price depends on the relative importance of the price and the selection effect which is determined by the relevance of the price for the review κ.

3.2. Dynamic Pricing - Infinite Horizon

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

The firms solves the problem

$$\max_{(p_t)_{t\geq 0}} \sum_{t=0}^{\infty} \delta^t p_t q(p_t, \overline{\psi}_t) \tag{20}$$

$$\text{s. t. } \overline{\psi}_t = \frac{m-1}{m}\overline{\psi}_t + \frac{1}{m}\psi(p_t, \overline{\psi}_t) \tag{21}$$

$$\overline{\psi}_0 = \overline{\psi}. \tag{22}$$

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

$$V(\overline{\psi}) = \max_p \left\{ pq(p, \overline{\psi}) + \delta V(\overline{\psi}') \right\} \tag{23}$$

$$\text{s. t. } \overline{\psi}' = \frac{m-1}{m}\overline{\psi} + \frac{1}{m}\psi(p, \overline{\psi}). \tag{24}$$

The review ψ in any period is linear in p given current aggregate rating $\bar{\psi}$ and given by

$$\psi = \alpha\theta + \beta - \overline{\psi} + p(1-2\kappa) \tag{25}$$

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

$$\overline{\psi}' = \frac{m-1}{m}\overline{\psi} + \frac{1}{m}\left(\alpha\theta + \beta - \overline{\psi} + p(1-2\kappa)\right) \tag{26}$$

$$\Leftrightarrow p = \frac{\alpha\theta + \beta + (m-2)\overline{\psi} - m\overline{\psi}'}{1-2\kappa}. \tag{27}$$

The value function $V(\cdot)$ is independent of t because both the rating system and consumer inference do not depend on it. The problem of the firm can be written as

$$V(\overline{\psi}) = \max_{\overline{\psi}'} \left\{ p(\overline{\psi}')q(\overline{\psi}', \overline{\psi}) + \delta V(\overline{\psi}') \right\}. \tag{28}$$

We solve the problem by guessing that the value function is of the form $V(\overline{\psi}) = c + d\overline{\psi} + e\overline{\psi}^2$, 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 firms of type $\theta > -\frac{\beta}{\alpha}$, there is a unique stationary equilibrium that is characterized by*

$$\tilde{p} = \frac{(\alpha\theta + \beta)(m(1-\delta) + 2\delta)}{4\delta + m(1-\delta)(3-2\kappa)} \tag{29}$$

$$\Psi = \frac{(\alpha\theta + \beta)(2m(1-\delta)(1-\kappa) + \delta(3-2\kappa))}{4\delta + m(1-\delta)(3-2\kappa)}. \tag{30}$$

The rating system is effective and consumers learn the quality $\tilde{\mu} = \theta$.

Proof. For the derivation, see Appendix A. ■

Proposition 1 shows that there is a 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 m.

Moreover, firms with qualities that are so low that they should not sell under full information ($\theta < -\frac{\beta}{\alpha}$; 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 firms 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 m. 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 m on the long run profits

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

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

(a) Prices are decreasing in m whenever the direct price effect in the reviews is not too large, $\frac{d\tilde{p}}{dm} < 0$ when $\kappa < \frac{1}{2}$ and $\frac{d\tilde{p}}{dm} > 0$, otherwise.

(b) Consumer surplus is increasing in m when the direct price effect in the rating is not too large, otherwise it is decreasing, $\frac{d\tilde{CS}}{dm} > 0$ when $\kappa < \frac{1}{2}$ and $\frac{d\tilde{CS}}{dm} < 0$, otherwise.

(c) Long-run profits are decreasing in m, $\frac{d\tilde{\pi}}{dm} < 0$.

Proof. Differentiating (29) gives

$$\frac{\partial \tilde{p}}{\partial m} = \overbrace{\frac{2(1-\delta)\delta(\alpha\theta + \beta)}{(4\delta + (1-\delta)(3-2\kappa)m)^2}}^{>0} \cdot (2\kappa - 1), \tag{32}$$

so that the sign depends on the sign of $2\kappa - 1$. This gives (a) and via the relation to CS (b). For (c), we differentiate (31) and get

$$\frac{\partial \tilde{\pi}}{\partial m} = \overbrace{-\frac{2(1-\delta)^2\delta m(\alpha\theta + \beta)^2}{\beta(4\delta + (1-\delta)(3-2\kappa)m)^3}}^{<0} \cdot (2\kappa - 1)^2, \tag{33}$$

which is unambiguously weakly (strictly for $\kappa \neq \frac{1}{2}$) negative. ■

Corollary 1 contains the results which characterize the main implications of the paper. In the stationary equilibrium, the firm 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 larger m, the more the firm would have to "invest" by deviating from the myopically optimal price to obtain a given next-period rating. Therefore, a lower sensitivity (higher m) has an unambiguously negative effect on the firm's profits.

