Alone, Together: A Model of Social (Mis)Learning from Consumer Reviews

Abstract. We develop a dynamic model of naïve social learning from consumer reviews. In our model, consumers decide if and what to buy based on both the product(s) expected quality and their idiosyncratic taste for them. Products' qualities are initially unknown, and are (mis)learned from reviews. At the heart of the model lies a dynamic feedback loop between reviews, beliefs, and choices: period t reviews influence $t + 1$ consumers' beliefs, and thus choices; these determine the average of $t+1$ reviews, which in turn influences $t+2$ beliefs, choices, and reviews. We show that in the long-run ($t = \infty$), reviews are systematically biased, leading some consumers astray. In particular, in both monopoly and duopoly, reviews relatively advantage lower quality and more polarizing products, since these products induce stronger taste-based consumer self-selection. Thus, in stark contrast with the winner-takes-all dynamics of classic observational learning models, in which consumers learn from the *choices* of their predecessors, social learning from *opinions* generates excessive choice fragmentation. Our findings have implications for interpreting the variance and number of reviews; pricing in the presence of reviews; and the short- and long-term effectiveness of fake reviews.

1. Introduction

Digitization has brought a substantial increase in variety in virtually all markets: music, books, movies and TV shows, for instance, are being produced at an unprecedented scale. In such a competitive landscape, in which thousands of products are fighting for consumers' attention (and money), it is of fundamental importance to understand how consumers sift through the large variety of products they are presented with.

The Internet has also had a significant impact on how consumers discover and evaluate products, in particular by means of consumer reviews. Learning from reviews, however, is made difficult by the fact that, to some extent, reviews measure idiosyncratic consumer-product fit, and not just objective quality. In this context, how and what can consumers learn from peer-generated information?

This paper studies the nature and impact of consumer reviews in horizontally differentiated products markets. At the heart of the model lies a dynamic feedback loop between reviews, beliefs, and choices: period t reviews influence $t+1$ consumers' beliefs, and thus the $t+1$ set of buyers of each product; this set, in turn, determines the nature of the product's $t+1$ reviews, which will influence $t+2$ beliefs, choices and reviews.
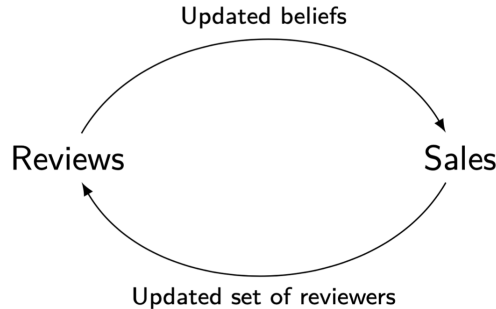
Figure 1
Dynamic feedback loop between reviews, beliefs and choices.

First, we characterize biases in the mean of reviews in both duopoly and monopoly settings. The first bias we identify is that differences in average reviews understate differences in objective quality. This is due to the fact that high-quality products end up inducing purchases (and thus reviews) even by buyers for whom the fit component of consumer satisfaction[1] is relatively lower (*"the curse of the best-seller"*).

Many industry players (and scholars) have recognized that firms might face a quality-quantity trade-off in reviews, reaching contradicting conclusions on how to

---

1. Throughout the paper, we will use the expressions "taste for the product" and "consumer-product fit" interchangeably.

optimally solve it.[2] This paper is, to the best of our knowledge, the first to carefully formalize and quantify this trade-off in a fully dynamic setting.

The review dynamics observed in a variety of highly differentiated product markets display patterns that are in line with our prediction. For example, looking at prize winning books and Academy Awards winning movies respectively, Kovács and Sharkey (2014) and Rossi (2021) show that average reviews decline right after the awards are announced. Conversely, both the winners' number and variance of ratings go up. By studying individual reviewers' behavior over time, both papers trace these effects back to an expanded consumer pool: as beliefs about quality go up, consumers are willing to experiment with genres they don't usually love. This is exactly what our model predicts.

We also show that consumer reviews favor "polarizing" products, that is, products whose fit component has very high variance (which is the case, for instance, when consumers either love or hate the product). The idea is that, because of consumer self-selection, the fit component of reviews is very high: consumers for whom the fit component is low do not purchase the product and thus do not review it. In other words, for a given level of objective quality, polarizing products receive higher average reviews than general-interest products.

A testable implication of this fact is that, for instance, books about politics that take an extreme stance are overrated compared to books that take a centrist one. The same is true for other polarizing attributes: really spicy food in restaurants, graphic violence in movies, and so forth.

Second, and closely related, we apply our model to characterize properties of the variance of reviews. We show that when consumers are aware of their taste for each product – for instance because they observe a restaurant's style of cuisine, or a movie's or book's genre – inferring which products are polarizing from the variance of their reviews is often incorrect.

This is because, since *ex-ante* highly polarizing products tend to be purchased by a very homogeneous set of buyers, their *ex-post* reviews often display low dispersion. In fact, we show that this dispersion can be lower than that of their less polarizing alternatives, which attract a much more diverse crowd. In other words: polarizing products split consumers into two very distinct camps, one who loves them, and one who loathes them. But since only the first camp purchases (and reviews) the product, much of this "between camps" diversity of opinions goes silent, and we are left with only the "within fans" variance – which is often very small.

Third, we discuss the robustness of our findings by looking at the full dynamics of reviews. Because consumers trade-off quality and fit, overrated products attract

---

2. See for instance `https://qiigo.com/blog/quality-or-quantity-whats-your-online-reviews-strategy/` favoring quantity and `https://www.stringcaninteractive.com/quality-or-quantity-whats-more-important-for-reviews/` favoring quality (specificity).

many poorly matched consumers. Therefore, shouldn't the biases highlighted above be short-lived, and naturally self-correcting?

We show that the aforementioned self-correction, while occurring, is only partial: long-run reviews display the same qualitative biases as short-run ones. More precisely, we show that at $t = \infty$ the amount of bias in reviews and consumer mistakes reinforce each other: biased reviews cause some consumers to purchase the wrong product given their subjective taste; and the resulting choice – and thus review – patterns confirm the bias.

This self-correction motive is in sharp contrast with the recent work of Park et al. (2021), who look at consumer electronics and show that the "fateful first consumer review" carries a disproportionate importance in determining both the valence and the number of future reviews. Importantly, taste-based self-selection – the key mechanism in our model – is arguably a non-factor in this market, in which vertical differentiation alone drives the vast majority of choices.

The opposite nature of our results and those in Park et al. (2021) illustrates the opposite issues plaguing reviews-based learning in different types of markets. In markets with large product differentiation, reviews are fundamentally biased but also self-correcting, and thus robust to both initial conditions and external manipulation. In contrast, in markets without taste-based differentiation, initial conditions are disproportionately important, giving rise to "winner-takes-all" dynamics similar to the ones in the classic *observational learning* models of Banerjee (1992) and Bikhchandani et al. (1992).[3]

Fourth, the model we have described so far is one of *naïve* social learning: consumers take reviews at face value, not being fully aware of the underlying self-selection patterns generating them. Would the biases disappear if consumers learned from reviews in a Bayesian fashion, that is, by correctly internalizing their informational content?

We show that, perversely, the presence of Bayesian consumers *worsens* the long-run biases we have identified. The reason is that because Bayesians can identify and choose their subjectively preferred products, they leave higher reviews, on average, than their "naive" counterparts. In doing so, they mitigate the self-correction of reviews described above, leading a greater fraction of their naïve peers to make incorrect choices, in both the short and the long run.

Our findings have important implications for consumers, sellers and platforms alike. For consumers, the key takeaway is to be wary of simple aggregate statistics such as the average and the variance of reviews. While this correction is made cognitively complex by the fact that its extent is highly product-specific, we show that a variety of heuristics – such as rewarding products receiving more reviews (*"the love for large numbers"*, introduced by Powell et al. (2017) and Watson et al. (2018)), and discounting the reviews of highly polarizing products – can (but need not) help.

---

3. See Zhang (2010), Tucker and Zhang (2011) and Tucker et al. (2013) for empirical applications in marketing.

For sellers, our paper highlights the importance of targeting as a reputation management tool: by targeting consumers with a high taste for their products, sellers can not only improve present revenue, but also future reviews. Moreover, sellers should consider the non-trivial trade-offs of different design and pricing strategies, often opting for counterintuitive strategies. When thinking about product design, one might assume that a mainstream design minimizes the chances of negative word of mouth. However, a mainstream product might fail to attract a targeted, passionate crowd, and as a result obtain mediocre reviews, to the detriment of its long-term success. When thinking about prices, the most natural assumption is that reviews might reflect the "quality of the deal" more so than quality *per se*.[4] However, high prices effectively function as a *matching device*: only consumers with a strong taste for the product will buy it.[5] These findings offer new insights compared to the classic marketing implications in Johnson and Myatt (2006), and are in line with empirical evidence on the perverse reputational effects of deep discounts, see for instance Byers et al. (2012), Liu et al. (2019).

We stress how platforms can mitigate this bias by offering targeted recommendations to consumers: provided similar groups of consumers review each of two products of the same genre, their relative reviews are unbiased. We also suggest an important channel through which reviews influence the nature of new product discovery and evaluation. By relatively overrating more polarizing and lower quality options, these platforms might contribute to excessive choice fragmentation. To counter it, the reviews of high quality and more mainstream products should be inflated, given the higher "burden of proof" faced by these products. Last, the fact that reviews, while biased, are naturally self-correcting over time has a crucial implication for platforms and sellers alike: the effects of reviews manipulation on the part of sellers are short-lived. Inducing excessively high expectations about product quality causes an increase in sales in the subsequent period, but in doing so attracts a set of buyers for whom the product will be, on average, a worse match. Thus, reviews will naturally decrease following an artificial boost. This is in line with recent research by He et al. (2022).

The rest of the paper is structured as follows. Section 2 reviews the related literature. Section 3 provides an extended example illustrating the main results. Section 4 introduces the model, and Section 5 presents the main results. Section 6 discusses additional connections to the literature. We conclude in Section 7.

2. Related Literature

There is currently a very active and interdisciplinary literature studying both the causal effects (Chevalier and Mayzlin (2006), Luca (2016)) and the informational

---

4. E.g., see Luca and Reshef (2021) in the context of *Yelp* restaurant reviews.
5. In Online Appendix A, we formalize this argument by analytically showing that optimal long-run duopoly prices are indeed higher in the presence of reviews; in equilibrium, high prices and biased reviews are self-reinforcing.

content of consumer reviews.

This latter strand of the literature has documented a variety of biases, due to social influence (Muchnik et al. (2013), Jacobsen (2015)), reciprocity towards sellers (Filippas et al., 2022), a desire to persuade future consumers (Chakraborty et al. (2022)), fake reviews (Luca and Zervas (2016), He et al. (2022)), or – most closely related to this paper – consumer self-selection (Li and Hitt (2008), Besbes and Scarsini (2018), Le Mens et al. (2018), Park et al. (2021), Acemoglu et al. (2022), Bondi et al. (2024)).

Li and Hitt (2008) study the dynamic pattern of reviews: they conjecture that for a variety of products, die-hard fans buy and (enthusiastically) review first, and consumers for whom the product is not as good a match naïvely follow. They show empirically that for this reason, many products' reviews follow a decreasing trend over time. This paper expands on theirs by fully characterizing the dynamic interplay of reviews and choices, as well as their long-run equilibrium behavior, and qualifying the subset of products for which reviews do, indeed, exhibit a decreasing pattern.

Chakraborty et al. (2022) shows, theoretically and empirically, that consumers who are motivated by their impact on future generations are more likely to leave positive reviews for weak brands (e.g., newcomers to the market), and more likely to leave negative reviews for strong brands (e.g., Starbucks or Chipotle). In their model, each consumer cares about their influence on future consumers, and is thus persuaded to leave a (honest) review only if it is impactful enough. Their findings relate to ours in that they also find that, in equilibrium, reviews penalize the stronger brands. In our model, everyone leaves a (honest) review, and a compression (or, worse, reversion) in reviews occurs as a function of product-specific patterns of taste-based consumer self-selection.

Acemoglu et al. (2022) show that despite time-varying self-selection in reviews (a common theme with our paper), a Bayesian learner can correctly infer the product's quality in a monopolistic setting in which he observes the whole sequence of reviews. Using a similar sequential model, Besbes and Scarsini (2018) show that accurate learning can be achieved under relatively weak assumptions of consumer sophistication. However, they also point out that simply using the mean review as a proxy of quality leads to an incorrect long-run estimate.

Notably, the aforementioned papers, as well as much of the theoretical literature on this topic, study the monopoly case, in which heterogeneous consumers decide between a product of unknown quality and an outside option. In contrast, we primarily focus on a competitive setting.[6] We believe this to be a key innovation, because reviews are usually employed to alleviate choice overload problems, that is, to aid consumers in deciding *what* – not (just) *if* – to buy. Thus, if all products' reviews were equally biased, *relative* reviews would be unbiased, leading to correct choices *between* products. In line with this, in the majority of this paper we study biases in relative, not absolute, reviews. These are much harder for consumers to correct

---

6. For a recent paper also studying the impact of reviews on competition, see Koh and Li (2023).

for than systematic platform-wide biases such as reputation inflation (Filippas et al., 2022).[7]

Another key contribution of this paper is that we model, and solve, the *full* learning process using an infinite horizon problem. Learning from reviews has been usually modeled as a two-periods game in which uninformed first period consumers leave reviews, and second period consumers learn from them (*e.g.*, Sun (2012), Papanastasiou and Savva (2016), Besbes and Scarsini (2018), Fainmesser et al. (2021), Lee et al. (2023)).[8] We show that studying the full dynamic evolution of reviews, while complex, is crucial, since it allows us to elucidate many phenomena that the study of two-period models does not. In doing so, our paper helps bridge the gap between the theoretical and empirical literature; the latter has investigated the dynamic evolution of reviews for over a decade (*e.g.*, Godes and Silva (2012)). Our model studies these issues in a parsimonious and tractable framework, while still accommodating for several dimensions of product heterogeneity, such as quality, design, and price.

While we study both naïve and Bayesian learning, we think the former is both more empirically realistic and behaviorally plausible. De Langhe et al. (2015) document that consumers lack sophistication when interpreting reviews, and "navigate by the stars", even when they are likely to lead them astray. In our case, internalizing the bias of reviews is made harder by the fact that these are both *relative* and *product specific*, hinging upon the interplay of several product characteristics, at least some of which are likely unknown to consumers (or else, they would not need to consult reviews in the first place). While likely pervasive, naïve learning is an understudied phenomenon in this context, as recently acknowledged by Acemoglu et al. (2022): *"It is important to move beyond Bayesian learning and investigate what types of review systems robustly aggregate information when agents use simple learning rules."*[9] It is also worth noting that, by regressing demand on reviews – and not on posterior beliefs of quality given reviews – this is implicitly the type of consumers the empirical literature on this topic (*e.g.*, Chevalier and Mayzlin (2006), Luca (2016) and more recently Reimers and Waldfogel (2021)) has focused on when trying to estimate the causal effects of reviews on demand.

3. Example

Before presenting the more general model in Section 4, we start with a simple, extended example that aims to illuminate the main elements of our theory.

Consider two products $i = 1, 2$ of equal quality: $Q_1 = Q_2$.[10] Without loss of

---

7. For instance, it is well known that an average review of 4.5 out of 5 is below average for an Uber driver, but also above average for a restaurant on Yelp.

8. One exception is represented by Acemoglu et al. (2022), who, however, study a different context from ours.

9. Outside of the context of online reviews, there is a rich literature in naïve social learning – for three classic examples, see DeGroot (1974), Ellison and Fudenberg (1995), and Golub and Jackson (2010).

10. We will consider the case of quality asymmetries $Q := Q_1 - Q_2 \neq 0$ later on.

generality, assume $Q_1 = Q_2 = 0$. A continuum of consumers (indexed by $j \in \mathcal{J}$) each has taste θ_{ij} for product $i = 1, 2$: we assume $\theta_{1j} \sim N(0,1)$, $\theta_{2j} \sim N(0,4)$. In line with Johnson and Myatt (2006), we refer to the first product as "mainstream" and the second product as "niche".[11] Without loss of generality (given the symmetric role of prices and qualities), we assume both prices are zero: $P_1 = P_2 = 0$. Each consumer knows their fit with (which we will also refer to as "taste for") each product, but do not know the products' qualities, and employs reviews to learn them.

After consumption, gross utility for consumer j and product i equally reflects product i's quality and consumer j's subjective taste for it: $\mathcal{U}_{ij} = Q_i + \theta_{ij}$.

We make two assumptions regarding the reviews generating process. First, each consumer buys exactly one product, and all consumers who buy a product review it (and no one else does).[12] Second, each of them does so in a subjectively honest way, by reporting their utility: $\mathcal{R}_{ij} := \mathcal{U}_{ij}$.

Denote by $\mathcal{J}_i^t$ the period t buyers of product i, by $\mathcal{R}_i^t = \{\mathcal{R}_{ij}^t\}_{j \in \mathcal{J}_i^t}$ the entire set of period t reviews for product i, and by $\mathbb{E}_t(\mathcal{R})$ and $Var_t(\mathcal{R})$ its expected value and variance respectively.

Because consumers use reviews to decide *what*, not (just) *if* to buy, our concept of unbiasedness is really one of *relative unbiasedness*: that is, we say that reviews are unbiased if and only if their average is the same for the two products (since $Q_1 = Q_2$), and if the variance of the reviews of product 2 is four times larger than that of product 1 (since $\theta_{1j} \sim N(0,1)$, $\theta_{2j} \sim N(0,4)$). Throughout this example, we define the bias in reviews at time t as $\mathcal{B}_t := (\mathbb{E}_t(\mathcal{R}_1) - \mathbb{E}_t(\mathcal{R}_2)) - (Q_1 - Q_2) =: \Delta_t(\mathcal{R}) - \Delta(Q)$. Since in our initial specification $Q_1 = Q_2$ and thus $\Delta(Q) = 0$, $\mathcal{B}_t := \Delta_t(\mathcal{R})$.

At each $t = 0, 1, ...$, a new generation of consumers arrives. We start by assuming that all consumers learn naïvely from reviews, by taking differences in average reviews as indicative of quality differences. That is, denoting by $\tilde{\Delta}_t(Q)$ the relative quality beliefs of generation t, we have $\tilde{\Delta}_t(Q) := \Delta_t(\mathcal{R})$. Thus, generation t beliefs understate quality differences whenever $\Delta_t(\mathcal{R}) < \Delta(Q)$, or $\mathcal{B}_t < 0$.

We first show that unbiased learning is not an equilibrium, that is, the long-run reviews are biased. To see this, notice that $Q_1 = Q_2$ implies that taste self-selection is the sole driver of consumers' choices, and thus reviews. Now assume $\tilde{\Delta}_t(Q) = 0$. Then, in the following period, we have

$$\mathbb{E}_{t+1}(\mathcal{R}_1) = \underbrace{\mathbb{E}(\theta_{1j} \mid \tilde{\Delta}_t(Q) + \theta_{1j} > \theta_{2j})}_{\text{Taste for Product 1, conditional on choosing it}} = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j}) \approx 0.35.$$

---

11. Notice that mainstream and niche do not necessarily have a market share interpretation, but rather simply refer to the (*ex-ante*) variance in consumers' taste for each product. Throughout the paper, we sometimes refer to niche products as "polarizing".

12. Because we have a continuum of consumers, the results are unchanged if only a positive proportion of consumers review the product, as long as the individual probability of reviewing is uncorrelated with the nature of the review. We discuss extensions to scenarios in which this is not the case in Appendix B.4.

$$\mathbb{E}_{t+1}(\mathcal{R}_2) = \underbrace{\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \tilde{\Delta}_t(Q) + \theta_{1j})}_{\text{Taste for Product 2, conditional on choosing it}} = \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j}) \approx 1.42.$$

Thus: if in any given period t we have $\mathbb{E}_t(\mathcal{R}_1) = \mathbb{E}_t(\mathcal{R}_2)$, in $t+1$ we obtain

$$\Delta_{t+1}(\mathcal{R}) = \mathbb{E}_{t+1}(\mathcal{R}_1) - \mathbb{E}_{t+1}(\mathcal{R}_2) \approx 0.35 - 1.42 = -1.07. \tag{1}$$

Therefore, naïve period $t+2$ consumers will believe $\tilde{\Delta}_{t+1}(Q) = \Delta_{t+1}(\mathcal{R}) = -1.07$. This implies that unbiased learning $(\tilde{\Delta}_{\infty}(Q) = 0)$ can not be the long-run fixed point of this process.

Notice that this also holds in percentile terms: $P(\theta_{1j} > 0.35) = 0.36$, $P(\theta_{2j} > 1.42) = 0.23$. That is, if only one representative consumer for each product were to leave reviews, the consumer of the mainstream (niche) product would like it more than 64% (77%) of their peers. The latter is roughly twice as removed from the median (and mean, by symmetry) consumer as the former. In other words: while both reviews are upward-biased, the bias is larger for the niche product.[13]

Last,

$$Var_{t+1}(\mathcal{R}_1) = \underbrace{Var(\theta_{1j}|\theta_{1j} > \theta_{2j})}_{\text{Variance in taste for product 1, conditional on choosing it}} \approx 0.87,$$

$$Var_{t+1}(\mathcal{R}_2) = \underbrace{Var(\theta_{2j}|\theta_{2j} > \theta_{1j})}_{\text{Variance in taste for product 2, conditional on choosing it}} \approx 1.96.$$

Both variances in reviews understate variances in the underlying taste distributions – a natural consequence of the fact that their most negative reviews are "missing", as the consumers who would have left them buy the other product instead – but the effect is stronger for the more niche product 2:

$$\frac{Var(\theta_{2j} \mid \theta_{2j} > \theta_{1j}) - Var(\theta_{2j})}{Var(\theta_{2j})} = \frac{4 - 1.96}{4} = 0.51$$

$$> 0.13 = \frac{1 - 0.87}{1} = \frac{Var(\theta_{1j} \mid \theta_{1j} > \theta_{2j}) - Var(\theta_{1j})}{Var(\theta_{1j})}.$$

That is, while consumer self-selection essentially cuts the *ex-ante* variance in half for product 2, the decrease is only 13% for product 1.

What happens in the next period? First notice that, since product 1 is now relatively underrated, some period 1 consumers will *wrongly* (given their taste for each

---

13. More generally, given $\theta_1 \sim N(0, \sigma_1^2)$ and $\theta_2 \sim N(0, \sigma_2^2)$ with $\sigma_1^2 < \sigma_2^2$, the following are true: i) $\mathbb{E}(\theta_1|\theta_1 > \theta_2) < \mathbb{E}(\theta_2|\theta_2 > \theta_1)$, ii) $P(\theta_1 > \mathbb{E}(\theta_1|\theta_1 > \theta_2)) > P(\theta_2 > \mathbb{E}(\theta_2|\theta_2 > \theta_1))$, iii) $\frac{\partial \mathbb{E}(\theta_i|\theta_i > \theta_{-i})}{\partial \sigma_{-i}} < 0$, $\frac{\partial \mathbb{E}(\theta_i|\theta_i > \theta_{-i})}{\partial \sigma_i} > 0$ for $i = 1, 2$ and iv) $\frac{\partial P(\theta_i > \mathbb{E}(\theta_i|\theta_i > \theta_{-i}))}{\partial \sigma_{-i}} > 0$, $\frac{\partial P(\theta_i > \mathbb{E}(\theta_i|\theta_i > \theta_{-i}))}{\partial \sigma_i} < 0$ for $i = 1, 2$.

product) purchase product 2 instead. Analytically, this subset of period 1 consumers is given by $\{j \mid \theta_{1j} < \theta_{2j} - \tilde{\Delta}_{t+1}(Q) \ \& \ \theta_{1j} > \theta_{2j}\}$. That is, given the bias in reviews favoring product 2, period $t + 2$ consumers are choosing product 2 (since $\theta_{1j} < \theta_{2j} - \tilde{\Delta}_{t+1}Q$) and $\tilde{\Delta}_{t+1}(Q) < 0$ they have a *perceived* preference for product 2), despite the fact that they should have chosen product 1 instead (since $\theta_{1j} > \theta_{2j}$). Numerically, roughly 18% of consumers fall into this interval, and are thus being misled by reviews.

As, upon purchasing it, these consumers' utility from product 2 is relatively low, their reviews for it will also be lower. Thus, by its very nature, this bias is at least partly self-correcting: as product 2 attracts too many consumers (and, thus, poorer taste matches) in period $t + 1$, its reviews will suffer in period $t + 2$. The opposite is true for product 1. We have

$$\mathbb{E}_{t+2}(\mathcal{R}_1) = \underbrace{\mathbb{E}\big(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \tilde{\Delta}_{t+1}(Q)\big)}_{\substack{\text{Average conditional taste for} \\ \text{product 1, given bias } \tilde{\Delta}_{t+1}(Q) \text{ in reviews}}} \approx 0.49$$

$$\mathbb{E}_{t+2}(\mathcal{R}_2) = \underbrace{\mathbb{E}\big(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \tilde{\Delta}_{t+1}(Q)\big)}_{\substack{\text{Average conditional taste for} \\ \text{product 2, given bias } \tilde{\Delta}_{t+1}(Q) \text{ in reviews}}} \approx 0.95$$

While the reviews of product 2 remain higher, the difference in review has considerably decreased compared to period $t + 1$: $\tilde{\Delta}_{t+2}(Q) = 0.49 - 0.95 = -0.46 > -1.07 = \tilde{\Delta}_{t+1}(Q)$. Product 2's $t + 1$ initial success is also its curse.

