Consumer Rating Dynamics*

André Stenzel†   Christoph Wolf‡

July 4, 2016

– preliminary and incomplete –

Abstract:
We consider dynamic price-setting by firms in the presence of rating systems and asymmetric information about product quality. The price a firm charges determines the characteristics of purchasing consumers and thereby affects future ratings and continuation profits. We outline the effects of prices on consumers' inference as well as their review upon purchase. We provide a characterization of the firm's pricing decision in the presence of consumers who conduct inference based on an observable aggregate rating and current price: Inference is characterized by a pair of inferred quality and purchasing consumers' tastes such that the aggregate rating is consistent and purchase decisions are individually rational. Sufficient conditions such that this inference is uniquely determined are obtained and the impact on the incentives of a strategic firm are outlined. We show that the strategic incentives may lead to the firm over- or undercharging relative to the myopically optimal price, depending on fundamentals. Moreover, we provide evidence in form of a numerical implementation which suggests that rating systems with limited memory, as recently suggested by e.g. Amazon.com, may come with adverse effects as regards consumers' inference.

JEL: D21, D82, L15
Keywords: Rating Systems, Dynamic Pricing, Asymmetric Information

*We are grateful to Luis Cabral, Martin Peitz, Johannes Schneider, Nicolas Schutz, Konrad Stahl and Stefan Weiergräber for helpful comments and suggestions, and appreciate input from participants at the MaCCI annual conference 2016.
†University of Leicester. E-Mail: as898[at]le.ac.uk.
‡University of Mannheim and MaCCI. E-Mail: christoph.wolf[at]gess.uni-mannheim.de.

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

It is a well-established fact that prices may serve as an effective means to signal 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)). At the same time, online platforms such as Amazon.com or eBay put extensive effort into the design of their review and rating systems, claiming to provide consumers with additional information to ameliorate the asymmetric information problem. This calls into question the effectiveness of price signaling in large, anonymous online markets – if prices were perfectly informative of product quality there would be no need for an effective rating system.[1] Moreover, the economic literature has documented ample empirical evidence that ratings have significant effects on demand (see e.g. Anderson and Magruder (2012), Chevalier and Mayzlin (2006), Luca (2011), Moreno and Terwiesch (2014) and Cabral and Hortacsu (2010)), as well as that signaling via price or advertising and informative rating systems are substitutes in terms of information provision regarding product quality (Bhargava and Feng (2015)).

Sellers in online markets hence need to take the effect of their strategic decisions on the reviews and ratings subsequently given by customers into account: Ratings reflect the experience of consumers in prior periods and shape the inference and thus demand and profits in the future. However, while the literature has addressed issues such as investment/disinvestment in quality (Board and Meyer-ter Vehn (2013)) or certification of quality (Marinovic et al. (2015)), we are not aware of articles which explicitly assess the strategic use of prices to affect the composition of purchasing consumers and hence future (expected) review scores.[2] We contribute to the understanding of price-setting in online markets by analyzing a model where a long-lived seller sells a good of fixed quality to short-lived consumers. Consumers value quality but do not directly observe it prior to consumption. Moreover, they differ in their taste for the product, i.e. exhibit horizontal differentiation. A consumer contemplating purchase of the good forms her belief about the product's quality using 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

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

[2] We are aware of ongoing work by Martin and Shelegia (2015) who consider strategic price-setting in a framework where information may be transmitted between periods via ratings. However, their focus is on the existence of equilibria where price still serves as a perfect signaling device. By contrast, we are interested in how price is able to affect reviews and hence indirectly impacts future quality inference and profits when prices are not able to perfectly signal quality.

were relative to the purchase price or quality expectation upon purchase. If 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. 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. Moreover, the assumption of imperfect observability of review- and rating-relevant information is validated by noting that even if individual reviews are reported in addition to aggregate ratings (e.g. on Amazon.com), there is no information obtainable about the price at which the transaction leading to the review occured.

As an example for the mechanism we have in mind, consider a monopolistic seller who has to decide on the current-period price of a good of fixed quality. All else, in particular beliefs about the inherent quality of the good, equal, a higher price charged by the seller increases the threshold level of horizontal taste for the good above which consumers choose to buy. In expectation, a higher price thus leads to a higher average gross utility of purchasing consumers. If this gross utility is (part of) the basis for consumer reviews, a high price may thus serve to increase (expected) review scores and hence future profits – potentially at the cost of current period flow profits. We thus emphasize the role of prices not as a signaling device, but as a selection device amongst potential consumers to strategically influence the induced reviews and ratings. Similarly, the price may also affect future reviews via inducing a different expectation of purchasing consumers, or directly if satisfaction is assessed relative to price.

A defining characteristic of platforms such as Amazon is that it allows for sellers to frequently and costlessly adjust prices, a feature which firms use extensively. A natural explanation for price changes is that they capture responses to fluctuating demand and new information about demand conditions for the seller. However, these types of adjustments seem unlikely to be able to explain the full extent of price variation. The view of prices as a strategic device to influence ratings and thus future profits, however, may help to explain the substantial variation prices exhibit over time in online markets. If prices indirectly affect future quality perceptions and hence profits via an exogenously given review and rating system, both the frequent price changes as a response to the rating score, as well as the indirect learning of consumers via ratings instead of direct inference from price, can be rationalized.

Preliminary Findings We provide sufficient conditions such that inference is uniquely determined. We assess the comparative statics of the inference and the impact of prices on inference and induced reviews. Taking into account future profits, a firm may deviate from myopically optimal pricing both by charging a higher or a lower price, which depends on fundamentals. In the case of linear additive utility and review functions, a firm has an incentive to price over the myopic optimum if the review is relatively insensitive to the purchase price, and relatively insensitive to expectation compared to actual quality, or if the review is relatively sensitive to the purchase price, and relatively sensitive to expectation compared to actual quality. In these cases, a higher price induces a higher review, while otherwise a lower price has the same effect. A numerical implementation of our setup shows that the recent emphasis by e.g. Amazon.com to make more recent reviews more important for the displayed aggregate rating may have adverse effects on consumers – it allows firms to engage in reputation-cycles and may lead to significantly higher price levels.

Structure of the Article Section 1.1 briefly surveys the related literature. Section 2 sets up the model and discusses the inference specification in detail. Section 3 provides conditions such that inference is uniquely determined, which is used in Section 4 to assess the comparative statics. Section 5 discusses the strategic price-setting effects, which lead to an assessment of dynamic pricing incentives in Section 6. We consider a special linear additive specification in Section 7 and illustrate our results with a numerical implementation in Section 8. Section 9 concludes.

1.1. Related Literature

We relate to several strands of literature. While we focus on signaling via ratings and the indirect effect prices have on future quality inference, there is a vast literature on direct price signaling. Starting with the seminal article by Wolinsky (1983), studies such as Bagwell and Riordan (1991) have focused on identifying the conditions such that separating equilibria are obtained – a theme of this strand of literature is the required markup by a high-quality (and high-cost) firm to separate from a low-quality (and low-cost) competitor. This markup prevents the low-quality competitor from imitating as it decrases sales which is comparatively more costly for a low-quality and low-cost firm. A central prediction is that prices should decrease over time as the required markup is lowered when more information is already known to consumers, e.g. due to learning in an experience good context (Bagwell and Riordan (1991)). This mechanism also plays a role in our preliminary results: As the information provided via ratings becomes less sensitive to new reviews over time, the value of strategically influencing reviews is lowered, which, in the context of naive consumers reporting gross utility, leads to a lower markup over the myopically optimal price. Building on these contributions, 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.

We also relate to articles assessing strategic firm behavior in the presence of review or reputation concerns. 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, the price-setting game is not explicitly assessed; firms simply 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. (2015) extend the analysis by considering a framework in which quality can be certified (and certification is costly). They show that benefits of reputation may be adversely affected by excessive certification, but show that first-best investment may arise if certification decisions are made based on the time since last certification. 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.

Finally, our analysis draws from and complements the empirical literature on the impact of review and rating systems on online platforms. 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.[3]

2. Model Setup

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 and covers $T < \infty$ periods, $t = 1, 2, 3, \ldots, T$. The economy consists of a single long-lived seller and a 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 known cdf F. The realization of θ is private information to the seller. In each period, the seller decides on the price p_t at which she is willing to sell. Marginal costs of production are independent of quality and normalized to 0. The seller is risk-neutral and discounts future profits at a rate $\delta \in [0, 1]$.[4]

[3]However, increased information provision via external sources may increase price distortions if it affects the signaling mix between advertising and price signaling a firm employs.

[4]For $\delta = 0$, the firm does not take the future into account and hence prices myopically optimal in every period.

Consumers In each period t, there is a mass one 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 $\omega_i \sim [\underline{\omega}, \bar{\omega}]$. The gross utility of a consumer may depend on the actual quality of the good, θ, the consumer's taste for the product ω, the price paid for the product p and the quality expectation μ. A consumer's net utility u, derived as the gross consumption utility $u(\cdot)$ net of price, is characterized by

$$u = u(\theta, \omega, p, \mu) - p. \tag{1}$$

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 conditional on purchase is characterized by

$$\psi = \psi\left(\theta, \omega, p, \mu\right). \tag{2}$$

Similar to the utility, the review provided by a consumer may depend not only on the actual quality θ and the taste ω, but also on the belief μ about quality θ prior to purchase, and the price paid for it. We assume that both a consumer's gross utility and the induced review are increasing in both the actual quality of the good and the individual taste for it. Moreover, we assume that a higher expectation weakly decreases both enjoyment and rating.

Assumption 1 (Gross utility & review) *A consumer's gross utility $u(\theta, \omega, p, \mu)$, as well as a consumer's rating $\psi\left(\theta, \omega, p, \mu\right)$, are continuously differentiable in all variables, and*

(i) strictly increasing in quality θ: $\frac{\partial u}{\partial \theta} > 0$, $\frac{\partial \psi}{\partial \theta} > 0$

(ii) strictly increasing in taste ω: $\frac{\partial u}{\partial \omega} > 0$, $\frac{\partial \psi}{\partial \omega} > 0$.

(iii) weakly decreasing in expectation μ: $\frac{\partial u}{\partial \mu} \leq 0$, $\frac{\partial \psi}{\partial \mu} \leq 0$

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).

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. The rating system is hence characterized by the mapping from current aggregate rating $\bar{\psi}_t$ and current average review ψ_t into next period's aggregate rating $\bar{\psi}_{t+1}$. We denote this mapping by ρ and let it potentially depend on t. This allows to capture different rating systems: For example, aggregate ratings may become less sensitive to new reviews over time if the average of all ratings is reported, or may 'forget' reviews sufficiently far in the past and only reflect more recent reviews. This is particularly relevant given the recent pushes by online platform such as Amazon.com to make more recent reviews matter more for the displayed aggregate rating.[5]

$$\bar{\psi}_{t+1} = \rho^t \left(\bar{\psi}_t, \psi_t\right). \tag{3}$$

Discussion of the Setup There are several assumptions inherent in the setup which warrant further discussion. First, we assume that consumers only observe the current aggregate rating and the current price. Both assumptions are unlikely to be perfectly satisfied in reality. Amazon.com, for example, lists all reviews written for a particular product, while price tracking websites such as camelcamelcamel.com provide at least partial access to historical price data. Nonetheless, we feel that at least a sizable number of potential consumers is primarily guided by the aggregate rating displayed in a 0 to 5 star format on Amazon and is moreover unlikely to use price tracking websites for more than price watches (consisting of alerts when the price falls below a threshold, which is the main service these websites provide), if even that. As such, we feel comfortable with this assumption as a first step of the analysis of price-setting behavior in online markets. In particular, it reflects an important feature of rating systems in online markets – a given review cannot be matched with the price at which the transaction occured.

The assumption that every purchasing consumer rates is more restrictive. As only the average review enters the updated rating, the fact that all consumers rate and the law of large numbers ensure that a strategic firm can perfectly forecast next period's rating, which greatly facilitates the analysis. While this is unlikely to hold true in reality, it nonetheless helps in isolating the qualitative tradeoffs faced by the firm in the presence of the strategic pricing incentives. Related to this consideration, the number of reviews in a given period has no impact in the current formulation: Only the average review in a given period affects the update process of the aggregate rating. It stands to reason that in reality, the more consumers purchase the good in a given period, the more reviews are obtained. This in turn affects the future responsiveness of the aggregate rating statistic. We plan on relaxing this assumption in the next steps – intuitively, however, the qualitative predictions in terms of price path and rating path should not be altered by incorporating this.