For consumers, the price level determines the long run consumer surplus as inference is correct. Whether consumers benefit from a higher or lower sensitivity therefore depends on whether the firm 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. This depends on the reactiveness of the review function to the price, κ. For low κ, the price has only a small direct effect on the induced reviews. The selection effect thus is more relevant than the direct price effect and the firm has an incentive to price higher than the myopic optimum to induce higher future ratings (see also the two-period case in (19)). This incentive is mitigated by a higher m as an individual period has a smaller effect on the rating. The price level hence is decreasing in m and consumers benefit from having the rating reflect past purchases (close to) equally instead of putting more weight on more recent reviews.

The converse is true when κ is large. When the direct price effect dominates, future profit considerations incentivize the firm to price *below* the myopically optimal price and a low m benefits consumers as it increases the relevance of these future considerations in the firm's optimization problem. Whenever reviews respond strongly to the purchase price, an emphasis on more recent reviews as e.g. implemented by 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 $\bar{\psi}_t = a + b\bar{\psi}_{t-1}$ where 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 $m \geq 2$, $b > 0$ and $\frac{\partial b}{\partial m} \geq 0$.*

Proof. See Appendix A. ■

The main takeaway from this is the comparative statics with respect to m. As convergence is quicker the lower the rate of convergence, a lower sensitivity (higher m) actually lowers 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.

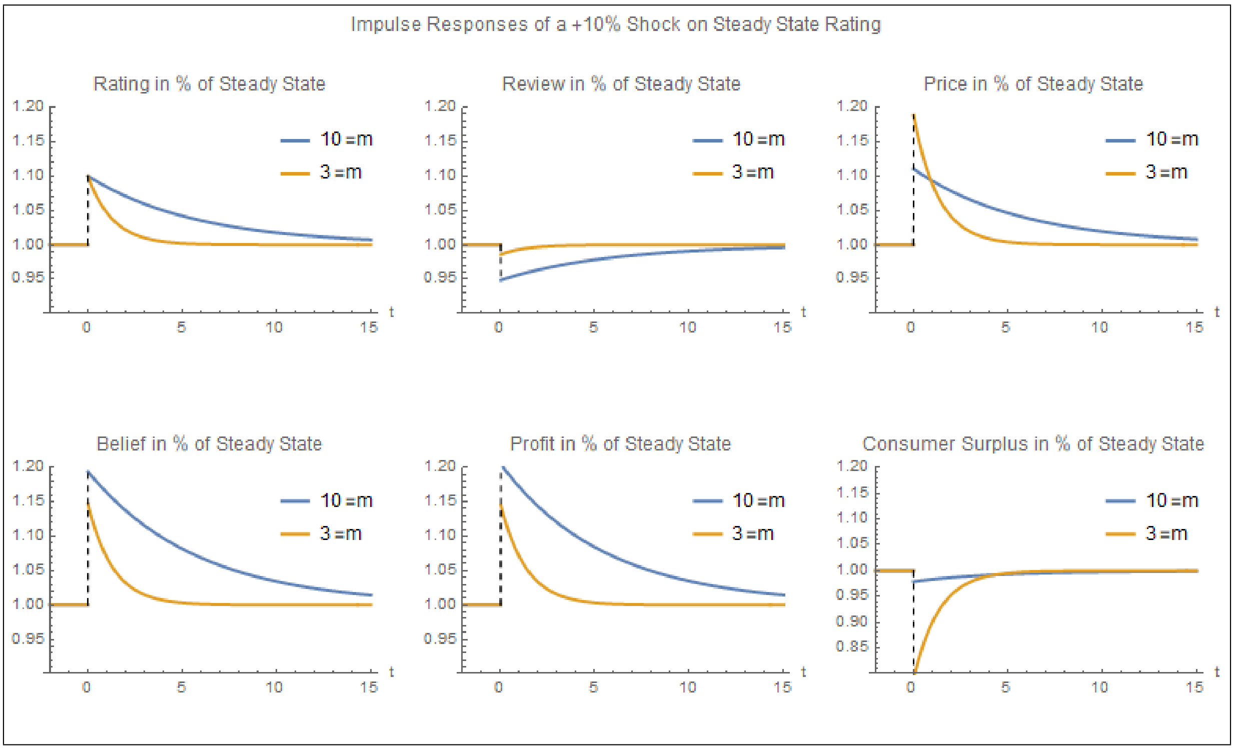

Figure 1: Impulse Response Functions given a 10% shock to the steady state rating