Where does this process converge? It is easy to show that $\Delta_{t+2}(\mathcal{R})$ is monotonically decreasing in $\tilde{\Delta}_{t+1}(Q)$; given this, the long-run bias in reviews $\tilde{\Delta}_{\infty}(Q)$ falls strictly between -1.07 and 0. To solve for it, notice that for reviews to stabilize, it must be the case that

$$\mathbb{E}\big(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta_{\infty}(\mathcal{R})\big) - \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta_{\infty}(\mathcal{R})\big) = \Delta_{\infty}(\mathcal{R}) \qquad (2)$$

In other words, equilibrium choices and reviews are mutually reinforcing: consumer choices are "optimal" given their naïve beliefs about the two products' relative qualities, which are shaped by the amount of bias in the system, and this amount of bias is generated exactly by consumers' choices.

Numerically, $\Delta_{\infty}(\mathcal{R}) = -0.66$. In equilibrium, product 2 sells more (62% of the market, which implies 12% of consumers should have bought product 1 instead) *while* enjoying better reviews ($\mathbb{E}_{\infty}(\mathcal{R}_2) = 1.1 > 0.44 = \mathbb{E}_{\infty}(\mathcal{R}_1)$). That is, while the initial overreaction is attenuated over time ($-0.66 > -1.07$), this correction is only partial. Long-run biases have the same sign as those in period 1.

9
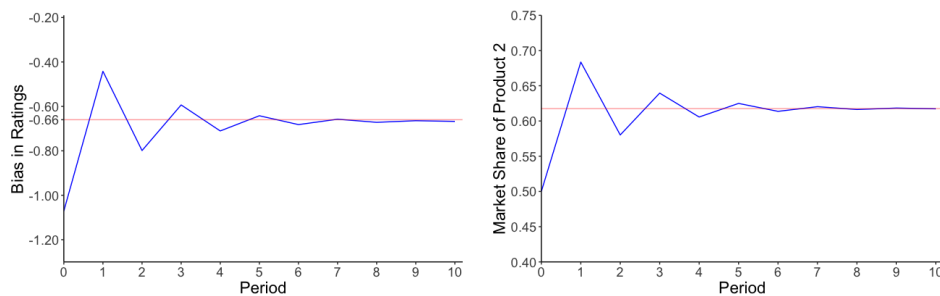
Figure 2

The bias in reviews starts at $\mathcal{B}_0 \approx -1.07$ and then converges to its long-run average, $\mathcal{B}_\infty \approx -0.66$. The number (or share) of reviews, or equivalently the market share, of product 2 starts at $\mathcal{N}_0(\mathcal{R}) = 0.50$ and then converges to its steady state, $\mathcal{N}_\infty(\mathcal{R}) \approx 0.62$.

Now assume that the two products do not just differ in the variance of their tastes distributions, but also in their qualities. Following Johnson and Myatt (2006), assume that the higher quality product is the more mainstream one: $Q := Q_1 - Q_2 > 0$.[14] Eq. (2) becomes

$$Q + \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta_\infty(\mathcal{R})\big) - \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta_\infty(\mathcal{R})\big) = \Delta_\infty(\mathcal{R}).$$

Let $Q = 1$. Numerically, we obtain $\Delta_\infty(\mathcal{R}) = 0.16$. That is, the bias in reviews, *relative to qualities*, got worse: $\mathcal{B}_\infty = \Delta_\infty(\mathcal{R}) - Q = 0.16 - 1 = -0.84 < -0.66$. The reason for this is that the lower quality product 2 induces even stronger taste self-selection (since a higher taste is required to choose the lower quality product 2 when a superior alternative is available), increasing its relative reviews.

---

14. Again, given the symmetric role of qualities and prices, we can assume $P_1 = P_2$ without loss of generality.
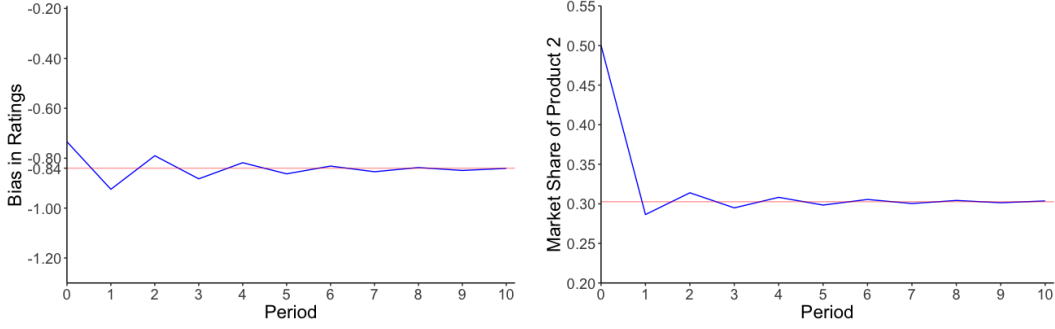
Figure 3

The bias in reviews starts at $\mathcal{B}_0 \approx -0.73$ and then converges to its long-run average, $\mathcal{B}_\infty \approx -0.84$. The reason why the bias gets worse over time is that it is at least partly driven by the quality asymmetry, Q, which is initially unknown and (mis)learned over time. The market share of product 2 starts at $\mathcal{N}_0(\mathcal{R}) = 0.50$ (period 0 consumers are unaware that $Q > 0$). Then, as period $t = 1, 2, \ldots$ consumers become aware of quality differences, it goes down to its steady state, $\mathcal{N}_\infty(\mathcal{R}) \approx 0.30$. This is still higher than if reviews were unbiased: $P(\theta_{2j} > \theta_{1j} + 1) \approx 0.18$.

So far, we have assumed *all* consumers learn naïvely from reviews by taking them at face value. We now return to the symmetric case $Q_1 = Q_2$ and introduce a share $1 - \alpha \in [0, 1]$ of Bayesian consumers, who are able to correct for the reviews biases and form correct beliefs: $\tilde{\Delta}(Q) = 0$. Denoting by $\Delta_\infty^\alpha(\mathcal{R})$ the long-run bias in reviews when only $\alpha \in (0, 1)$ of consumers are naïve, a similar reasoning to the one employed above implies that $\Delta_\infty^\alpha(\mathcal{R})$ satisfies

$$\underbrace{\alpha \cdot \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta_\infty^\alpha(\mathcal{R})\big) + (1 - \alpha) \cdot \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} > \theta_{2j}\big)}_{\substack{\text{Average conditional taste} \\ \text{for Product 1, given bias}}}$$

$$\underbrace{- \alpha \cdot \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta_\infty^\alpha(\mathcal{R})\big) - (1 - \alpha) \cdot \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} > \theta_{1j}\big)}_{\substack{\text{Average conditional taste} \\ \text{for Product 2, given bias}}} = \underbrace{\Delta_\infty^\alpha(\mathcal{R})}_{\text{Bias in reviews}}$$

Both types of consumers contribute to a bias in reviews (in fact, we will see in a second that Bayesian consumers contribute *more*), but only the naïve ones' choices respond to it. Fix $\alpha = \frac{1}{2}$. Numerically, we obtain $\Delta_\infty^{1/2}(\mathcal{R}) = -0.82 < -0.66$. Interestingly, the presence of Bayesian consumer made the bias *worse*. And, therefore, their naïve peers worse off. This is an interesting result: instead of steering their naïve peers towards better learning, Bayesian consumers actually impose a negative externality on them (Section 5.5). The number of reviews is now given by approximately 0.57 for product 2 and 0.43 for product 1. This is less unbalanced than the 0.61 vs 0.39 split we obtained with less biased reviews and $\alpha = 1$. These figures result from the average market shares for Bayesian consumers (who correctly split $50\% - 50\%$, given $Q_1 = Q_2$ and $P(\theta_{1j} > \theta_{2j}) = 1/2$), and those of naïve consumers, which have indeed become more unequal ($35\% - 65\%$), a consequence of the larger equilibrium bias.

11

Last, to verify the robustness of the above findings, we study three extensions: what happens when i) consumers also have an outside option, ii) there are many products, or iii) consumers learn from cumulative, not period by period, reviews? For simplicity of exposition, we go back to the case of naïve consumers and symmetric qualities: $\alpha = 1$, $Q = 0$.

First, let $c \in \mathbb{R}$ be the quality of the outside option. It is easy to see that Eq. (2) becomes

$$\mathbb{E}\big(\theta_{1j} \mid \theta_{1j} \geq max(\theta_{2j} - \Delta_\infty(\mathcal{R}), c)\big) - \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} \geq max(\theta_{1j} + \Delta_\infty(\mathcal{R}), c)\big) = \Delta_\infty(\mathcal{R}).$$

Numerically, we find that when $c = 0$, $\Delta_\infty(\mathcal{R}) = -0.73$. That is, the presence of an outside option with an *ex-ante* equal expected quality to each product has *worsened* the bias. Is this a general finding? Numerical simulations show that the presence and direction of the bias are unchanged,[15] and that c has a non-monotonic effect on its magnitude. When c is very small, the outside option is not important, and we recover the initial long-run bias, $\Delta_\infty(\mathcal{R}) - 0.66$. We then see a worsening of the bias as c grows, followed by an improvement when c becomes large. Importantly, in all of these cases the bias is quantitatively important: in our simulations, $\mathcal{B}_\infty(c) \leq -0.52$, $\forall c \in [-3, 3]$.[16]
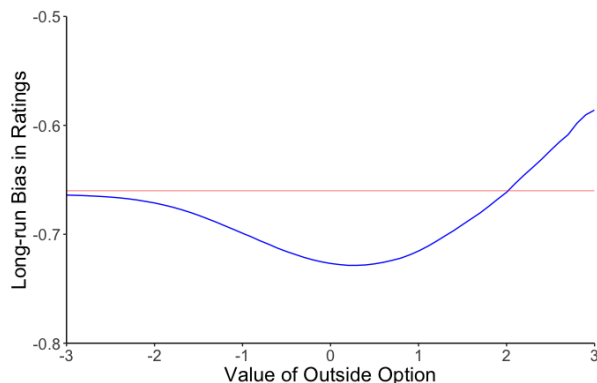
Figure 4
Long-run bias in reviews, $\mathcal{B}_\infty(c)$, as a function of the outside option's quality c.

Is the bias alleviated by more intense competition? Consider the case of $2N$ products, $N \geq 2$, half of which are mainstream like product 1 at the beginning of the example, half of which are niche like product 2.[17] Let $N = 3$, so that there are

---

15. We also confirm these findings analytically in Appendix B.1.
16. As shown in Figure 4, a high c has two effects: first, because the conditional expected values of both θ_1 and θ_2 grow, we have that reviews go up, and importantly, that they get (slightly) closer to each other. Second, when c is very high, the two products' market shares quickly decrease, adding some noise to our estimates at the right end of Figure 4.
17. When N is large, the presence of an outside option that is not exceedingly high is inconsequential, since with high probability there will be at least one product whose subjective value to each consumer will exceed it.

six products in total, three mainstream ($i = 1, 2, 3$), three niche ($i = 4, 5, 6$). By symmetry, we can focus on the first and the fourth products. Eq. (2) now becomes

$$\mathbb{E}\big(\theta_{1j} \mid \theta_{1j} \geq max(\theta_{2j}, \theta_{3j}, \theta_{4j} - \Delta_\infty(\mathcal{R}), \theta_{5j} - \Delta_\infty(\mathcal{R}), \theta_{6j} - \Delta_\infty(\mathcal{R})\big)$$
$$-\mathbb{E}\big(\theta_{4j} \mid \theta_{4j} \geq max(\theta_{5j}, \theta_{6j}, \theta_{1j} + \Delta_\infty(\mathcal{R}), \theta_{2j} + \Delta_\infty(\mathcal{R}), \theta_3 + \Delta_\infty(\mathcal{R})\big)$$
$$= \Delta_\infty(\mathcal{R}).$$

where we have used that, by symmetry, $\mathbb{E}_\infty(\mathcal{R}_1) = \mathbb{E}_\infty(\mathcal{R}_2) = \mathbb{E}_\infty(\mathcal{R}_3)$ and $\mathbb{E}_\infty(\mathcal{R}_4) = \mathbb{E}_\infty(\mathcal{R}_5) = \mathbb{E}_\infty(\mathcal{R}_6)$, and $\Delta_\infty(\mathcal{R})$ denotes, as before, the difference between the reviews of the mainstream product(s) and those of the niche one(s), that is, $\Delta_\infty(\mathcal{R}) := \mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_4)$. Numerically, we find that $\mathbb{E}_\infty(\mathcal{R}_1) = 1.38$, $\mathbb{E}_\infty(\mathcal{R}_4) = 2.09$, and thus $\Delta_\infty(\mathcal{R}) = -0.71$. The number of reviews is given by $\mathcal{N}_\infty(\mathcal{R}_1) = \mathcal{N}_\infty(\mathcal{R}_2) = \mathcal{N}_\infty(\mathcal{R}_3) = 0.06$ and $\mathcal{N}_\infty(\mathcal{R}_4) = \mathcal{N}_\infty(\mathcal{R}_5) = \mathcal{N}_\infty(\mathcal{R}_6) = 0.27$. We highlight three facts. First, all reviews have gone up compared to the duopoly case: this is natural given the stronger taste self-selection induced by each product conditional on choice among six, not two, options.[18] Second, long-run reviews are more biased that in two products case. Third, the number of reviews has become more unequal compared to the two duopoly case: with many options and symmetric qualities, niche products are disproportionately likely to be chosen.[19]

Last, consider now the (empirically relevant) case in which generation $t + 1$ consumers observe the average of all reviews received by each product between period 0 and t, weighted by the number of reviews in each period. The next figure shows the evolution of cumulative reviews over time.
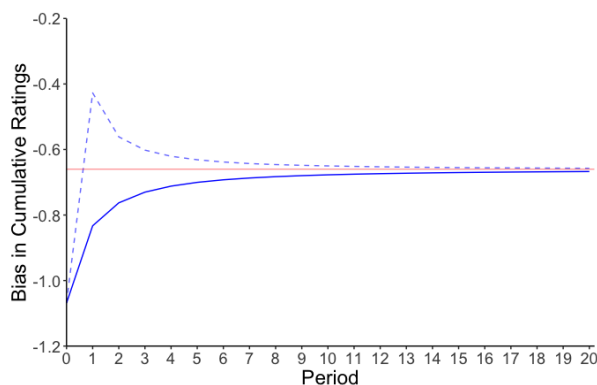
Figure 5
The solid blue line shows the evolution of cumulative reviews. The dashed blue line shows the evolution of period by period reviews when consumers learn from cumulative reviews.

---

18. The same remains true if we instead consider the duopoly plus outside option case.
19. This extends further: for example, if, instead, we had $N = 10$, so that each consumer would have to choose between 20 products, we would have $\Delta_\infty(\mathcal{R}) = -0.78$, $\mathbb{E}_\infty(\mathcal{R}_1) = 2.36$, $\mathbb{E}_\infty(\mathcal{R}_{11}) = 3.14$, and very unequal number of reviews of 0.003 for each of the 10 mainstream products and 0.097 for each of the 10 niche products.

Comparing with Figure 3, we see that $\mathcal{B}_\infty = -0.66$ is unchanged, despite the completely different reviews and, thus, learning dynamics.[20]

We are now ready to present a general model that demonstrates the much wider scope of the facts highlighted by this example.

4. General Theoretical Framework

4.1. Products

Our baseline model features two competing products, $i = \{1, 2\}$, and a continuum set of buyers, $\mathcal{J}$, totaling mass 1. (Considering the large number of different variables used, Table 1 lists the main notation used in the paper.)

The products are both vertically and horizontally differentiated. Following Johnson and Myatt (2006), we model vertical differentiation in terms of quality Q and price P, and horizontal differentation in terms of product design, s. Design measures how polarizing (or, following Johnson and Myatt (2006), "niche"[21]) the product is: a mainstream (high mass appeal, $s = s_H$) design is inoffensive to all consumers, while a niche, or polarizing (low mass appeal, $s = s_L$) design polarizes consumers, who will either love it or hate it.

Consumer j's utility for product i is given by

$$\mathcal{U}_{ij} = Q_i + \theta_{ij} - P_i.$$

θ_{ij} represents idiosyncratic consumer-product match, and is drawn from a continuous and smooth cumulative distribution $F_{s_i}(\cdot)$ with mean 0, *iid* across consumers. In other words, product design $s_i \in \{s_L, s_H\}$ influences the shape of $F_{s_i}(\cdot)$, subject to the constraint that its mean be fixed at 0. More specifically, following Johnson and Myatt (2006), designs are ranked in terms of second order stochastic dominance.[22] In this setting, we can think of the cumulative distributions of consumers' taste in terms of demand rotations – something we will use extensively in proving our results.

Definition 1 (Johnson and Myatt (2006)) *We say that $F_{s'_i}(\cdot)$ is a rotation of $F_{s_i}(\cdot)$ if there exists a $\theta^\dagger_{s_i}$ such that*

$$F_{s_i}(\theta) < (>)F_{s'_i}(\theta) \quad \Longleftrightarrow \quad \theta < (>)\theta^\dagger_{s_i}.$$

Intuitively, $F_{s'_i}(\cdot)$ concentrates more mass around $\theta^\dagger_{s_i}$ than $F_{s_i}(\cdot)$ does. In economic terms, this measures the difference between more mainstream designs, which are

---

20. This is a general property, which we formally prove in Appendix B.2.
21. Note that "niche" here should **not** be interpreted as having a lower market share: in fact, in our main example (Section 3) the niche product obtains a higher market share, and thus a higher number of reviews, in equilibrium, in light of its inflated reviews.
22. Given the two distributions both have 0 mean, this notion simplifies to that of a mean preserving spread.

moderately appealing to most consumers and offensive to none, and more niche ones, which will be loved by some, loathed by others.

One classic specialization to this is the case of variance-ordered distributions, where s governs the spread of the distribution, $F_s(\cdot) := F(\theta/\sigma(s))$ for $\sigma(s) > 0$ and $\sigma'(s) > 0$, but our definitions apply more broadly, for instance to a mean preserving spread that transform a single-peaked distribution into a double-peaked one, as in Figure 6.

Throughout the paper, we assume that prices and taste are always observable to consumers before purchasing,[23] while vertical quality is not, and is inferred from reviews.

4.2. Choice and Reviews

Each consumer picks the product promising the higher expected utility (given beliefs about quality, which as we will see depend on reviews). Upon choosing a product, each buyer reviews it honestly, but *subjectively*, by reporting their own experienced utility.[24] Emphasizing the self-selection in choices, we have:

$$\mathcal{R}_{ij} := \begin{cases} \mathcal{U}_{ij} = Q_i + \theta_{ij} - P_i & \text{if } i \in argmax\{\mathbb{E}(\mathcal{U}_{1j}), \mathbb{E}(\mathcal{U}_{2j})\}, \\ \varnothing & \text{otherwise.} \end{cases} \tag{3}$$

Truth-telling is an especially sensible assumption on platforms – like *Amazon, Goodreads* or *Netflix* – that motivate consumers to leave product reviews at least partly to receive future personalized recommendations.

Denote by $\mathcal{J}_1$ and $\mathcal{J}_2$ the sets of buyers of product 1 and 2 respectively (we are omitting the t subscript for notational simplicity). That is,

$$\mathcal{J}_1 = \{j \in \mathcal{J} \mid \mathbb{E}(Q_1) + \theta_{1j} - P_1 \geq \mathbb{E}(Q_2) + \theta_{2j} - P_2\}, \tag{4}$$

and similarly for $\mathcal{J}_2$. The expectations depend on consumers' subjective beliefs, as discussed below. Denote by $\mathcal{R}_i = \{\mathcal{R}_{ij}\}_{j \in \mathcal{J}_i}$ the entire set of product i's reviews. Last, denote by $F_{s_i}^{\mathcal{J}_i}$ the conditional distributions of θ_i, $i = 1, 2$:

$$F_{s_i}^{\mathcal{J}_i}(\theta_i) = \int_{\mathcal{J}_i} f_{s_i}(\theta_{ij}) dj, \quad j = 1, 2.$$

---

23. Assuming that prices are observable is straightforward, while the same is not necessarily true for product fit. Nevertheless, in a large variety of applications, the main determinants of taste are observable: *e.g.*, genre, setting, author and year for books, movies or TV shows; cuisine, vegetarian-friendliness and atmosphere for restaurants; positioning for non-fiction books; and so forth.

24. Note that this buries two assumptions: first, each buyer posts reviews; second, they do so truthfully. Motivated by a large body of empirical research on extremity bias, motivated reviews, and social influence, among others, we discuss the robustness of our main findings to these assumptions in Appendix B.4.

We are interested in characterizing biases in the mean and variance of $\mathcal{R}_i$, $i = 1, 2$. These are given by

$$\mathbb{E}(\mathcal{R}_i) = \int_{\mathcal{J}_i} \mathcal{R}_{ij} dj = Q_i + \mathbb{E}_{F_{s_i}^{\mathcal{J}_i}}(\theta_i) - P_i, \qquad Var(\mathcal{R}_i) = Var_{F_{s_i}^{\mathcal{J}_i}}(\theta_i).$$

A few remarks are in order. *Prima facie*, Q_i and P_i only shift the distribution of reviews up and down respectively, and do not enter the variance of reviews $Var_{F_{s_i}^{\mathcal{J}_i}}(\theta_i)$. However, they do so indirectly, through consumer self-selection, since

$$\mathcal{J}_i = \mathcal{J}_i(Q_i, Q_{-i}, s_i, s_{-i}, P_i, P_{-i}).$$

This also sheds light on the interplay between a product's (absolute and relative) reviews and its alternative's characteristics, Q_{-i}, s_{-i} and P_{-i}.

Second, Q_i and P_i play symmetric roles. That is, changes in Q_i and P_i that leave $Q_i - P_i$ unchanged do not influence our results. Therefore, we will generally ignore the role of prices (by assuming $P_1 = P_2 = 0$). Moreover, because prices are observable, consumers can correct for them. Therefore, we can equivalently assume that reviews represent gross utilities: $\mathcal{R}_{ij} = Q_i + \theta_{ij}$, and thus shape future consumers' beliefs about Q, not $Q - P$. We discuss the role of pricing in greater detail in Section A.

Third, due to self-selection, consumers usually buy products for which they have positive taste, the more so the higher their relative prices:

$$\mathbb{E}\big(\theta_{1j} \mid \mathbb{E}(Q_1) + \theta_{1,j} - P_1 \geq \mathbb{E}(Q_2) + \theta_{2,j} - P_2\big) \geq 0.$$

Thus, we generally have $\mathbb{E}(\mathcal{R}_i) > Q$ and $Var(\mathcal{R}_i) < Var(\theta_i)$ for $i = 1, 2$. In other words, *all* reviews are upward-biased, and *all* variances in taste distributions are downward-biased. This is because buyers of each product have a stronger taste for it than the average consumer. This is realistic: horror fans are more likely to watch and rate horror movies, lovers of spicy food are more likely to visit and review Szechuan restaurants, and so forth. Moreover, the two go hand in hand: the more upward-biased average reviews are, the more downward-biased the variances, as a natural result of the most negative reviews going "missing" (or, more precisely, going to the other product). The magnitude of these biases depends on the extent of truncation in the taste distribution stemming from consumer self-selection.

However, the key observation is that these biases are highly product-specific: for some products characteristics (which we characterize in our main propositions), we observe a drastic increase in average reviews and decrease in variance; for others, both changes are much smaller. Thus, the reviews of some products are *relatively biased* compared to those of their alternatives. This is crucial, because – in the majority of our model, and overwhelmingly in empirical applications, too – reviews inform choices

between products[25] – that is, choices of *what*, not (just) *if* to buy. We start with the following:

Definition 2. *Average reviews are biased in favor of product 1 (2) whenever*

$$\mathcal{B}(Q_1, Q_2, s_1, s_2, P_1, P_2) := (\mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2)) - (Q_1 - Q_2) > 0 \;\; (< 0)$$

Suppressing the dependence of $\mathcal{B}$ on products features for notational simplicity, we have that $\mathcal{B} > 0$ whenever product 1's reviews are higher than product 2's, *relative to their qualities*. Put differently, $\mathcal{B} > 0$ is equivalent to product 1's reviews being more upward-biased relative to qualities than product 2's: $\mathbb{E}(\mathcal{R}_1) - Q_1 > \mathbb{E}(\mathcal{R}_2) - Q_2$.

4.3. Learning

Modelling learning from reviews is not straightforward. A majority of the literature (e.g., Sun (2012), Papanastasiou and Savva (2016), Fainmesser et al. (2021)) models learning from reviews in a two-period setting, in which period-0 consumers leave reviews, and period-1 consumers learn from them.[26]

In our paper, it is crucial to consider an infinite horizon learning problem, since as we will see short-run and long-run dynamics are quantitatively different. We are mostly interested in the long-run, equilibrium properties of reviews.

Each period $t = 0, 1, 2, \ldots$, a continuum of consumers arrives to the market, and observes the reviews of the previous generation. Each generation is the same, that is, there is no exogenous taste evolution over time. Despite this, choices evolve endogenously as each generation of consumers' beliefs responds to their predecessors' reviews.

In order to describe the long-term dynamics of reviews, we need to specify how consumers internalize the information contained in them. Motivated by empirical realism, we will consider the case of naïve consumers,[27] and then discuss the consequence of relaxing this assumption in Section 5.5. Naïve consumers simply take reviews at face value, that is, they treat average reviews as the quality of each product. Denoting by $\mathbb{E}^N(Q_i)$ their expectations about quality, we have:

---

25. The fact that the real object of social learning here is $\Delta(Q) := Q_1 - Q_2$, more so than individual product qualities Q_1 and Q_2, has a parallel in classic herding models (Banerjee, 1992, Bikhchandani et al., 1992). There, it is the gap in qualities that matters (as proxied by the precision of the quality signal, ρ), not qualities *per se*. Note how, combined with the normalization of prices to 0 (which, as explained earlier, comes without loss of generality), this allows us to reduce a four dimensional problem (Q_1, Q_2, P_1, P_2) to a one dimensional one, $\Delta(Q)$.

26. An important exception is Acemoglu et al. (2022), who look at a the full dynamic of reviews – albeit in a monopolistic, Bayesian setting very different from ours.