2.1. Consumer Inference

A central requirement is to specify how consumers conduct quality inference given their observables. In principle, quality inference could be based on direct price-signaling. In Appendix A, we characterize the equilibria of the static game where consumer inference is based purely on observed prices. While both pooling and separating equilibria arise, pooling in fact weakly dominates separating equilibria in terms of firm profits. Absent ratings, any combination whereby a stage-game equilibrium (either pooling or separating) is played in every period would hence

[5] Preliminary findings from a numerical implementation of the model given different rating systems, i.e. different ρ^t, are presented in Section 8.

constitute an equilibrium.[6] However, if ratings are present, extending the analysis with price-signaling to multiple periods is not straightforward: Given that ratings carry information, the implications of any deviation from a conjectured equilibrium are often ambiguous and it is – without a substantial set of restrictive assumptions on the behavior of $u(\cdot)$ and $\psi(\cdot)$ – not clear in which direction the most profitable deviation takes place.[7]

Crucially, empirical evidence suggests that asymmetric information concerns prevail even in the presence of rating systems: Uncertainty about product quality is a major concern. This speaks against fully rational consumers even given the limited observables in the form of current aggregate rating and price only.[8] In principle, fully rational consumers could 'solve' the firm's maximization problem at any point in time for any quality θ and use this to infer the quality which leads to an aggregate rating of $\bar{\psi}_t$.[9] Full rationality thus has two drawbacks: On the one hand, it imposes high computational requirements on the consumers. Given that consumers only observe the aggregate rating and the current price, they would need to compute the optimal response by the firm for all possible situations and use this to draw inference along the unobserved path of prices and reviews. On the other hand, even if these computational requirements and associated tractability issues are met, it would allow for ostensibly perfect inference, which, as previously argued, is unlikely to hold in reality.

Given the tractability issues associated with perfectly rational consumers, as well as the observation that asymmetric information concerns prevail despite the presence of informative rating systems, we abstract from direct price signaling considerations. Moreover, we assume that consumers use a heuristic: They conduct quality inference from the observables (consisting of aggregate rating $\bar{\psi}_t$ and current price p_t) under the hypothetical scenario that all past consumers were faced with the same price-rating-combination. Note that this still lets price have an indirect effect on quality inference; however, there is no direct informational content associated with a given price.

[6] If there are no ratings, the fact that consumers are short-lived implies that no information is transmitted across periods.

[7] To illustrate the difficulties, consider the case of $T = 2$ and a conjectured pooling equilibrium in the first period. Under reasonable assumptions, this would lead to perfect revelation of the quality in period 2 via the observed rating. If a firm deviates from the equilibrium in the first period, it sacrifices flow profits (due to pessimistic off-path beliefs) but may induce a higher rating and hence higher profits at $t = 2$. Whether this higher rating is attained by charging a higher or lower price at $t = 1$, however, depends on the relative impact of θ, p, μ, and ω on the average review ψ and requires substantial assumptions to be evaluated. Similarly, the degree to which flow profits are sacrificed and the resolution of the tradeoff required to verify whether a conjectured equilibrium can be rationalized depends on the sensitivity of $u(\cdot)$ to these inputs. Even for relatively simple utility functions such as $u(\theta, \omega, p, \mu) = \theta + \rho\omega$, i.e. additively separable gross utility independent of price and belief, the analysis is rendered intractable for reasonably general forms of $\psi(\cdot)$.

[8] If consumers were to observe more than current aggregate rating and price, rational inference would pin down the firm's quality in a variety of different settings. Any review which is based on a known price would immediately identify product quality θ given a specified inference, as would knowledge of the full review or price path (where the fact that the firm sets prices optimally allows for perfect inference).

[9] Note that this is feasible if t is known, i.e. if consumers are aware of the number of periods which have already passed. While online platforms typically depict since when a particular product has been sold, it seems unreasonable to map this directly into the number of periods of the game. This is particularly the case once one takes into account that while the number of reviews may provide a good proxy for the number of sales, there is no information provided about the price at which a given sale took place. Thus, there is no way for consumers to infer how many reviews (purchases) occured at one price, i.e. in one period. The product's online presence timestamp only applies to exactly that particular listing – if the seller sold the same product with a different listing before, or changed the specifics of a particular listing (such as selling a new combination of products), this would be hidden from consumers.

Assumption 2 (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*

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

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

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 essentially reduces quality 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 rating (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 alternative way of interpreting the assumption is that consumers conduct quality inference by treating the game as if it were already 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 restrictive, we do not consider it to be too far removed from reality. As discussed previously, past prices are not directly observable on online platforms. While individual reviews typically are available, 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). As they are moreover noisy 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 'quasi-stationary' equilibrium seems a reasonable approximation in light of the difficulty of imposing full rationality – consumers are often unclear about how many periods have passed and how often the firm changed prices in the past.

The biggest drawback is admittedly that it rules out explicit price-signaling. In light of the discussion about the sustainability of direct price-signaling in this setup, as well as to isolate the strategic aspect of prices in regards to determining the set of purchasing consumers, we think that this is a sensible first step. A combination of direct price-signaling as in e.g. Bagwell and Riordan (1991) and our mechanism is, however, an interesting avenue for future work.

The next section provides additional assumptions and sufficient conditions such that the inference given Assumption 2 is uniquely determined.

3. Existence & uniqueness of inference

We want to show under which conditions a solution $(\mu, \omega) = (\mu^*, \omega^*)$ to the equation system

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

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

exists. To do so, we proceed in several steps. We first impose an additional restriction on the utility and review functions which is that they depend more on changes in quality than on changes in expectation. This ensures that the slope of the solution pairs (μ^*, ω^*) of (CONS) and (RAT) respectively in the $\mu - \omega$-space is monotone. Next, we require that the price is such that solutions to (CONS) and (RAT) exist. Finally, we show that provided the slopes of the solutions to (CONS) and (RAT) exhibit a monotonicity property, the solution is uniquely determined and the implicit function theorem (henceforth IFT) can be applied to assess the impact of price changes on inference and future reviews. Note that these conditions are not necessary but sufficient. By studying several examples we conclude that they are not overly restrictive and cover several intuitive utility and review function combinations.

Assumption 3 *In addition to Assumption 1 it holds that for all $\theta \in [\underline{\theta}, \bar{\theta}]$*

$$\frac{\partial u}{\partial \theta} > \left|\frac{\partial u}{\partial \mu}\right| \tag{4}$$

$$\frac{\partial \psi}{\partial \theta} > \left|\frac{\partial \psi}{\partial \mu}\right|, \tag{5}$$

whenever belief and quality coincide, that is, for $\theta = \mu$.

Denote by $\mu_i^*(\omega^*)$ for $i \in \{C, R\}$ the solution for (CONS) and (RAT), respectively. Assumption 3 implies that both $\mu^*(\omega^*)^i$ are decreasing in ω^*. To see this, note that both (CONS) and (RAT) are evaluated at $\theta = \mu$. Under Assumption 3, a higher μ^* hence induces a higher $u(\cdot)$ and $\psi(\cdot)$, respectively, necessitating a lower ω^* to satisfy the equation.

The monotonicity of the solutions implies that if at the boundaries $\underline{\omega}$ and $\overline{\omega}$ the corresponding solution $\mu_i^*(\underline{\omega}), \mu_i^*(\overline{\omega}) \in [\underline{\theta}, \overline{\theta}]$, then so is any μ^* in the interior. Moreover, the $\mu_i^*(\cdot)$ are continuous due to the continuous differentiability of $u(\cdot)$ and $\psi(\cdot)$. To ensure that a solution for (CONS) and (RAT) exists (separately) over the entire support of ω, i.e. over $\Omega = [\underline{\omega}, \overline{\omega}]$, it suffices that the price p observed by consumers is such that at the boundaries the corresponding μ^* is in $[\underline{\theta}, \overline{\theta}]$.[10]

Assumption 4 *The current price p is such that the solution $\mu_C^*(\underline{\omega}), \mu_R^*(\underline{\omega})$ to each of the equations*

$$\psi(\mu_C^*(\underline{\omega}), \omega^e(\underline{\omega}), p, \mu_C^*(\underline{\omega})) = \overline{\psi} \tag{6}$$

$$u(\mu_R^*(\underline{\omega}), \underline{\omega}, p, \mu_R^*(\underline{\omega})) = p \tag{7}$$

and $\mu_C^*(\overline{\omega}), \mu_R^*(\overline{\omega})$ to each of the equations

$$\psi(\mu_C^*(\overline{\omega}), \omega^e(\overline{\omega}), p, \mu_C^*(\overline{\omega})) = \overline{\psi} \tag{8}$$

$$u(\mu_R^*(\overline{\omega}), \overline{\omega}, p, \mu_R^*(\overline{\omega})) = p \tag{9}$$

are in the set of possible quality levels $[\underline{\theta}, \overline{\theta}]$.

This implicitly gives bounds on p such that (CONS) and (RAT) separately exhibit solutions $\mu^*(\omega^*) \in [\underline{\theta}, \bar{\theta}]$ for all $\omega^* \in \Omega$. The final step is to provide a sufficient condition which allows us to conclude existence of a (μ^*, ω^*)-pair which jointly solves (CONS) and (RAT).

Assumption 5 *We assume that $\mu_i^*(\underline{\omega}) \geq \mu_j^*(\underline{\omega})$ and $\mu_i^*(\overline{\omega}) \leq \mu_j^*(\overline{\omega})$ for $i, j \in \{C, R\}$, $i \neq j$.*

If Assumption 5 holds, existence of a joint solution to (CONS) and (RAT) immediately follows from the intermediate value theorem as $\mu_{i,j}^*(\omega)$ are continuous on Ω. This allows us to conclude the following Lemma.

Lemma 1 *Suppose Assumption 3 to 5 hold. Then, the consumer-inference problem described by (CONS) and (RAT) has at least one solution.*

3.1. Uniqueness

While Assumption 3 to 5 ensure that inference can be conducted, it is still possible that multiple solution pairs (μ^*, ω^*) to (CONS) and (RAT) exist. To rule out this case, it suffices to ensure that the solution functions $\mu_i^*(\cdot)$ are such that one is steeper than the other over the entire support. This is captured by the following condition.

Assumption 6 *For the entire range of $\tilde{\omega} \in [\underline{\omega}, \bar{\omega}]$ and $\tilde{\mu} \in [\underline{\theta}, \bar{\theta}]$, and for any combination $(\bar{\psi}, \tilde{p})$ which satisfies Assumption 4, it holds that*

$$\frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} > \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}} \quad \textit{or} \quad \frac{\frac{\partial \psi}{\partial \omega}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} < \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}, \tag{STP}$$

[10] As discussed in more detail subsequently, this is particularly relevant to avoid the case of unbounded flow profits – inference may be such that higher prices actually decrease the threshold taste for purchase. If this is the case, the implicit bounds due to a valid inference provide upper bounds on prices and thus flow profits.

where partial derivatives are evaluated at $\theta = \tilde{\mu}$, $\omega = \tilde{\omega}$, $p = \tilde{p}$ and $\mu = \tilde{\mu}$.

Note that as $\psi(\cdot)$ is evaluated at $\omega^e(\omega)$ for a given cutoff ω, we denote by $\frac{\partial \psi}{\partial \tilde{\omega}}$ not the partial derivative of ψ with respect to ω, but the partial derivative with respect to a change in the cutoff, i.e. formally

$$\frac{\partial \psi}{\partial \tilde{\omega}} = \left.\frac{\partial \psi}{\partial \omega}\right|_{\omega=\omega^e(\tilde{\omega})} \cdot \frac{\partial \omega^e}{\partial \tilde{\omega}},$$

where $\frac{\partial \psi}{\partial \omega}$ is the derivative with respect to the second input of $\psi(\theta, \omega, p, \mu)$ and $\frac{\partial \omega^e}{\partial \tilde{\omega}}$ is the derivative of the induced average taste ω^e with respect to the cutoff $\tilde{\omega}$ – $\frac{\partial \omega^e}{\partial \tilde{\omega}}$ measures how much the average consumer's taste changes given a change in the cutoff.

(STP) itself assesses the relative impact of a change in the cutoff ω compared to a simultaneous change in belief μ across $\psi(\cdot)$ and $u(\cdot)$. Note that a change in μ affects $\psi(\cdot)$ both via the actual belief ($\frac{\partial \psi}{\partial \mu}$) and via the quality ($\frac{\partial \psi}{\partial \theta}$) as inference is conducted treating the inferred quality as the actual one, i.e. imposing $\theta = \mu = \mu^*$. The same also applies to $u(\cdot)$. To illustrate this, suppose that

$$\frac{\frac{\partial \psi}{\partial \omega}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} < \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}$$

and that $\mu = \theta$ increases while ω decreases in such a fashion as to leave $\psi(\cdot)$ overall unchanged. In that case, we can expect $u(\cdot)$ – given the same changes – to decrease: It reacts comparatively more strongly to the change in the cutoff than $\psi(\cdot)$ and this negative effect thus outweighs the positive overall effect of an increase in $\mu = \theta$. The sign of (STP) at the solution pair (μ^*, ω^*) to (CONS) and (RAT) has important implications for the comparative statics of the inference. This is because the relative strength of these effects determines whether the utility (and hence (RAT)) or review (and hence (CONS)) function reacts more strongly to changes in the aggregate rating ($\bar{\psi}$) and price (p), respectively.