3.3. Extensions

In this subsection, we discuss some 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 consumer 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

$$\overline{\psi}_{t+1} = \frac{m - q_t}{m}\overline{\psi}_t + \frac{q_t}{m}\psi_t, \tag{34}$$

where q_t is the number of consumers in period t. In this sense, the rating reacts more strongly when more consumers purchase. 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 firm's first-order condition for prices

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

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

$$= \underbrace{q_t(p_t) + \frac{dq_t(p_t)}{dp_t}p_t}_{\text{static monopoly pricing}} + \delta \underbrace{\frac{dV_{t+1}}{d\overline{\psi}_{t+1}}}_{\text{effect of reviews on CV}} \left(\underbrace{\frac{q_t(p_t)}{m}\frac{d\psi_t(p_t)}{dp_t}}_{\text{higher reviews}\rightarrow\text{higher ratings}} + \underbrace{\frac{dq_t(p_t)}{dp_t}(\overline{\psi_t} - \psi_t(p_t))}_{\text{number of reviews effect}} \right).$$

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 firm 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. Similar to the number of reviews not mattering, we made this assumption for expositional and tractability purposes. In practice, 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, and that, if the number of purchasing consumers decreases ($\tilde{\omega}$ increases), the average reviewing consumer, $w^e(\tilde{\omega}) \equiv \int_{\tilde{\omega}} w \frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(w,\tilde{\omega})} d\omega$, increases, but not too much, $\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} \in (0,1)$. Under these assumptions the consumers' inference and demand is given by the solution to the equation system

$$\alpha\mu + \beta\tilde{\omega} = p \tag{35}$$

$$\alpha\mu + \beta\int_{\tilde{\omega}}^1 \omega \frac{f_\psi(\omega;\tilde{\omega})}{\int_{\tilde{\omega}}^1 f_\psi(\omega,\tilde{\omega})} d\omega - \kappa p = \overline{\psi}. \tag{36}$$

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

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = -\begin{pmatrix} \frac{\partial u}{\partial \theta} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} & \frac{\partial \psi}{\partial \tilde{\omega}} \end{pmatrix}^{-1} \begin{pmatrix} -1 \\ \frac{\partial \psi}{\partial p} \end{pmatrix}. \tag{37}$$

[13] The 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.

Inverting the matrix and plugging in the partial derivatives yields

$$\begin{pmatrix} \frac{d\mu}{dp} \\ \frac{d\tilde{\omega}}{dp} \end{pmatrix} = \begin{pmatrix} \frac{1}{\alpha}\frac{\kappa - \frac{d}{d\tilde{\omega}}\int_{\tilde{\omega}}^{1}\omega\frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^{1}f_\psi(\omega,\tilde{\omega})d\omega}d\omega}{1-\frac{d}{d\tilde{\omega}}\int_{\tilde{\omega}}^{1}\omega\frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^{1}f_\psi(\omega,\tilde{\omega})d\omega}d\omega} \\ \frac{1}{\beta}\frac{1-\kappa}{1-\frac{d}{d\tilde{\omega}}\int_{\tilde{\omega}}^{1}\omega\frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^{1}f_\psi(\omega,\tilde{\omega})d\omega}d\omega} \end{pmatrix}, \tag{38}$$

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 firm'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 firms' first-order condition

$$\frac{dV(\overline{\psi}_t)}{dp_t} = q_t(p_t) + \frac{dq_t(p_t)}{dp_t}p_t + \delta\frac{dV_{t+1}}{dp_t} \tag{39}$$

$$= 1-\tilde{\omega}(p_t) - \frac{1}{\beta}\frac{1-\kappa}{1-\frac{d}{d\tilde{\omega}}\int_{\tilde{\omega}}^{1}\omega\frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^{1}f_\psi(\omega,\tilde{\omega})d\omega}d\omega}p_t \tag{40}$$

$$+ \delta\frac{dV_{t+1}}{d\overline{\psi}_{t+1}}\frac{1}{m}\left(\frac{d}{d\tilde{\omega}}\int_{\tilde{\omega}}^{1}\omega\frac{f_\psi(\omega,\tilde{\omega})}{\int_{\tilde{\omega}}^{1}f_\psi(\omega,\tilde{\omega})d\omega}d\omega - \kappa\right) \tag{41}$$

Note that the main term is again the change in the average *reviewing* consumer. The relevant consideration for the firm 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 higher reviews and the firm 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 $[\underline{\omega}, \overline{\omega}]$ 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 $\omega^e(\tilde{\omega}) \equiv \int_{\tilde{\omega}}^{\overline{\omega}}\omega\frac{g(\omega)}{\int_{\tilde{\omega}}^{\overline{\omega}}g(\omega)d\omega}d\omega$. Then, consumer inference is given by the solution to the equation system

$$\alpha\mu + \beta\tilde{\omega} = p \tag{42}$$

$$\alpha\mu + \beta\int_{\tilde{\omega}}^{1}\omega\frac{g(\omega;\tilde{\omega})}{\int_{\tilde{\omega}}^{1}g(\omega,\tilde{\omega})}d\omega - \kappa p = \overline{\psi}. \tag{43}$$