27. For a discussion of "cognitively simple" – and behaviorally realistic – decision rules in marketing see Lin et al. (2015), as well as Payne et al. (1993), and citations therein.

$$\mathbb{E}^N(Q_i) = \mathbb{E}(\mathcal{R}_i) \quad i = 1, 2.^{28} \tag{5}$$

(We will omit the N superscript whenever it is obvious from context.)

Before presenting our main results, we reiterate that all biases in reviews we identify are solely predicated on consumers' taste-based self-selection: absent that, all reviews would be perfectly informative of products' features, independently on Q_i and s_i, $i = 1, 2$, as shown in the following:

Lemma 1. *Assume consumers are uninformed about their taste for each product:* $\mathbb{E}(\theta_{1j}) = \mathbb{E}(\theta_{2j}) = 0 \quad \forall j \in \mathcal{J}$. *Then,* $\mathbb{E}(\mathcal{R}_i) = Q_i$ *and* $Var(\mathcal{R}_i) = Var_{F_{s_i}}(\theta_i)$, *for all* Q_i *and* s_i, $i = 1, 2$.

Table 1
Main notation used in the paper (subscripts t omitted for simplicity)

Variable	Description
Q_i	quality of product i
$\Delta(Q)$	quality difference in favor of product 1
P_i	price of product i (normalized, wlog, to 0 until Section 6)
s_i	design of product i
$\mathcal{J}$	set of consumers
$\mathcal{J}_i$	subset of consumers who buy product i
θ_{ij}	consumer j's taste for product i
$f_s(\theta_{ij}), F_s(\theta_{ij})$	design dependent PDF and CDF of taste distribution for product i
$\mathcal{R}_{ij}$	consumer j's review for product i
$\mathcal{R}_i$	set of all reviews $\{\mathcal{R}_{ij}\}_{j \in \mathcal{J}_i}$ for product i
$\mathbb{E}(\mathcal{R}_i)$	mean of reviews for product i
$var(\mathcal{R}_i)$	variance of reviews for product i
$\mathcal{N}(\mathcal{R}_i)$	number of reviews of product i
$\mathcal{B}$	relative bias in reviews in favor of product 1

---

28. We use this definition as it also applies to the monopoly case studied in Section 5.3. In the context of duopoly, we stress that this condition is sufficient, but not necessary, for our results to hold. (This is also apparent in our leading example.) With two products, it is enough that consumers are naïve about the relationship between relative qualities and relative reviews, that is,

$$\mathbb{E}^N(Q_1 - Q_2) = \mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2).$$

This definition implies that mislearning is independent of consumer beliefs about platform-wide review (dis)inflation (Filippas et al. (2022)): even if consumers believed that $\mathbb{E}^N(Q_i) = \mathbb{E}(\mathcal{R}_i) + c$ for some constant $c \neq 0$, this would not affect their *relative* qualities beliefs, which drive choices *between* products.

4.4. Discussion of the Model's Assumptions

Before moving on to presenting our results, we stress the role of our assumptions, as well as the similarities and differences from the existing literature.

- **Everyone Reviews Honestly.** *All* consumers in our model rate the products they buy by reporting their *honest* (but, crucially, *subjective*) opinion. A large body of empirical literature has shown a variety of biases in who reviews conditional on choice (*e.g.*, extremeness bias, review trolls, ...) and how (*e.g.*, social influence). We discuss incorporating some of these extensions in our model in Appendix B.4. Nevertheless, we stress that studying a relative frictionless environment is a deliberate choice: if systematic mislearning arises here, then even more severe mislearning is likely to arise when we include additional biases on the review formation process.

- **No Noise in Reviews.** A variety of papers in the aforementioned literature consider the case in which each individual review contains, on top of a subjective element θ_{ij}, and additional term ϵ_{ij}, drawn from a distribution $H(\cdot)$ with 0 expected value, which measures noise, or individual variability of reviews *even for consumers with the same taste*. For instance, a chef can have a bad day, spoiling the quality of a meal even for consumers who normally love a certain type of cuisine.[29] There are two reasons for not including ϵ in our model. First, we primarily focus on products without variability over time (e.g. books or movies, not restaurants or hotels). Second and more importantly, we study the aggregation of large number of reviews, more so than individual ones. When enough opinions are aggregated, the law of large number applies: $E(\theta_{ij} + \epsilon_{ij} \mid Choice) = E(\theta_{ij} \mid Choice)$. Thus, this model applies even to goods of variable quality, as long as opinions aggregate fast. To summarize, this is a model of *bias*, not *noise*.

- **No Learning about Taste.** In our model, each consumer knows her taste for both products, θ_{ij} for $i = 1, 2$, and employs reviews to try and learn about quality. Clearly, this is not always the case in reality (Sun, 2012). For example, a consumer might not know her taste for a restaurant in advance, even for fixed quality. Nevertheless, we think that quality being the object of social learning is a good approximation in many settings: the genre of a book or a movie are usually known before consumption, and so is a restaurant's style of cuisine, while their quality is not.[30] In Appendix B.3, we propose a simple way to incorporate learning about taste into the model.

- **Cumulative vs One Period Stock of Reviews.** In our model, generation t of consumers learn from the opinions of generation $t-1$. In reality, generation t usually (but not always) observes the (backward discounted) average of reviews

---

29. See Lee et al. (2023) for a decomposition of the variance of reviews into a taste and a quality-variability component, and its implications for consumer learning.

30. On top of this, it should be noted that, from a purely theoretical perspective, it is not straightforward how to model *social* learning about an *idiosyncratic* component.

of generation $t-2, t-3, \ldots, 1$. We chose this approach for analytical tractability and ease of exposition. One can think of our model as on in which the platform quickly discount past opinions as new ones arrive, or one in which consumers only read the most recent reviews. In Appendix B.2, we formally show that none of our major results would be affected if we assumed, instead, that period t consumers observed the (discounted or undiscounted) average of all reviews up to time $t-1$.

5. Main Results

We now turn to study the long-run reviews and learning dynamics in the presence of taste-based self-selection.

We assume that period 0 choices reflect the lack of any information about relative qualities (or qualities per dollar), and thus solely reflect taste: $\mathcal{J}_i = \{j \in \mathcal{J} \mid \theta_{ij} \geq \theta_{-ij}\}$, $i = 1, 2$. (As discussed at length, long-run outcomes are robust to initial conditions.) Thus, assumptions regarding period 0 beliefs about quality, and thus choices, do not play a role in our analysis.

Period 1 choices incorporate information contained in period 0 reviews, which are given by $\mathbb{E}_0(\mathcal{R}_i) = Q_i + \mathbb{E}\big(\mathbb{E}(\theta_{ij} \mid \theta_{ij} \geq \theta_{-ij})\big)$, $i = 1, 2$. More generally, given period t reviews $\mathbb{E}_t(\mathcal{R}_1)$ and $\mathbb{E}_t(\mathcal{R}_2)$, we have that a naïve consumer in period $t+1$ chooses product 1 if and only if $\mathbb{E}_t(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_t(\mathcal{R}_2) + \theta_{2j}$. Thus, product 1's period $t+1$ reviews from naïve consumers are given by

$$\mathbb{E}_{t+1}(\mathcal{R}_1) = Q_1 + \mathbb{E}\big(\mathbb{E}\big(\theta_{1j} \mid \mathbb{E}_t(\mathcal{R}_1) + \theta_{1j} \geq \mathbb{E}_t(\mathcal{R}_2) + \theta_{2j}\big)\big),$$

and similarly for product 2.

To rule out pathological cases in the dynamics of reviews, throughout this Section we also impose a weak technical condition, namely

Assumption 1.
$$\frac{\partial \mathbb{E}\big(\theta_{ij} \mid \theta_{ij} > \theta_{-ij} + k\big)}{\partial k} < \frac{1}{2} \quad \forall s_i, s_{-i}.$$

Intuitively, the condition ensures that the conditional taste for a product does not change too sharply with respect to consumers' beliefs about its relative quality. If this condition were to fail, we could have divergent dynamics in long-run reviews.[31]

To characterize the long-run properties of averages reviews $\big(\mathbb{E}_\infty(\mathcal{R}_1), \mathbb{E}_\infty(\mathcal{R}_2)\big)$, we thus have to solve for the following system of fixed-point equations:

$$\begin{cases} \mathbb{E}_\infty(\mathcal{R}_1) = Q_1 + \mathbb{E}\big(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} \geq \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j}\big) \\ \mathbb{E}_\infty(\mathcal{R}_2) = Q_2 + \mathbb{E}\big(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} \geq \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j}\big) \end{cases}. \tag{6}$$

---

31. This condition is satisfied by the normal, uniform, and beta distributions among many others.

Clearly, this is independent on initial conditions. Furthermore, notice how that the left hand side of both equations is increasing in $\mathbb{E}(\mathcal{R}_i)$, while the right hand side is decreasing in $\mathbb{E}(\mathcal{R}_i)$ and increasing in $\mathbb{E}(\mathcal{R}_{-i})$. This follows straightforwardly from properties of conditional expectations. Two important facts follow: first, for a fixed $\mathbb{E}(\mathcal{R}_{-i})$, there is at most one $\mathbb{E}(\mathcal{R}_i)$ solving each equation. Second, the long-run reviews of each product are *increasing* in the other product's reviews. This is because competing with a better (perceived) alternative strengthens the self-selection of a product's buyers, inflating their average taste for it, and thus its average reviews.

We now characterize the (unique) solution of the above system. In particular, we are interested in properties of $\mathcal{B}_\infty$ as a function of the two products' (relative) characteristics. As we will see, these are both intuitive and empirically realistic, and rationalize a variety of phenomena that have been documented by researchers over the last few years, for instance the perverse effect of prizes such as literary or Academy Awards on reviews (Kovács and Sharkey (2014), Rossi (2021)), the short-lived effects of fake reviews (He et al., 2022), the backfiring of deep discounts (Byers et al. (2012), Liu et al. (2019)), the bunching of reviews around high but not stellar averages[32], and the relatively high average reviews (and, relatedly, low variance in reviews) observed for highly polarizing options, among others.

5.1. Reviews and Product Design

We start with the case of two products having the same quality, but differing in their design. We have the following:

Proposition 1 (More Polarizing Products Are Relatively Overrated) *Let the two products differ only in their design: $Q_1 = Q_2$, $s_1 = H$, $s_2 = L$. Assume that s_1 and s_2 are symmetric.[33] Then, in the long-run ($t = \infty$):*

- *Reviews converge, and the more polarizing product 2 is relatively overrated: $\mathcal{B}_\infty < 0$ or, equivalently since $\Delta(Q) = 0$, $\mathbb{E}_\infty(\mathcal{R}_2) > \mathbb{E}_\infty(\mathcal{R}_1)$,*
- *and thus obtains a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.*
- *Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_0 < \mathcal{B}_\infty < 0$, $\mathcal{N}_0(\mathcal{R}_2) > \mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.*

Proof: The proof for this and all other results can be found in Appendix A. ∎

To gather some intuition for this result, assume we knew that the mainstream product was chosen. Since its valuation among consumers is fairly concentrated (low variance in θ_{1j} given $s_1 = H$), this was likely caused by a distaste for the niche alternative more so than a (statistically rare) strong taste for product 1. So, product

---

32. See *e.g.*, `https://www.researchgate.net/figure/`
    `Histogram-of-venues-review-for-restaurants-fast-foods-and-bars_fig4_`
    `314160022`
33. Symmetry is a sufficient, but not necessary, condition.

1's reviews will not be particularly upward-biased. On the flip side, when observing a consumer choosing product 2, it is more likely that this is due to a strong taste for it than to a (again, statistically rare) strong distaste for the mainstream product.

Put differently, reviews of niche products reflect the opinions of their fans, while those of mainstream products reflect the opinions of anyone who is not a fan of the available alternative(s). This implies a higher conditional taste – and thus relatively more upward-biased reviews – for the niche product.

Another way to grasp the logic behind this result is in terms of *missing* reviews – reviews that are never posted because the (potential) consumer ended up choosing the alternative. While the missing reviews of the less polarizing products are only very slightly negative, by virtue of its taste distribution being concentrated around 0, the ones for the polarizing product are much more negative. In other words: since the two products' have the same *unconditional* quality, product 2 having higher reviews is equivalent to it having lower *missing* reviews.
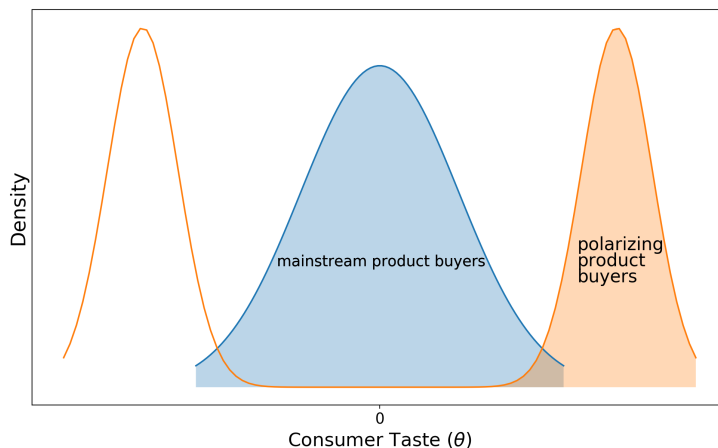
Figure 6
More polarizing products are overrated: intuition with a bimodal (e.g., far-right book) and a unimodal (e.g., centrist book) distribution.

Notice, importantly, that Proposition 1 is *not* saying "The niche product sells only to a small set of super fans, so it obtains fewer but better reviews". In fact, with symmetric qualities, product 2 obtains a *larger* number of reviews ($\mathcal{N}_\infty(\mathcal{R}_2) > 1/2$) *and* obtains a higher average review in equilibrium. We have seen a concrete instance of this in Section 3: there, $Prob(\theta_H \geq \theta_L) = 0.50$, $\mathcal{N}_2(\mathcal{R}_1) \approx 0.67$ and $\mathcal{N}_\infty(\mathcal{R}_1) \approx 0.62$.

If anything, the fact that in equilibrium the more polarizing product obtains higher reviews *while* obtaining a higher number of reviews is surprising, especially in light of Proposition 2 and Corollary 1 in the next Section. Again, the number and average of reviews reinforce each other at $t = \infty$: the polarizing product's reviews are just upward-biased enough to allow it to capture additional consumers who would have preferred product 1, but not enough of them so as to lower its mean review.

5.2. Reviews and Quality: "The Curse of the Best Seller"

We now turn to our second central result: can we trust reviews to accurately reflect quality differences when products have the same design, and are only vertically differentiated? One conjecture is that, since quality is agreed upon by all consumers, quality differences should not bias relative reviews. However, this intuition turns out to be incorrect, since at the heart of consumers' choice is the interplay between vertical quality and horizontal fit: the higher the former, the lower the latter can be while still justifying a purchase. Building on this intuition, we have the following:

Proposition 2 (High Quality Products Are Relatively Underrated.) *Let the two products differ only in their qualities: $Q_1 > Q_2$, $s_1 = s_2$. Then, in the long-run ($t = \infty$):*

- *Reviews converge, and the higher quality product 1 has higher reviews: $\mathbb{E}_\infty(\mathcal{R}_1) > \mathbb{E}_\infty(\mathcal{R}_2)$,*
- *Despite being relatively underrated: $\mathcal{B}_\infty < 0$.*
- *It thus obtains a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_1) > 1/2$, but less than it would if reviews were unbiased.*
- *Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $\mathcal{N}_\infty(\mathcal{R}_1) > \mathcal{N}_2(\mathcal{R}_1)$.*
- *Despite these distortions, reviews unambiguously increase consumer welfare.*

To gather some intuition for this result, assume that reviews are unbiased at period t: $\mathcal{B}_t = 0$. Then, because product 1 is of higher quality, the marginal consumer has a stronger taste for product 2 than it does for product 1:

$$\mathbb{E}\big(\theta_1 \mid Q_1 + \theta_1 > Q_2 + \theta_2\big) < \mathbb{E}\big(\theta_2 \mid Q_2 + \theta_2 > Q_1 + \theta_1\big),$$

where the inequality follows straightforwardly from properties of conditional expectations, and the fact that $s_1 = s_2$. This implies that $t + 1$ reviews will reflect a higher average taste for product 2 than for product 1, leading to $\mathcal{B}_{t+1} < 0$. Thus, it cannot be the case that $\mathcal{B}_\infty = 0$.

An alternative – but equivalent – way to state this is that higher vertical quality can persuade buyers to choose a product even though *it is not the best match for them.* That is, there exists a non-empty set $\underline{\mathcal{J}_1} \subset \mathcal{J}_1$ such that for every $j \in \underline{\mathcal{J}_1}$ we have $Q_1 + \theta_{1j} > Q_2 + \theta_{2j}$ even though $\theta_{1j} < \theta_{2j}$. Though individually rational, these decisions imply that, on average, consumers are better matched with the vertically inferior product, inflating its relative review (while decreasing the number of its reviews) compared to the case in which the two products are vertically indistinguishable.

Notice that, while in Proposition 1 the short- vs long-run result was obtained by comparing $\mathcal{B}_0$ and $\mathcal{B}_\infty$, here we use $\mathcal{B}_1$ and $\mathcal{B}_\infty$ instead.[34] The reason is that, since qualities are initially unknown, and the bias in Proposition 2 comes from the trade-off

---

34. The same is true for the number of reviews, which always lag one period compared to $\mathcal{B}$.

between qualities and taste, here we have $\mathcal{B}_0 = 0$. The bias first appear in period 1, as consumers start to (mis)learn qualities. In Proposition 1, on the other hand, it is product fit alone that is responsible for the bias, and therefore initial reviews are already biased, even though consumers initially hold symmetric beliefs about quality.

Proposition 2 has a host of interesting corollaries. First, it implies the bunching of reviews at the platform level around fairly high, but not stellar, average scores, something that anecdotally resonates when looking at Amazon, Yelp, IMDb, and many more. Whenever a product approach a stellar review, this will attract a high number of (relatively) poorly matched consumers to purchase it, decreasing its future reviews.

Building on the above fact, Proposition 2 also implies a fairly flat relationship between the average and the number of reviews. For instance, on several platforms we see that the average review of a product with fewer than 100 reviews is almost as that of the average product with over 1000 reviews. At first sight, this might be puzzling: shouldn't commercial success be more correlated with quality? Our model rationalizes this fact: the "burden of proof" faced by the higher quality (and thus more popular) products is dramatically higher than that faced by their lower quality alternatives. Large differences in quality coexist with both small differences in reviews and with large differences in the number of reviews.

Proposition 2 suggests that – for a fixed mean of reviews and symmetric designs – products with a high number of reviews should be of higher quality. Thus, in this case rational consumers should have a preference for them. This offers a potential rationalization for the "love of large numbers" (Powell et al., 2017, Watson et al., 2018) commonly observed on consumer reviews platforms. Interestingly, this rationalization of adopting "observational learning" even when consumers' opinions are available is orthogonal to the classic ones of Banerjee (1992), Bikhchandani et al. (1992) and Caminal and Vives (1996).
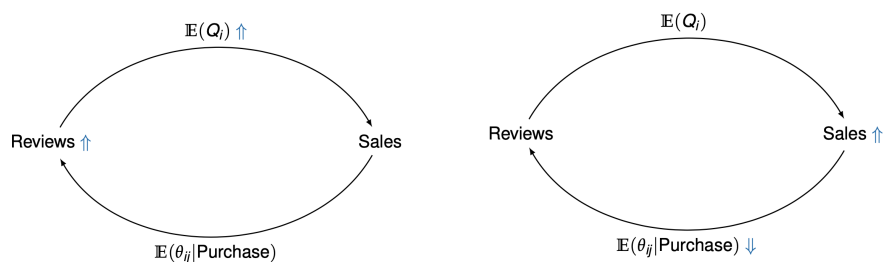
Figure 7
The dynamic feedback loop between reviews, beliefs and choices behind Proposition 2 and Corollary 1. A period t increase in reviews (left) – whether fraudulent or not – leads to an increase in $t + 1$ beliefs about quality. This in turn increases $t + 1$ sales (right), and decreases $t + 1$ product-consumer matches, which leads to lower $t + 1$ reviews.

The next Corollary follows straightforwardly from Proposition 2.

Corollary 1 (High Reviews Are Self-Defeating.) $\mathbb{E}_{t+1}(\mathcal{R}_i)$ *is decreasing in* $\mathbb{E}_t(\mathcal{R}_i)$ *and increasing in* $\mathbb{E}_t(\mathcal{R}_{-i})$. *Jointly, these imply that the same is true,* a fortiori, *for relative reviews:* $\mathcal{B}_{t+1}$ *is decreasing in* $\mathcal{B}_t$.

Corollary 1 formalizes how, contrary to models of social influence (Muchnik et al., 2013) or path dependence more generally (Le Mens et al. (2018), Park et al. (2021)) in which initial reputational advantages are reinforced over time, there is a negative correlation between (both absolute and relative) reviews in consecutive periods: higher reviews today lead to higher number of reviews, and thus lower average matches, and a worse average of reviews, tomorrow.

Last, it is interesting to contrast the case of consumers learning from the *opinions* of their predecessors with that of observational learning, in which they learn from their *actions* (Banerjee (1992) and Bikhchandani et al. (1992)). Models of observational learning are characterized by informational cascades, generating a "winner-take-all" dynamic for sellers and an almost immediate breakdown in the aggregation of information. Here, the opposite happens: niche and lower quality options are overvalued at the expenses of more popular ones, increasing market fragmentation.

An important consequence of this fact is that, while systematically *biased*, reviews are *robust* to manipulation, by their very self-defeating nature. This suggests that, in markets in which consumer taste-based self-selection is prevalent, the impact of fake reviews might be short-lived.

Corollary 2 (Short- and Long-Term Impact of Fake Reviews.) *Let product* 1*'s relative reviews be artificially inflated due to seller* 1*'s manipulation in period* t: $\mathbb{E}_t(\mathcal{R}_\infty^{\mathcal{F}}) > \mathbb{E}_t(\mathcal{R}_1)$, *where* $\mathcal{R}^{\mathcal{F}}$ *indicates the presence of fake reviews. Then:*

- *Fake reviews are effective in the short-run: product* 1*'s number of reviews will be higher in period* $t + 1$: $\mathcal{N}_{t+1}(\mathcal{R}_1^{\mathcal{F}}) > \mathcal{N}_{t+1}(\mathcal{R}_1)$.
- *But absent additional manipulation, they will backfire in the next period: product* 1*'s average reviews in* $t + 1$, *and thus its number of reviews in* $t + 2$, *will be lower:* $\mathbb{E}_{t+1}(\mathcal{R}^{\mathcal{F}}) < \mathbb{E}_{t+1}(\mathcal{R})$, $\mathcal{N}_{t+2}(\mathcal{R}_1^{\mathcal{F}}) < \mathcal{N}_{t+2}(\mathcal{R}_1)$.
- *Both the long-run (*$t = \infty$*) average and the number of reviews are unaffected:* $\mathbb{E}_\infty(\mathcal{R}_1^F) = \mathbb{E}_\infty(\mathcal{R}_1)$, $\mathcal{N}_\infty(\mathcal{R}_1^{\mathcal{F}}) = \mathcal{N}_\infty(\mathcal{R}_1)$.

Notice that the symmetry in our model implies that this result can be immediately restated in terms of negative fake reviews that business post to their competitors, as in Mayzlin et al. (2014). That is, an increase in $\mathbb{E}_t(\mathcal{R}_i)$ is conceptually equivalent to a decrease in $\mathbb{E}_t(\mathcal{R}_{-i})$, as formalized in Corollary 1: future relative reviews are decreasing in today's relative reviews.

Second, and importantly, this result does not just imply that reviews immediately regress to the mean following fake reviews in period t. It goes further, and suggests that fake reviews in period t actively *decrease* $t + 1$ reviews. Then, due to the robustness properties discussed above, reviews eventually converge (back) to their equilibrium levels, as characterized in Proposition 2. Both of these are strikingly in

line with recent empirical work by He et al. (2022), who document a fall in average *Amazon* reviews immediately after brands purchase fake reviews on *Facebook*, and no long-term effects of reviews manipulation, whether positive or negative.

Last, we stress that, just like the rest of our other results, the short-lived effects of fake reviews we document in Corollary 2 are predicated on taste-based self-selection. Absent the correction we highlight, the long-term dynamics could be completely different, and even amplify the initial manipulation. Therefore, we are not suggesting that fake reviews *always* have limited use. In markets that are not as horizontally differentiated as the ones we focus on here, fake reviews might have a disproportionate impact due to either social influence, information cascades, or both (Muchnik et al. (2013), Jacobsen (2015), Chen and Papanastasiou (2021), Park et al. (2021)).
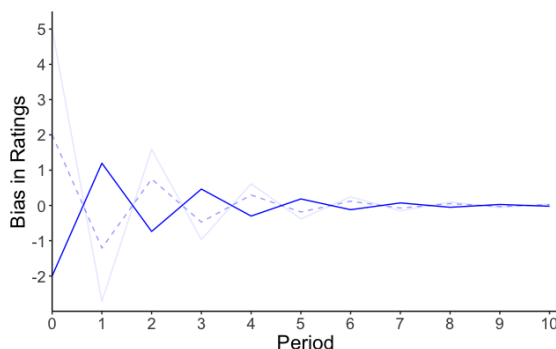
Figure 8
The impact of fake reviews is short-lived. Let $Q_1 = Q_2$ and $\theta_{1j} = \theta_{2j} \sim N(0,1)$. Following an initial inflation (deflation) of product 1's reviews equal to 5 (light blue line), 2 (dotted line) and -2 (dark blue line), relative reviews rebound in period 1, and quickly converge to the unconditional level $\mathcal{B}_\infty = 0$.