To see that the steepness condition (STP) is sufficient for at most one joint solution to (CONS) and (RAT) to exist, note the following. The implicit function theorem yields that the solution functions $\mu_i^*(\cdot)$ exhibit

$$\frac{\partial \mu_C^*}{\partial \omega^*} = \frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} \qquad \text{and} \qquad \frac{\partial \mu_R^*}{\partial \omega^*} = \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}, \tag{10}$$

and hence Assumption 6 implies that $\mu_C^*(\cdot)$ and $\mu_R^*(\cdot)$ can intersect at most once on Ω.

Lemma 2 *Suppose that Assumption 3 to 6 hold. Then, a solution to the consumer-inference problem described by* (CONS) *and* (RAT) *exists and is unique.*

Given Lemma 2, we know that consumer inference according to Assumption 2 can be conducted and yields to a uniquely determined consumer behavior. Thus, the consumer problem is well-defined and can the firm's problem can be studied.

3.2. Discussion of the assumptions

The main purpose of this section was to provide assumptions under which the inference conducted by the consumer according to the quasi-stationarity Assumption 2 is uniquely determined. We will see that it is straightforward to verify the assumptions once utility and review functions are specified and that they are in fact satisfied by a wide range of functions determining utility and review behavior.

The assumptions themselves are relatively innocuous. Assumption 3 states that consumers respond more strongly to changes in the actual quality than to changes in the expectation. Even if utility and subsequent reviews are determined relative to expectation, the actual quality of the good is required to have a stronger effect. It seems sensible to assume that actual quality is more important for a consumer than the expectation of quality. Assumption 4 is a constraint on the price charged by the firm and can easily be rationalized once the full game is assessed: Denoting the implicit bounds on p by $\underline{p}$ and $\bar{p}$, we can simply close the inference for the full game by imposing that no consumer is willing to purchase if the price falls outside $[\underline{p}(\bar{\psi}_t), \bar{p}(\bar{\psi}_t)]$. This captures the fact that consumers are not willing to purchase if rational and consistent inference according to (RAT) and (CONS) is not possible: No consumer is willing to purchase if the deal offered for the product seems off, and thereby ensures that the firm only charges prices such that valid nference is feasible. Finally, the boundary and steepness assumptions Assumption 5 and Assumption 6 can both be rationalized once one takes into account that ψ is evaluated at $\omega^e(\tilde{\omega})$, the expected horizontal component conditional on purchase, while u is evaluated at $\tilde{\omega}$, which is the lowest purchasing consumer's horizontal component. This difference leaves plenty of scope for the assumptions to be satisfied, as we will see in detail when discussing examples.

It should also be noted that even if (STP) in Assumption 6 does not hold on the entire support, it may be the case that the inference according to Assumption 2 is uniquely determined. Moreover, provided that (STP) holds locally at such a solution, the comparative statics can be assessed as in Section 4 as the IFT is applicable.

4. Inference: comparative statics

Given a unique inference (μ^*, ω^*), this section assesses the comparative statics. Note that inference is affected by the current aggregate rating $\bar{\psi}$, as well as the price charged, p. Recall that the IFT is applicable as the determinant of the Jacobian is nonzero provided that the steepness condition (STP) in Assumption 6 is satisfied. We apply the IFT to the equation system

$$\begin{aligned} 0 &= \psi\left(\tilde{\mu}, \omega^e(\tilde{\omega}), p, \tilde{\mu}\right) - \bar{\psi} \\ 0 &= u\left(\tilde{\mu}, \tilde{\omega}, p, \tilde{\mu}\right) - p. \end{aligned} \tag{11}$$

around the solution (μ^*, ω^*) which exists and is unique if Assumption 1 to 6 hold. In interpreting the effects, note that a higher cutoff is synonymous with an in-period reduction in the quantity sold. This is particularly important to bear in mind once considering the tradeoff faced by the monopolist in determining the optimal price which balances current period profit considerations with those pertaining to future profits via changes in the induced rating.

4.1. Effect of aggregate rating on belief and cutoff

The comparative statics of how the present rating affects the induced consumers' belief and cutoff (for a given price p) are characterized by:

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial \tilde{\psi}} \\ \frac{\partial \omega^*}{\partial \tilde{\psi}} \end{pmatrix} = -\begin{pmatrix} \frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu} & \frac{\partial \psi}{\partial \tilde{\omega}} \end{pmatrix}^{-1} \begin{pmatrix} 0 \\ -1 \end{pmatrix}. \tag{12}$$

For the inverse of the Jacobian, we have

$$\begin{pmatrix} \frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu} & \frac{\partial \psi}{\partial \tilde{\omega}} \end{pmatrix}^{-1} = \frac{1}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \begin{pmatrix} \frac{\partial \psi}{\partial \tilde{\omega}} & -\frac{\partial u}{\partial \omega} \\ -\frac{\partial \psi}{\partial \theta} - \frac{\partial \psi}{\partial \mu} & \frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu} \end{pmatrix}. \tag{13}$$

Hence, we obtain that

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial \tilde{\psi}} \\ \frac{\partial \omega^*}{\partial \tilde{\psi}} \end{pmatrix} = -\frac{1}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \begin{pmatrix} \frac{\partial \psi}{\partial \tilde{\omega}} & -\frac{\partial u}{\partial \omega} \\ -\frac{\partial \psi}{\partial \theta} - \frac{\partial \psi}{\partial \mu} & \frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu} \end{pmatrix} \begin{pmatrix} 0 \\ -1 \end{pmatrix} \tag{14}$$

$$= -\frac{1}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \begin{pmatrix} \frac{\partial u}{\partial \omega} \\ -(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}) \end{pmatrix} \tag{15}$$

$$= \begin{pmatrix} \frac{-\frac{\partial u}{\partial \omega}}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \\ \frac{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \end{pmatrix}. \tag{16}$$

Given Assumption 1, we have that $-\frac{\partial u}{\partial \omega} < 0$ and hence belief increases (decreases) in the current rating if

$$\begin{aligned} &\left(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}\right)\frac{\partial \psi}{\partial \tilde{\omega}} < (>) \frac{\partial u}{\partial \omega}\left(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}\right) \\ \Leftrightarrow\ & \frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} < (>) \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}, \end{aligned} \tag{17}$$

while the cutoff increases (decreases) if the belief decreases (increases) as the numerator is positive by Assumption 3 and the denominator the same as for the belief. Note that the steepness condition in Assumption 6 thus pins down whether the inference pair (μ^*, ω^*) increases or decreases in the current rating. Moreover, a change in $\bar{\psi}_t$ always has opposing effects on the induced belief μ^* and the induced cutoff ω^*: Either more consumers purchase the product ($\omega^* \downarrow$) while holding a higher belief ($\mu^* \uparrow$), or less consumers purchase and hold a lower belief ($\omega^* \uparrow$, $\mu^* \downarrow$).

Suppose for example that (STP) is satisfied in that we have on the full support

$$\frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} < \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}, \tag{18}$$

i.e. that a cutoff change is relatively less important for $\psi(\cdot)$ than for $u(\cdot)$. This implies that the consistency of the rating (CONS) reacts less strongly to the change in $\bar{\psi}$ than the rationality condition (RAT). In that case, the denominator is negative and the quality inference increases in the aggregate rating, while the cutoff decreases.

To illustrate this, consider consumer inference given an increase in the aggregate rating $\bar{\psi}$. The inference (μ^*, ω^*) has to adjust such that the higher rating is matched. In principle, this can be facilitated either by a higher inferred quality μ^*, or a higher cutoff type ω^*. However, the cutoff type still needs to be indifferent and the change in rating does not matter for (RAT). Thus, the changes in μ^* and ω^* must exactly offset each other as it regards the cutoff type's utility. Because of (18), the only way to facilitate this is via an increased inferred quality μ^* and decreased cutoff ω^*: $\psi(\cdot)$ reacts relatively less to a change in ω^* than μ^* compared to $u(\cdot)$ and (RAT) – by keeping the cutoff type indifferent, the lower responsiveness to the decreasing ω^* than the increasing inferred quality μ^* thus allows to match the higher rating $\bar{\psi}$ keeping the cutoff type indifferent.

Intuitively, the consumer's inference adjusts by adjusting the relatively more responsive dimension in the same direction as the change in aggregate rating: If $\bar{\psi}$ increases and reviews are more responsive to changes in μ^* than ω^*, inferred quality μ^* increases and the cutoff ω^* decreases. Similarly, if $\bar{\psi}$ were to decrease, μ^* would decrease under the same condition: The higher responsiveness to μ^* allows to match the aggregate rating consistently, while the cutoff type is exactly indifferent.

4.2. Effect of price on belief and cutoff

Similar to the previous section, we can also assess how a price change affects the consumer's belief about quality μ^* and the cutoff type ω^*. The implicit function theorem yields

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial p} \\ \frac{\partial \omega^*}{\partial p} \end{pmatrix} = - \begin{pmatrix} \frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu} & \frac{\partial u}{\partial \omega} \\ \frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu} & \frac{\partial \psi}{\partial \tilde{\omega}} \end{pmatrix}^{-1} \begin{pmatrix} \frac{\partial u}{\partial p} - 1 \\ \frac{\partial \psi}{\partial p} \end{pmatrix}. \tag{19}$$

The inverse of the Jacobian is identical to the previous subsection and we obtain for the effect of a price change on belief and cutoff

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial p} \\ \frac{\partial \omega^*}{\partial p} \end{pmatrix} = -\frac{1}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} \begin{pmatrix} \frac{\partial \psi}{\partial \tilde{\omega}}(\frac{\partial u}{\partial p}-1) - \frac{\partial u}{\partial \omega}\frac{\partial \psi}{\partial p} \\ (-\frac{\partial \psi}{\partial \theta}-\frac{\partial \psi}{\partial \mu})(\frac{\partial u}{\partial p}-1) + (\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial p} \end{pmatrix} \tag{20}$$

$$= \begin{pmatrix} \frac{\frac{\partial u}{\partial \omega}\frac{\partial \psi}{\partial p} - \frac{\partial \psi}{\partial \tilde{\omega}}(\frac{\partial u}{\partial p}-1)}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} \\ \frac{(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})(\frac{\partial u}{\partial p}-1) - (\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial p}}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} \end{pmatrix}. \tag{21}$$

A marginal increase in price therefore increases the belief if

$$\frac{\partial \mu^*}{\partial p} = \frac{\frac{\partial u}{\partial \omega}\frac{\partial \psi}{\partial p} - \frac{\partial \psi}{\partial \tilde{\omega}}(\frac{\partial u}{\partial p}-1)}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} > 0 \tag{22}$$

and decreases it otherwise. Similarly, the cutoff is increased by a marginal increase in the price if

$$\frac{\partial \omega^*}{\partial p} = \frac{(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})\left(\frac{\partial u}{\partial p}-1\right) - (\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial p}}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} > 0. \tag{23}$$

Note that the denominator in both cases again corresponds to the steepness condition (STP). For the change in the induced belief ($\frac{\partial \mu^*}{\partial p}$), the numerator assesses relative impact of a simultaneous change in the cutoff to a simultaneous change in price across (CONS) and (RAT). For the change in the induced cutoff ($\frac{\partial \omega^*}{\partial p}$), it is the relative impact of a simultaneous change in belief (which corresponds to θ for the inference) to a simultaneous change in price which determines the sign.

To illustrate this, consider consumer inference when faced with a price increase. Suppose again that

$$\frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu}} < \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu}}, \tag{24}$$

holds, i.e. that a cutoff change is relatively less important for $\psi(\cdot)$ than for $u(\cdot)$. In addition, let

$$\frac{\partial u}{\partial \omega}\frac{\partial \psi}{\partial p} - \frac{\partial \psi}{\partial \tilde{\omega}}\left(\frac{\partial u}{\partial p}-1\right) < 0 \Leftrightarrow \frac{\frac{\partial \psi}{\partial p}}{\frac{\partial \psi}{\partial \tilde{\omega}}} < \frac{\frac{\partial u}{\partial p}-1}{\frac{\partial u}{\partial \omega}} \overset{\frac{\partial \psi}{\partial p}, \frac{\partial u}{\partial p}-1<0}{\Leftrightarrow} \left|\frac{\frac{\partial \psi}{\partial p}}{\frac{\partial \psi}{\partial \tilde{\omega}}}\right| > \left|\frac{\frac{\partial u}{\partial p}-1}{\frac{\partial u}{\partial \omega}}\right|, \tag{25}$$

that is let $\psi(\cdot)$ (and thus (RAT)) be relatively more sensitive to a change in the cutoff than to a change in price compared with $u(\cdot)$ (and thus (CONS)).[11] The consumers' inferred belief and cutoff have to adjust following the price increase such that (RAT) and (CONS) are satisfied. As price increases, the gross utility of the marginal consumer has to increase as well. This can be accomplished either via a higher inferred quality equal to belief, μ^*, or a higher cutoff type, ω^*. If μ^* decreases, ω^* has to increase significantly such that the aggregate rating is consistent:

[11] Note that $(\frac{\partial u}{\partial p}-1)$ addresses the fact that for (CONS), p enters both via the utility and that the cutoff type has to be exactly indifferent and that both sides in (25) are negative due to the negative numerators.