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

$$\frac{dV(\overline{\psi}_t)}{dp_t} = 1-\tilde{\omega}(p_t) - \frac{1}{\beta}\frac{1-\kappa}{1-\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}}}p_t + \delta\frac{dV_{t+1}}{d\overline{\psi}_{t+1}}\frac{1}{m}\left(\frac{d\omega^e(\tilde{\omega})}{d\tilde{\omega}} - \kappa\right) \tag{44}$$

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 necessarily correlated with the observables. In particular, it is not necessarily the case that reviews in a given period represent the average consumer in that period. 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

$$\psi_t = \alpha\theta + \beta\tilde{\omega}_t^e - \kappa p_t + \varepsilon_t \tag{45}$$

where $\varepsilon \in \{-\epsilon, \epsilon\}$ 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 can be obtained from the authors upon request.

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

$$\psi_t = \alpha\theta + \beta\tilde{\omega}_t^e - \kappa(\overline{p} - p_t)^2. \tag{46}$$

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

$$\frac{d\psi(p, \overline{\psi})}{d(\overline{p} - p)} = 4\kappa(\overline{p} - p) - 1 \tag{47}$$

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 higher 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 (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 non-linear utility and review functions.

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. 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^t such that all consumers up to x^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

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

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

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

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 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 added effect that the own price affects the selection of purchasing consumers of *the rival firm* which exerts influence on that firm's review 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 from the rival firm and closer to oneself as it requires a higher taste for the own product to still 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

$$V_i = \max_{p_i^0} p_i^0 \cdot q_i^0 + \pi_i^1(\bar{\psi}_i^1, \bar{\psi}_{-i}^1). \tag{48}$$

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

$$\frac{dV_i}{dp_i^0} = \overbrace{\frac{dq_i^0}{dp_i^0} + p_i^0 \frac{dq_i^0}{dp_i^0}}^{\text{flow profit effect}} + \overbrace{\underbrace{\frac{d\pi_i^1}{d\bar{\psi}_i^1}}_{>0} \cdot \underbrace{\frac{d\bar{\psi}_i^1}{\psi_i^0}}_{>0} \underbrace{\frac{d\psi_i^0}{dp_i^0}}_{\gtrless 0}}^{\text{dynamic effect via } \textit{own} \text{ rating}} + \overbrace{\underbrace{\frac{d\pi_i^1}{d\bar{\psi}_{-i}^1}}_{<0} \cdot \underbrace{\frac{d\bar{\psi}_{-i}^1}{\psi_{-i}^0}}_{>0} \underbrace{\frac{d\psi_{-i}^0}{dp_i^0}}_{<0}}^{\text{dynamic effect via } \textit{rival's} \text{ rating}} \tag{49}$$

[14] We implicitly assume full market coverage which facilitates the exposition of the effects. We emphasize that we do not want to use this model to draw qualitative predictions but only illustrate that our main channels are also present in a competitive setting. This is important as the full market coverage assumption can in principle lead to equilibria where the cutoff consumer would obtain a negative utility given the inference.

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 own rating, but in addition an incentive to *increase* prices due to the dynamic effect via the induced reviews and rating for the opponent.

4. Platform Incentives

Given the closed-form expressions for the long-run consumer surplus and seller profits, we are also 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 m, to maximize

$$\pi^P = w_c \cdot \tilde{CS} + (1 - w_c) \cdot \tilde{\pi}, \tag{50}$$

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.

It is straightforward from Corollary 1 that the lowest possible m maximizes π^P whenever the direct price effect is large, i.e. $\kappa > \frac{1}{2}$. This is because in this case both consumers and the seller benefit from a low m – the firm always prefers the rating system most responsive to recent reviews, while consumers in this case want the firm to place a high weight of the current price on future ratings as this leads to a downward pressure on prices so as to manage the rating. However, when the direct price effect is small, the firm's and consumers' interests diverge. The firm prefers a responsive rating system, i.e. a low m, while consumers are better off whenever m is large.

Proposition 3 (Platform Incentives) *A platform maximizing π^P as given in (50) always chooses the longest possible memory unless*

(i) the direct price effect in the review is weak, $\kappa < \frac{1}{2}$, and

(ii) the weight the platform attaches to long run consumer surplus is sufficiently low, $w_c < \frac{1-2\kappa}{3-4\kappa}$.