5.3. Monopoly

Before moving on to studying consumer welfare implications, we emphasize that the intuitions behind Propositions 1 and 2 do not depend on (duopoly) competition, so that both results are readily applicable to the much more studied setting of online reviews in monopoly (e.g., Papanastasiou et al. (2017), Besbes and Scarsini (2018), Acemoglu et al. (2022)).[35]

In our setting, choice between a product and an outside option can be studied very similarly to choice between two products. The only difference is that, in this case, the outside option has the same value for all consumers, so that it does not induce any taste-based self-selection. Thus, the outside option can be interpreted as a product i of quality c and such that $\theta_{ij} \equiv 0 \;\; \forall j$.

The next Corollary is the monopolistic version of Proposition 1.

---

35. The same is true for all of our next results as well. We omit those Corollaries for brevity.

Corollary 3. *Let each consumer choose between a product of quality Q and design s and an outside option, c. Then, for $c > c^{*}$,*[36]

- *The presence of reviews lead consumers to excessive consumption: compared to the outside option, the product is always overrated in equilibrium.*

- *More polarizing products are more overrated in equilibrium:*

$$\mathbb{E}_{\infty}(\mathcal{R}(s_L)) > \mathbb{E}_{\infty}(\mathcal{R}(s_H)).$$

Corollary 3 formalizes a simple result: product reviews are always inflated, and the degree of inflation is proportional to the extent of taste-based self-selection, which in turns increases the more polarizing the product is.[37]

Similarly, Proposition 2 also has a natural monopolistic interpretation, as shown in the following:

Corollary 4. *Let each consumer choose between a product of quality Q and design s and an outside option, c. Then, for every $s \in \{s_L, s_H\}$ and $c \in \mathbb{R}$,*

- *The presence of reviews lead consumers to excessive consumption: compared to the outside option, the product is always overrated.*

- *Higher quality products are less overrated, while obtaining a higher number of reviews:*

$$\frac{\partial \mathcal{B}_{\infty}(Q)}{\partial Q} < 0, \quad \frac{\partial \mathcal{N}_{\infty}(\mathcal{R})(Q)}{\partial Q} > 0.$$

First of all, notice that in a monopolistic setting $\mathcal{B}_{\infty} = \mathbb{E}_{\infty}(\mathcal{R}) - Q$. The intuition for this result is that, when the product is of very high quality compared to the outside option, most consumers purchase it, leading to a high number of reviews but a weak (but always positive) upward-bias in reviews. The opposite is true when the relative quality of the outside option is high: in this case, only the few consumers with a very high taste for the product buy it, leading to a low number of reviews and a very upward-biased average of reviews.

5.4. Consumer Welfare

Notably, we have analyzed the effects of products' features on their reviews one by one. In reality, products simultaneously differ in both quality and design, and their (relative) reviews will reflect differences in both dimensions. How do product characteristics interact in a competitive market? That is, are high or low quality products more likely to have a polarizing design?

---

36. Assuming symmetry of s, it is easy to show that c^{*} is negative.

37. The requirement that c is high enough is a sufficient condition. The example in Section 3 shows that it is not necessary: for $X \sim N(0, \sigma^2)$, we have that $\mathbb{E}(X \mid X > c) = \sigma^2 \cdot \frac{\phi(c)}{1 - \Phi(c)}$, which is increasing in $\sigma \ \forall c \in \mathbb{R}$.

Johnson and Myatt (2006), Bar-Isaac et al. (2012), Sun (2012) show that sellers benefit from mainstream designs if and only if their quality is relatively high. This important fact has an intuitive explanation: a polarizing design (loved by some consumers, loathed by others) acts as a differentiation tool when competing on the vertical dimension is not possible. In other words: the seller adopts a polarizing design to appeal to at least some consumers (only) if appealing to all consumers is not feasible due to quality deficiencies.

This fact implies that the biases described in Proposition 1 and 2 compound each other: on many platforms, reviews systems display a clear dichotomy. High quality, mainstream products face a very high burden of proof and thus obtain a high volume of relatively low reviews; lower quality and polarizing products, on the other hand, attract extremely strong matches and thus obtain fewer, but very upward-biased reviews. We have seen a numerical example of this compounding effect in Section 3.

How is consumer welfare affected by the interaction of these two biases? The answer is ambiguous. Proposition 1 shows that the presence of reviews can mislead consumers enough to make them worse off compared to the case of no social learning. Proposition 2, on the other hand, shows that when consumers fall prey to the *"curse of the best-seller"*, they fail to achieve first best but still increase their welfare. Can the former effect dominate? Yes, as shown in the following:

Corollary 5. *(Mis)learning from online reviews can be welfare reducing. This is more likely to occur if either quality differences are small, or design differences are large.*

5.5. Bayesian Consumers

We now discuss how the presence of Bayesian consumers affects our results. The main reason for doing so is that, as mentioned in the introduction, one could suspect that the biases we have highlighted so far arise simply due to consumers' naïvete.

That is, if consumers could internalize the biases in reviews, correct for them, and choose accordingly, would the biases we have seen in Proposition 1 and 2 disappear? In this Section, we show that this is not the case, and highlight some additional surprising findings.

We start by showing that Bayesian learning, in this context, takes a rather dramatic form: because reviews contain no noise, Bayesian consumers can correctly back up quality differences in any period (but the initial one, since no reviews are available then). To see this, notice that a period $t+1$ Bayesian consumer observes reviews that are given by

$$\mathbb{E}_{t+1}(\mathcal{R}_i) = Q_i + \mathbb{E}(\theta_{ij} \mid \theta_{ij} + \mathbb{E}_t(\mathcal{R}_i) \geq \theta_{-ij} + \mathbb{E}_t(\mathcal{R}_{-i})), \quad i = 1, 2.$$

Then, she could invert both equations to back up qualities and, thus, *a fortiori*, quality differences:

$$Q_i = \mathbb{E}_{t+1}(\mathcal{R}_i) - \mathbb{E}(\theta_{ij} \mid \theta_{ij} + \mathbb{E}_t(\mathcal{R}_i) \geq \theta_{-ij} + \mathbb{E}_t(\mathcal{R}_{-i})), \quad i = 1, 2. \quad (7)$$

Notice, however, the inherent complexity of such an inversion: besides observing current reviews $\mathbb{E}_{t+1}(\mathcal{R}_i)$ for each product, the consumer should know the entire distributions of θ_{1j} and θ_{2j} (not just her taste for each of the two products), as well as past reviews (which determine self-selection patterns in t, and hence $t+1$ reviews), and be able to compute the double expected value in parentheses. Then, she should take the difference of the two expressions on the right hand side and, in light of her individual tastes for each product, make her subjectively correct purchase. This inference becomes even more complex when more Bayesians are present, since they respond to informational differently from naïves: in that case, the consumer should also know the percentages of each consumer type.

We believe that, in general, such level of inference is very unlikely to happen in reality (see De Langhe et al. (2015) for empirical evidence). Nevertheless, it is interesting to study what happens if it does, and to contrast naïve and Bayesian learning rules.

We are interested in two specific questions: first, how do Bayesian and naïve consumers interact? And second, do the biases highlighted in the previous two sections disappear when a fraction of (potentially all) consumers are Bayesian? In other words, do Bayesian consumers exert a positive externality on their naïve peers, leading them to make their subjectively correct decision?

To this end, assume a measure $\alpha \in [0,1]$ of consumers are as previously described, while the remaining $1-\alpha$ are Bayesian - that is, they are able to perform the inversion in Eq. (7) and back up quality differences.

Denote by $\mathcal{B}_\infty(\alpha)$ the amount of equilibrium bias in reviews when a fraction α of consumers are naïves. (Clearly, $\mathcal{B}_\infty(1)$ corresponds to the case we have studied up to this Section.)

The system of Equations (6) which governs the long-run behavior of reviews now becomes

$$
\begin{cases}
\mathbb{E}_\infty(\mathcal{R}_1) = Q_1 + \alpha\mathbb{E}\big(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} \geq \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j}\big) \\
\qquad\qquad + (1-\alpha)\mathbb{E}(\mathbb{E}(\theta_{1j} \mid Q_1 + \theta_{1j} \geq Q_2 + \theta_{2j})) \\
\mathbb{E}_\infty(\mathcal{R}_2) = Q_2 + \alpha\mathbb{E}\big(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} \geq \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j}\big) \\
\qquad\qquad + (1-\alpha)\mathbb{E}(\mathbb{E}(\theta_{2j} \mid Q_2 + \theta_{2j} \geq Q_1 + \theta_{1j}))
\end{cases}
\tag{8}
$$

Both types of consumers contribute to the average of reviews, but only naïves respond to it period by period. Unbiased reviews ($\mathcal{B}_\infty = 0$) would imply that naïve and Bayesian consumers with the same idiosyncratic taste make the same choices in the long-run. We will see this is generally not the case.

Interestingly, and surprisingly, we show that naïve and Bayesian consumers' beliefs (and thus choices) differ *more* the higher the share of Bayesian consumers, $1-\alpha$. In other words, not only do Bayesian consumers not nudge their naïve peers towards subjectively better choices, they impose a negative externality on them.

More formally, we have the following:

Proposition 3. *Increasing the share of Bayesian consumers, $1-\alpha \in [0, 1]$:*

- *Worsens the amount of equilibrium bias, $\mathcal{B}_\infty(\alpha)$, in Propositions 1 and 2,*
- *Thus making naïve consumers strictly worse off.*

Proposition 3 highlights a key fact: all the biases worsen with a higher share of Bayesian consumers, and thus not only do not disappear, but are in fact largest when only Bayesian consumers – who internalize, and mentally correct for, the biases – are present ($\alpha = 0$).

To gather some intuition for why this is the case, notice that, whenever product 2 is relatively overrated ($\mathcal{B}_t < 0$), too many naïves choose it over product 1 compared to their Bayesian peers (and to the normative optimum). This means that naïve consumers are, on average, a worse match for the product and leave lower reviews. Put differently, by internalizing the bias, Bayesian consumers are less likely to buy product 2 and thus increase its relative review, leading a higher share of $t+1$ naïves to mislearn from reviews and purchase product 2 even when it is the subjectively inferior option for them. Naturally, this effect is stronger the more rational consumers there are.

To sum up, we emphasize that it is not consumer naïvete that leads reviews to be biased. Rather, biases in reviews are solely caused by consumer self-selection patterns into buying products of different qualities and design. Naïve consumers, by taking reviews at face value, are more likely to buy overrated products compared to their Bayesian peers: crucially, this leads to a (partial) self-correction in reviews. Conversely, when more (all) consumers are Bayesian, a smaller degree of (no) self-correction takes place and the biases' magnitude is maximized.

5.6. The Variance of Reviews

To conclude this Section, we apply our model to study the nature, and role, of the variance of reviews. Here, we ask the same question about the variance as we have asked about the mean of reviews in the previous sections, namely: How informative are reviews about product characteristics – in this case, product design? A natural conjecture is that, since product design determines the variance of taste shocks, the reviews of more polarizing products should have a higher variance.

In widely influential work, Clemons et al. (2006) and Sun (2012) study the interplay between product design and reviews variance: Sun (2012) shows, theoretically and empirically, that a high variance increases sales only when the average review is low. Clemons et al. (2006) show that reviews can help the most niche brands expand their market shares.

Our model provides an alternative way to interpret these results: in our formulation, it is easy to see that the first and second moments of reviews are inversely related. As consumers' strategies are always determined by intertwined cutoffs for θ_{1j} and θ_{2j}, reviews become more dispersed if and only if they become lower on average.

Moreover, the dynamic feedback loop between reviews and choices in our model extends to the variance of reviews: not only does the variance affect demand, but the converse is also true. In particular, higher demand results in lower matches on

average, which is equivalent to a higher variance in match qualities. This suggests caution in the causal estimation of the effects of reviews' variance on demand.

In fact, a stronger negative result holds. Given self-selection on taste, a product's reviews' variance need not be indicative of its design. This can translate into a complete reversal of *ex-ante* and *ex-post* variances, as shown in the following:

Proposition 4 (The Variance of Reviews Needs Not Proxy Product Design.)
Let the support for θ_1 and θ_2 be bounded above. Then, there exists a quality gap $Q := Q_1 - Q_2 > 0$ such that $Var_\infty(\mathcal{R}_1) > Var_\infty(\mathcal{R}_2)$ for all s_1 and s_2.

Proposition 4 formalizes the following idea: when product 1 is of much higher quality than product 2, it will attract a much wider – and thus much more diverse taste-wise – audience, while reducing the audience of product 2 to its die-hard fans. When the quality gap is large enough, eventually all buyers of product 2 will have a very strong taste for it: θ_2 will be close to $\bar{\theta}$ for each consumer in $\mathcal{J}_2$. Thus, the observed variance of reviews of θ_2 will get arbitrarily small – and eventually smaller than $Var_\infty(\mathcal{R}_1)$, independently on *ex-ante* designs.

To illustrate the logic further, Figure 6 in Section 5.1 provides an example without quality asymmetries. In Figure 6, the high dispersion in valuations for the orange book comes from the presence of two opposite, but fairly homogeneous, taste groups. Self-selection eliminates reviews from the left one, and thus the observed reviews are all coming from the homogeneous right group, resulting in low variance. Clearly, this *between* groups – *within* group variance gap can be arbitrarily large, even when $Q_1 \leq Q_2$.

Put differently, this finding results from the combination of two forces: on one hand, the niche product has higher unconditional variance. On the other hand, we can make more inference on $\theta_{ij} \sim F_{s_L}(\cdot)$ conditional on choice than we can on $\theta_{ij} \sim F_{s_H}(\cdot)$, so that the *reduction* in variance is larger for the polarizing product. The latter effect can dominate.

To sum up, using the variance of *observed* reviews to infer products designs is appropriate when consumer have no *ex-ante* information and match with products randomly (as is the case in Sun (2012), in which the taste mismatch cost is unknown), but might lead to misclassification when products' horizontal attributes are known to consumers.[38] This is the case, for instance, with political books (in most cases, their titles and covers reveal the books' political stance quite clearly), but also movie genres, restaurants cuisine types, and so forth.

6. Discussion

We now present a number of direct corollaries of our main propositions that help illustrate the relationship between our paper and the existing literature, as well as

---

38. Goodreads' ranking of the most polarizing books of all times clearly shows the risks of wrongly interpreting the information contained in the variance of reviews in this case: Goodreads: The Most Polarizing Books Of All Time.

highlight some original (to the best of our knowledge) predictions.

6.1. Consumption Segregation Goes Up

Contrary to classic models of observational learning, learning from reviews occur from negative, as well as positive, opinions. Thus, the implications for learning are completely different: our model displays excessive dispersion in choices, with fewer consumers purchasing the higher quality product than normatively optimal (Proposition 2). This decreases the probability that two consumers purchase the same product (which is given by $\mathcal{N}(\mathcal{R}_1)^2 + \mathcal{N}(\mathcal{R}_2)^2$ and thus maximized when either $\mathcal{N}(\mathcal{R}_1) = 1$ or $\mathcal{N}(\mathcal{R}_2) = 1$), increasing consumption segregation.

That is, in stark contrast with models of learning from others as conformity, (naïve) consumers in our model end up being less alike, despite the social nature of their learning (*alone, together*). Social learning here fragments market, aiding niche and lower quality products to the expense of their superior, less polarizing alternatives. As we have shown, seeding the model with Bayesian consumers only make matters worse, exacerbating these biases.

6.2. Exploration vs Exploitation

An influential recent stream of research (*e.g.,* Kremer et al. (2014), Papanastasiou and Savva (2016), Che and Hörner (2017), Vellodi (2021)) analyzes the role of platform's reviews design in incentivizing experimentation by consumers, as opposed to simply rewarding myopic exploitation.

Kremer et al. (2014), Papanastasiou and Savva (2016) and Che and Hörner (2017) show that, when platforms want to incentivize exploration – say, Netflix wants to persuade its users to watch a lesser known series that it thinks has promise, – they optimally inflate the reviews of such options at the expense of more established ones. In a similar vein, but focusing on supply-side implications, Vellodi (2021) argues that platforms offer an advantage to incumbent firms, and proposes a reviews design solution: depressing the reviews of the most successful incumbent firms in order to facilitate entry of newcomers.

There are several important differences between our model and theirs. Nevertheless, it seems noteworthy that in each of these four papers, reviews simply proxy qualities. When reviews reflect consumers' idiosyncratic product fit as well as qualities, we show, equilibrium reviews endogenously display biases that are in line with the aforementioned platform design recommendations, even absent platform interventions.

6.3. (Short-Term) Increased Returns to Targeted Advertising

By attracting precisely the type of consumers the firm believes to have a higher taste for its product, targeted advertising offers future reputational benefits, on top of immediate revenue ones. Fainmesser et al. (2021) provide a two-period analysis

formalizing these points. In particular, they show that the firm restricts advertising in period 1 to bump up its reviews in period 2. Our model builds on this by showing that, due to the self-correcting nature of reviews, the magnitude of these effects diminishes over time. However, our results also suggest that the qualitative impact is likely to persist.

7. Conclusion

We study social learning from consumer reviews in a horizontally differentiated market. We do so by modeling the dynamic feedback loop between reviews, beliefs and choices: reviews today influence consumer beliefs and choices tomorrow, but these in turn influence tomorrow's reviews.

We first build a simple model in which the dynamics of reviews can be traced back to the time-varying patterns of taste-based self-selection for each product. We then characterize the fixed-point ($t = \infty$) of this process, and show that its features parsimoniously rationalize a variety of findings from the literature, both theoretical and empirical.

Reviews distort market outcomes in systematic and sizable ways, relatively advantaging lower quality and more polarizing products to the expense of their higher quality and more mainstream alternatives. This is because higher quality and more mainstream product fail to attract a highly self-selected crowd. Similarly, the variance of reviews need not be informative of products' designs: more polarizing products attract a uniform set of buyers, potentially leading their reviews to display lower dispersion than their less polarizing alternatives.

These findings have a large number of immediate (and testable) implications, and rationalize some disparate findings from the literature. For instance: there is a weak relationship between the number and the average of reviews, and the platform could benefit consumers by inflating the reviews of the most rated products; high reviews are self-defeating, and thus fake reviews might not be as impactful; deep discounts can backfire; and sellers (including, contrary to the classic insight of Johnson and Myatt (2006), high quality ones) have an incentive to market more polarizing products, and to target their advertising (or even "demarket", Kotler and Levy (1971)).

Clearly, our study is not without limitations, and in seeking a tractable model, we made several modelling choices (some of which are discussed in Appendix B, on top of the ones already highlighted in Section 4.4). Partly informed by these, we see multiple avenues for future work, both theoretical and empirical, which builds on the model presented in this paper.

First, in our model *all* consumers review the product they buy, and they do so *honestly*. Future work should further investigate the more complex incentives behind information sharing, from social image concerns (*"Will sharing this opinion enhance my reputation?"*) to motivated reviews (*"Will it meaningfully change my peers' opinions about this product?"*) and extremity bias (*"Do I feel strongly enough about this product to take the time to leave a review?"*), as well as their interaction, in

the context of our model. While we believe our results to be largely robust to these and other extensions, we also think these are interesting in their own right, and likely to generate a wealth of additional insights.

Second, despite the crucial role played by taste self-selection, our model is one of social learning about *quality*. What happens when consumers simultaneously learn about quality *and* fit? Or primarily about the latter (reviews as a matching device)? Moreover, since fit is subjective, how can we model consumers learning about it from their peers? We see instances of learning about fit everyday. For instance, *"I recommend this restaurant if you love Mexican food"* or *"Fans of Cormac McCarthy will enjoy this novel"* are examples of statements that are informative about product fit: they will read positive to some, and neutral, or even negative, to others. We believe this would be an exciting avenue for future inquiry.

Despite these caveats, our model can provide a fruitful toolkit for (theoretical and empirical) marketing researchers and managers alike, rationalizing a variety of empirical facts by providing a parsimonious, dynamic theory, and highlighting several directions of future research.

Funding and Competing Interests

References

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Learning from reviews: The selection effect and the speed of learning. *Econometrica*, 90(6):2857–2899, 2022.

Abhijit V Banerjee. A simple model of herd behavior. *The quarterly journal of economics*, 107(3):797–817, 1992.

Heski Bar-Isaac, Guillermo Caruana, and Vicente Cuñat. Search, design, and market structure. *American Economic Review*, 102(2):1140–60, 2012.

Omar Besbes and Marco Scarsini. On information distortions in online ratings. *Operations Research*, 2018.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026, 1992.

Tommaso Bondi, Michelangelo Rossi, and Ryan Stevens. The good, the bad and the picky: Consumer heterogeneity and the reversal of product ratings. *Management Science (forthcoming)*, 2024.

Leif Brandes, David Godes, and Dina Mayzlin. Extremity bias in online reviews: a field experiment. 2018.

John W Byers, Michael Mitzenmacher, and Georgios Zervas. The groupon effect on yelp ratings: a root cause analysis. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 248–265. ACM, 2012.

Ramon Caminal and Xavier Vives. Why market shares matter: An information-based theory. *The RAND Journal of Economics*, pages 221–239, 1996.

Christoph Carnehl, André Stenzel, and Peter Schmidt. Pricing for the stars: Dynamic pricing in the presence of rating systems. *Management Science*, 2023.

Ishita Chakraborty, Joyee Deb, and Aniko Oery. When do consumers talk? *Available at SSRN 4155523*, 2022.

Yeon-Koo Che and Johannes Hörner. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2017.

Li Chen and Yiangos Papanastasiou. Seeding the herd: Pricing and welfare effects of social learning manipulation. *Management Science*, 67(11):6734–6750, 2021.

Judith A Chevalier and Dina Mayzlin. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3):345–354, 2006.

Eric K Clemons, Guodong Gordon Gao, and Lorin M Hitt. When online reviews meet hyperdifferentiation: A study of the craft beer industry. *Journal of management information systems*, 23(2):149–171, 2006.

Bart De Langhe, Philip M Fernbach, and Donald R Lichtenstein. Navigating by the stars: Investigating the actual and perceived validity of online user ratings. *Journal of Consumer Research*, 42(6):817–833, 2015.

Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical association*, 69(345):118–121, 1974.

Glenn Ellison and Drew Fudenberg. Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, 110(1):93–125, 1995.

Itay P Fainmesser, Dominique Olié Lauga, and Elie Ofek. Ratings, reviews, and the marketing of new products. *Management Science*, 67(11):7023–7045, 2021.

Apostolos Filippas, John Horton, and Joseph Golden. Reputation inflation: Evidence from an online labor market. *Marketing Science*, 2022.

David Godes and José C Silva. Sequential and temporal dynamics of online opinion. *Marketing Science*, 31(3):448–473, 2012.

Benjamin Golub and Matthew O Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.

Sherry He, Brett Hollenbeck, and Davide Proserpio. The market for fake reviews. *Marketing Science*, 2022.

Grant D Jacobsen. Consumers, experts, and online product evaluations: Evidence from the brewing industry. *Journal of Public Economics*, 126:114–123, 2015.

Justin P Johnson and David P Myatt. On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3):756–784, 2006.

Andrew Koh and Ricky Li. Competition and social learning. *Available at SSRN*, 2023.

Philip Kotler and Sidney J Levy. Demarketing, yes, demarketing. *Harvard Business Review*, 79:74–80, 1971.

Balázs Kovács and Amanda J Sharkey. The paradox of publicity: How awards can negatively affect the evaluation of quality. *Administrative science quarterly*, 59(1): 1–33, 2014.

Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the wisdom of the crowd. *Journal of Political Economy*, 122(5):988–1012, 2014.

Gaël Le Mens, Balázs Kovács, Judith Avrahami, and Yaakov Kareev. How endogenous crowd formation undermines the wisdom of the crowd in online ratings. *Psychological science*, 29(9):1475–1490, 2018.

Nah Lee, Bryan Bollinger, and Richard Staelin. Vertical versus horizontal variance in online reviews and their impact on demand. *Journal of Marketing Research*, 60 (1):130–154, 2023.

Xinxin Li and Lorin M Hitt. Self-selection and information role of online product reviews. *Information Systems Research*, 19(4):456–474, 2008.

Song Lin, Juanjuan Zhang, and John R Hauser. Learning from experience, simply. *Marketing Science*, 34(1):1–19, 2015.

Yongdong Liu, Anja Lambrecht, and Yiting Deng. Price promotions and online product evaluations. *Available at SSRN 3506009*, 2019.

Michael Luca. Reviews, reputation, and revenue: The case of yelp. com. 2016.

Michael Luca and Oren Reshef. What are consumers rating? evidence from *Yelp! Management Science*, 2021.

Michael Luca and Georgios Zervas. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science*, 62(12):3412–3427, 2016.

Dina Mayzlin, Yaniv Dover, and Judith Chevalier. Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review*, 104(8): 2421–55, 2014.

Lev Muchnik, Sinan Aral, and Sean J Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.

Yiangos Papanastasiou and Nicos Savva. Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*, 63(4):919–939, 2016.

Yiangos Papanastasiou, Kostas Bimpikis, and Nicos Savva. Crowdsourcing exploration. *Management Science*, 64(4):1727–1746, 2017.

Sungsik Park, Woochoel Shin, and Jinhong Xie. The fateful first consumer review. *Marketing Science*, 40(3):481–507, 2021.

John W Payne, James R Bettman, and Eric J Johnson. *The adaptive decision maker*. Cambridge university press, 1993.

Tiberiu Popoviciu. Sur l'approximation des fonctions convexes d'ordre supérieur. *Mathematica (Cluj)*, 10:49–54, 1935.

Derek Powell, Jingqi Yu, Melissa DeWolf, and Keith J Holyoak. The love of large numbers: a popularity bias in consumer choice. *Psychological science*, 28(10):1432–1442, 2017.

Imke Reimers and Joel Waldfogel. Digitization and pre-purchase information: the causal and welfare impacts of reviews and crowd ratings. *American Economic Review*, 111(6):1944–1971, 2021.

Michelangelo Rossi. Quality disclosures and disappointment: Evidence from the academy awards. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 790–791, 2021.