Overall, the price-increase must be offset such that $\bar{\psi}_t$ is consistent. But because the consistency requirement reacts less strong to a change in the cutoff than a change in belief (given (24)) *and* the change in price (given (25)) than the rationality requirement, the increase in ω^* necessary to satisfy (CONS) will lead to (RAT) being violated. Hence, the only way that the consumers inference is both consistent and rational is for μ^* to increase following the price increase, and ω^* adjusting in such a fashion as to balance (RAT) and (CONS). Whether this is facilitated by an increase or decrease in ω^* depends in addition on the relative reactiveness of (RAT) and (CONS) to changes in belief and price (see the numerator of (23)).

Implication of $\frac{\partial \omega^*}{\partial p} < 0$. One aspect worthy of discussion is that inference according to Assumption 2 may lead to the induced cutoff ω^* being decreasing in price. In principle, this may seem to suggest that a strategic firm could obtain unbounded flow profits by charging an infinite price if $\frac{\partial \omega^*}{\partial p} < 0$ were to hold for all parameters: A higher price would be accompanied by a higher quantity and hence strictly increase profits. However, this is avoided given that the price is bounded by Assumption 4: Inference needs to satisfy both consistency and rationality and be valid in the sense that the induced inference is consistent with the bounds on ω and θ. To illustrate this further, note that $\frac{\partial \omega^*}{\partial p} < 0$ necessarily implies $\frac{\partial \mu^*}{\partial p} > 0$ (as otherwise (RAT) would be violated). Given that the induced cutoff is bounded from below by $\underline{\omega}$, and the induced belief bounded from above by $\bar{\theta}$, there will hence be a price $\hat{p}(\bar{\psi})$ such that for $p > \hat{p}$, either $\omega^*(\bar{\psi}, p) < \underline{\omega}$ or $\mu^*(\bar{\psi}, p) > \bar{\theta}$ – the inference would need to violate at least one of the bounds to still satisfy (RAT) and (CONS) simultaneously. This provides a clear rationale for avoiding unbounded flow profits: If inference is not *within the bounds* and simultaneously satisfying (CONS) and (RAT), a consumer is unwilling to purchase. This can easily be rationalized by considering that faced with a price-aggregate-rating combination $(\bar{\psi}_t, p_t)$ leading to such an invalid inference being considered as a likely scam – if there is no realistic way this can be rationalized, a sceptic consumer does not purchase.[12] We will revisit and expand on this issue when considering a linearly additive specification in Section 7.

5. Effect of price on review

The core contribution of this paper is to analyze the dynamic pricing incentives generated by the linkage of periods through the rating system. Hence, it is not only relevant to understand how the current price affects the current belief and cutoff, but also how the current price affects the review and hence future profits via the updated rating. Overall, there are three ways in which the price affects the average review in a given period: first, price directly affects the induced average review through the effect of price on the review itself, $\frac{\partial \psi}{\partial p}$. Second, it indirectly affects the expectation of consumers and therefore their review as this conditions on the expectation. Third, the expectation (which is affected by the price) and the price itself determine the marginal consumer who is indifferent between buying and not buying. This selection effect of the price determines who purchases the product and, via the impact of the taste on the review, the induced

[12] The reasoning with the bounds applies irrespective of the sign of $\frac{\partial \omega^*}{\partial p}$; however, it is relevant to constrain flow profits only in the case of $\frac{\partial \omega^*}{\partial p} < 0$.

average review.

The total effect of the price on the review is given by the following equation

$$\frac{d\psi_t}{dp_t} = \underbrace{\frac{\partial\psi_t}{\partial p_t}}_{\substack{<0 \\ \text{direct effect of price} \\ \text{on review}}} + \underbrace{\frac{\partial\psi_t}{\partial\tilde{\omega}_t}}_{\substack{>0 \\ \text{effect of selection} \\ \text{on review}}} \underbrace{\frac{d\omega_t^*}{dp_t}}_{\substack{\gtrless 0 \\ \text{effect of price} \\ \text{on selection}}} + \underbrace{\frac{\partial\psi_t}{\partial\mu_t}}_{\substack{<0 \\ \text{effect of belief} \\ \text{on review}}} \underbrace{\frac{d\mu_t^*}{dp_t}}_{\substack{\gtrless 0 \\ \text{effect of price} \\ \text{on expectation}}} \gtrless 0 \tag{26}$$

Plugging in the previously obtained equations for $\frac{\partial\omega^*}{\partial p}$ and $\frac{\partial\mu^*}{\partial p}$ yields[13]

$$\begin{aligned}
\frac{d\psi_t}{dp_t} &= \frac{\partial\psi_t}{\partial p_t} + \frac{\frac{\partial\psi}{\partial\tilde{\omega}}\frac{\partial\psi}{\partial\theta}\left(\frac{\partial u}{\partial p}-1\right)+\frac{\partial\psi}{\partial p}\left(\frac{\partial\psi}{\partial\mu}\frac{\partial u}{\partial\omega}-\left(\frac{\partial u}{\partial\theta}+\frac{\partial u}{\partial\mu}\right)\frac{\partial\psi}{\partial\tilde{\omega}}\right)}{\left(\frac{\partial u}{\partial\theta}+\frac{\partial u}{\partial\mu}\right)\frac{\partial\psi}{\partial\tilde{\omega}}-\frac{\partial u}{\partial\omega}\left(\frac{\partial\psi}{\partial\theta}+\frac{\partial\psi}{\partial\mu}\right)} \\
&= \frac{\frac{\partial\psi}{\partial\tilde{\omega}}\frac{\partial\psi}{\partial\theta}\left(\frac{\partial u}{\partial p}-1\right)-\frac{\partial\psi}{\partial p}\frac{\partial\psi}{\partial\theta}\frac{\partial u}{\partial\omega}}{\left(\frac{\partial u}{\partial\theta}+\frac{\partial u}{\partial\mu}\right)\frac{\partial\psi}{\partial\tilde{\omega}}-\frac{\partial u}{\partial\omega}\left(\frac{\partial\psi}{\partial\theta}+\frac{\partial\psi}{\partial\mu}\right)} = \frac{\frac{\partial\psi}{\partial\theta}\left[\frac{\partial\psi}{\partial\tilde{\omega}}\left(\frac{\partial u}{\partial p}-1\right)-\frac{\partial\psi}{\partial p}\frac{\partial u}{\partial\omega}\right]}{\left(\frac{\partial u}{\partial\theta}+\frac{\partial u}{\partial\mu}\right)\frac{\partial\psi}{\partial\tilde{\omega}}-\frac{\partial u}{\partial\omega}\left(\frac{\partial\psi}{\partial\theta}+\frac{\partial\psi}{\partial\mu}\right)} \\
&= -\frac{\partial\psi}{\partial\theta}\frac{d\mu^*}{dp}.
\end{aligned} \tag{27}$$

There are three points to take away from this equation: First, the firm is able to manipulate induced reviews strategically via its pricing. This is because inference is conducted using the current price as the best guess for past prices (given the quasi-stationary inference). Second, given that $\frac{\partial\psi}{\partial\theta} > 0$, the effect of a price change on the current inference ($\frac{d\mu^*}{dp}$) always goes in the opposite direction than that of a price change on the induced rating in the current period after purchase. Third, the degree to which the rating is affected depends particularly on the reactiveness of the review function ψ to the actual quality θ. This is notable because one might expect that it is the reactiveness of $\psi(\cdot)$ to the belief itself which affects the induced rating.

To rationalize these observations, consider the two cases where the belief decreases in price ($\frac{d\mu^*}{dp} < 0$) and increases it ($\frac{d\mu^*}{dp} > 0$), respectively. If the inferred belief μ^* decreases (locally) following a price increase, it immediately follows that as (μ^*, ω^*) has to satisfy (CONS) and (RAT), ω^* has to increase significantly. This in turn induces a higher rating in the future due to the selection effect. If μ^* increases, this selection effect is either mitigated (if $\frac{d\mu^*}{dp} > 0$ still holds) or even reversed (if $\frac{d\omega^*}{dp} < 0$). Crucially, in both cases, a change in the belief enters the quality inference both via the effect it has directly ($\frac{\partial\psi}{\partial\mu}$ and $\frac{\partial u}{\partial\mu}$) and via the effect of quality on rating and utility (i.e. via $\frac{\partial\psi}{\partial\theta}$ and $\frac{\partial\psi}{\partial\theta}$) – this is because inference is conducted treating the actual quality as identical to the belief. However, for the actual review it is the true quality θ (which is exogenously fixed) which matters.

This can also be seen in the following manner: Denote $\mu^*(p, \bar{\psi}_t)$ and $\omega^e(\omega^*(p, \bar{\psi}_t)) \equiv \hat{\omega}$. Recall that $\bar{\psi}_t$ is a state variable and cannot be affected by the firm. As such, how the induced review

[13] The below derivation implicitly assumes that partial derivatives are identical irrespective of whether they are evaluated at $(\theta, \omega^e(\omega^*), p, \mu^*)$ (for the actual review) or $(\mu^*, \omega^e(\omega^*), p, \mu^*)$ (for inference). If this is not the case, we have equality only for correct inference, i.e. $\mu^* = \theta$ and an approximation otherwise. Note that the linear additive specification in Section 7 satisfies this.

$\hat{\psi} \equiv \psi(\theta, \hat{\omega}, p, \mu^*)$ changes in response to a price change is identical to how $\hat{\psi} - \bar{\psi}_t$ changes. But as $\bar{\psi}_t = \psi(\mu^*, \hat{\omega}, p, \mu^*)$, we have

$$\begin{aligned} \frac{\partial \hat{\psi} - \bar{\psi}_t}{\partial p} &= \underbrace{\frac{\partial \psi}{\partial \omega}\frac{\partial \hat{\omega}}{\partial p} + \frac{\partial \psi}{\partial p} + \frac{\partial \psi}{\partial \mu}\frac{\partial \mu^*}{\partial p}}_{\text{evaluated at } (\theta, \hat{\omega}, p, \mu^*)} - \underbrace{\left(\frac{\partial \psi}{\partial \theta}\frac{\partial \mu^*}{\partial p} + \frac{\partial \psi}{\partial \omega}\frac{\partial \hat{\omega}}{\partial p} + \frac{\partial \psi}{\partial p} + \frac{\partial \psi}{\partial \mu}\frac{\partial \mu^*}{\partial p}\right)}_{\text{evaluated at } (\mu^*, \hat{\omega}, p, \mu^*)} \\ &\approx -\frac{\partial \psi}{\partial \theta}\frac{\partial \mu^*}{\partial p}, \end{aligned} \tag{28}$$

where we have equality provided that $\mu^* = \theta$, i.e. that inference is correct, or if partial derivatives are identical for the two points of evaluations. The latter is given in the case of a linear additive specification which we discuss subsequently to illustrate our findings.

Given that even after the price change, the same $\bar{\psi}_t$ has to be matched for the inference to be consistent. The only difference between $\bar{\psi}_t$ and the induced review $\hat{\psi}$ is that the review function $\psi(\cdot)$ is evaluated at the inferred quality μ^* to match $\bar{\psi}_t$, while the true quality θ enters it for the induced review. Given that the inference changes following a price change are exactly such that $\bar{\psi}_t$ is still matched when evaluated at inferred quality μ^*, the induced review moves in the opposite direction of the induced belief change – it is still evaluated at the unaffected true quality θ. For example, if $\frac{\partial \mu^*}{\partial p} > 0$, a price change leads to a higher inferred quality. $\bar{\psi}_t$ is still matched given this positive effect on $\psi(\mu^*, \hat{\omega}, p, \mu^*)$, which in turn implies that the effect of the change in inference and price via the induced cutoff, the price itself, and the belief itself is negative. As these are the effects which also carry over to the induced review, while the quality θ remains unchanged, the induced review is negatively affected.

Thus, if the sensitivity of the rating to quality is large, i.e. for $\frac{\partial \psi}{\partial \theta} >> 0$, strategic pricing of the firm can be significantly biased because inference treats the inferred quality as the true one in looking for the fixed-point (μ^*, ω^*). A small change in the inferred μ^* thus leaves scope for larger corresponding changes in ω^* and hence a larger effect on the actual review via the selection effect.