If (i) and (ii) are satisfied, the platform sets $m^* = \max\{1, \tilde{m}\}$ where

$$\tilde{m} = \frac{2w_c\delta}{(1-\delta)\,[1 - 2\kappa - (3 - 4\kappa)w_c]}. \tag{51}$$

Proof. See Appendix A ■

Proposition 3 is intuitive in that whenever the incentives vis-a-vis maximizing firm profits and consumer surplus are misaligned, the platform either balances the two aspects by choosing an interior m, 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, $\kappa < \frac{1}{2}$.

Even in this case, however, the platform needs to care significantly more about the firm than about consumers to deviate from choosing the memory which is best for consumers, i.e. from setting the longest possible memory. Specifically, as $\frac{1-2\kappa}{3-4\kappa} < \frac{1}{3}$ for all $\kappa < \frac{1}{2}$, the platform sets the longest possible memory even if it puts twice the weight on the firm than it does on consumers ($w_c = \frac{1}{3}$). This hence also extends to the case where the platform maximizes total surplus, i.e. the equally weighted sum of $\tilde{CS}$ and $\tilde{\pi}$ ($w_c = \frac{1}{2}$).

Finally, it is important to note that the threshold weight $\frac{1-2\kappa}{3-4\kappa}$ below which the platform does not choose the consumer-optimal memory for $\kappa < \frac{1}{2}$ is decreasing in κ. For κ close to $\frac{1}{2}$, even a small weight w_c on consumer surplus is sufficient for the platform to choose the consumer-optimal longest possible memory. Once κ is sufficiently low, however, this weight needs to be larger.

5. Conclusion

This paper develops a model which allows to assess strategic pricing incentives of firms in the presence of rating systems. We incorporate two key effects: a *selection effect* where a higher price induces only consumers who are more positively inclined towards the product to purchase and hence increases reviews, and a *direct price effect* where a higher price directly lowers reviews as consumers evaluate a good relative to its purchase price. To the best of our knowledge, this is the first tractable analysis of a firm's strategic pricing decision when it holds an informational advantage over consumers in the presence of rating systems.

We propose an inference method whereby consumers treat the game as stationary and use the current price as the best predictor for past prices. They look for a combination of purchasing consumers' tastes and quality which matches individual rationality conditions and the aggregate rating. For this inference method, we show how the firm's pricing is affected once it takes into account the impact on future profits via reviews. Crucially, whether a firm has an incentive to over- or underprice relative to the myopically optimal price depends on the relative strength of the direct price effect and indirect selection effect. If the direct price effect is large, that is, if a small price change has a large direct effect on the induced reviews because consumers' reviews have a strong dependence on the price level, the firm has an incentive to underprice. By contrast, if this effect is small, the selection effect – a larger price on average leads to higher tastes of purchasing consumers and thus a higher induced review – dominates and the firm has an incentive to overprice.

We show that the game always converges to a stationary equilibrium in which the firm charges the same price every period, the rating is stable, and consumers correctly infer the quality of the good. However, the price level, consumer surplus, and profits in the long run equilibrium depend on how much past reviews matter for the aggregate rating. Firms' profits are always maximized by a rating system which is sensitive to new reviews as it makes it cheaper for them to strategically manipulate the aggregate rating displayed to future consumers. In this sense, the recent pushes by Amazon and Steam for ratings putting more weight on more recent reviews can be seen as beneficial for sellers. In contrast, the effect of the rating system's sensitivity on long run prices and consumer surplus depends on the relative strengths of the direct price and selection effect. A longer memory always lessens the firm's incentive to strategically deviate from the myopically optimal price. If the direct price effect is stronger than the selection effect, the firm has an incentive to underprice relative to the myopic optimum and a high-sensitivity rating system benefits consumers. The reverse is true when the selection effect dominates.

A platform which maximizes a weighted sum of long-run consumer surplus and firm profits by choosing the optimal sensitivity is inclined to choose the consumer-optimal sensitivity unless it cares sufficiently little about them. To the extent that situations captured by our model are concerned, i.e. situations in which a good of time-invariant quality is sold, the recent rating system changes can be interpreted in different lights. It is possible that they benefit consumers as most products feature a weak selection and strong price effect, i.e. are standardized homogenous products with little taste idiosyncrasy. However, our results should serve as a warning sign as the alternative implication is that platforms choose a high sensitivity precisely because they care much more about sellers than about buyers. The latter interpretation is particularly likely for products where the empirical evidence suggests a positive relationship between prices and reviews, such as books or non-casual video games.