Amin Sayedi. Pricing in a duopoly with observational learning. *Available at SSRN 3131561*, 2018.

Monic Sun. How does the variance of product ratings matter? *Management Science*, 58(4):696–707, 2012.

Catherine Tucker and Juanjuan Zhang. How does popularity information affect choices? a field experiment. *Management Science*, 57(5):828–842, 2011.

Catherine Tucker, Juanjuan Zhang, and Ting Zhu. Days on market and home sales. *The RAND Journal of Economics*, 44(2):337–360, 2013.

Nikhil Vellodi. Ratings design and barriers to entry. *Available at SSRN 3267061*, 2021.

Jared Watson, Anastasiya Pocheptsova Ghosh, and Michael Trusov. Swayed by the numbers: the consequences of displaying product review attributes. *Journal of Marketing*, 82(6):109–131, 2018.

Juanjuan Zhang. The sound of silence: Observational learning in the us kidney market. *Marketing Science*, 29(2):315–335, 2010.

Appendix

A. Proofs

Proof of Proposition 1

This proof consists of several steps. We proof each of them individually:

Claim 1: Reviews converge, and the more polarizing product 2 is relatively overrated: $\mathcal{B}_\infty < 0$ or, equivalently since $\Delta(Q) = 0$, $\mathbb{E}_\infty(\mathcal{R}_2) > \mathbb{E}_\infty(\mathcal{R}_1)$

Proof of Claim 1:

We start by showing that reviews – and thus, *a fortiori*, their difference $\mathcal{B}_t$ – converge.

Because consumers are initially uniformed about quality differences, it is easy to see that $\mathbb{E}_0(\mathcal{R}) = Q_1 + \mathbb{E}(\theta_1|\theta_1 > \theta_2)$ and $\mathbb{E}_0(\mathcal{R}) = Q_2 + \mathbb{E}(\theta_2|\theta_2 > \theta_1)$. Because $\Delta(Q) = 0$ and s_1 is less polarizing than s_2, we have $\Delta_0(\mathcal{R}) < 0$.

Thus,

$$\Delta_1(\mathcal{R}) = \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_0(\mathcal{R})) - \mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta_0(\mathcal{R})).$$

Clearly, $\Delta_0(\mathcal{R}) < 0 \Rightarrow \Delta_1(\mathcal{R}) > \Delta_0(\mathcal{R})$.

Moreover, we have that

$$
\begin{aligned}
\Delta_1(\mathcal{R}) &= \mathbb{E}\Big(\theta_1|\theta_1 > \theta_2 - \Delta_0(\mathcal{R})\Big) - \mathbb{E}\Big(\theta_2|\theta_2 > \theta_1 + \Delta_0(\mathcal{R})\Big) \\
&< \mathbb{E}(\theta_1|\theta_1 > \theta_2) - \frac{\Delta_0(\mathcal{R})}{2} - \Big(\mathbb{E}(\theta_2|\theta_2 > \theta_1) + \frac{\Delta_0(\mathcal{R})}{2}\Big) \\
&= \Big(\mathbb{E}(\theta_1|\theta_1 > \theta_2) - \mathbb{E}(\theta_2|\theta_2 > \theta_1)\Big) - \frac{\Delta_0(\mathcal{R})}{2} - \frac{\Delta_0(\mathcal{R})}{2} \\
&= \Delta_0(\mathcal{R}) - \frac{\Delta_0(\mathcal{R})}{2} - \frac{\Delta_0(\mathcal{R})}{2} \\
&= 0
\end{aligned}
$$

where the inequality follows from Assumption 1, which implies that $\mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_0(\mathcal{R})) \leq \mathbb{E}(\theta_1|\theta_1 > \theta_2) - \frac{\Delta_0(\mathcal{R})}{2}$ and $\mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta_0(\mathcal{R})) \geq \mathbb{E}(\theta_2|\theta_2 > \theta_1) + \frac{\Delta_0(\mathcal{R})}{2}$.

Thus, $\Delta_1(\mathcal{R}) \in (\Delta_0(\mathcal{R}), 0)$. Because $\Delta_{t+1}(\mathcal{R})$ is monotonically decreasing in $\Delta_t(\mathcal{R})$, this implies that $\Delta_2(\mathcal{R}) \in (\Delta_0(\mathcal{R}), \Delta_1(\mathcal{R}))$.

By induction, it is easy to see that

$$\Delta_{2k} \in (\Delta_{2k-2}(\mathcal{R}), \Delta_{2k-1}(\mathcal{R})), \quad \Delta_{2k+1} \in (\Delta_{2k}(\mathcal{R}), \Delta_{2k-1}(\mathcal{R})) \quad \forall k \geq 0. \tag{9}$$

Notice how this implies that the sequences $\{\Delta_{2k}(\mathcal{R})\}_{k=0}^\infty$ is increasing, while $\{\Delta_{2k+1}(\mathcal{R})\}_{k=0}^\infty$ is decreasing.

Moreover, Equation 9 implies that

$$\Delta_{2k}(\mathcal{R}) - \Delta_{2k+1}(\mathcal{R}) < \Delta_{2k-2}(\mathcal{R}) - \Delta_{2k-1}(\mathcal{R}).$$

Thus, $\{\Delta_{2k}(\mathcal{R}) - \Delta_{2k+1}(\mathcal{R})\}_{k=0}^{\infty} \nearrow 0$, proving that reviews to converge to $\Delta_{\infty}(\mathcal{R})$, which is defined by the unique solution of Equation 6. Because the right hand side of $\mathbb{E}_{t+1}(\mathcal{R}_1)$ and $\mathbb{E}_{t+1}(\mathcal{R}_2)$ solely depends on $\Delta_t(\mathcal{R})$, by continuity this implies that $\mathbb{E}_t(\mathcal{R}_1)$ and $\mathbb{E}_t(\mathcal{R}_2)$ also converge.

Subtracting the second line in Eq. 6 from the first we get:

$$\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2) = (Q_1 - Q_2) + \mathbb{E}(\theta_{1j} \mid \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j})) \\ - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j})). \tag{10}$$

Using the definition of

$$\mathcal{B}_{\infty} = (\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2)) - (Q_1 - Q_2),$$

we can simplify this expression to

$$\mathcal{B}_{\infty} = \mathbb{E}(\theta_{1j} \mid \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j})) - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j})).$$

Now note that $Q_1 = Q_2$ by assumption, and thus $\mathcal{B}_{\infty} = (\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2))$. Therefore, the above Eq. can be rewritten as

$$\mathcal{B}_{\infty} = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_{\infty})) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_{\infty})). \tag{11}$$

We now have one Eq. in one variable, $\mathcal{B}_{\infty}$. To show that a solution exists and that it is unique, first notice that the LHS of Eq. 11 is (trivially) increasing in $\mathcal{B}_{\infty}$. The RHS, on the other hand, is decreasing in $\mathcal{B}_{\infty}$: this follows from the fact that $\mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_{\infty})$ is decreasing in $\mathcal{B}_{\infty}$, due to basic properties of conditional expectations, while the opposite is true for $\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_{\infty})$.

To show that the (only) solution $\mathcal{B}_{\infty}$ is negative, therefore, we have to show that *i)* if $\mathcal{B}_{\infty} = 0$, the RHS is negative and *ii)* if $\mathcal{B}_{\infty}$ becomes small (in a way that will be defined later in the proof), the RHS is larger than the LHS.

Let's start with *i)*. Notice that whenever $\mathcal{B}_{\infty} = 0$, the RHS becomes

$$\mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}).$$

Therefore, we have to show the following:

Lemma 2. *Let the assumptions of Proposition 1 hold. Then*

$$\mathbb{E}(\theta_2 \mid \theta_2 \geq \theta_1) > \mathbb{E}(\mathbb{E}(\theta_1 \mid \theta_1 \geq \theta_2).$$

Proof:

We have that

$$\mathbb{E}(\theta_1 | \theta_1 \geq \theta_2) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\theta_1} x_1 f_1(\theta_1) f_2(\theta_2)\, d\theta_2\, d\theta_1}{\int_{-\infty}^{\infty} \int_{-\infty}^{\theta_1} f_1(\theta_1) f_2(\theta_2)\, d\theta_2\, d\theta_1} = \frac{\mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta))}{1/2}$$

and similarly

$$\mathbb{E}(\theta_2 | \theta_2 \geq \theta_1) = \frac{\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta))}{1/2},$$

where the denominators simplify to $1/2 = Prob(\theta_2 > \theta_1) = Prob(\theta_1 > \theta_2)$ given symmetry.

Thus, we are left with having to show that

$$\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta)).$$

Let us start by focusing on $\theta > 0$. Notice that symmetry implies that $F_{s_2}(\theta) < F_{s_1}(\theta)$ if and only if $\theta > 0$. In other words, for $\theta > 0$ we have that s_2 FOSD s_1. To show that the above inequality holds, we show that

$$\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta)).$$

The proof that $\mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta))$ is immediate, since $F_{s_2}(\theta) < F_{s_1}(\theta)$ for every $\theta > 0$. To show that $\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta))$, notice that this is a property of FOSD (which holds for $\theta > 0$): for every non-negative function, the expected value under F_{s_2} is higher than under F_{s_1}.

To show the same for $\theta < 0$, notice that here $F_{s_2}(\theta) > F_{s_1}(\theta)$.

Once again, to show that the above inequality holds, we show that

$$\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta)).$$

The proof that $\mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_2}(\theta))$ is immediate, since $F_{s_2}(\theta) > F_{s_1}(\theta)$ for every $\theta < 0$, which implies $\theta F_{s_2}(\theta) < \theta F_{s_1}(\theta)$. To show that $\mathbb{E}_{F_{s_2}}(\theta F_{s_1}(\theta)) > \mathbb{E}_{F_{s_1}}(\theta F_{s_1}(\theta))$, notice that this is a property of FOSD (which holds for $\theta < 0$): for every negative function, the expected value under F_{s_1} is higher (less negative) than under F_{s_2}. This concludes the proof of Lemma 5. ∎

Conversely, denote by $\bar{\theta}_1$ and $\underline{\theta}_2$ the maximum possible value for θ_1 and minimum possible value for θ_2 respectively. Whenever $\mathcal{B}_\infty < \underline{\theta}_2 - \bar{\theta}_1$, we have $\mathcal{B}_\infty + \theta_1 \leq \underline{\theta}_2$ and thus

$$\mathbb{E}\left(\theta_2 \mid \theta_2 > \theta_1 + \mathcal{B}_\infty\right) \leq \mathbb{E}\left(\theta_2 \mid \theta_2 > \underline{\theta}_2\right) = \mathbb{E}\left(\theta_2\right) = 0.$$

On the other hand, it is straightforward to see that the first conditional expected value on the RHS is positive.

Therefore, we have that $\mathcal{B}_\infty < 0$, as desired. ∎

Claim 2: Product 2 thus obtains a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.

Proof of Claim 2:
We have that

$$\mathcal{N}_\infty(\mathcal{R}_2) = Prob(\theta_{2j} > \theta_{1j} + \mathcal{B}_\infty).$$

Since $\mathcal{B}_\infty < 0$, we have $\mathcal{N}_\infty(\mathcal{R}_2) > 1/2$. ∎

Claim 3: Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_0 < \mathcal{B}_\infty < 0$, $\mathcal{N}_0(\mathcal{R}_2) > \mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.

Proof of Claim 3:
We have that

$$\mathbb{E}_0(\mathcal{R}_1) = Q_1 + \mathbb{E}(\theta_{1j}|\theta_{1j} > \theta_{2j})) \tag{12}$$

and similarly for product 2. Thus, $\mathcal{B}_0 = \mathbb{E}_1(\mathcal{R}_1) - \mathbb{E}_1(\mathcal{R}_2)$ implies

$$\mathcal{B}_0 = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j})) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j})), \tag{13}$$

where we have simplified the RHS using the fact that $Q_1 = Q_2$ by assumption. To show that $\mathcal{B}_0 < \mathcal{B}_\infty$, assume by contradiction $\mathcal{B}_0 = \mathcal{B}_\infty$. But then, we obtain

$$\mathcal{B}_0 = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_0)) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_0)). \tag{14}$$

Therefore,

$$\begin{aligned}
\mathcal{B}_0 &= \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j})) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j})) \\
&< \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_\infty)) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_\infty)) \\
&= \mathcal{B}_\infty,
\end{aligned}$$

where the inequality follows from the fact that $\mathcal{B}_\infty < 0$, as established in the Proof of **Claim 1**. The conclusions hold *a fortiori* if $\mathcal{B}_0 > \mathcal{B}_\infty$. This proves that $\mathcal{B}_0 < \mathcal{B}_\infty$. The fact that $\mathcal{N}_1(\mathcal{R}_2) > \mathcal{N}_\infty(\mathcal{R}_2)$ follows straightforwardly from $\mathcal{B}_0 < \mathcal{B}_\infty$ using the same argument as in the Proof of **Claim 2**. ∎

Proof of Proposition 2

Claim 1: Reviews converge, and the higher quality product 1 has higher reviews: $\mathbb{E}_\infty(\mathcal{R}_1) > \mathbb{E}_\infty(\mathcal{R}_2)$,

Proof of Claim 1:
We start by showing that reviews – and thus, *a fortiori*, $\mathcal{B}_t$ – converge.

Because consumers are initially uniformed about quality differences, it is easy to see that $\mathbb{E}_0(\mathcal{R}) = Q_1 + \mathbb{E}(\theta_1|\theta_1 > \theta_2)$ and $\mathbb{E}_0(\mathcal{R}) = Q_2 + \mathbb{E}(\theta_2|\theta_2 > \theta_1)$. Because $\Delta(Q) > 0$ and $s_1 = s_2$, we have $\Delta_0(\mathcal{R}) = \Delta(Q)$. Thus,

$$\Delta_1(\mathcal{R}) = \Delta(Q) + \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_0(\mathcal{R})) - \mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta_0(\mathcal{R})).$$

Clearly, $\Delta_0(\mathcal{R}) > 0 \Rightarrow \Delta_1(\mathcal{R}) < \Delta_0(\mathcal{R})$.

Moreover, we have that

$$\begin{aligned}
\Delta_1(\mathcal{R}) &= \Delta(Q) + \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_0(\mathcal{R})) - \mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta_0(\mathcal{R})) \\
&= \Delta(Q) + \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta(Q)) - \mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta(Q)) \\
&> \Delta(Q) + \mathbb{E}(\theta_1|\theta_1 > \theta_2) - \frac{\Delta(Q)}{2} - \left(\mathbb{E}(\theta_2|\theta_2 > \theta_1) + \frac{\Delta(Q)}{2}\right) \\
&= 0,
\end{aligned}$$

where the last equality follows from the fact that $s_1 = s_2$ (and, therefore, $\mathbb{E}(\theta_1|\theta_1 > \theta_2) = \mathbb{E}(\theta_2|\theta_2 > \theta_1)$), while the inequality follows from our regularity assumption **??**, which implies that $\mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta(Q)) \geq \mathbb{E}(\theta_1|\theta_1 > \theta_2) - \frac{\Delta(Q)}{2}$ and $\mathbb{E}(\theta_2|\theta_2 > \theta_1 + \Delta(Q)) \leq \mathbb{E}(\theta_2|\theta_2 > \theta_1) + \frac{\Delta(Q)}{2}$.

Thus, $\Delta_1(\mathcal{R}) \in (\Delta_0(\mathcal{R}), 0)$. Because $\Delta_{t+1}(\mathcal{R})$ is monotonically decreasing in $\Delta_t(\mathcal{R})$, this implies that $\Delta_2(\mathcal{R}) \in (\Delta_0(\mathcal{R}), \Delta_1(\mathcal{R}))$.

By induction, it is easy to see that

$$\Delta_{2k} \in (\Delta_{2k-2}(\mathcal{R}), \Delta_{2k-1}(\mathcal{R})), \quad \Delta_{2k+1} \in (\Delta_{2k}(\mathcal{R}), \Delta_{2k-1}(\mathcal{R})) \quad \forall k \geq 0. \tag{15}$$

Notice how this implies that the sequences $\{\Delta_{2k}(\mathcal{R})\}_{k=0}^{\infty}$ is increasing, while $\{\Delta_{2k+1}(\mathcal{R})\}_{k=0}^{\infty}$ is decreasing.

Moreover, Equation 15 implies that

$$\Delta_{2k+1}(\mathcal{R}) - \Delta_{2k}(\mathcal{R}) < \Delta_{2k-1}(\mathcal{R}) - \Delta_{2k-2}(\mathcal{R}).$$

Thus, $\{\Delta_{2k+1}(\mathcal{R}) - \Delta_{2k}(\mathcal{R})\}_{k=0}^{\infty} \searrow 0$, proving that reviews converge to $\Delta_{\infty}(\mathcal{R})$, which is defined by the unique solution of Equation 6. Because the right hand side of $\mathbb{E}_{t+1}(\mathcal{R}_1)$ and $\mathbb{E}_{t+1}(\mathcal{R}_2)$ solely depends on $\Delta_t(\mathcal{R})$, by continuity this implies that $\mathbb{E}_t(\mathcal{R}_1)$ and $\mathbb{E}_t(\mathcal{R}_2)$ also converge.

Denote $\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2)$ by $\Delta(\mathcal{R})$. Notice that, by definition of $\mathcal{B}_{\infty}$, the claim is equivalent to $\mathcal{B}_{\infty} > -(Q_1 - Q_2)$.

To show that this is indeed the case, suppose by contradiction that $\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2) < 0$ and, thus, $\mathcal{B}_{\infty} < -(Q_1 - Q_2)(< 0)$.

But then, we have

$$\begin{aligned}
\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2) = (Q_1 - Q_2) &+ \mathbb{E}(\theta_{1j} \mid \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j})) \\
&- \mathbb{E}(\theta_{2j} \mid \mathbb{E}_{\infty}(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_{\infty}(\mathcal{R}_1) + \theta_{1j})).
\end{aligned}$$

which, using the definition of

$$\mathcal{B}_{\infty} = (\mathbb{E}_{\infty}(\mathcal{R}_1) - \mathbb{E}_{\infty}(\mathcal{R}_2)) - (Q_1 - Q_2),$$

simplifies to

$$
\begin{aligned}
\mathcal{B}_\infty = {} & \mathbb{E}(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j})) \\
& - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j})).
\end{aligned}
\tag{16}
$$

Since the LHS is negative, it is enough to show that the RHS is positive to reach a contradiction. To see that this is indeed the case, notice that

$$
\begin{aligned}
& \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} + \mathbb{E}_\infty(\mathcal{R}_2) - \mathbb{E}_\infty(\mathcal{R}_1)) \\
& - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2)) \\
& = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(\mathcal{R}))) \\
& - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(\mathcal{R}))) \\
& > 0
\end{aligned}
\tag{17}
$$

Where the inequality follows from $s_1 = s_2$ ($\in \{s_L, \ s_H\}$) and the fact that $\Delta_\infty(\mathcal{R}) < 0$ by assumption. Thus, we have that $\mathbb{E}_\infty(\mathcal{R}_1)) \geq \mathbb{E}_\infty(\mathcal{R}_2))$.

To rule out equality, notice that if $\mathbb{E}_\infty(\mathcal{R}_1)) = \mathbb{E}_\infty(\mathcal{R}_2))$ the RHS is Eq. (16) is 0 by symmetry of designs, while the LHS is negative since $\mathcal{B}_\infty = -Q_1 + Q_2 < 0$. We have thus proved that $\mathbb{E}_\infty(\mathcal{R}_1) > \mathbb{E}_\infty(\mathcal{R}_2)$. $\blacksquare$

Claim 2: Despite being relatively underrated: $\mathcal{B}_\infty < 0$.

Proof of Claim 2:

Subtracting the second line in Eq. 6 from the first – just like we did in the Proof of Proposition 1 – we get:

$$
\begin{aligned}
\mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2) = {} & (Q_1 - Q_2) + \mathbb{E}(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j})) \\
& - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j})).
\end{aligned}
$$

Using the definition of

$$
\mathcal{B}_\infty = (\mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2)) - (Q_1 - Q_2),
$$

we can simplify this expression to

$$
\begin{aligned}
\mathcal{B}_\infty = {} & \mathbb{E}(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j}) \\
& - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j}).
\end{aligned}
\tag{18}
$$

Now assume $\mathcal{B}_\infty = 0$. Then, Eq. (18) becomes

$$
\mathbb{E}(\theta_{1j} \mid Q_1 + \theta_{1j} > Q_2 + \theta_{2j})) = \mathbb{E}(\theta_{2j} \mid Q_2 + \theta_{2j} > Q_1 + \theta_{1j})).
\tag{19}
$$

Denoting by $\Delta(Q) = Q_1 - Q_2$, and noticing that $\Delta(Q) > 0$ by assumption, this is equivalent to

$$
\mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(Q))) = \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(Q)))).
\tag{20}
$$

But this is a contradiction, since:

$$\begin{aligned}
\mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(Q))) &= \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} - \Delta(Q))) \\
&< \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j})) \\
&< \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(Q)))).
\end{aligned} \tag{21}$$

Where the equality follows from $s_1 = s_2 (\in \{s_L, \ s_H\})$ and the two inequalities follow from progressively increasing the lower bound of integration in the conditional expected value of θ_{2j}.

The case $\mathcal{B}_\infty > 0$ can be handled similarly. Therefore, in equilibrium we have $\mathcal{B}_\infty < 0$.   ∎

Claim 3: It thus obtains a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_1) > 1/2$, but less than it would if reviews were unbiased.

Proof of Claim 3: The Proof is a natural consequence of Claim 1 and 2. We have that

$$\begin{aligned}
\mathcal{N}_\infty(\mathcal{R}_1) &= Prob(\theta_{1j} + \mathbb{E}_\infty(\mathcal{R}_1) > \theta_{2j} + \mathbb{E}_\infty(\mathcal{R}_2)) \\
&> Prob(\theta_{1j} > \theta_{2j}) \\
&= \frac{1}{2},
\end{aligned}$$

where the inequality follows from the fact that $\mathbb{E}_\infty(\mathcal{R}_1) > \mathbb{E}_\infty(\mathcal{R}_2)$ as shown in Claim 1, and the last equality follows from symmetry in designs.   ∎

Claim 4: Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $\mathcal{N}_\infty(\mathcal{R}_1) > \mathcal{N}_2(\mathcal{R}_1)$.

Proof of Claim 4:

We have already shown that $\mathcal{B}_0 = 0$, that is, $\mathbb{E}_0(\mathcal{R}_1) - \mathbb{E}_0(\mathcal{R}_2) = Q_1 - Q_2 > 0$. Thus, we have

$$\mathbb{E}_1(\mathcal{R}_1) = Q_1 + \mathbb{E}(\theta_1 \mid \theta_1 + \mathbb{E}_0(\mathcal{R}_1)) > \theta_2 + \mathbb{E}_0(\mathcal{R}_2)) = Q_1 + \mathbb{E}(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q))$$

and similarly

$$\mathbb{E}_1(\mathcal{R}_2) = Q_2 + \mathbb{E}(\theta_2 \mid \theta_2 + \mathbb{E}_0(\mathcal{R}_2) > \theta_1 + \mathbb{E}_0(\mathcal{R}_1)) = Q_2 + \mathbb{E}(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q))).$$

Jointly, these two imply that

$$\begin{aligned}
\mathcal{B}_1 &= (Q_1 + \mathbb{E}(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q))) - (Q_2 + \mathbb{E}(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q))) - Q_1 + Q_2 \\
&= \mathbb{E}(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q))) - \mathbb{E}(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q))).
\end{aligned} \tag{22}$$

But then, comparing the expressions in Eq. (18) and (54), we obtain $\mathcal{B}_1 < \mathcal{B}_\infty$ if and only if $\Delta(Q) > \mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2)$. But this is equivalent to $\mathcal{B}_\infty < 0$, which we have shown to be true in the Proof of **Claim 1**.

As in the previous cases, the fact that $\mathcal{N}_2(\mathcal{R}_1) < \mathcal{N}_\infty(\mathcal{R}_1)$ follows straightforwardly from $\mathcal{B}_1 < \mathcal{B}_\infty$. ∎

Claim 5: Despite these distortions, reviews unambiguously increase consumer welfare.

Proof of Claim 5:
To show that reviews improve consumer welfare, define

$$\mathcal{J}_{BR} := \{j \mid \theta_{1j} > \theta_{2j} - \Delta(Q) \quad \& \quad \theta_{1j} < \theta_{2j} - \Delta(R)\}$$

and

$$\mathcal{J}_{NR} = \{j \mid \theta_{1j} > \theta_{2j} - \Delta(Q) \quad \& \quad \theta_{1j} < \theta_{2j}\}.$$

In words, $\mathcal{J}_{BR}$ is the set of consumers who would prefer product 1 but end up choosing their *subjectively* less preferred option given the presence of biased reviews (BR). That is, every consumer $j \in \mathcal{J}_{BR}$ prefers product 1 in light of the quality difference between the two product and their taste for each of the two products, but end up choosing product 2 because they are misled by reviews. Noice that $\Delta_\infty(R) < \Delta(Q)$ – or, equivalently, $\mathcal{B}_\infty$, which we have shown in **Claim 2** – implies this set is non-empty.

Similarly, $\mathcal{J}_{NR}$ is the set of consumers who would prefer product 1 but choose their *subjectively* less preferred option given no reviews (NR): these consumers choose purely based on taste, without accounting for the fact that $Q_1 > Q_2$.

We want to show that $\mathcal{J}_{BR} \subset \mathcal{J}_{NR}$. But this is immediate, since

$$\theta_{1j} < \theta_{2j} - \Delta(R) \quad \Rightarrow \quad \theta_{1j} < \theta_{2j}$$

given that $\Delta(R) > 0$ as we have shown in **Claim 1**. Therefore, $Prob(j \in \mathcal{J}_{NR}) > Prob(j \in \mathcal{J}_{BR})$.