6. Dynamic Pricing

Consider the pricing problem of a firm in a given period t: It aims to maximize the sum of the flow profits and discounted future profits. Future profits in turn depend on two state variables: The period, t, which cannot be affected by the firm, and the aggregate rating, which can be strategically influenced. Denote the value of future profits as $V(\bar{\psi}_{t+1})$ and we obtain for the firm's maximization problem:

$$V_t(\bar{\psi}_t) = \max_{p_t} p_t \cdot q_t(p_t) + \delta V_{t+1}(\bar{\psi}_{t+1}), \tag{29}$$

where

$$q_t = \frac{\bar{\omega} - \omega^*(p_t, \bar{\psi}_t)}{\bar{\omega} - \underline{\omega}} \tag{30}$$

$$\bar{\psi}_{t+1} = \rho^t(\bar{\psi}_t, \psi_t) \tag{31}$$

$$\psi_t = \psi(\theta, \omega^e(\omega^*(p_t, \bar{\psi}_t)), p_t, \mu^*(p_t, \bar{\psi}_t)). \tag{32}$$

The first order effect of a price change is hence given by

$$\frac{\bar{\omega} - \omega^*(p, \bar{\psi}_t)}{\bar{\omega} - \underline{\omega}} + p \cdot -\frac{\partial \omega^*}{\partial p}\frac{1}{\bar{\omega} - \underline{\omega}} + \delta V'_{t+1}(\bar{\psi}_{t+1}) \cdot \frac{\partial \rho^t}{\partial \psi_t} \cdot \underbrace{-\frac{\partial \psi}{\partial \theta}\frac{d\mu^*}{dp}}_{\frac{\partial \psi_t}{\partial p_t}}. \tag{33}$$

We thus have three effects: First, flow profits are affected by the increase in the price. This is captured by $\frac{\bar{\omega} - \omega^*(p, \bar{\psi}_t)}{\bar{\omega} - \underline{\omega}} = q(p_t, \bar{\psi}_t)$. Second, flow profits may either decrease or increase depending on the impact the change in price has on the induced cutoff and hence the quantity. This is reflected by $p \cdot \left(-\frac{\partial \omega^*}{\partial p}\frac{1}{\bar{\omega} - \underline{\omega}}\right) = p \cdot \frac{\partial q}{\partial p}$. It is possible that a price increase has a (locally) unambiguously positive effect on flow profits if the price increase is associated with an increase in the quantity due to the inference.[14] Third, 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 ($\frac{\partial \psi_t}{\partial p} = -\frac{\partial \psi}{\partial \theta}\frac{\partial \mu^*}{\partial p}$).

It follows immediately that over time, the incentive to price strategically decreases as the responsiveness of the future profits to the aggregate rating decreases over time (as $t \to T$, less periods remain in which the high aggregate rating can be milked) *provided that the aggregate rating becomes less sensitive over time*. Moreover, the way the aggregate rating is computed matters for the incentives to price away from the myopic optimum. For example, if the aggregate rating is simply the average over all past ratings, $\bar{\psi}_t = \frac{\sum_{i=0}^t \psi_t}{t}$, we have $\frac{\partial \rho^t}{\partial \psi_t} \overset{t \to \infty}{\to} 0$ and the incentive to price away from the myopic optimum lessens over time. However, it is not immediately clear whether this speaks for or against rating systems which weigh recent reviews more heavily than older ones. While this would ensure that the responsiveness of the aggregate rating to the induced review is high, leading to persistent incentives to price away from the myopic optimum, it also affects the rating- and price-paths (along with the induced inference) such that a clear evaluation is not possible. We will illustrate this in more detail with numerical examples in Section 8.

Finally, if incentives to deviate from myopically optimal prices obtain, they may come both in the form of incentives to increase the price (if $\frac{\partial \psi_t}{\partial p} > 0$) or decrease the price ($\frac{\partial \psi_t}{\partial p} < 0$), which in turn depends on the specific sensitivity of utility and review functions to their inputs. We illustrate different cases in the linear additive specification below.

[14] Recall from the previous discussion that this does not imply that flow profits are unbounded – given that inference is required to be valid, there still remains an upper bound on myopic profits even in this case.

7. Linear Additive Specification

In this section, we consider a simple linear additive specification for both the utility and the review function. This parametrization allows us to discuss various implications of the model in reasonable generality while still giving a more detailed illustration of the implications gained from the generalized framework in the previous sections. Specifically, we consider a (gross) utility function where consumers only care about the true quality of the good, θ, and their taste for it. Moreover, the two components are additively separable. We hence have

$$u(\theta, \omega, p, \mu) = \alpha\theta + \beta\omega, \tag{34}$$

where α , $\beta > 0$ parametrize the relative relevance of quality and taste. Regarding the review a consumer gives, we let the review resemble the gross utility, but also be negatively affected by the price paid for the good, as well as the initial quality belief. We again impose an additively separable specification and obtain

$$\psi(\theta, \omega, p, \mu) = \alpha\theta + \beta\omega - \kappa p - \gamma\mu, \tag{35}$$

where $\kappa, \gamma \geq 0$. Given the restrictions on $\alpha, \beta, \gamma, \kappa$, it is clear that Assumption 1 is satisfied. Moreover, we impose that actual quality matters more than belief for the review, i.e. $\alpha > \gamma$, in line with Assumption 3.[15] Moreover, we let $\omega \sim U[\underline{\omega}, \bar{\omega}]$ which implies $\frac{\partial \omega^e(\omega)}{\partial \omega} = \frac{1}{2}$. For now ignoring the issue of price-bounds (see Assumption 4; we revisit this later on), this implies that the steepness condition (STP) in Assumption 6 simplifies to

$$\begin{aligned} \frac{\frac{\partial \psi}{\partial \tilde{\omega}}}{\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu}} = \frac{\frac{1}{2}\beta}{\alpha - \gamma} \quad &\gtreqless \quad \frac{\beta}{\alpha} = \frac{\frac{\partial u}{\partial \omega}}{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}} \\ \Leftrightarrow \frac{1}{2}\alpha - \gamma \quad &\lesseqgtr \quad 0. \end{aligned} \tag{36}$$

Note that if a unique inference exists for the price-rating-pair $(\bar{\psi}_t, p_t)$, we know that the comparative statics with respect to price and rating are given by (16) and (21) and we have

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial \bar{\psi}} \\ \frac{\partial \omega^*}{\partial \bar{\psi}} \end{pmatrix} = \begin{pmatrix} \frac{-\frac{\partial u}{\partial \omega}}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \\ \frac{\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu}}{(\frac{\partial u}{\partial \theta} + \frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \tilde{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta} + \frac{\partial \psi}{\partial \mu})} \end{pmatrix} = \begin{pmatrix} -\frac{1}{\gamma - \frac{1}{2}\alpha} \\ \frac{\alpha}{\beta(\gamma - \frac{1}{2}\alpha)} \end{pmatrix},$$

[15] For the utility, this is satisfied by our specification as $\frac{\partial u}{\partial \theta} = \alpha > 0 = \frac{\partial u}{\partial \mu}$.

as well as

$$\begin{pmatrix} \frac{\partial \mu^*}{\partial p} \\ \frac{\partial \omega^*}{\partial p} \end{pmatrix} = \begin{pmatrix} \frac{\frac{\partial u}{\partial \omega}\frac{\partial \psi}{\partial p} - \frac{\partial \psi}{\partial \bar{\omega}}(\frac{\partial u}{\partial p}-1)}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \bar{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} \\ \frac{(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})(\frac{\partial u}{\partial p}-1)-(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial p}}{(\frac{\partial u}{\partial \theta}+\frac{\partial u}{\partial \mu})\frac{\partial \psi}{\partial \bar{\omega}} - \frac{\partial u}{\partial \omega}(\frac{\partial \psi}{\partial \theta}+\frac{\partial \psi}{\partial \mu})} \end{pmatrix} = \begin{pmatrix} \frac{\frac{1}{2}-\kappa}{\gamma-\frac{1}{2}\alpha} \\ \frac{\gamma-\alpha(1-\kappa)}{\beta(\gamma-\frac{1}{2}\alpha)} \end{pmatrix}.$$

7.1. Explicit Inference

We can also characterize the inference explicitly by solving the equation system

$$\begin{aligned} \bar{\psi}_t = \quad & (\alpha-\gamma)\mu^* + \beta\tfrac{\bar{\omega}+\omega^*}{2} - \kappa p && = \psi(\mu^*, \omega^e(\omega^*), p, \mu^*) && (37) \\ p = \quad & \alpha\mu^* + \beta\omega^* && = u(\mu^*, \omega^*, p, \mu^*). && (38) \end{aligned}$$

This gives

$$\omega^* = \frac{2\alpha(\bar{\psi}_t - p) + 2p(\gamma+\alpha\kappa) - \alpha\beta\bar{\omega}}{\beta(2\gamma-\alpha)}, \; \mu^* = \frac{(1-2\kappa)p - 2\bar{\psi}_t + \beta\bar{\omega}}{2\gamma-\alpha} \tag{39}$$

which is consistent with

$$\frac{\partial \omega^*}{\partial \bar{\psi}} = \frac{\alpha}{\beta(\gamma-\frac{1}{2}\alpha)}, \; \frac{\partial \omega^*}{\partial p} = \frac{\gamma-(1-\kappa)\alpha}{\beta(\gamma-\frac{1}{2}\alpha)}$$

and

$$\frac{\partial \mu^*}{\partial \bar{\psi}} = -\frac{1}{\gamma-\frac{1}{2}\alpha}, \; \frac{\partial \mu^*}{\partial p} = \frac{\frac{1}{2}-\kappa}{\gamma-\frac{1}{2}\alpha}$$

as the general analysis suggests. Regarding the – until now omitted – price-bounds, we can obtain them implicitly from (39): We know that $\omega^* \in [\underline{\omega}, \bar{\omega}]$ and $\mu^* \in [\underline{\theta}, \bar{\theta}]$ have to be satisfied for inference to be sensible. Depending on the signs of $\frac{\partial \mu^*}{\partial p}$ and $\frac{\partial \omega^*}{\partial p}$, solving $\omega^* = \underline{\omega}$ ($\omega^* = \bar{\omega}$) for p and $\underline{\theta} = \bar{\omega}$ ($\mu^* = \underline{\theta}$) for p, we obtain the lower and upper bounds, respectively. Consider for example the case of $\frac{\partial \omega^*}{\partial p} < 0$ and $\frac{\partial \mu^*}{\partial p} > 0$. Then the upper bounds on p are obtained from $\omega^* = \underline{\omega}$ and $\mu^* = \bar{\theta}$ and we have

$$\underline{\omega} = \omega^* = \frac{2\alpha(\bar{\psi}_t - p) + 2p(\gamma+\alpha\kappa) - \alpha\beta\bar{\omega}}{\beta(2\gamma-\alpha)} \Leftrightarrow p = \frac{2\beta\gamma\underline{\omega} - \alpha\beta(\underline{\omega}+\bar{\omega}) - 2\alpha\bar{\psi}_t}{2(\gamma-(1-\kappa)\alpha)}$$

and

$$\bar{\theta} = \mu^* = \frac{(1-2\kappa)p - 2\bar{\psi}_t + \beta\bar{\omega}}{2\gamma-\alpha} \Leftrightarrow p = \frac{2\bar{\psi}_t - \beta\bar{\omega} - 2(\alpha-\gamma)\bar{\theta}}{(1-2\kappa)}$$

which implies that for inference to be valid, p must satisfy

$$p \geq \max\left\{\frac{2\beta\gamma\underline{\omega} - \alpha\beta(\underline{\omega}+\bar{\omega}) - 2\alpha\bar{\psi}_t}{2(\gamma-(1-\kappa)\alpha)}, \frac{2\bar{\psi}_t - \beta\bar{\omega} - 2(\alpha-\gamma)\bar{\theta}}{(1-2\kappa)}\right\} \tag{40}$$

and for the upper bound that

$$p \leq \min \left\{ \frac{2\beta\gamma\underline{\omega} - 2\alpha\beta\bar{\omega} - 2\alpha\bar{\psi}_t}{2(\gamma - (1-\kappa)\alpha)}, \frac{2\bar{\psi}_t - \beta\bar{\omega} - 2(\alpha-\gamma)\underline{\theta}}{(1-2\kappa)} \right\}. \tag{41}$$

7.2. Comparative Statics & Effect of Price on Induced Review

Consider first $\frac{\partial\omega^*}{\partial\bar{\psi}}$ and $\frac{\partial\mu^*}{\partial\bar{\psi}}$ as given by

$$\frac{\partial\omega^*}{\partial\bar{\psi}} = \frac{\alpha}{\beta(\gamma - \frac{1}{2}\alpha)}, \frac{\partial\mu^*}{\partial\bar{\psi}} = -\frac{1}{\gamma - \frac{1}{2}\alpha}.$$

Note that given the assumptions, the signs are uniquely determined by the sign of $(\gamma - \frac{1}{2}\alpha)$: If the belief matters relatively little compared to quality ($\gamma < \frac{1}{2}\alpha$), an increase in the aggregate rating has a negative effect on the induced cutoff and a positive effect on the induced belief. This is because (CONS) is evaluated at $\theta = \mu = \mu^*$: Inference is conducted by treating the belief as equal to the quality and hence $(\alpha - \gamma)$ gives the net effect of an increase in the belief of the rating. If this is high (i.e. γ low), a higher rating can only be rationalized by a higher belief (while keeping (RAT) satisfied). By contrast, if this net effect is low (γ high relative to α), a higher rating can only be rationalized by having come from people with higher tastes – in this case, we would have $\frac{\partial\omega^*}{\partial\bar{\psi}} > 0$, where $\frac{\partial\mu^*}{\partial\bar{\psi}} < 0$ would follow to keep (RAT) satisfied.[16]

A similar natural explanation arises for $\frac{\partial\omega^*}{\partial p}$ and $\frac{\partial\mu^*}{\partial p}$ given by

$$\frac{\partial\omega^*}{\partial p} = \frac{\gamma - (1-\kappa)\alpha}{\beta(\gamma - \frac{1}{2}\alpha)}, \frac{\partial\mu^*}{\partial p} = \frac{\frac{1}{2} - \kappa}{\gamma - \frac{1}{2}\alpha}.$$

If the review function is relatively sensitive to belief compared to actual quality ($\gamma > \frac{1}{2}\alpha$), an increase in the price increases the induced belief if and only if the review function is not very sensitive to the price ($\kappa < \frac{1}{2}$). If the review function were very sensitive to price, the inferred quality would have to be lower as inferred quality overall has a small effect and a high rating can be reached with a high quality/belief given a high price only if the price-effect on reviews is small. By contrast, if the sensitivity to the belief is low, an (anticipated) higher quality has a strong effect on the review/rating and quality increases in price only if the review function is very sensitive to price – a given rating can only be rationalized at a high price if the quality is also high.