Our results can also be interpreted in a different fashion to inform the debate about the coarseness of rating systems.[15] The strategic pricing incentives in our setup arise precisely because consumers only observe an aggregate statistic which reflects multiple different components entering a consumer's review. This is a prominent feature of many platforms which feature a uni-dimensional rating system (e.g. Amazon, Steam). In this sense, we can interpret the coarseness to be beneficial for firms (as they can strategically affect consumers' quality perceptions) and, situationally also for consumers (whenever the strategic incentives exert a downward pressure on firms' pricing decisions). In particular, we find that prices are below the full information benchmark when the price effect dominates. Whether this is the case depends on the characteristics of products being sold, i.e. whether price or selection effect are more likely to dominate. This may explain why certain platform markets such 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.

[15] We thank Heski Bar-Isaac for making this observation.

References

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2019): "Learning From Reviews: The Selection Effect and the Speed of Learning," *working paper.*

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

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.

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.*

Cabral, L. and A. Hortacsu (2004): "The dynamics of seller reputation: Theory and evidence from eBay," *working paper.*

——— (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.

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.

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.

Ho, Y.-C., J. Wu, and Y. Tan (2017): "Disconfirmation effect on online rating behavior: A structural model," *Information Systems Research*, 28, 626–642.

Juan, F., L. Xin, and M. X. ZHANG (2018): "Online Product Reviews-Triggered Dynamic Pricing: Theory and Evidence," *Information Systems Research.*

Kovbasyuk, S. and G. Spagnolo (2018): "Memory and markets," *Available at SSRN 2756540.*

Li, X. and L. M. Hitt (2008): "Self-selection and information role of online product reviews," *Information Systems Research*, 19, 456–474.

Luca, M. (2011): "Reviews, reputation, and revenue: The case of Yelp. com," *working paper.*

Marinovic, I., A. Skrzypacz, and F. Varas (2018): "Dynamic certification and reputation for quality," *American Economic Journal: Microeconomics*, 10, 58–82.

Moreno, A. and C. Terwiesch (2014): "Doing business with strangers: Reputation in online service marketplaces," *Information Systems Research*, 25, 865–886.

Osborne, M. and A. H. Shapiro (2014): "A dynamic model of price signaling, consumer learning, and price adjustment," Tech. rep.

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.

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.*

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 $V(\overline{\psi}) = c + d\overline{\psi} + e\overline{\psi}^2$. As discussed, we replace the price as a control by the rating tomorrow such that

$$p = \frac{\alpha\theta + \beta - m\overline{\psi}' + (m-2)\overline{\psi}}{2\kappa - 1} \tag{52}$$

$$q = \frac{2}{\beta}\frac{(\alpha\theta + \beta - m\overline{\psi}')(\kappa - 1) + (1 + m(\kappa - 1))\overline{\psi}}{(2\kappa - 1)} \tag{53}$$

and the Bellman equation becomes

$$V(\overline{\psi}) = \max_{\overline{\psi}'} \frac{\alpha\theta + \beta - m\overline{\psi}' + (m-2)\overline{\psi}}{2\kappa - 1} \cdot \frac{2}{\beta}\frac{(\alpha\theta + \beta - m\overline{\psi}')(\kappa - 1) + (1 + m(\kappa - 1))\overline{\psi}}{(2\kappa - 1)} + \delta V(\overline{\psi}'). \tag{54}$$

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

$$V'(\overline{\psi}) = d + 2e\overline{\psi}. \tag{55}$$

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

$$\overline{\psi}' = \frac{(1-\kappa)4m(\alpha\theta + \beta) + \delta\beta d(1-2\kappa)^2}{(1-\kappa)(m^2 + 2e\delta\beta\kappa) - 2\delta\beta\kappa e} - \frac{2m(3-2\kappa) - 4m^2(1-\kappa)}{(1-\kappa)(m^2 + 2e\delta\beta\kappa) - 2\delta\beta\kappa e}\overline{\psi} \tag{56}$$

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.

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

$$\begin{aligned}\overline{\psi}' = &\frac{(\alpha\theta + \beta)(\delta(3-2\kappa) + 2(\delta-1)(\kappa-1)m)}{2\delta + (\delta-1)(\kappa-1)m^2 + \sqrt{(\delta(m-2)^2 - m^2)\left(\delta + (\kappa-1)m(2\delta + (\delta-1)(\kappa-1)m)\right)}}\\ &+ \frac{m(2\kappa + 2(1-\kappa)m - 3)}{\sqrt{(\delta(m-2)^2 - m^2)\left(\delta + (\kappa-1)m(2\delta + (\delta-1)(\kappa-1)m)\right)} + (1-\kappa)m^2 + \delta(m-2)((1-\kappa)m - 1)}\bar{\psi}.\end{aligned} \tag{57}$$