It is immediate to realize that, in this context, welfare is proportional to the number of consumers who choose their subjectively optimal product. We have that this probability is 1 in the first best case, $1 - Prob(j \in \mathcal{J}_{BR})$ in the biased reviews case, and $1 - Prob(j \in \mathcal{J}_{NR})$ in the no reviews case.

Clearly, the above reasoning implies that

$$1 - Prob(j \in \mathcal{J}_{NR}) < 1 - Prob(j \in \mathcal{J}_{BR}) < 1, \tag{23}$$

which concludes the proof. ∎

Proof of Proposition 3

Claim 1: Increasing the share of informed consumers, $1 - \alpha$, worsens the amount of equilibrium bias, $\mathcal{B}_\infty(\alpha)$, in Propositions 1 and 2

Proof of Claim 1:

We have seen that in this case, Eq. (6) becomes

$$\begin{cases} \mathbb{E}_\infty(\mathcal{R}_1) = Q_1 + \alpha \cdot \mathbb{E}\left(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} \geq \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j}\right) \\ \qquad\qquad + (1 - \alpha) \cdot \mathbb{E}\left(\theta_{1j} \mid Q_1 + \theta_{1j} \geq Q_2 + \theta_{2j}\right) \\ \mathbb{E}_\infty(\mathcal{R}_2) = Q_2 + \alpha \cdot \mathbb{E}\left(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} \geq \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j}\right) \\ \qquad\qquad + (1 - \alpha) \cdot \mathbb{E}\left(\theta_{2j} \mid Q_2 + \theta_{2j} \geq Q_1 + \theta_{1j}\right) \end{cases} \tag{24}$$

Denote by $\mathcal{B}_\infty(\alpha)$ the amount of bias in reviews as a function of the fraction of naïve consumers.

The proof proceeds in two step. First, we show that $\mathcal{B}_\infty(0) < \mathcal{B}_\infty(1)$. Then, we show that $\mathcal{B}_\infty(\alpha)$ is monotonic in $[0, 1]$. Each of these two steps must be performed for both the case in Proposition 1 and that in Proposition 2.

Case 1: Extension of Proposition 1

Subtracting the second Eq. from the first in (25) and noting that here, $Q_1 = Q_2$ (and, thus, $B_\infty = \mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2) < 0$), we get:

$$\begin{aligned} \mathcal{B}_\infty = {} & \alpha \cdot \mathbb{E}\left(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}_\infty\right) \\ & - \alpha \cdot \mathbb{E}\left(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}_\infty\right) \\ & + (1 - \alpha) \cdot \mathbb{E}\left(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}\right) \\ & - (1 - \alpha) \cdot \mathbb{E}\left(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}\right). \end{aligned} \tag{25}$$

First, we want to show that $\mathcal{B}_\infty(1) < \mathcal{B}_\infty(0)$. We have

$$\begin{aligned} \mathcal{B}_\infty(1) = {} & \mathbb{E}\left(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}_\infty(1)\right) \\ & - \mathbb{E}\left(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}_\infty(1)\right). \end{aligned}$$

and

$$\begin{aligned} \mathcal{B}_\infty(0) = {} & \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}) \\ & - \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}). \end{aligned}$$

Recall that $\mathcal{B}_\infty(1) < 0$ (Proposition 1).

But then,

$$\begin{aligned} \mathcal{B}_\infty(0) = {} & \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}) \\ \leq {} & \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}_\infty(1)) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}_\infty(1)) \\ = {} & \mathcal{B}_\infty(1). \end{aligned}$$

Thus, $\mathcal{B}_\infty(0) < \mathcal{B}_\infty(1) < 0$: the equilibrium bias got *worse* if only Bayesian are present ($\alpha = 0$), compared to only naïves ($\alpha = 1$).

To show that we have monotonicity in $\alpha \in [0,\ 1]$, denote by

$$
\begin{aligned}
G(\alpha, \mathcal{B}) = \ &\mathcal{B} - \alpha \cdot \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}) \\
&+ \alpha \cdot \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}) \\
&- (1 - \alpha) \cdot \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}) \\
&+ (1 - \alpha) \cdot \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}).
\end{aligned}
\tag{26}
$$

Then, for every $\alpha \in [0,1]$, $\mathcal{B}_\infty(\alpha)$ solves $G(\alpha, \mathcal{B}) = 0$. By the Implicit Function Theorem, we have that

$$
\frac{\partial \mathcal{B}(\alpha)}{\partial \alpha} = - \frac{\frac{\partial G(\alpha, \mathcal{B})}{\partial \alpha}}{\frac{\partial G(\alpha, \mathcal{B})}{\partial \mathcal{B}}}
\tag{27}
$$

The numerator is given by

$$
\begin{aligned}
\frac{\partial G(\alpha, \mathcal{B})}{\partial \alpha} = \ &- \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}) \\
&+ \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}) + \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j}) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}),
\end{aligned}
$$

which is negative, because $-\mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B}) + \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j})$ is since $\mathcal{B} < 0$, and so is $\mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B}) - \mathbb{E}((\theta_{2j} \mid \theta_{2j} \geq \theta_{1j}))$.

The denominator is given by

$$
\frac{\partial G(\alpha, \mathcal{B})}{\partial \mathcal{B}} = 1 - \alpha \cdot \frac{\mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B})}{\partial \mathcal{B}} + \alpha \cdot \frac{\mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B})}{\partial \mathcal{B}}
$$

which is positive because the first and third term also are, while the second one is negative.

So, overall we have that $\mathcal{B}_\infty(\alpha)$ is increasing for $\alpha \in [0,\ 1]$. An increase in the share of Bayesian consumers make the reviews more biased, that is, decreases $\mathcal{B}_\infty(\alpha)$ further away from 0.

Case 2: Extension of Proposition 2

The proof follows very similar steps to that of **Case 1: Extension of Proposition 1**. Nevertheless, there are some differences, so we also report this one in its entirety.

Subtracting the second Eq. from the first in (25), we get:

$$
\begin{aligned}
\mathcal{B}_\infty = \ &\alpha \cdot \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathbb{E}(\mathcal{R}_1) + \mathbb{E}(\mathcal{R}_2) \\
&- \alpha \cdot \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2) \\
&+ (1 - \alpha) \cdot \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - Q_1 + Q_2) \\
&- (1 - \alpha) \cdot \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + Q_1 - Q_2).
\end{aligned}
\tag{28}
$$

First, we want to show that $\mathcal{B}_\infty(1) < \mathcal{B}_\infty(0)$. We have

$$
\begin{aligned}
\mathcal{B}_\infty(1) = \ &\mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathbb{E}(\mathcal{R}_1) + \mathbb{E}(\mathcal{R}_2)) \\
&- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2)).
\end{aligned}
$$

and

$$\mathcal{B}_\infty(0) = \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - Q_1 + Q_2)$$
$$- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + Q_1 - Q_2).$$

First notice that Proposition 2 implies that $\mathcal{B}_\infty(1) < 0$.

But then,

$$\mathcal{B}_\infty(0) = \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - Q_1 + Q_2)$$
$$- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + Q_1 - Q_2)$$
$$< \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathbb{E}(\mathcal{R}_1) + \mathbb{E}(\mathcal{R}_2))$$
$$- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2))$$
$$= \mathcal{B}_\infty(1).$$

where the inequality uses the fact that $\mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2)) - Q_1 + Q_2 = \mathcal{B}_\infty(1)$ is negative. Thus, $\mathcal{B}_\infty(0) < \mathcal{B}_\infty(1) < 0$: the equilibrium bias got *worse* if only Bayesian are present, compared to only naïves.

To show that we have monotonicity in $\alpha \in [0, \ 1]$, denote by

$$G(\alpha, \mathcal{B}) = \mathcal{B} - \alpha \cdot \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - (\mathcal{B} + Q_1 - Q_2)$$
$$+ \alpha \cdot \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B} + Q_1 - Q_2\big)$$
$$+ (1 - \alpha) \cdot \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - Q_1 + Q_2)$$
$$- (1 - \alpha) \cdot \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + Q_1 - Q_2). \tag{29}$$

where we have substituted $\mathbb{E}(\mathcal{R}_1) - \mathbb{E}(\mathcal{R}_2)) = \mathcal{B} + Q_1 - Q_2$.

Then, for every $\alpha \in [0, 1]$, $\mathcal{B}_\infty(\alpha)$ solves $G(\alpha, \mathcal{B}) = 0$. By the Implicit Function Theorem, we have that

$$\frac{\partial \mathcal{B}(\alpha)}{\partial \alpha} = -\frac{\frac{\partial G(\alpha, \mathcal{B})}{\partial \alpha}}{\frac{\partial G(\alpha, \mathcal{B})}{\partial \mathcal{B}}} \tag{30}$$

The numerator is given by

$$\frac{\partial G(\alpha, \mathcal{B})}{\partial \alpha} = - \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - (\mathcal{B} + Q_1 - Q_2)$$
$$- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B} + Q_1 - Q_2)$$
$$+ \mathbb{E}(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - Q_1 + Q_2)$$
$$- \mathbb{E}(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + Q_1 - Q_2),$$

which is negative, because the difference between its 1st and 3rd terms is, and the same is true for the 2nd and 4th.

The denominator is given by

$$\frac{\partial G(\alpha, \mathcal{B})}{\partial \mathcal{B}} = 1 - \alpha \cdot \frac{\partial \mathbb{E}\big(\theta_{1j} \mid \theta_{1j} \geq \theta_{2j} - \mathcal{B} - Q_1 + Q_2\big)}{\partial \mathcal{B}} + \alpha \cdot \frac{\partial \mathbb{E}\big(\theta_{2j} \mid \theta_{2j} \geq \theta_{1j} + \mathcal{B} + Q_1 - Q_2\big)}{\partial \mathcal{B}}$$

which is positive.

So, overall we have that $\mathcal{B}_\infty(\alpha)$ is decreasing for $\alpha \in [0, 1]$, as desired. An increase in the share of Bayesian consumers worsens the bias in reviews, that is, decreases $\mathcal{B}_\infty(\alpha)$ further away from 0. ■

Claim 2: Thus making naïve consumers strictly worse off.

Proof of Claim 2: Our welfare results in Proposition 5 show that welfare and bias go hand in hand: the more biased the reviews, the larger the welfare losses for naïve consumers. This result follows straightforwardly. ■

Proof of Proposition 4

We want to show that

$$Var\big(\theta_L \mid \theta_L > \theta_H + \Delta(Q)\big) < Var\big(\theta_H \mid \theta_H > \theta_H - \Delta(Q)\big), \quad \forall \Delta(Q) > \Delta^*(Q).$$

First notice that, when $\Delta(Q)$ approaches $\bar{\theta} - \underline{\theta}$, we have

$$Var\big(\theta_H \mid \theta_H > \theta_H - \Delta(Q)\big) \to Var(\theta_H).$$

On the other hand, the fact that $\theta_L > \theta_H + \Delta(Q)$ implies $\theta_L \in (\underline{\theta} + \Delta(Q), \, \bar{\theta})$. Therefore, Popoviciu's Inequality (Popoviciu, 1935) implies that

$$Var\big(\theta_L \mid \theta_L > \theta_H + \Delta(Q)\big) \leq \frac{1}{4}\big(\bar{\theta} - \underline{\theta} - \Delta(Q)\big)^2.$$

Notice that the right hand side gets arbitrarily small as $\Delta(Q) \to \bar{\theta} - \underline{\theta}$, implying the existence of a $\Delta^*(Q)$ such that $Var\big(\theta_L | \theta_L > \theta_H + \Delta(Q)\big) < Var(\theta_H)$ for every $\Delta(Q) > \Delta^*(Q)$. ■

B. Extensions

B.1. Duopoly with Outside Option

Assume now that consumers are not only decide *what*, but also *if* to buy. That is, they also have an outside option, of quality c. Without loss of generality, we assume that the outside option is non trivial, that is, it is chosen by at least some consumers. In particular, given that each consumer's taste shocks for the two products are independent, this is equivalent to requiring that a consumer with the lowest possible taste for both product 1 and product 2 would choose the outside option instead. In other words, this requires

$$c \geq \max(Q_1 + \underline{\theta}_1, Q_2 + \underline{\theta}_2).$$

For example, when $s_1 = s_2$ and $Q_1 > Q_2$, this amounts to requiring that $c > Q_2 + \underline{\theta}$. This condition is trivially satisfied for unbounded distribution (that is, in our leading example, the outside option is chosen by some consumers, even for very low values of c).

As before, upon choosing a product, each buyer reviews it honestly, but *subjectively*, by reporting their own experienced utility. Thus, in this case, Eq. (3) becomes

$$\mathcal{R}_{ij} := \begin{cases} \mathcal{U}_{ij} = Q_i + \theta_{ij} - P_i & \text{if } \mathbb{E}(\mathcal{U}_{ij}) \geq \max(\mathbb{E}(\mathcal{U}_{-ij}), c), \\ \varnothing & \text{otherwise.} \end{cases}$$

Denote by $\mathcal{J}_1^c$ and $\mathcal{J}_2^c$ the sets of buyers of product 1 and 2 respectively (we are omitting the t subscript for notational simplicity). That is,

$$\mathcal{J}_1^c = \{j \in \mathcal{J} \mid \mathbb{E}(Q_1) + \theta_{1j} - P_1 \geq \max(\mathbb{E}(Q_2) + \theta_{2j} - P_2, c)\},$$

and similarly for $\mathcal{J}_2^c$. It is immediate to see that, given the $\mathcal{J}_1$ and $\mathcal{J}_2$ defined in Eq. (4), we have $\mathcal{J}_1^c \subseteq \mathcal{J}_1$, $\mathcal{J}_2^c \subseteq \mathcal{J}_2$, with the inclusion being strict whenever the value of the outside option, c, is non-trivial. Moreover, for $c' > c$, we have $\mathcal{J}_1^{c'} \subseteq \mathcal{J}_1^c$ and $\mathcal{J}_2^{c'} \subseteq \mathcal{J}_2^c$. These simple observations form the basis for our first result: the presence of an outside option increases the reviews of each product, and decreases the number of reviews for each product.

We are interested in studying the robustness of our main results – in particular, Proposition 1, 2 and 4 – to the inclusion of an outside option.

We start from Proposition 1. We have the following result:

Proposition 5 (More Polarizing Products Are Relatively Overrated) *Let the quality of the outside option be c, and let the two products differ only in their design: $Q_1 = Q_2$, $s_1 = H$, $s_2 = L$. Assume that s_1 and s_2 are symmetric.[39] Then, in the long-run ($t = \infty$):*

---

39. Symmetry is a sufficient, but not necessary, condition.

- *Both products' reviews are higher than they would be absent an outside option, and increasing in the outside option quality, c.*
- *The more polarizing product 2 is relatively overrated: $\mathcal{B}_\infty < 0$,*
- *and thus captures a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.*
- *Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_0 < \mathcal{B}_\infty < 0$, $\mathcal{N}_1(\mathcal{R}_2) > \mathcal{N}_\infty(\mathcal{R}_2) > 1/2$.*

Proposition 6 (High Quality Products Are Relatively Underrated.) *Let the quality of the outside option be c, and let the two products differ only in their qualities: $Q_1 > Q_2$, $s_1 = s_2$. Then, in equilibrium ($t = \infty$):*

- *The higher quality product 1 has higher reviews: $\mathbb{E}_\infty(\mathcal{R}_1) > \mathbb{E}_\infty(\mathcal{R}_2)$,*
- *Despite being relatively underrated: $\mathcal{B}_\infty < 0$.*
- *It thus obtains a higher number of reviews: $\mathcal{N}_\infty(\mathcal{R}_1) > 1/2$, but less than it would if consumers were fully informed.*
- *Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $\mathcal{N}_\infty(\mathcal{R}_1) > \mathcal{N}_2(\mathcal{R}_1)$.*
- *Despite these distortions, reviews unambiguously increase consumer welfare.*

Proposition 7 (The Variance of Reviews Needs Not Proxy Product Design) *Let the quality of the outside option be c, and let the support for θ_1 and θ_2 be bounded above. Then, there exists a quality gap $Q := Q_1 - Q_2 > 0$ such that $Var_\infty(\mathcal{R}_1) > Var_\infty(\mathcal{R}_2)$ for all s_1 and s_2.*

Now that we have established robustness of our main results to the introduction of an outside option, we turn to a different question: quantitatively, does the outside option mitigate or worsen the biases? We show, directly, that the answer to this question depends on the outside option's quality, c. That is, one can find examples such that an outside option of quality c increases $\mathcal{B}_\infty$, while an outside option of quality c' decreases $\mathcal{B}_\infty$.

One such example can be found in Figure 4. Here, we can clearly see that low quality outside options worsen the bias in reviews, while higher quality ones mitigate (but fail to erase) it. Of course, this is a byproduct of the normal distribution. It is plausible that other distributions would cause the dependence of $\mathcal{B}_\infty^c$ on c to change. However, what the last three Propositions show is that the qualitative nature of our main results is robust to the introduction of an outside option.

B.2. Learning from Cumulative Reviews

Throughout our paper, period t consumers learn from the reviews of their predecessors, generation $t - 1$. We now turn to studying the (empirically realistic) case of consumers learning from a (possibly weighted) average of reviews up to time t. To

this end, we first formalize the average of cumulative reviews as follows. For $\beta \in [0,1]$, denote by $\mathbb{E}_t(\mathcal{R}^c)$ the average of cumulative reviews at time t:

$$\mathbb{E}_t(\mathcal{R}^C) = \frac{\sum_{\tau=0}^{\tau=t} \beta^{t-\tau} \mathcal{N}_\tau(\mathcal{R}) \mathbb{E}_\tau(\mathcal{R})}{\sum_{\tau=0}^{\tau=t} \beta^{t-\tau} \mathcal{N}_\tau(\mathcal{R})}. \tag{31}$$

As the formula makes apparent, $\mathbb{E}_t(\mathcal{R}^C)$ depends on both the average and the number of reviews that the product received in each period $\tau = 0, 1, ..., t$. The (backward) discount factor $\beta^{t-\tau}$ measures how much the platform underweights past reviews compared to more recent ones.[40] If $\beta = 0$, $\mathbb{E}_t(\mathcal{R}^C) = \mathbb{E}_t(\mathcal{R})$, and we recover the case studied throughout this paper. Conversely, if $\beta = 1$,

$$\mathbb{E}_t(\mathcal{R}^C) = \frac{\sum_{\tau=0}^{\tau=t} \mathcal{N}_\tau(\mathcal{R}) \mathbb{E}_\tau(\mathcal{R})}{\sum_{\tau=0}^{\tau=t} \mathcal{N}_\tau(\mathcal{R})},$$

so that $\mathbb{E}_t(\mathcal{R}^C)$ is simply the weighted average of all reviews, with the weights being given solely by the number of reviews in each period. Building on this definition, we also define the cumulative reviews advantage of product 1 as $\Delta_t^C(\mathcal{R}) = \mathbb{E}_t(\mathcal{R}_1^C) - \mathbb{E}_t(\mathcal{R}_2^C)$.

Clearly, our definition of cumulative reviews implies that, for every β, $\mathbb{E}_0(\mathcal{R}^C) = \mathbb{E}_0(\mathcal{R})$. The two definitions start to differ in period 1. For example, if reviews for product i decrease in period 1, ($\mathbb{E}_1(\mathcal{R}_i) < \mathbb{E}_0(\mathcal{R}_i)$), then naturally cumulative reviews also do ($\mathbb{E}_1(\mathcal{R}_i^C) < \mathbb{E}_0(\mathcal{R}_i^C)$), but less so than period by period ones ($\mathbb{E}_1(\mathcal{R}_i^C) > \mathbb{E}_1(\mathcal{R}_i)$). As a result, period 3 consumers will form different beliefs and make different choices, resulting in different period 3 reviews and, *a fortiori*, a difference between $\mathbb{E}_2(\mathcal{R}^C)$ and $\mathbb{E}_2(\mathcal{R})$.

Therefore, the dynamics of consumers' self-selection patterns, and thus reviews, are different depending on whether $\beta = 0$ (as is the case in our main body of the paper) or $\beta > 0$. We emphasize that not only there is a difference between $\mathbb{E}_t(\mathcal{R}^C)$ and $\mathbb{E}_t(\mathcal{R})$, but also that period by period reviews are different in the two cases, since they result from different learning dynamics. Figure 5, for instance, shows that cumulative reviews are smoother than period by period ones and, as a natural consequence of this fact, that when consumers learn from cumulative reviews, their period by period reviews are also smoother.

Given the totally different learning dynamics in the two cases, it is then natural to ask whether the reviews' convergence in the long-run, as well as the long-run bias (if any), depend on β. The next Proposition offers a striking answer to this question: despite completely different learning and review dynamics, both sequences of reviews converge, and moreover, their long-run limit is unchanged. Thus, the assumption

---

40. Overweighting (or overemphasizing) recent reviews is a widespread practice on many online platforms. In our context, quality is fixed over time, so that, in principle, older reviews are just as informative as more recent ones. Ultimately, as we show explicitly, the results holds independently of $\beta \in [0,1]$, and therefore it is worth presenting this model extension in its most general form.

that $\beta = 0$ made throughout the paper is without loss of generality for the purpose of studying long-run biases in reviews.

We start with a straightforward Lemma:

Lemma 3. *For every $\beta \in [0,1]$ and every t, $\mathbb{E}_t(\mathcal{R}^C)$ satisfies the following recursive equation:*

$$\mathbb{E}_t(\mathcal{R}^C) = \frac{\beta \mathbb{E}_{t-1}(\mathcal{R}^C) + \mathcal{N}_t(\mathcal{R})\mathbb{E}_t(\mathcal{R})}{\beta + \mathcal{N}_t(\mathcal{R})}. \tag{32}$$

Similarly, $\Delta_t^C(\mathcal{R})$ satisfies

$$\Delta_t^C(\mathcal{R}) = \frac{\beta \mathbb{E}_{t-1}(\mathcal{R}_1^C) + \mathcal{N}_t(\mathcal{R}_1)\mathbb{E}_t(\mathcal{R}_1)}{\beta + \mathcal{N}_t(\mathcal{R}_1)} - \frac{\beta \mathbb{E}_{t-1}(\mathcal{R}_2^C) + \mathcal{N}_t(\mathcal{R}_2)\mathbb{E}_t(\mathcal{R}_2)}{\beta + \mathcal{N}_t(\mathcal{R}_2)}. \tag{33}$$

Lemma 3 follows immediately from the definition of $\mathbb{E}_t(\mathcal{R}^C)$ in Eq. (31), and, in our context, we show that it implies that $\Delta_\infty(\mathcal{R}^C)(\beta)$ converges whenever $\Delta_\infty(\mathcal{R})$ does. Moreover, the two convergence points are the same, as shown in the following Proposition.

Proposition 8. *$\Delta_t(\mathcal{R}^C)(\beta)$ converges whenever $\Delta_t(\mathcal{R})(\beta)$ does. Moreover, they converge to the same long-run outcome:*

$$\Delta_\infty(\mathcal{R}^C)(\beta) = \Delta_\infty(\mathcal{R}) \quad \forall \beta \in [0,1].$$

This result has important implications, since it shows that in the general context of our model, one can study learning from cumulative or period by period reviews interchangeably – at least when the focus in on the long-run properties of reviews. Clearly, if the goal were to study short-term changes in reviews (for example after a popularity shock like a major award, or after a seller obtains fake reviews), then the learning technology would quantitatively matter. We stress, however, that even in this case our results would qualitatively hold true. For example, in the case of fake reviews, the large pool of buyers immediately following would decrease period by period reviews and, albeit by a lesser amount, cumulative ones.

B.3. Learning About Taste

We now discuss the possibility of consumers employing reviews to simultaneously learn about both quality and fit. For instance, while some determinants of consumer-product fit can be easily observable by consumers even absent any reviews (e.g., the genre of a movie, or the cuisine of a restaurant), others might be more subtle, and thus be learned over time through reviews. For example, consumers might pick up more of a movie's characteristics over time, or get more precise information about the atmosphere of a restaurant.

As we highlight in our conclusions (Section 7), modeling social learning about the latter is not straightforward, because taste is *iid* across consumers. Nevertheless, one

could think of a simple model in which, in each period, consumers' perceived taste for each product, which we denote by $\tilde{\theta}_{ij}$, is a weighted average of their actual taste θ_{ij}, plus an uninformative signal $\xi_{ij} \sim H(\cdot)$, also with mean 0, which is uncorrelated with θ_{ij}:

$$\tilde{\theta}_{ij} = \rho(t)\theta_{ij} + (1 - \rho(t))\xi_{ij}. \tag{34}$$

$\rho(t)$, the weight assigned to the informative signal, can be assumed to be increasing over time ($\frac{\partial \rho(t)}{\partial t} > 0$) and bounded above by 1 ($\lim_{t \to \infty} \rho(t) \leq 1$).

This is essentially equivalent to assuming that, in each period, each consumer observes a signal of its match for each product, $s_{ij} \sim \mathcal{N}(\theta_{ij}, 1/\rho(t))$, where $\rho(t)$ is an increasing function of t. It is also essentially equivalent to assuming that a fraction $\rho(t)$ of consumers in each period are aware of their taste for each product, while a fraction $1 - \rho(t)$ is not.