Finally, we can assess the impact of a change in the price on the induced review – this is crucial for assessing whether a strategic firm wishes to price higher or lower compared to the myopically

[16] Note that (RAT) is unaffected by the change in $\bar{\psi}$.

optimal price which maximizes flow profits. Following (27), we obtain[17]

$$\frac{\partial \psi_t}{\partial p_t} = -\frac{\partial \psi}{\partial \theta} \cdot \frac{\partial \mu^*}{\partial p_t} = -\alpha \cdot \frac{\frac{1}{2} - \kappa}{\gamma - \frac{1}{2}\alpha}. \tag{42}$$

As previously established, the induced review is affected in the opposite direction of the induced belief – by (CONS), the review evaluated at the induced belief as the quality still matches $\bar{\psi}_t$, while for the induced review all other inputs are identically affected, but the true quality θ matters. If the belief increases (decreases), the effect of the changes via the direct effects of the changed cutoff, price, and belief following a price increase is negative (positive), as is thus the effect on the induced review. A price increase hence increases the induced review if and only if the belief is negatively affected, i.e. iff one of the following conditions holds.

$$\textbf{Case 1:}\ \frac{1}{2} - \kappa > 0 \wedge \gamma - \frac{1}{2}\alpha < 0. \tag{43}$$

$$\textbf{Case 2:}\ \frac{1}{2} - \kappa < 0 \wedge \gamma - \frac{1}{2}\alpha > 0. \tag{44}$$

There are hence two possible cases such that the firm has an incentive to price *above* the myopically optimal price to positively affect future profits via the induced rating. It does so if the price impact itself on a review is relatively minor ($\kappa < \frac{1}{2}$), while the quality aspect entering into a review is significantly more important than the belief ($\gamma < \frac{1}{2}\alpha$, **Case 1**). Alternatively, this is the case if there is a sizeable price impact on the review ($\kappa > \frac{1}{2}$) and a large impact of the belief on the review compared to the actual quality ($\gamma > \frac{1}{2}\alpha$, **Case 2**).

By contrast, if the prize impact is sizeable ($\kappa > \frac{1}{2}$) but the belief itself is of minor importance for reviews compared to the actual quality ($\gamma < \frac{1}{2}\alpha$), the reverse holds: The induced belief is positively affected following a price increase, which leads to a negative effect on the induced review. The firm thus has an incentive to underprice compared to the myopic optimum to positively affect future profits. We summarize as follows.

Observation 1 (Strategic pricing in the linear additive specification) *Consider the linear additive specification of utility and review functions given by* (34) *and* (35)*. A strategic firm has an incentive to charge a higher price compared to the myopic optimum if and only if:*

(i) The review is relatively insensitive to the purchase price ($\kappa < \frac{1}{2}$), and relatively insensitive to expectation compared to actual quality ($\gamma < \frac{1}{2}\alpha$)

or

(ii) the review is relatively sensitive to the purchase price ($\kappa > \frac{1}{2}$), and relatively sensitive to expectation compared to actual quality ($\gamma > \frac{1}{2}\alpha$).

Otherwise, the firm has an incentive to charge a lower price compared to the myopic optimum.

[17] Note that given the additive linear specification, (27) holds with equality as first-order effects are constant over the entire support.

8. Effectiveness of rating systems – numerical implementation

In this section, we outline observations gathered from numerically implementing the model in Mathematica. We stick to the linear additive specification as in the previous section and in particular parametrize the review and utility functions as follows:

$$\begin{aligned} u(\theta,\omega,p,\mu) &= 2\theta+1.5\omega \\ \psi(\theta,\omega,p,\mu) &= 2\theta+1.5\omega-0.2p-0.5\mu, \end{aligned}$$

that is, we choose $\alpha=2$, $\beta=1.5$, $\kappa=0.2$, and $\gamma=0.5$. Moreover, we let the taste be uniformly distributed, $\omega\sim U[-5,5]$, and consider a 12-period-version of the model: $T=12$. Throughout, we consider two different specifications for the rating function: First, we consider the case where the aggregate rating exhibits 'full memory', that is, the aggregate rating is simply the average of all past reviews. Second, we let the aggregate rating only reflect recent reviews and in particular in our implementation the average of the previous two ratings – $\bar{\psi}_t=\frac{\psi_{t-2}+\psi_{t-1}}{2}$. In our implementation, the prior of the consumers in the initial period is characterized by some initial rating $\bar{\psi}_0$

Strategic Firm, Full Memory Rating and Different Qualities

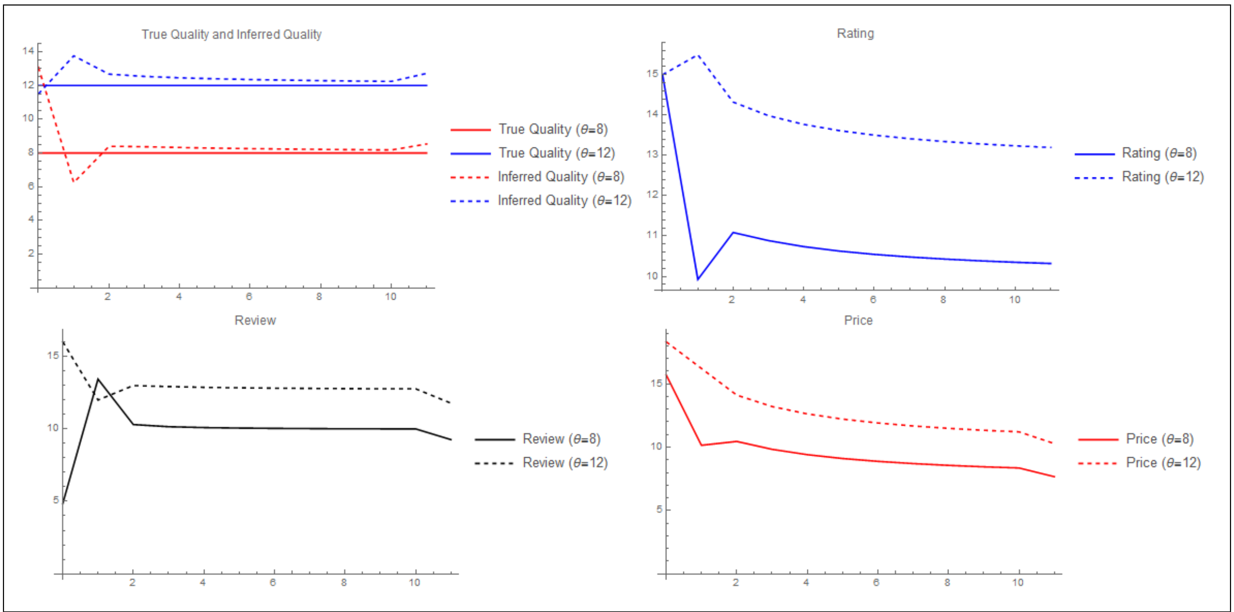

Figure 1: Strategic firm, $\bar{\psi}_t=\frac{\sum_{i=0}^{t-1}\psi_i}{t}$, $\theta=8$ vs. $\theta=12$

Figure 1 depicts the comparison between two firms of different qualities ($\theta_1=8$ and $\theta_2=12$) faced with the same initial condition characterized by the initial rating $\bar{\psi}_0$. For both qualites, the inferred qualities quickly converge close to the actual qualities with the transition period characterized by fluctuations in price. However, the path until period 3 markedly differs: For the firm with lower quality good ($\theta=8$), the initial rating $\bar{\psi}_0$ is such that the firm immediately exploits the (too high) rating. This is accompanied by a very low first-period review and associated sharp drop in the rating. In fact, the exploitation incentive is sufficiently strong initially that the firm needs to build its rating back up in the next period. By contrast, the firm with the higher quality good ($\theta = 12$) builds up its rating even further before starting to exploit it. Finally, note that despite the near-constant rating after period $t = 3$, the firms gradually lower their prices – this is precisely to keep the rating near-constant and balances flow profits with future profit considerations. Note also that in the final period, there is a price-drop: Given the parametrization, we have $\frac{\partial \psi_t}{\partial p_t} = -\frac{\partial \psi}{\partial \theta}\frac{\partial \mu_t^*}{\partial p_t} = -\alpha \cdot \frac{\frac{1}{2}-\kappa}{\gamma - \frac{1}{2}\alpha} = 1.2 > 0$ and hence an incentive to price above the myopic optimum when taking future profits into account.

8.1. Comparison strategic and myopic firm behavior under full memory

In Figure 2 we fix the quality at $\theta = 12$ but plot the resulting paths for prices, reviews, inferred qualities and ratings contrasting the case where the firm is indeed forward-looking with the case where it prices myopically. As previously discussed, given our parametrization we expect the strategic firm to price above the myopic one in the first period when both are faced with identical initial conditions. Moreover, while the strategic firm's pricing is such that inferred qualities are close to the actual quality, there still remains a small difference, while this is not the case for the myopic firm – prices, inferred qualities, reviews and ratings quickly stabilize at a quasi-stationary level. This speaks for the success of aggregate ratings reflecting all past reviews in transmitting information about the product's quality if the firm does not take into account the long-term impact of its pricing.

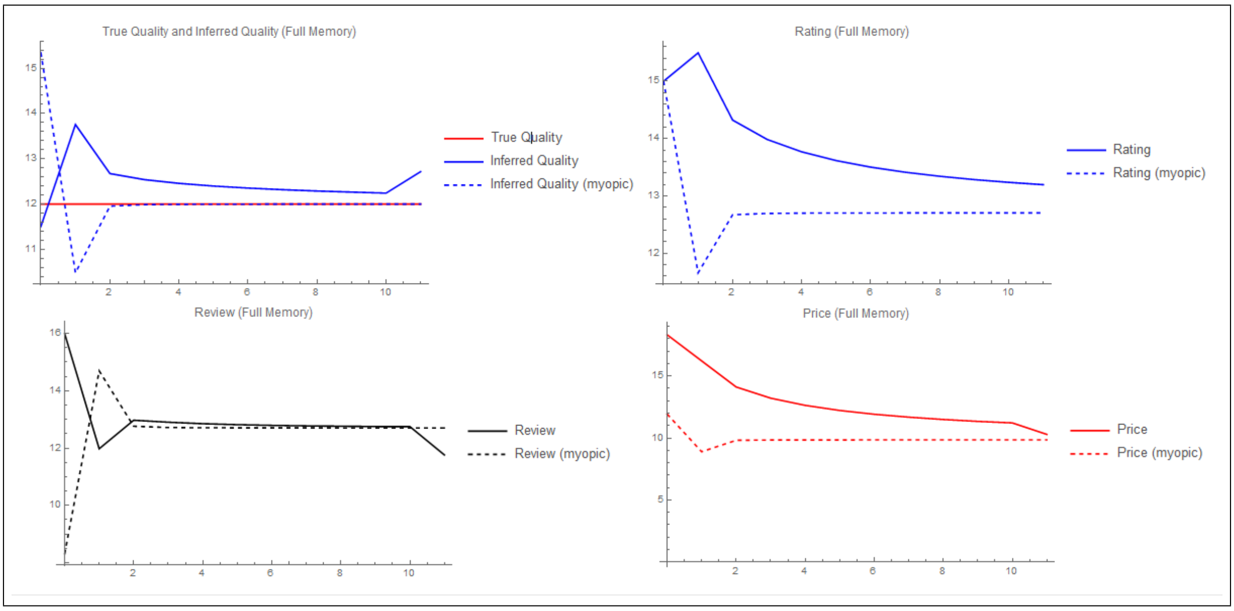

Figure 2: Strategic vs. myopic pricing with $\bar{\psi}_t = \frac{\sum_{i=0}^{t-1}\psi_i}{t}$

8.2. Comparison strategic and myopic firm behavior under limited memory

We repeat this exericse for the case of aggregate ratings displaying limited memory, that is, only reflecting the average of the two most recent reviews. The resulting plots are depicted in Figure 3.