Denote

$$a \equiv \frac{(\alpha\theta+\beta)(\delta(3-2\kappa)+2(\delta-1)(\kappa-1)m)}{2\delta+(\delta-1)(\kappa-1)m^2+\sqrt{(\delta(m-2)^2-m^2)\left(\delta+(\kappa-1)m(2\delta+(\delta-1)(\kappa-1)m)\right)}} \tag{58}$$

$$b \equiv \frac{m(2\kappa+2(1-\kappa)m-3)}{\sqrt{(\delta(m-2)^2-m^2)\left(\delta+(\kappa-1)m(2\delta+(\delta-1)(\kappa-1)m)\right)}+(1-\kappa)m^2+\delta(m-2)((1-\kappa)m-1)}. \tag{59}$$

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

$$\overline{\psi_\tau} = \left(a \cdot \sum_{i=0}^{\tau-2} b^i\right) + b^{\tau-1}\bar{\psi}_1 \tag{60}$$

and thus, using $|b| < 1$,[16]

$$\lim_{\tau\to\infty} \bar{\psi}_\tau = \frac{a}{1-b} = \frac{(\alpha\theta+\beta)(\delta(3-2\kappa)+2(\delta-1)(\kappa-1)m)}{4\delta+(\delta-1)(2\kappa-3)m} \equiv \Psi. \tag{61}$$

At this long-run rating, we can use (27) and obtain the long run price

$$\tilde{p} = p(\Psi,\Psi) = -\frac{((\delta-1)m-2\delta)(\alpha\theta+\beta)}{4\delta+(\delta-1)(2\kappa-3)m}. \tag{62}$$

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 firm's problem.

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

$$\begin{aligned}
\frac{|\bar{\psi}_t - \bar{\psi}|}{|\bar{\psi}_{t-1} - \bar{\psi}|} &= \frac{|a + b\bar{\psi}_{t-1} - \frac{a}{1-b}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} && (63)\\
&= \frac{|a \cdot \left(1 - \frac{1}{1-b}\right) + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} && (64)\\
&= \frac{|-\frac{ab}{1-b} + b\bar{\psi}_{t-1}|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} && (65)\\
&= \frac{|b \cdot \left(\bar{\psi}_{t-1} - \frac{a}{1-b}\right)|}{|\bar{\psi}_{t-1} - \frac{a}{1-b}|} && (66)\\
&= |b|. && (67)
\end{aligned}$$

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

We explicitly derived for b that

$$b = \frac{m(2m-3-(m-1)\kappa)}{\sqrt{(\delta(m-2)^2-m^2)\left(\delta+(\kappa-1)m(2\delta+(\delta-1)(\kappa-1)m)\right)}+(1-\kappa)m^2-\delta(m-2)(1-(1-\kappa)m)} \tag{68}$$

Recall that $\kappa < 1 - \frac{1}{m}$ was assumed, which ensures that $b > 0$ for $m \geq 2$. Under the same restrictions, it can be established that $\frac{\partial b}{\partial m} > 0$.[17]

Proof of Proposition 3 In line with the previous argument for optimality of a large m whenever $\kappa > \frac{1}{2}$, we restrict attention to $\kappa < \frac{1}{2}$. Using (29) to obtain $\tilde{CS}$ and plugging this together with (31) into (50), we obtain

$$\pi^P = \frac{2(\alpha\theta+\beta)^2}{\beta} \cdot \frac{(\delta+(1-\delta)(1-\kappa)m)\left((2-w_c)\delta+(1-\delta)m(1-\kappa w_c)\right)}{(4\delta+(1-\delta)(3-2\kappa)m)^2} \tag{69}$$

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

$$\frac{\partial \pi^P}{\partial m} > 0 \iff 2w_c\delta - m(1-\delta)\left[(1-2\kappa)-(3-4\kappa)w_c\right] > 0 \tag{70}$$

and analogously

$$\frac{\partial \pi^P}{\partial m} < 0 \iff 2w_c\delta - m(1-\delta)\left[(1-2\kappa)-(3-4\kappa)w_c\right] < 0. \tag{71}$$

Denote $A \equiv (1-2\kappa)-(3-4\kappa)w_c$ so that $A \leq 0 \iff w_c > \frac{1-2\kappa}{3-4\kappa} \equiv \overline{w_c}$. Clearly, for $A \leq 0$ we have that $\frac{\partial \pi^P}{\partial m}$ is positive for all positive m. It follows similarly that for $A > 0$, i.e. for $w_c < \frac{1-2\kappa}{3-4\kappa}$, that $m = \tilde{m} = \frac{2w_c\delta}{(1-\delta)[1-2\kappa-(3-4\kappa)w_c]} > 0$ is the unique maximizer of π^P given the first-order behavior in (70) and (71). As the lowest permitted memory in our setup is $m = 1$, the platform's optimal choice in this case is hence $m^* = \max\{1, \tilde{m}\}$. Specifically, we have that $\tilde{m} \geq 1 \iff w_c \geq \frac{(1-\delta)(1-2\kappa)}{3-\delta-4(1-\delta)\kappa} \equiv \underline{w_c}$.