It follows from Eq. (34) that all of the conditional taste distributions governing the dynamics of our model would change to

$$\begin{aligned}
&\mathbb{E}_t(\theta_{ij}|\tilde{\theta}_{ij} > \tilde{\theta}_{-ij} - \Delta_t(\mathcal{R})) \\
=&\mathbb{E}_{t+1}(\theta_{ij} \mid \rho(t)\theta_{ij} + (1 - \rho(t))\xi_{ij} > \rho(t)\theta_{-ij} + (1 - \rho(t))\xi_{-ij} - \Delta_t(\mathcal{R})).
\end{aligned} \tag{35}$$

When $t = 0$, it is easy to see that Eq. (35) implies

$$\begin{aligned}
&\mathbb{E}_0(\theta_{ij} \mid \tilde{\theta}_{ij} > \tilde{\theta}_{-ij} - \Delta_t(\mathcal{R})) \\
=&\mathbb{E}_0(\theta_{ij} \mid \xi_{ij} > \xi_{-ij} - \Delta_t(\mathcal{R})) = 0,
\end{aligned}$$

due to the fact that ξ_{ij} is uncorrelated with actual taste, θ_{ij}. As a consequence of this fact, as Lemma 1 shows, period 0 reviews are unbiased, since they do not reflect any form of taste-based self-selection. However, as t increases, so does $\rho(t)$. If $\rho(t)$ eventually reaches 1 (say, for $t = \bar{t}$), we recover our original model for every $t \geq \bar{t}$.

None of our findings would be qualitatively affected by this (admittedly simple) modification. What the above formula for the conditional expectation of θ_{ij} implies is that the biases would get stronger over time, as consumers become more aware of their match with each product. This can be seen as a continuous equivalent (or extension) of Lemma 2, which states that absent taste-based self-selection, all biases disappear. In this slightly modified model, the stronger taste-based self-selection, the larger the biases.

B.4. Alternative Reviewing Behavior

We now briefly discuss our model's robustness to changes in its core assumptions. Particularly, we have made three key assumptions for our analysis: i) reviews are subjectively honest, that is, each consumer reports their subjective utility upon purchasing a product, ii) no self-selection at the writing stage, conditional on purchase:

everyone purchasing a product reviews it and *iii*) consumers are choosing between two options.

Inspired by both empirical realism and the sizable existing literature already presented in Section 2 (and further discussed here), we consider the following extensions.

Self-Selection Into Leaving Reviews. In our model, every consumer leaves a review upon purchasing a product. In reality, very few consumers leave reviews: a variety of surveys estimate this percentage to lie between 1% and 5%, depending on the market.

It is important to stress that, especially in a model (like ours) in which reviews are not subject to noise (see discussion in Section 4.4), this fact *per se* would be inconsequential for our findings whenever self-selection into review conditional on choice is orthogonal to the nature of the review.

However, this need not be the case. Perhaps the most common form of self-selection on writing conditional on choice documented in this context is *extremity bias* (see *e.g.* Brandes et al. (2018) and citations therein). Put simply, consumers with strong feelings towards the product – whether positive or negative – are more likely to express them compared to their peers that feel neutral towards it.

It is interesting to spell out how extremity bias would affect our results. To this end, assume that consumers in both tails (say, consumers that are either below the $10th$ percentile or above the $90th$ in their idiosyncratic taste for the product) are the only ones to leave reviews.[41] Denote by $\mathcal{J}_i^{10-}$ and $\mathcal{J}_i^{90+}$ these two camps of buyers for product i. Then, the average conditional taste for the product as reflected by reviews will be given by

$$\frac{1}{2}\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i^{10-}) + \frac{1}{2}\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i^{90+}).$$

How does this compare to the case without extremity bias, $\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i)$? It is immediate to see that the two are equal for symmetric distributions. So, for instance, all of the numerical results in our Section 3 would be unaffected by this change.

Our conclusions become less sharp whenever the skew of the distribution changes. In this case, one can imagine two products with the same quality, same variance in θ_{ij}, same prices, and yet different reviews resulting from asymmetries in $\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i^{10-})$ and $\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i^{90+})$.

In this case, for instance, a product that is loved by few and mildly (dis)liked by many might do better than one that is appreciated – but not loved – by most, in line with Proposition 1.

It is *a priori* unclear how this dimension of heterogeneity would interact with the other bias we discuss in this paper, and particularly in Propositions 2. A more in depth analysis of the nature (and dynamics) of reviews in light of this bias is beyond the scope of this paper, and seems like a noteworthy research question.

---

41. One could also assume that these consumers are simply more likely to post reviews, and not the only ones to do so. This would not affect any of our reasoning below.

Another interesting case is the one in which it is the absolute – not relative – levels of love or hate for the products that shapes self-selection into reviewing. That is, consumer j leaves a review for product i when either $U_{ij} > \bar{U}$ or $U_{ij} < \underline{U}$, for two consumer- and product-independent thresholds $\underline{U} < \bar{U}$.

Under these assumptions, the average reviews of low quality products would be downward biased, while the opposite is true for products of high quality, contrary to Proposition 2 and somewhat similarly (though with slightly different drivers) to Park et al. (2021).

When niche products are also of lower quality – which has been shown to be the case in a variety of contexts, see Johnson and Myatt (2006), Bar-Isaac et al. (2012) and Sun (2012) – the conclusions are ambiguous. Again, spelling these out in greater detail seems like a promising avenue for future research.

Reviewing to Persuade. In our model, consumers are not strategic in their review behavior. They simply report their subjective opinion regarding the chosen option, irrespective of the impact of their reviews on their successors. This assumption is psychologically realistic, and additionally justified by the consumers' desire to receive future personalized recommendations, which is an important driver of review behavior on *Netflix, Yelp* and *Goodreads*, among other platforms.

Nevertheless, it is interesting to briefly discuss the case of consumers leaving reviews with the explicit desire to be persuasive (as is the case, for instance, in Chakraborty et al. (2022)). Generally, consumers motivated by persuading their peers will not rate truthfully. To see this, consider a consumer who believes that a product is of good quality (say, 4 out of 5), and before posting, notices that the product currently has an average review of 3.5. Then, her best response is to inflate her review to 5, to get the *ex-post* average review closer to her subjective quality assessment, 4.

That is, for a product of quality Q_i for which she has taste θ_{ij}, a period $t+1$ consumer reacting to period t reviews would seek to minimize the strategic (S) loss function
$$L^S(\mathcal{R}_{ij} \mid Q_i, \theta_{ij}, \mathcal{R}_i) := -\big(Q_i + \theta_{ij} - \mathbb{E}_{t+1}(\mathcal{R}_i|\mathcal{R}_{ij})\big)^2$$
instead of the purely individual (I) one
$$L^I(\mathcal{R}_{ij} \mid Q_i, \theta_{ij}, \mathcal{R}_i) := -\big(Q_i + \theta_{ij} - \mathcal{R}_{ij}\big)^2.$$

An in-depth study of social learning with strategic review behavior is beyond the scope of this paper, and seems a promising area for future research (as also suggested by Acemoglu et al. (2022)). Here, we will only add two observations that mitigate concerns regarding the possibility (and impact) of strategic review in this context.

The first one is that $\mathbb{E}_{t+1}(\mathcal{R}_i|\mathcal{R}_{ij}) \approx \mathbb{E}_{t+1}(\mathcal{R}_i)$ whenever the number of reviews that the product had already received is large. In other words, the ability to move the average is limited when such average is built on a high number of reviews. Thus, $L^S(Q_i - \theta_{ij} - P_i \mid Q_i, \theta_{ij}, P_i, \mathcal{R}_i) \approx \max_{\mathcal{R}_{ij}} L^S(\mathcal{R}_{ij} \mid Q_i, \theta_{ij}, \mathcal{R}_i)$. This is usually the

case on many online platforms such as *Netflix, Goodreads* and *IMDb*, in which every product has several thousands (and often millions) of reviews.

Second, notice that for each $j^* \in \mathcal{J}_i$, it is straightforward to sign the difference between individual and strategic reviews, $\mathcal{R}_{ij}^I - \mathcal{R}_{ij}^S$:

$$\mathcal{R}_{ij}^I < (>)\mathcal{R}_{ij}^S \;\Leftrightarrow\; \mathbb{E}(\mathcal{R}_i) < (>)Q_i + \theta_{ij^*} \;\Leftrightarrow\; \mathbb{E}(\theta_{ij}|j \in \mathcal{J}_i^t) < (>)\theta_{ij^*}.$$

In other words, consumer j^* strategic review is lower than the truthful one if and only if consumer j^* has a lower taste for the product than the average period t rater.

Much like we have seen in Proposition 2 and Corollary 1, this also gives rise to self-defeating review dynamics: products with very high reviews will motivate future strategic consumers to skew their reviews down in order to have an impact, and the opposite is true for products with low reviews. Therefore, assuming strategic motives strengthen our conclusions that reviews are pushed to the middle, understating quality differences and thus penalizing higher quality products.

Social Influence. The deviation from truthful review behavior that we have just highlighted is not the only possible one. Contrary to the contrarian-like behavior of a reviewer who has a desire to sway future consumers towards her preferred options, one can imagine at least some reviewers' opinions are at least partly reflective of (that is, anchored to) those of their predecessors.

This phenomenon is an example of social influence (see Muchnik et al. (2013) and citations therein) and can be conceptualized as "biasing the judgement of an experience – and, thus, adapting one's review – in the direction of what previous consumers have reported".

For instance, if every consumer in the previous generation has left a product glowing reviews, future consumers will rate the product higher if they were to consume it in isolation. That is,

$$\frac{\partial \mathbb{E}_{t+1}(\mathcal{R})}{\partial \mathbb{E}_t(\mathcal{R})} > 0.^{42}$$

Social influence is an important force in the digital world. For example, Muchnik et al. (2013) demonstrate, using a large scale field experiment, that randomly manipulating the first upvote or downvote received by a user post on a popular online forum influences the post's long-term upvotes to downvotes ratio. Similarly, Jacobsen (2015) shows that when famous beer bloggers review a beer more positively or negatively than the average of consumers, future consumer reviews shift in the direction of the bloggers' opinion.

This type of review behavior is often opposite to the one described in the subsection above. There, consumers effectively look as contrarians (despite their lack of

---

42. While beyong the scope of our paper, it is interesting to notice that this could be either because the *perceived* consumption utility went up, $\frac{\partial U_{ij}^{t+1}\left(Q_i, \theta_{ij}, \mathbb{E}_t(\mathcal{R}_i)\right)}{\partial \mathbb{E}_t(\mathcal{R}_i)} > 0$, or because reviews went up for a given U_{ij}, reflecting the rater's desire to conform to the raters in the previous period.

social image concerns), since that is what is required to affect the average review. Here, consumers have a desire to conform (or they perceive products differently depending on the previous reviews), and thus they conform to the crowd preceding them. From a learning standpoint, conformity is dangerous in this setting, because – much like in the classic work of Banerjee (1992) and Bikhchandani et al. (1992) – it leads to a halt in the aggregation of information.

Our model is *not* robust to social influence, and in fact generates prediction that are to it, as discussed at length in both Section 1 and Section 5.2. Clearly, the presence of social influence leads to *winners-take-all* dynamics: better reviews today translate into (more and) better ones tomorrow. Particularly, the opinions of particularly influential members should sway not only readers' choices, but also their very perceptions conditional on that.

We believe that a variety of empirical findings – including those of Kovács and Sharkey (2014), Rossi (2021), and He et al. (2022) – offer substantial evidence that social influence is not prevalent in this context and, if anything, high reviews (and thus sales) end up *hurting* a product's future review, as described in our Proposition 2 and Corollaries 1 and 2.

Understanding when social influence is the dominant force, and when, on the contrary, taste-based self-selection leads to robust (but potentially biased) reviews seems like a promising research question moving forward.

Online Appendix for *Alone, Together*: A Model of Social (Mis)Learning from Consumer Reviews

A. Optimal Pricing

Throughout our analysis up to this point, we have fully focused on the demand side, that is, we have studied consumer (mis)learning in a context in which firms were "passive". It is interesting to now consider a more strategic model in which firms react to – and, as we will see, influence – the information environment by setting their optimal prices. For analytical simplicity, we go back to the case of duopoly without an outside option – which also nests the case of monopoly with an outside options, if we assume that one of the two products, i, is such that $\theta_{ij} \equiv 0 \ \forall j$.

Our goal is two-fold: first, we aim to gather insights on the robustness (or lack thereof) of our findings to endogenous prices, that is, study how optimal prices influence properties of long-run reviews. Second, conversely, we want to study how the presence of reviews influence firms' pricing.

To this end, assume, as in the previous analysis, that firm i has quality Q_i, design s_i and set price P_i, $i = 1, 2$. Denote by $\Delta(Q) := Q_1 - Q_2$ the quality advantage of firm 1 and assume, without loss of generality, that it is non-negative. We define the *relative* taste for firm 1 by $\theta_j \equiv \theta_{1j} - \theta_{2j}$, with $\theta_{ij} \sim F_i(\cdot)$, and by $G(\cdot)$ and $g(\cdot)$ its cumulative and density respectively. We assume that $G(\cdot)$ is symmetric and that it satisfies the monotone hazard rate, that is, $h(\cdot) := \frac{g(\cdot)}{1-G(\cdot)}$ is non-decreasing. For brevity, we will drop the consumer subscript j from this point on.

To emphasize the role of subjective reviews on prices (and, conversely, that of optimal prices on reviews), we compare this case to two opposite benchmarks: no information and full information. Thus, we characterize pricing equilibria in three informational environments: 1) **no information** (N), in which consumers are unaware of quality differences, 2) **full information** (F) , in which the consumers are aware of the firms' qualities Q_1 and Q_2, and thus their difference $\Delta(Q)$, and 3) **subjective reviews** $(\mathcal{R})$, in which buyers of each products report their utility from it, as previously described in this paper. In this third case, consistent with the rest of our paper, we assume consumers take reviews at face value.[1]

Throughout the analysis, we denote consumers' beliefs about firm 1's quality advantage by $\tilde{\Delta}(Q)$. In particular, we have $\tilde{\Delta}(Q) = 0$, $\tilde{\Delta}(Q) = \Delta(Q)$ and $\tilde{\Delta}(Q) = \Delta_\infty(\mathcal{R})$ in the no information, full information and subjective reviews cases, respectively.[2]

---

1. The second case can be seen as one in which all consumers are Bayesian, and thus fully internalize, and correct for, potential biases in subjective reviews.

2. Throughout this Section, to ensure consistency between the three cases, reviews should be thought of as reflecting gross utilities: $\mathcal{R}_{ij} = Q_i + \theta_{ij}$. Equivalently, we assume that, since prices are observable, if reviews were reflecting net utilities instead, consumers could simply account for prices. We refer to Section 4.2 for a more detailed discussion of this point.

First, we fully characterize the properties of optimal prices in each of these three environments. Then, we compare prices across environments.

A.1. No Information

Given the lack of information regarding firms' (relative) qualities, each consumer simply trades off price and fit for each product. That is, a consumer buys from firm 1 if and only if

$$\theta_j \geq P_1 - P_2.$$

Thus, we have

$$\pi_1(P_1, P_2) = P_1 \cdot \big(1 - G(P_1 - P_2)\big)$$

and

$$\pi_1(P_1, P_2) = P_2 \cdot G(P_1 - P_2).$$

We have the following

Proposition 9. *In the unique equilibrium,*

$$P_1^N = P_2^N = \frac{G(0)}{g(0)} = \frac{1}{2g(0)}. \tag{36}$$

A.2. Full Information

We now turn to the case in which consumers are fully aware of the quality difference, $\Delta(Q)$. We have the following:

Proposition 10. *In the unique equilibrium, we have*

$$P_1^F = \frac{1 - G(\Delta(P) - \Delta(Q))}{g(\Delta(P) - \Delta(Q))} \tag{37}$$

and

$$P_2^F = \frac{G(\Delta(P) - \Delta(Q))}{g(\Delta(P) - \Delta(Q))}. \tag{38}$$

Comparing to the no-information case, we have $P_2^F < P_2^N = P_1^N < P_1^F$. Last, firm 1 captures more than half of the market.

A.3. Subjective Reviews

Now consider the case in which reviews are present, and consumers take them at face value, that is, equating $\tilde{\Delta}(Q) = \Delta_\infty(\mathcal{R})$.

We have

$$\mathbb{E}_\infty(\mathcal{R}_1) = Q_1 + \mathbb{E}(\theta_1 \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_1 - P_1 > \mathbb{E}_\infty(\mathcal{R}_2) + \theta_2 - P_2))$$

2

and

$$\mathbb{E}_\infty(\mathcal{R}_2) = Q_2 + \mathbb{E}(\theta_2 \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_2 - P_2 > \mathbb{E}_\infty(\mathcal{R}_1) + \theta_1 - P_1)).$$

This is in line with the rest of our paper. In this context, the key observation is that reviews do not *directly* reflect prices. While not a crucial assumption, this is justified by the fact that consumers observe prices directly, so that, if reviews reflected them, consumers could easily invert reviews to then employ them as (potentially biased) quality signals.[3]

It is useful to combine the previous two equations into

$$\Delta_\infty(\mathcal{R}) = \Delta(Q) + \mathbb{E}\big(\mathbb{E}\big(\theta_1 \mid \Delta\theta > \Delta(P) - \Delta_\infty(\mathcal{R})\big) \\ - \mathbb{E}\big(\theta_2 \mid \Delta\theta < \Delta(P) - \Delta_\infty(\mathcal{R})\big) \tag{39}$$

Crucially, notice how firm 1 can increase the RHS of Eq. 39 by increasing P_1, and the symmetric is true for firm 2 increasing P_2. By setting a high price a firm improves the self-selection of its consumers, while worsening the self-selection of its competitor's consumers.

We start with three Lemma's, each instrumental to proving our main Proposition:

Lemma 4. $0 < \frac{\partial \Delta_\infty(\mathcal{R})}{\partial \Delta(P)} < 1.$

Lemma 5. *If* $\Delta(P^\mathcal{R}) < \Delta(Q)$, *then* $\Delta(P^\mathcal{R}) < \Delta_\infty(\mathcal{R}) < \Delta(Q)$; *If* $\Delta(P^\mathcal{R}) > \Delta(Q)$, *then* $\Delta(Q) < \Delta_\infty(\mathcal{R}) < \Delta(P^\mathcal{R})$.

Lemma 6. *In equilibrium, we have* $0 < \Delta(P^\mathcal{R}) < \Delta(Q)$.

We are now ready to state the following:

Proposition 11. *Equilibrium prices are given by*

$$P_1^\mathcal{R} = \frac{1 - G(\Delta(P) - \Delta_\infty(\mathcal{R}))}{g(\Delta(P) - \Delta_\infty(\mathcal{R}))\left(1 - \frac{\partial \Delta_\infty(\mathcal{R})}{\partial \Delta(P)}\right)} \tag{40}$$

and

$$P_2^\mathcal{R} = \frac{G(\Delta(P) - \Delta_\infty(\mathcal{R}))}{g(\Delta(P) - \Delta_\infty(\mathcal{R}))\left(1 - \frac{\partial \Delta_\infty(\mathcal{R})}{\partial \Delta(P)}\right)}. \tag{41}$$

In equilibrium, $0 < \Delta(P^\mathcal{R}) < \Delta_\infty(\mathcal{R}) < \Delta(Q)$.

---

3. For an alternative modeling choice, in which consumers are unaware of whether reviews include prices, and thus fail to correct for them, see Carnehl et al. (2023).

First and foremost, Proposition 11 demonstrates the robustness of our main results to the introduction of optimal pricing by firms. This is an important insight, as one could have suspected prices to work as a corrective mechanism, undoing the reviews' biases. Proposition 11 also has a host of additional, interesting implications. We start by noticing that, because $\Delta(P^{\mathcal{R}}) - \Delta_\infty(\mathcal{R}) < 0$, we have

$$P_1^{\mathcal{R}} > \frac{1 - G(0)}{g(0)\left(1 - \frac{\partial \Delta \mathcal{R}}{\partial \Delta(P)}\right)} > \frac{1 - G(0)}{g(0)} = P_1^N. \tag{42}$$

The two inequalities emphasize the presence of two totally distinct forces, each pushing Firm 1's prices up. First, because $\Delta_\infty(\mathcal{R}) > 0$, the quality advantage of Firm 1 is now known to consumers. Second, both firms have additional incentives to increase their prices to improve the matches with their consumers ("price as a matching device"), while decreasing their competitor's matches, thus increasing their relative reviews: this is quantified by the $1/\left(1 - \frac{\partial \Delta_\infty(\mathcal{R})}{\partial \Delta(P)}\right)$ term appearing in both $P_1^{\mathcal{R}}$ and $P_2^{\mathcal{R}}$.

The analysis for firm 2 is more ambiguous, since the two aforementioned forces point in opposite directions: while (part of) the quality disadvantage of firm 2 is revealed, firm 2 also has the same incentive to increase prices to improve matches, and thus reviews, as firm 1.

A.4. Combining the Three Cases

Despite the aforementioned ambiguity, we find that, strikingly, when $\Delta(Q)$ is not too large, both firms' prices are highest in the case of subjective reviews.

Proposition 12. *Suppose that* $\frac{d\mathbb{E}(\theta_1|\Delta\theta > k)}{dk} + \frac{d\mathbb{E}(\theta_2|\Delta\theta < k)}{dk} > \epsilon > 0$ *in a neighborhood of* 0 *(for k).*[4] *Then, there exists a* $\Delta(Q^*)$ *such that whenever* $\Delta(Q) < \Delta(Q^*)$, *we have*

$$P_1^{\mathcal{R}} > P_1^F > P_1^N, \quad P_2^{\mathcal{R}} > P_2^N > P_2^F. \tag{43}$$

Taken together, our findings in Section A highlight several interesting facts. First and most obviously, the combination of subjective reviews and consumer naïvete can lead both firms to price higher compared to both the no information and the full information case. This happens despite the fact that, from an informational point of view, we have shown the subjective review case to always lie between the other two: $0 < \Delta_\infty(\mathcal{R}) < \Delta(Q)$. The specific price rankings highlighted in Proposition 12 crucially depend on $\Delta(Q)$ to be small – if this condition is violated, then firm 2 would be revealed to be at a large quality disadvantage, which would decrease its prices compared to the no information case, leading to $P_2^N > P_2^{\mathcal{R}}$, while firm 1's perceived quality advantage would be much smaller than in the full information case,

---

4. The condition that $\frac{d\mathbb{E}(\theta_1|\Delta\theta > k)}{dk} + \frac{d\mathbb{E}(\theta_2|\Delta\theta < k)}{dk} > \epsilon > 0$ is easily satisfied by many distributions. For example, it is satisfied if θ_1 and θ_2 are both uniformly distributed or both normally distributed.

leading to $P_1^F > P_1^{\mathcal{R}}$. But the more general point does not: with subjective reviews that depend on average conditional taste distributions, both firms have incentives to price higher than they would otherwise, since prices act as *matching devices*, or, put differently, as a reviews management tool: a higher relative taste is required to pick a product when its relative price is higher.

Perhaps more importantly for our analysis, Proposition 11 highlights that Proposition 1 and 2 are robust to the introduction of optimal pricing by both firms. It is interesting to justaxpose our findings with those of Sayedi (2018), who shows that the pathological outcomes in the classic observational learning models of Banerjee (1992) and Bikhchandani et al. (1992) disappear whenever the two firms are allowed to price optimally. This is because, with observational learning, the higher quality firm is incentivized to dramatically lower its prices whenever a cascade on the lower quality product starts. Here, on the other hand, when quality differences are not overwhelming, both firms are incentivized to increase their equilibrium prices to obtain better matches, and thus higher reviews, than they would if reviews were absent altogether.

Because equilibrium reviews reveal some information ($\Delta_\infty(\mathcal{R}) > 0$), firm 1 still partly cashes in on its revealed quality advantage, and thus sets higher prices than firm 2. Therefore, the number of reviews for the two products are closer to $0.50 - -0.50$ than they would be with symmetric prices. But the point remains that this correction is only partial, and all the biases highlighted in our analysis ($\Delta_\infty(\mathcal{R}) < \Delta(Q)$) are robust to optimal pricing by both firms.

B. Additional Proofs

B.1. Proofs for Main Text Results

Proof of Corollary 1

We have that the period $t+1$ set of consumers for product i is given by:

$$\mathcal{J}_i^{t+1} = \{j \mid \theta_{ij} + \mathbb{E}_t(\mathcal{R}_i) > \theta_{-ij} + \mathbb{E}_t(\mathcal{R}_{-i})\}$$

Now let $\mathbb{E}_t(\mathcal{R}_i)$ increase to $\mathbb{E}_t(\mathcal{R}_i) + \epsilon$. Clearly, this implies that $\mathcal{J}_i^{t+1,\epsilon} \supset \mathcal{J}_i^{t+1}$. It also implies that the crowd of product i buyers becomes less self-selected:

$$\mathbb{E}(\theta_{ij} \mid \theta_{ij} + \mathbb{E}_t(\mathcal{R}_i) + \epsilon > \theta_{-ij} + \mathbb{E}_t(\mathcal{R}_{-i}))) < \mathbb{E}(\theta_{ij} \mid \theta_{ij} + \mathbb{E}_t(\mathcal{R}_i) > \theta_{-ij} + \mathbb{E}_t(\mathcal{R}_{-i}))),$$

which in turns causes $\mathbb{E}_{t+1}(\mathcal{R}_i)$ to decline. The case of $\mathbb{E}_{t+1}(\mathcal{R}_{-i})$ can be handled symmetrically.

Clearly, because $\mathcal{B}_{t+1} := (\mathbb{E}_{t+1}(\mathcal{R}_1) - \mathbb{E}_{t+1}(\mathcal{R}_2)) - (Q_1 - Q_2)$, a decrease in $\mathbb{E}_{t+1}(\mathcal{R}_1)$ and an increase in $\mathbb{E}_{t+1}(\mathcal{R}_2)$ will compound to cause a larger decrease in $\mathcal{B}_{t+1}$. ∎

Proof of Corollary 2

The proof follows directly from our previous results. ∎

Proof of Corollary 3

Assume that product quality is Q, and, without loss of generality, that initial belief is 0, compared to an outside option $c \in \mathbb{R}$. We have

$$\mathbb{E}_0(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta > c) > Q.$$

But then,

$$\mathbb{E}_1(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta + \mathbb{E}_0(\mathcal{R}) > c).$$

Clearly, $\mathbb{E}_1(\mathcal{R}) \in (0, \mathbb{E}_0(\mathcal{R}))$. Thus, a similar argument to the one applied in the proof of Proposition 1 implies that reviews converge.