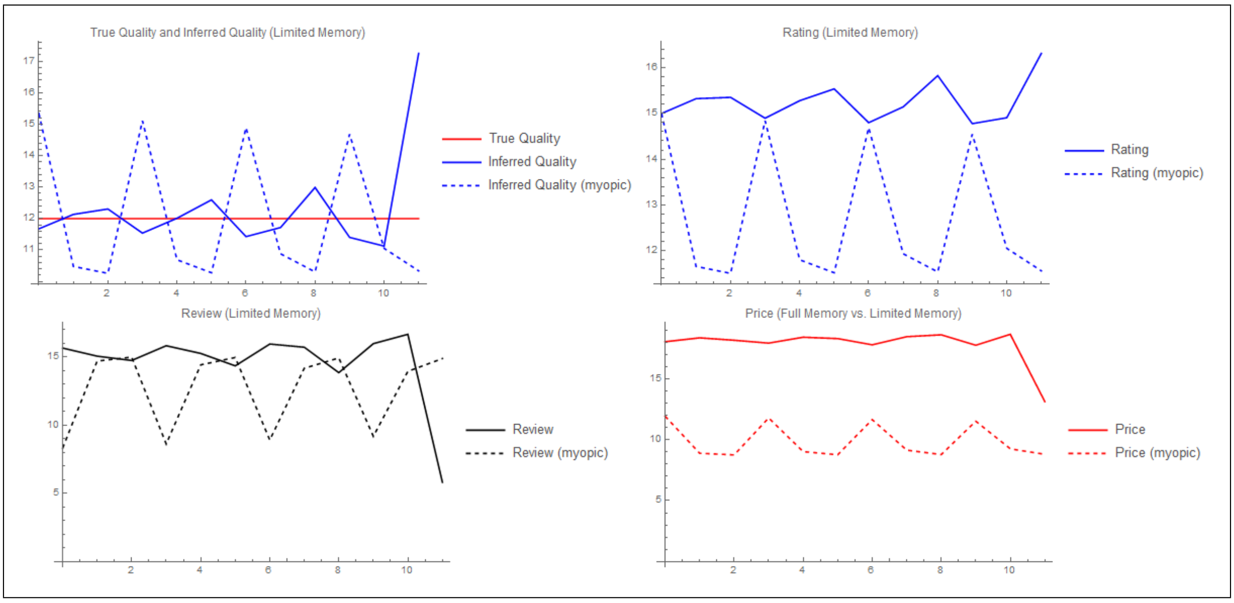

Figure 3: Strategic vs. myopic pricing with $\bar{\psi}_t = \frac{\psi_{t-2}+\psi_{t-1}}{2}$

There are several things worth noting: First, the price level charged by the strategic firm is markedly higher than that of the myopic one, as expected. Second, the limited memory induces fluctuations in the quality inference for both types of firms. However, while the fluctuations (slowly) shrink over time given the myopic pricing, which again reinforces that ratings serve as a means of information transmission if firms are non-strategic, they actually grow larger as time goes on if the firm is strategic: It deliberately enters a cycle of reputation-build-up and reputation-exploitation and even amplifies the effect of this cycle over time.

8.3. Comparison full vs. limited memory

The previous observations are reinforced when explicitly contrasting the cases of full and limited memory for the aggregate rating. The resulting plots can be seen in Figure 4 for both strategic (Panel (a)) and myopic (Panel(b)) firms. An important takeaway comes from the case of strategic firms: While, as previously discussed, a full memory leads to a reasonably quick stabilization of ratings, reviews, and quality inference close to the actual quality (along with a slowly decreasing price-path), there is a marked price level effect from ratings exhibiting limited memory. This, together with the previous observation of intentionally inducing reputation cycles given the limited memory, suggests that the recent emphasis of e.g. Amazon on 'improving' its rating systems by putting more weight on more recent reviews may not be in the consumer's interest after all. We intend to look into this and other considerations in more detail in the future, also focusing on the resulting consumer welfare which is not explicitly assessed yet.

In Appendix B we prove some of these observations analytically for a special case of the linearly additive model with $\kappa = \gamma = 0$ and full memory for the rating system, i.e. where the review reflects the gross utility of a consumer, and the rating reflects all past reviews. In that specification, a strategic firm always charges a markup over the myopically optimal price. Moreover, under myopic pricing, the rating system would be effective, that is, the quality inference is correct as t grows large. By contrast, a strategic firm always biases the inference, that is, even if in a given period quality inference is correct ($\mu_t^* = \theta$), the firm prices such that future inference will again be biased. The degree of this bias, however, vanishes as t grows large as the sensitivity of the aggregate rating declines.

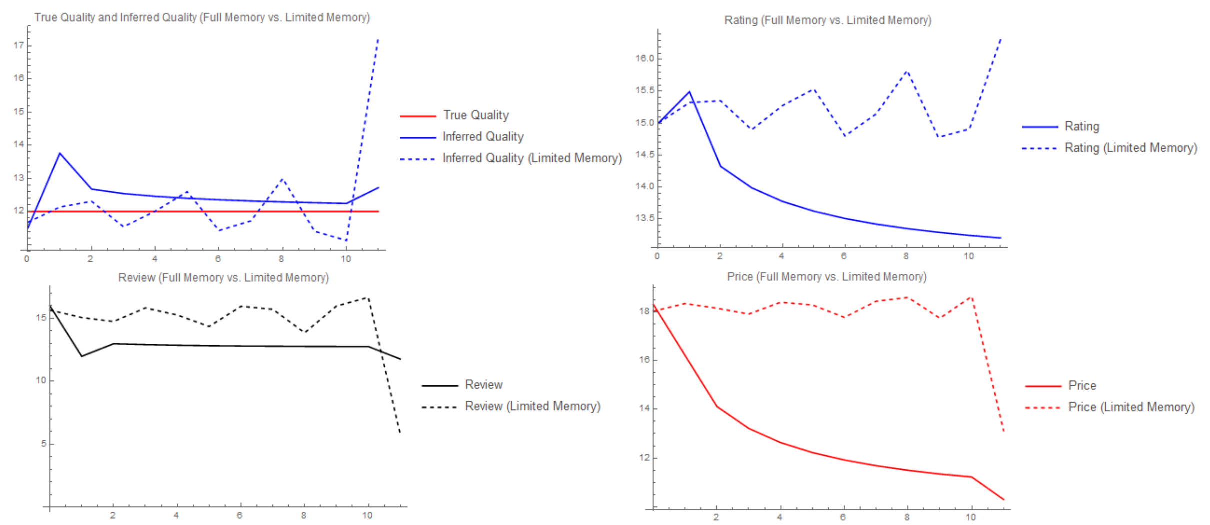

(a) Strategic firm

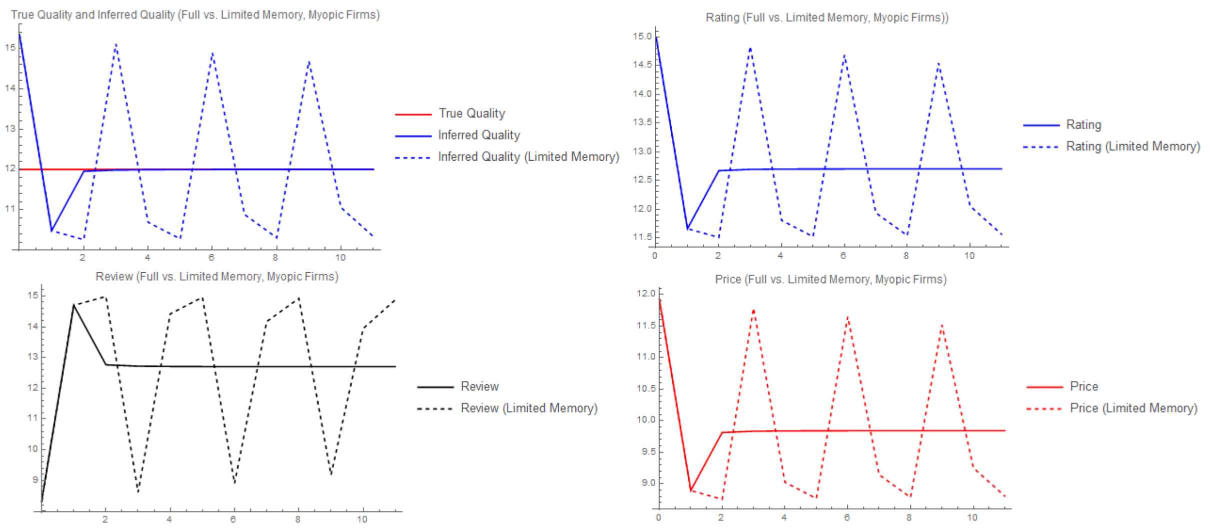

(b) Myopic firm

Figure 4: Comparison full & limited memory

9. Conclusion and outlook

This paper develops a framework for assessing strategic pricing incentives of firms in online markets. The price they charge is used as a selection device for purchasing consumers' characteristics which in turn affect the reviews they write and therefore future profits. We propose an inference method whereby consumers use the current price as the best predictor for past prices and look for a combination of purchasing consumers' tastes and quality which matches individual rationality conditions and the aggregate rating. For this inference method, we provide sufficient conditions such that it is unique and show how the firm's pricing is affected once it takes into account the impact on future profits via reviews.

We show that the effect of a price increase on the inferred quality and induced review depends on fundamentals: In the case of a linearly additive specification of review and utility functions, the effect of a price increase on the inferred quality and induced review go in opposite directions. For example, a price increase negatively affects the inferred quality but positively affects the induced review and thus future profits provided that the review itself is neither too sensitive to the price, nor to the induced belief. A strategic firm taking future profits into account in its price-setting thus has an incentive to overprice relative to the myopic optimum. By doing so, purchasing consumers' characteristics lead to a higher review and hence rating which the firm can exploit in the future.

We implement a variation of the linear additive specification numerically and document some interesting patterns. For example, the recent push by e.g. Amazon.com to have ratings reflect more recent reviews instead of all past reviews, ostensibly to improve the accuracy of the rating system, may actually incentivize strategic firms to engage in reputation build-up and exploitation cycles. Crucially, while inference with 'full memory' rating systems converges to the true quality, it fluctuates around the true quality when rating systems have 'limited memory'. Moreover, the fluctuations seem to amplify over time.

While there remains a lot of work to be done on generalizing these findings and providing sufficient conditions such that they hold (analytically) for generic utility and review functions, we think that our approach provides a toolbox for further research into the strategic price-setting of firms on online platforms.

References

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.

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.

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

——— (2010): "The dynamics of seller reputation: Evidence from ebay," *The Journal of Industrial Economics*, 58, 54–78.

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.*

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

Marinovic, I., A. Skrzypacz, and F. Varas (2015): "Dynamic Certification and Reputation for Quality," *working paper.*

Martin, S. and S. Shelegia (2015): "Under-promise and Over-deliver," *working paper.*

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.

Wolinsky, A. (1983): "Prices as signals of product quality," *The Review of Economic Studies*, 647–658.

A. Static Game Equilibrium with Price-Signaling

A.1. Monopoly Price under Perfect Information

Consider the static one-period game (i.e. no ratings) and a firm of a given quality θ. We assume that in addition to Assumption 1, a consumer's utility is weakly decreasing in p and μ, i.e. that a higher price or higher quality expectation weakly decreases the gross utility.[18] Formally, we assume that

$$\frac{\partial u}{\partial p} \leq 0\,, \;\frac{\partial u}{\partial \mu} \leq 0.$$

To characterize the monopoly price under perfect information, suppose that the quality θ is publicly known. Hence, a consumer's gross utility is $u(\theta, \omega_i, p, \theta)$ where $\omega_i \sim U[\underline{\omega}, \bar{\omega}]$ with $[\underline{\omega}, \bar{\omega}] \equiv \Omega$. A consumer buys iff

$$u(\theta, \omega_i, p, \theta) \geq p \Leftrightarrow \omega_i \geq \tilde{\omega}(\theta, p) \text{ where } u(\theta, \tilde{\omega}(\theta, p), p, \theta) = p$$

implicitly defines $\tilde{\omega}$ as a function of θ and p. Moreover, the assumptions ensure that $\tilde{\omega}(\theta, p)$ is unique, decreasing in θ and increasing and convex in p. This induces a quantity

$$q(\theta, p) = \frac{\bar{\omega} - \tilde{\omega}(\theta, p)}{\bar{\omega} - \underline{\omega}}$$

which is increasing in θ and decreasing and concave in p. As such, the monopoly price of a firm of quality θ under symmetric information is derived from

$$\max_p p \cdot q(\theta, p)$$

and is unique with $p^m(\theta)$ increasing in θ.

Throughout the subsequent analysis of the static game, we restrict attention to the most pessimistic off-path beliefs, i.e. upon observing a price $\tilde{p}$ not featured in equilibrium, consumers believe the quality to be $\underline{\theta}$: $\tilde{\theta}(\tilde{p}) = \underline{\theta}$.

[18] The idea being that a consumer is less likely to be satisfied the higher the price paid for the good, or the higher the initial quality expectation.