B. Calculations for Competition Setup

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

$$u_1(\theta_1, x) = \alpha\theta_1 + \beta(1-x) \tag{72}$$
$$u_2(\theta_2, x) = \alpha\theta_2 + \beta x \tag{73}$$
$$\psi_1(\theta_1, x, p_1) = \alpha\theta_1 + \beta(1-x) - \kappa p_1 \tag{74}$$
$$\psi_1(\theta_2, x, p_2) = \alpha\theta_2 + x - \kappa p_2. \tag{75}$$
$$\tag{76}$$

[17]The detailed calculations are extensive and are omitted for brevity. They are available upon request, as is a Mathematica file verifying the result.

We can solve the equation system characterized by (IND), ($CONS_1$) and ($CONS_2$) and obtain the inference (μ_1, μ_2, x_c) given $(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)$ as

$$\mu_1 = \frac{(6\kappa - 2)p_1 + 2(1-\kappa)p_2 + 6\bar{\psi}_1 - 2\bar{\psi}_2 - 3\beta}{4\alpha} \quad (77)$$

$$\mu_2 = \frac{(6\kappa - 2)p_2 + 2(1-\kappa)p_1 + 6\bar{\psi}_2 - 2\bar{\psi}_1 - 3\beta}{4\alpha} \quad (78)$$

$$x_c = \frac{1}{2} + \frac{(1-\kappa)(p_2 - p_1) + (\bar{\psi}_1 - \bar{\psi}_2)}{\beta}. \quad (79)$$

which induces quantities

$$q_1\left(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)\right) = x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)\,,\ q_2\left(x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2)\right) = 1 - x_c(p_1, p_2, \bar{\psi}_1, \bar{\psi}_2) \quad (80)$$

and reviews

$$\psi_1 = \frac{1}{4}\left(4\alpha\theta_1 + 3\beta + (2 - 6\kappa)p_1 - 2(1-\kappa)p_2 - 2\bar{\psi}_1 + 2\bar{\psi}_2\right) \quad (81)$$

$$\psi_2 = \frac{1}{4}\left(4\alpha\theta_2 + 3\beta - 2(1-\kappa)p_1 + (2 - 6\kappa)p_2 + 2\bar{\psi}_1 - 2\bar{\psi}_2^2\right). \quad (82)$$

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, in the final period, $t = 1$, firms simply solve

$$\max_{p_i^1} p_i^1 \cdot q_i\left(x_c(p_1^1, p_2^1, \bar{\psi}_1^1, \bar{\psi}_2^1)\right). \quad (83)$$

Solving the system obtained from the two firms' first order conditions yields

$$p_1^1 = \frac{3\beta + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1}{6 - 6\kappa}\,,\ p_2^1 = \frac{3\beta - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1}{6 - 6\kappa} \quad (84)$$

which induces profits

$$\pi_1^1 = \frac{\left(3\beta + 2\bar{\psi}_1^1 - 2\bar{\psi}_2^1\right)^2}{36\beta(1-\kappa)}\,,\ \pi_2^1 = \frac{\left(3\beta - 2\bar{\psi}_1^1 + 2\bar{\psi}_2^1\right)^2}{36\beta(1-\kappa)} \quad (85)$$

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

$$\max_{p_i^0} p_i^0 \cdot q_i^0 + \pi_i^1(\bar{\psi}_i^1, \bar{\psi}_{-i}^1), \quad (86)$$

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%.[18] Annual revenue of the platform in 2017 is estimated at $4.3 billion.[19]

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 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:

$$\mathrm{y}_{ig} = \lambda_g + \beta \cdot X_{ig} + \delta \cdot \mathrm{P}_{ig} + \epsilon_{ig} \tag{87}$$

In (87), 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 (87) 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.

[18] https://expandedramblings.com/index.php/steam-statistics/

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

Table 2 shows estimation results for regression (87). 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) recommend	(2) recommend	(3) recommend	(4) recommend
price	0.041	0.074*	0.115***	−0.418***
	(0.043)	(0.042)	(0.044)	(0.134)
game_age		−0.001**	−0.001**	−0.001
		(0.000)	(0.000)	(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***	−0.001***	−0.002***
		(0.000)	(0.000)	(0.001)
frac_helpful		0.596***	0.604***	0.502***
		(0.024)	(0.025)	(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 thu 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”.