A.2. Static Game – Pooling Equilibria

First note that in any pooling equilibrium of the static game, all firms charge the same price p^p and the associated quality inference is $\tilde{\theta}(p^p) = E[\theta] \equiv \bar{\theta}$. Next, note that given that production costs are independent of θ (and normalized to 0), profits at any price $\tilde{p}$ are independent of θ and only depend on the quality inference.

It follows immediately that pooling can be sustained in equilibrium at any price p^p such that all firms earn at least as much as the monopoly profits of the lowest-quality firm under symmetric information. Crucially, this allows for pooling at the monopoly price under symmetric information charged by the lowest type firm, in which case profits are strictly larger than the monopoly profits of the firm with quality $\underline{\theta}$ under symmetric information.

Lemma 3 (Pooling Equilibria of the Static Game) *Any p^p such that*

$$\pi(p^p) \geq p^m(\underline{\theta}) \cdot q(\underline{\theta}, p^m(\underline{\theta})) \equiv \pi^m(\underline{\theta})$$

can be sustained as a pooling equilibrium. $p^m(\underline{\theta})$ is one price at which pooling can be sustained.

Proof.

Any deviation to $\tilde{p} \neq p^m(\underline{\theta})$ gives quality inference $\tilde{\theta}(\tilde{p}) = \underline{\theta}$ and hence yields profits $\pi(\tilde{p}) = \tilde{p} \cdot q(\underline{\theta}, \tilde{p}) \leq p^m(\underline{\theta})$ so no firm wants to deviate. If $p^p = p^m(\underline{\theta})$, $\bar{\theta} > \underline{\theta}$ ensures that $\tilde{\theta}(p^m(\underline{\theta})) > \underline{\theta}$ and hence that $\pi(p^m(\underline{\theta})) > \pi^m(\underline{\theta})$. ■

A.3. Static Game – Separating Equilibria

In any separating equilibrium, each firm is uniquely identified through its price. Denote by $p^s(\theta)$ the equilibrium price schedule (which depends on θ) and by $\pi^s(\theta) = p^s(\theta) \cdot q(\theta, p^s(\theta))$ the equilibrium profits. Again recall that a firm's profit only depends on the price it charges and thus the inferred θ by consumers, but not directly on its own quality. This directly implies that all firms have to make identical profits.

Given the most pessimistic off-path beliefs, it is straightforward that the lowest quality firm ($\underline{\theta}$) has to exactly charge its monopoly price: As it is identified as the lowest-quality firm in equilibrium, even most pessimistic off-path beliefs would cause it to deviate to $p^m(\underline{\theta})$ – if $p^m(\underline{\theta})$ is not part of the equilibrium pricing of any firm, this yields monopoly profits, while if $p^m(\underline{\theta})$ is associated with a quality $\tilde{\theta} > \underline{\theta}$, profits are even higher. Both cases would imply a profitable deviation.

The fact that firms' profits have to be identical also directly pins down all other prices: It has to hold that for all $\theta \in (\underline{\theta}, \bar{\theta}]$:

$$p^s(\theta) \cdot q(\theta, p^s(\theta)) = \pi^m(\underline{\theta}) < p^m(\theta). \tag{45}$$

This is intuitive: Any firm with $\theta \neq \underline{\theta}$ has to charge a price in equilibrium distorted from its monopoly price under symmetric information to prevent the lowest type firm to imitate it. Because of the previously derived properties of $q(\theta)$ under symmetric information (which prevails in equilibrium), there are two candidate prices $\underline{p}^s(\theta)$ and $\bar{p}^s(\theta)$ with $\underline{p}^s(\theta) < p^m(\theta) < \bar{p}^s(\theta)$ such that the profit conditon (45) is satisfied.

Lemma 4 (Separating Equilibria of Static Game) *There are infinitely many separating equilibria. In each separating equilibrium, the lowest quality firm charges its monopoly price $p^m(\underline{\theta})$, while all other firms of quality $\theta \neq \underline{\theta}$ charge either $\bar{p}^s(\theta)$ or $\underline{p}^s(\theta)$.*

A.4. Pooling vs. Separating

Note that pooling profits are bounded from below to be $p^m(\underline{\theta})$ for each firm (and strictly higher for all but one of the equilibria), while in a separating firm each firm receives $p^m(\underline{\theta})$.

B. Results with consumers reporting linear additive gross utility

Let the utility and review functions be given by

$$u(\theta, \omega, p, \mu) == \alpha\theta + \beta\omega = \psi(\theta, \omega, p, \mu). \tag{46}$$

It follows immediately that, as the review and hence rating is independent of price and inferred quality μ_i, price has no direct effect on rating but still indirectly affects it via the composition of purchasing consumers. In terms of the aggregate rating, it is simply computed as the weighted average of the previous aggregate rating and the current average review:

$$\bar{\psi}_t = \frac{t\bar{\psi}_{t-1} + \psi_{t-1}}{t}, \tag{47}$$

We let the initial rating $\bar{\psi}_0 = \bar{\mu}$ be equal to the initial belief of consumers: In the first period, consumers do not observe a rating but have access to the shared prior belief about the product's quality $\bar{\mu}$. It should be noted that the first period quality inference is still affected by the initial price p_0 set by the firm – nonetheless, $\bar{\psi}_0$ matters and is positively related to the first-period quality inference.

As this is a special case of the linear additive model described in Section 7 (notably with $\kappa = \gamma = 0$), inference is characterized as in (39) and we obtain

$$\omega^* = \frac{2\alpha(\bar{\psi}_t - p) + 2p(\gamma + \alpha\kappa) - \alpha\beta\bar{\omega}}{\beta(2\gamma - \alpha)} = \bar{\omega} - \frac{2(\bar{\psi}_t - p)}{\beta}, \; \mu^* = \frac{(1 - 2\kappa)p - 2\bar{\psi}_t + \beta\bar{\omega}}{2\gamma - \alpha} = \frac{2\bar{\psi}_t - p - \beta\bar{\omega}}{\alpha}. \tag{48}$$

Given the quality inference described above, it is straightforward to compute demand and flow profit for the firm for a given price p_t and given rating $\bar{\psi}_t$, where we for simplicity let $\bar{\omega} - \underline{\omega} = 1$.

Lemma 5 (Demand, Flow Profit & Induced Review) *Demand $q_t(p_t, \bar{\psi}_t)$ and flow profit $\pi_t(p_t, \bar{\psi}_t)$ are characterized as follows:*

$$q_t(p_t, \bar{\psi}_t) = \bar{\omega} - \tilde{\omega}_t = \frac{2(\bar{\psi}_t - p_t)}{\beta} \tag{49}$$

$$\pi_t(p_t, \bar{\psi}_t) = p_t \cdot q_t(p_t, \bar{\psi}_t) = \frac{2p_t(\bar{\psi}_t - p_t)}{\beta}. \tag{50}$$

The induced average review is then given by

$$\psi_t(p_t, \bar{\psi}_t) = \alpha\theta + \beta E[\omega|\omega > \omega^*] = \alpha\theta + \beta\bar{\omega} + (p_t - \bar{\psi}_t). \tag{51}$$

Proof. Quantity and profits follow immediately from the inference. Finally, note that

$$E[\omega|\omega > \omega^*] = \frac{\bar{\omega} + \omega^*}{2} = \frac{\bar{\omega} + \frac{2(p_t - \bar{\psi}_t)}{\beta} + \bar{\omega}}{2} = \bar{\omega} + \frac{p_t - \bar{\psi}_t}{\beta}.$$

■

Suppose as a benchmark that the firm prices myopically – it simply maximizes flow profits π_t. Differentiating (50) with respect to p_t and solving the FOC immediately yields the myopically optimal price.

Lemma 6 (Price if Firm is Myopic) *If the firm is myopic and simply maximizes flow profits, it sets a period-t-price of*

$$p_t = \frac{1}{2}\bar{\psi}_t.$$

Corollary 1 *If the firm prices myopically, the quality perception μ_t^* always moves towards θ. Rating systems are thus effective at communicating information.*

Proof. $p_t = \frac{1}{2}\bar{\psi}_t$ induces a period-t-average review of (see (51))

$$\psi_t(p_t, \bar{\psi}_t) = \alpha\theta + \beta\bar{\omega} + (p_t - \bar{\psi}_t) = \alpha\theta + \beta\bar{\omega} - \frac{1}{2}\bar{\psi}_t. \tag{52}$$

Thus,

$$\bar{\psi}_{t+1} = \bar{\psi}_t + \frac{\alpha\theta + \beta\bar{\omega} - \frac{3}{2}\bar{\psi}_t}{t},$$

With $\mu_t^* = \frac{2\bar{\psi}_t - \beta\bar{\omega} - p_t}{\alpha} = \frac{\frac{3}{2}\bar{\psi}_t - \beta\bar{\omega}}{\alpha}$, we can rewrite this as

$$\bar{\psi}_{t+1} = \bar{\psi}_t + \frac{\alpha\left(\theta - \mu_t^*\right)}{t}$$

to see that the rating always moves in the direction of the current bias of the inference. Next, note that for consumers to draw correct quality inference, we would require a rating $\bar{\psi} = \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right)$ – in that case, the myopic firm would charge $p = \frac{1}{3}\left(\alpha\theta + \beta\bar{\omega}\right)$ and induce quality inference

$$\mu^* = \frac{2 \cdot \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right) - \beta\bar{\omega} - \frac{1}{3}\left(\alpha\theta + \beta\bar{\omega}\right)}{\alpha} = \theta.$$

It is thus sufficient to show that $\bar{\psi}_{t+1}$ moves towards $\bar{\psi} = \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right)$. This is easily established by writing

$$\epsilon \equiv \bar{\psi}_t - \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right).$$

This yields

$$\bar{\psi}_{t+1} = \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right) + \epsilon - \frac{\frac{3}{2}\epsilon}{t} = \frac{2}{3}\left(\alpha\theta + \beta\bar{\omega}\right) + \epsilon\left(1 - \frac{3}{t}\right),$$

where with $t \geq 1$ we immediately have that $\left(1 - \frac{3}{t}\right) \in (0, 1)$ and are done. ■

A second potential benchmark case is that of consumers who base their purchase decision solely on some prior belief $\tilde{\mu}$, i.e. a world without ratings where price does not serve as a signaling device. In that case, it is straightforward that the firm simply sets the quasi-monopoly price $p_m(\tilde{\mu}) = \frac{\alpha\tilde{\mu}+\beta\bar{\omega}}{2}$ given the quality belief $\tilde{\mu}$. Price is constant over time, consumers do not learn anything and the resulting loss in terms of welfare depends on the extent and direction of the consumers' misperception: If $\theta > \tilde{\mu}$, too few consumers buy compared to the true quality (and the firm loses out on profits as it sets a lower price than it could if θ were known; if $\theta < \tilde{\mu}$ this is reversed.

Consider thus the case where the firm is strategic and incorporates the effect of price on reviews and hence future ratings into its decision. In period t, it hence maximizes

$$\pi_t + \delta V^{t+1}(\bar{\psi}_{t+1}),$$

where $V^{t+1}(\cdot)$ is the continuation value from a rating $\bar{\psi}_{t+1}$ entering period $t+1$. First note that irrespective of t, we have that

$$\frac{\partial V^{t+1}}{\partial \bar{\psi}_{t+1}} > 0,$$

that is, that the continuation value is increasing in rating. This is because profits in a given period are increasing in $\bar{\psi}_t$. This immediately leads to the following conclusion.

Lemma 7 (Markup with Strategic Pricing) *If the firm is strategic, it will always charge a price p_t above the myopically optimal price $\frac{1}{2}\bar{\psi}_t$ which maximizes flow profits. It thus sacrifices flow profits to induce a higher average review and hence rating in the next period.*

Proof. We have already argued that $\frac{\partial V^{t+1}}{\partial \bar{\psi}_{t+1}} > 0$. Moreover, we have $\frac{\partial \bar{\psi}_{t+1}}{\psi_t} > 0$ from (47) and we can see from Lemma 5 that $\frac{\partial \psi_t}{\partial p_t} > 0$. Thus, $\frac{\partial V^{t+1}}{\partial p_t} > 0$, which together with the analysis of myopically optimal pricing yields the Lemma. ■

Corollary 2 *If the firm is strategic, even if a correct inference $\mu_t^* = \theta$ is induced in a given period t, the next period's induced rating $\bar{\psi}_{t+1}$ is distorted upward and future inference will again be biased.*

A strategic firm hence sacrifices flow profits by charging a markup over the myopically optimal price which serves to increase future profits by inducing a higher current-period average review and hence a higher future rating. As sensitivity of the rating to additional reviews decreases over time, pricing will thus become closer to the myopic optimum even absent the endgame effect (i.e. for T very large).[19] Moreover, it is interesting to note that a high price p_t decreases the current period quality inference μ^*. As such, the firm's price increase decreases demand not only by requiring a higher horizontal taste ω due to the higher price, but also because the inferred quality μ^* decreases which further pushes the induced cutoff ω^* upward.

[19] In the final period, pricing is myopic in any case.