The convergence point is given by the fixed point of the above equation:

$$\mathbb{E}_\infty(\mathcal{R}) = Q + \mathbb{E}(\theta \mid \theta + \mathbb{E}_\infty(\mathcal{R}) > c).$$

Assume by contradiction that $\mathbb{E}_\infty(\mathcal{R}_H) \geq \mathbb{E}_\infty(\mathcal{R}_L)$. This implies

$$\mathbb{E}(\theta_H \mid \mathbb{E}_\infty(\mathcal{R}_H) + \theta_H > c) \geq \mathbb{E}(\theta_L \mid \mathbb{E}_\infty(\mathcal{R}_L) + \theta_L > c).$$

But then, *a fortiori*, we have

$$\mathbb{E}(\theta_H \mid \theta_H > c) \geq \mathbb{E}(\theta_L \mid \theta_L > c).$$

Therefore, to reach a contradiction we have to show that, when c is high enough,

$$\mathbb{E}(\theta_L \mid \theta_L > c) > \mathbb{E}(\theta_H \mid \theta_H > c)$$

We have that

$$\mathbb{E}(\theta_L \mid \theta_L > c) = \frac{\int_c^{\bar{\theta}} \left(1 - F_{s_L}(\theta)\right) d\theta}{1 - F_{s_L}(c)} \tag{44}$$

and similarly

$$\mathbb{E}(\theta_1 \mid \theta_1 > c) = \frac{\int_c^{\bar{\theta}} \left(1 - F_{s_H}(\theta)\right) d\theta}{1 - F_{s_H}(c)}. \tag{45}$$

Note that, by the definition of demand rotations, there exists a unique $\theta^\dagger$ such that $F_{s_L}(\theta^\dagger) = F_{s_H}(\theta^\dagger)$ and $F_{s_L}(\theta) < F_{s_H}(\theta)$ for every $\theta > \theta^\dagger$ (in particular, $\theta^\dagger = 0$ whenever s is symmetric).

As a result,

$$\int_{\theta^\dagger}^{\bar{\theta}} \left(1 - F_{s_H}(\theta)\right) d\theta < \int_{\theta^\dagger}^{\bar{\theta}} \left(1 - F_{s_L}(\theta)\right) d\theta$$

and thus, substituting $c = \theta^\dagger$ in Equations (44) and (45), we have

$$\mathbb{E}(\theta_L \mid \theta_L > \theta^\dagger) > \mathbb{E}(\theta_H \mid \theta_H > \theta^\dagger).$$

The result is true *a fortiori* for every $c > \theta^\dagger$. By continuity of both conditional expected values – which follows from the smoothness of both $F_L(\cdot)$ and $F_H(\cdot)$ – we have that there exists a $c^* \in [\underline{\theta}, \theta^\dagger)$ such that the result is true for every $c > c^*$. Thus, $\mathbb{E}_\infty(\mathcal{R}(s_L)) > \mathbb{E}_\infty(\mathcal{R}(s_H))$ as desired. This concludes the proof. ∎

Proof of Corollary 4

Assume product quality is Q. Initial belief 0 (wlog) compared to outside option $c \in \mathbb{R}$. We have

$$\mathbb{E}_0(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta > c) > Q.$$

But then,

$$\mathbb{E}_1(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta + \mathbb{E}_0(\mathcal{R}) > c).$$

Clearly, $\mathbb{E}_1(\mathcal{R}) \in (0, \mathbb{E}_1(\mathcal{R}))$. A similar logic implies $\mathbb{E}_2(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta + \mathbb{E}_1(\mathcal{R}) > c) \in (\mathbb{E}_1(\mathcal{R})), \mathbb{E}_0(\mathcal{R})))$. Thus, a similar argument to the one applied in the proof of Proposition 2 implies that reviews converge.

The convergence point is given by the fixed point of the above equation:

$$\mathbb{E}_\infty(\mathcal{R}) = Q + \mathbb{E}(\theta|\theta + \mathbb{E}_\infty(\mathcal{R}) > c). \tag{46}$$

We want to show that $\mathcal{B}_\infty = \mathbb{E}_\infty(\mathcal{R}) - Q$ is decreasing in Q.

To show this in a straightforward fashion, notice that as Q increases by ϵ, so does the RHS of Eq. (46).

Now suppose $\mathbb{E}_\infty(\mathcal{R})$ also increases by ϵ: then, first, the LHS goes up by ϵ; and second, the RHS decreases, since $\mathbb{E}(\theta \mid \theta + \mathbb{E}_\infty(\mathcal{R}) + \epsilon > c)$ is decreasing in ϵ. Thus, if Q and $\mathbb{E}_\infty(\mathcal{R})$ were to increase equally, the LHS would exceed the RHS.

At the same time, if Q increased by ϵ and $\mathbb{E}_\infty(\mathcal{R})$ were unchaged, the RHS would exceed the LHS.

Thus, we have

$$\frac{\partial \mathbb{E}_\infty(\mathcal{R})}{\partial Q} \in (0,1) \quad \Rightarrow \quad \frac{\partial \mathcal{B}_\infty(Q)}{\partial Q} = \frac{\partial \mathbb{E}_\infty(\mathcal{R})}{\partial Q} - 1 < 0.$$

The number of reviews is given by

$$\mathcal{N}_\infty(\mathcal{R})(Q) = Prob(\theta + \mathbb{E}_\infty(\mathcal{R}) > c),$$

which is increasing in Q because $\mathbb{E}_\infty(\mathcal{R})$ is. This concludes the proof. $\blacksquare$

Proof of Corollary 5

We show a direct example of welfare reducing social learning from reviews, and then extend it.

We have seen in Proposition 2, Claim 5 that whenever $s_1 = s_2$ learning from reviews is welfare enhancing. Therefore, assume now that the products differ in their designs $s_1 = L$, $s_2 = H$. Assume furthermore that $Q_1 = Q_2$. Then, we have that

$$\mathcal{J}_1^{NR} = \emptyset \tag{47}$$

whereas in the presence of biased reviews we have

$$\mathcal{J}_1^{BR} = \{j \mid \theta_{1j} < \theta_{2j} - \Delta_\infty(\mathcal{R}) \ \ \& \ \ \theta_{2j} < \theta_{1j}\}. \tag{48}$$

Clearly, $\mathcal{J}_1^{BR}$ is non-empty because we know from Proposition 1 that in this case, long-run reviews are biased in favor of the more polarizing product: $\mathcal{B}_\infty = \Delta_\infty(R) < 0$.

Welfare in this case is given by the probability of a consumer making the correct choice. This is given by $1 - Prob(j \in \mathcal{J}_1^{BR})$ in the case of biased reviews and 1 otherwise. Thus, in this case reviews reduce welfare.

Clearly, the above example is not particularly surprising: when quality differences are 0 to begin with, and reviews help consumers make inference about quality differences, no improvement over the prior $Q_1 = Q_2$ is possible.

Thus, we extend this result to the case of quality asymmetries: $\Delta(Q) = Q_1 - Q_2 > 0$. Now define by

$$\mathcal{J}_1^{NR}(\Delta(Q)) = \{j \mid \theta_{1j} > \theta_{2j} \ \ \& \ \ \theta_{1j} > \theta_{2j} - \Delta(Q)\},$$

and

$$\mathcal{J}_1^{BR}(\Delta(Q)) = \{j \mid \theta_{1j} > \theta_{2j} - \Delta_\infty(\mathcal{R})(\Delta(Q)) \ \& \ \theta_{2j} > \theta_{1j}\}.$$

We have $\mathcal{J}_1^{BR}(0) \subset \mathcal{J}_1^{NR}(0)$. By continuity, there exists a $\Delta^*(Q)$ such that the result holds for any $\Delta^*(Q) < \Delta(Q)$. (We have suppressed the dependence of $\Delta^*(Q)$ on $\mathcal{B}_\infty$.)

But then,

$$\mathcal{J}_{NR}(\Delta^*(Q)) \subset \mathcal{J}_{QR}(\Delta^*(Q)) \quad \forall \Delta^*(Q) < \Delta(Q). \tag{49}$$

Thus,

$$1 - Prob(\mathcal{J}_{BR}(\Delta^*(Q))) < 1 - Prob(\mathcal{J}_{NR}(\Delta^*(Q))) \quad \forall \Delta^*(Q) < \Delta(Q). \tag{50}$$

Thus, reviews are welfare reducing whenever differences in designs are large, and quality differences are small. ∎

B.2. Proofs for Appendix B Results

Proof of Proposition 5

Proof:

The proof of Proposition 5 essentially replicates the steps of that of Proposition 1.

- The fact that both reviews are higher than they would have been without an outside option follows from the fact that

$$\mathbb{E}(\theta_1|\theta_1 > (\max \theta_2 - \Delta_\infty(\mathcal{R}), c)) > \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_\infty(\mathcal{R})),$$

and similarly for product 2, as the presence of a non-trivial outside option c increases the lower bound of integration for at least some consumers – the more so the higher c.

- Following the reasoning in the proof of Proposition 1, this is equivalent to showing that

$$\mathbb{E}(\theta_2|\theta_2 > \max(\theta_1, c)) > \mathbb{E}(\theta_1|\theta_1 > \max(\theta_2, c)).$$

The LHS can be rewritten as

$$\mathbb{E}(\theta_2|\theta_2 > \max(\theta_1, c)) = P(\theta_1 > c)\mathbb{E}(\theta_2|\theta_2 > \theta_1) + P(\theta_1 < c)\mathbb{E}(\theta_2|\theta_2 > c).$$

We know from our Proofs of Propositions 1 and Corollary 3 that $\mathbb{E}(\theta_2|\theta_2 > \theta_1) > \mathbb{E}(\theta_1|\theta_1 > \theta_2)$ and $\mathbb{E}(\theta_2|\theta_2 > c) > \mathbb{E}(\theta_1|\theta_1 > c)$. The result follows.

- The proof for this result is straightforward, and can be found in the proof of Proposition 1.
- We have that

$$\mathbb{E}_0(\mathcal{R}_1) = Q_1 + \mathbb{E}(\theta_{1j}|\theta_{1j} > max(\theta_{2j}, 0)) \tag{51}$$

and similarly for product 2. Thus, $\mathcal{B}_0 = \mathbb{E}_1(\mathcal{R}_1) - \mathbb{E}_1(\mathcal{R}_2)$ implies

$$\mathcal{B}_0 = \mathbb{E}(\theta_{1j} \mid \theta_{1j} > max(\theta_{2j}, c))) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > max(\theta_{1j}, c))), \tag{52}$$

where we have simplified the RHS using the fact that $Q_1 = Q_2$ by assumption. To show that $\mathcal{B}_0 < \mathcal{B}_\infty$, assume by contradiction $\mathcal{B}_0 = \mathcal{B}_\infty$. But then, we obtain

$$\begin{aligned} \mathcal{B}_0 =& \mathbb{E}(\theta_{1j} \mid \theta_{1j} > max(\theta_{2j} - \mathcal{B}_0, c - \mathbb{E}_\infty(\mathcal{R}_1)) - \\ & \mathbb{E}(\theta_{2j} \mid \theta_{2j} > max(\theta_{1j} + \mathcal{B}_0, c - \mathbb{E}_\infty(\mathcal{R}_2)). \end{aligned} \tag{53}$$

Therefore,

$$\begin{aligned} \mathcal{B}_0 =& \mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j}) - \mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j}) \\ <& \mathbb{E}(\theta_{1j} \mid \theta_{1j} > max(\theta_{2j} - \mathcal{B}_0, c - \mathbb{E}_\infty(\mathcal{R}_1)) \\ &- \mathbb{E}(\theta_{2j} \mid \theta_{2j} > max(\theta_{1j} + \mathcal{B}_0, c - \mathbb{E}_\infty(\mathcal{R}_2)) \\ =& \mathcal{B}_\infty, \end{aligned}$$

where the inequality follows from the fact that $\mathcal{B}_\infty < 0$ (equivalently, $\mathbb{E}_\infty(\mathcal{R}_1) < \mathbb{E}_\infty(\mathcal{R}_2)$, as established in the Proof of **Claim 1**. The conclusions hold *a fortiori* if $\mathcal{B}_0 > \mathcal{B}_\infty$. This proves that $\mathcal{B}_0 < \mathcal{B}_\infty$. The fact that $\mathcal{N}_1(\mathcal{R}_2) > \mathcal{N}_\infty(\mathcal{R}_2)$ follows straightforwardly from $\mathcal{B}_0 < \mathcal{B}_\infty$ using the same argument as in the Proof of **Claim 2**. ■

■

Proof of Proposition 6

Proof:

The proof of Proposition 6 essentially replicates the steps of that of Proposition 2, with some minor modifications.

- The fact that both reviews are higher than they would have been without an outside option follows from the fact that

$$\mathbb{E}(\theta_1|\theta_1 > max(\theta_2 - \Delta_\infty(\mathcal{R}), c)) > \mathbb{E}(\theta_1|\theta_1 > \theta_2 - \Delta_\infty(\mathcal{R})),$$

as the presence of a non-trivial outside option c increases the lower bound of integration for at least some consumers – the more so the higher c.

10

- To see that product 1 has higher long-run reviews, assume by contradiction that $\mathbb{E}(\mathcal{R}_1) = \mathbb{E}(\mathcal{R}_2)$. Then, we have that

$$0 > -(Q_1 - Q_2) = \mathcal{B}_\infty = \mathbb{E}(\theta_1|\theta_1 > \max(\theta_2, c)) - \mathbb{E}(\theta_2|\theta_2 > \max(\theta_1, c)) = 0$$

  where the last equality comes from the fact that $s_1 = s_2$. We have reached a contradiction. Just like in the case of Proposition 2, one can show that if $\mathbb{E}(\mathcal{R}_1) < \mathbb{E}(\mathcal{R}_2)$, the LHS decreases, while the RHS increases.

- To show that $\mathcal{B}_\infty < 0$, assume by contradiction that $\mathcal{B}_\infty \geq 0$. This is equivalent to
$$\mathbb{E}(\theta_1|\theta_1 > \max(\theta_2 - \Delta_\infty(\mathcal{R}), c - \mathbb{E}_\infty(\mathcal{R}_1))) -$$
$$\mathbb{E}(\theta_2|\theta_2 > \max(\theta_1 + \Delta_\infty(\mathcal{R}), c - \mathbb{E}_\infty(\mathcal{R}_2))) > 0.$$

  But this is a contradiction, because $s_1 = s_2$, $\Delta_\infty(\mathcal{R}) > \Delta(Q) > 0$ and thus $c - \mathbb{E}_\infty(\mathcal{R}_1) < c - \mathbb{E}_\infty(\mathcal{R}_2)$, implying that the second conditional expected value exceeds the first.

- The proof for this result is straightforward, and can be found in the proof of Proposition 2.

  We have that $\mathcal{B}_0 = 0$ since $\mathbb{E}(\theta_1|\theta_1 > max(\theta_2, c)) = \mathbb{E}(\theta_2|\theta_2 > max(\theta_1, c))$. That is, $\mathbb{E}_0(\mathcal{R}_1) - \mathbb{E}_0(\mathcal{R}_2) = Q_1 - Q_2 > 0$. Therefore,

$$\mathbb{E}_1(\mathcal{R}_1) = Q_1 + \mathbb{E}(\theta_1 \mid \theta_1 + \mathbb{E}_0(\mathcal{R}_1) > max(\theta_2 + \mathbb{E}_0(\mathcal{R}_2), c)$$
$$= Q_1 + \mathbb{E}(\theta_1 \mid \theta_1 > max(\theta_2 - \Delta(Q), c))$$

  and similarly

$$\mathbb{E}_1(\mathcal{R}_2) = Q_1 + \mathbb{E}(\theta_2 \mid \theta_2 + \mathbb{E}_0(\mathcal{R}_2) > max(\theta_1 + \mathbb{E}_0(\mathcal{R}_1), c)$$
$$= Q_2 + \mathbb{E}(\theta_2 \mid \theta_2 > max(\theta_1 - \Delta(Q), c))$$

  Jointly, these two imply that

$$\mathcal{B}_1 = \mathbb{E}(\theta_1 \mid \theta_1 > max(\theta_2 - \Delta(Q), c)) - \mathbb{E}(\theta_2 \mid \theta_2 > max(\theta_1 - \Delta(Q), c)).$$
$$(54)$$

  But then, $\mathcal{B}_1 < \mathcal{B}_\infty$ if and only if $\Delta(Q) > \mathbb{E}_\infty(\mathcal{R}_1) - \mathbb{E}_\infty(\mathcal{R}_2)$. But this is equivalent to $\mathcal{B}_\infty < 0$, which we have shown to be true in the Proof of **Claim 1**.

  As in the previous cases, the fact that $\mathcal{N}_2(\mathcal{R}_1) < \mathcal{N}_\infty(\mathcal{R}_1)$ follows straightforwardly from $\mathcal{B}_1 < \mathcal{B}_\infty$.

- The proof for this claim is exactly the same as the one for the corresponding claim in Proposition 2.

■

Proof of Proposition 7

Proof:

We want to show that

$$Var\big(\theta_L \mid \theta_L > max(\theta_H + \Delta(Q), c)\big) < Var\big(\theta_H \mid \theta_H > max(\theta_H - \Delta(Q), c)\big), \quad \forall \Delta(Q) > \Delta^*(Q).$$

First notice that, when $\Delta(Q)$ approaches $\bar{\theta} - \underline{\theta}$, we have

$$Var\big(\theta_H \mid \theta_H > max(\theta_H - \Delta(Q), c)\big) \to Var(\theta_H | \theta_H > c).$$

On the other hand, the fact that $\theta_L > \theta_H + \Delta(Q)$ implies $\theta_L \in (\underline{\theta} + \Delta(Q), \bar{\theta})$. Therefore, Popoviciu's Inequality Popoviciu (1935) implies that

$$Var\big(\theta_L \mid \theta_L > \theta_H + \Delta(Q)\big) \leq \frac{1}{4}\big(\bar{\theta} - \underline{\theta} - \Delta(Q)\big)^2.$$

Notice that the right hand side gets arbitrarily small as $\Delta(Q) \to \bar{\theta} - \underline{\theta}$, implying the existence of a $\Delta^*(Q)$ such that $Var\big(\theta_L | \theta_L > \theta_H + \Delta(Q)\big) < Var(\theta_H)$ for every $\Delta(Q) > \Delta^*(Q)$. A fortiori, this is true for $Var\big(\theta_L \mid \theta_L > max(\theta_H + \Delta(Q), c)\big)$. ■

Proof of Proposition 8

Proof: Assume, for now, that $\Delta(Q) > 0$ and $s_1 = s_2$. We start by showing uniqueness. The proofs follows straightforwardly from the uniqueness of a solution for the equation

$$\Delta(\mathcal{R}) = \Delta(Q) + \mathbb{E}(\theta_1 | \theta_1 > \theta_2 - \Delta(\mathcal{R})) - \mathbb{E}(\theta_2 | \theta_2 > \theta_1 + \Delta(\mathcal{R})). \qquad (55)$$

As discussed in our Proofs of Propositions 2, 1 and 11, this uniqueness is guaranteed by the fact i) the LHS is increasing in $\Delta(\mathcal{R})$, while the RHS is decreasing, ii) the RHS exceeds the LHS when $\Delta(\mathcal{R}) = 0$ and iii) the LHS exceeds the RHS when $\Delta(\mathcal{R}) = \Delta(Q)$, and similarly if the products differ, instead, in their design.

Next, we want to show that $\Delta_t(\mathcal{R}^C)(\beta)$ converges whenever $\Delta_t(\mathcal{R})$ does. To see this, we start by showing that $\Delta_2(\mathcal{R}^C)(\beta) \in (\Delta_1(\mathcal{R}^C)(\beta), \Delta_0(\mathcal{R}^C)(\beta))$.

But this is immediate since (as shown in the proof of Proposition 2), when $\Delta(Q) > 0$ we have $\Delta_0(\mathcal{R}) = \Delta(Q)$, $0 < \Delta_1(\mathcal{R}) < \Delta(Q)$ and $\Delta_2(\mathcal{R}) \in (\Delta_1(\mathcal{R}), \Delta_0(\mathcal{R}))$. Thus, in the cumulative case, $\Delta_1(\mathcal{R}^C)(\beta) \in (0, \Delta_0(\mathcal{R}^C)(\beta))$, as $\Delta_1(\mathcal{R}^C)(\beta)$ is a weighted average of $\Delta_0(\mathcal{R})$ and $\Delta_1(\mathcal{R})$ for every β. This immediately implies that $\Delta_2(\mathcal{R}^C)(\beta) < \Delta_0(\mathcal{R}^C)(\beta) = \Delta(Q)$. Thus, a similar reasoning to the one employed in the proof of Proposition 2 implies that the sequence $\{\Delta_t(\mathcal{R}^C)(\beta)\}_{t=0}^{\infty}$ converges for every $\beta \in [0, 1]$.

The proof for the case for $s_1 \neq s_2$ follows very similar steps, and be easily derived by adapting the previous case and combining it with the proof of Proposition 1.

■

C. Proofs for Online Appendix A: Optimal Pricing

Proof of Proposition 9

First order conditions for firm 1 and 2 are given by, respectively,

$$1 - F(P_1^N - P_2^N) - P_1^N f(P_1^N - P_2^N) = 0 \tag{56}$$

and

$$F(P_1^N - P_2^N) - P_2^N f(P_1^N - P_2^N) = 0 \tag{57}$$

Jointly, these imply $P_1^N = \frac{1 - F(P_1^N - P_2^N)}{f(P_2^N - P_1^N)}$ and $P_2^N = \frac{F(P_1^N - P_2^N)}{f(P_1^N - P_2^N)}$. Now, define $\Delta(P) := P_1^N - P_2^N$. Subtracting the two expressions for P_1^N and P_2^N we have just found, we get

$$\Delta(P)^N f(\Delta(P)^N) = 1 - 2F(\Delta(P)^N). \tag{58}$$

Now, notice that the LHS is positive if and only if $\Delta(P)^N$ is. Conversely, the RHS is positive whenever $F(\Delta(P)^N) < 1/2$, and negative afterwards. Because $F(0) = 1/2$, this implies $\Delta(P)^N = 0$. Thus, the two only intersect at 0, which implies that $P_1^N = P_2^N$ in every equilibrium.

Now, plugging this back into Eq. (57), we get that in equilibrium

$$P_1^N = P_2^N = \frac{F(0)}{f(0)} = \frac{1}{2f(0)}, \tag{59}$$

when the equality comes from the symmetry of $F(\cdot)$. ∎

Proof of Proposition 10

The first order conditions are given by

$$1 - F(\Delta(P)^F - \Delta(Q)) - P_1 f(\Delta(P)^F - \Delta(Q)) = 0 \tag{60}$$

and

$$F(\Delta(P)^F - \Delta(Q)) - P_2 f(\Delta(P)^F - \Delta(Q)) = 0 \tag{61}$$

which can be combined into

$$\Delta(P)^F f(\Delta(P)^F - \Delta(Q)) = 1 - 2F(\Delta(P)^F - \Delta(Q)). \tag{62}$$

In equilibrium, we have $0 < \Delta(P)^F < \Delta(Q)$ whenever $\Delta(Q) > 0$. This is because the RHS is weakly negative whenever $\Delta(P)^F \geq \Delta(Q)$, while the LHS is always positive when $\Delta(P)^F > 0$. Thus, $\Delta(P)^F < \Delta(Q)$. A similar argument can rule out $\Delta(P)^F < 0$.

Moreover, explicitly solving for P_1 and P_2 we obtain

$$P_1^F = \frac{1 - F(\Delta(P)^F - \Delta(Q))}{f(\Delta(P)^F - \Delta(Q))} \tag{63}$$

and

$$P_2^F = \frac{F(\Delta(P)^F - \Delta(Q))}{f(\Delta(P)^F - \Delta(Q))}. \tag{64}$$

Comparing to the no-information case, we have:

$$P_2^F < P_2^N = P_1^N < P_1^F. \tag{65}$$

To see this, notice that the monotone hazard rate implies

$$P_2^F = \frac{F(\Delta(P)^F - \Delta(Q))}{f(\Delta(P)^F - \Delta(Q))} < \frac{F(0)}{f(0)} = \frac{1}{2f(0)} = P_2^N,$$

and similarly $P_1^F > P_1^N$. ∎

Proof of Lemma 4

First, to show that $\frac{\partial \Delta \mathcal{R}}{\partial \Delta(P)} > 0$, assume by contradiction that $\Delta(P^{\mathcal{R}})$ increases and $\Delta \mathcal{R}$ does not. Then, the LHS of Eq. 39 does not change, while the RHS increases, violating their equality.

To show that $\frac{\partial \Delta \mathcal{R}}{\partial \Delta(P^{\mathcal{R}})} < 1$, assume by contradiction it is not. Then, $\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R}$ goes down (or remains unchanged) following an increase in $\Delta(P^{\mathcal{R}})$. Thus, as a whole, the RHS of Eq. 39 decreases. On the contrary, the LHS increases, again reaching a contradiction. ∎

Proof of Lemma 5

From Eq. 39, we see that if $\Delta(P^{\mathcal{R}}) = \Delta(Q)$, then $\Delta \mathcal{R} = \Delta(Q)$. The result then follows immediately from Lemma 4. ∎

Proof of Lemma 6

Suppose (by contradiction) that $\Delta(P^{\mathcal{R}}) \leq 0$. Then $\Delta(P^{\mathcal{R}}) < \Delta(Q)$, and by Lemma 5, $\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R} < 0$. In this case, the LHS of Eq. (68) is positive but the RHS is weakly negative, a contradiction, so $\Delta(P^{\mathcal{R}}) > 0$. ∎

Proof of Proposition 11

The first order conditions are given by

$$\frac{\partial \pi_1}{\partial P_1^{\mathcal{R}}} = 1 - F(\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R}) - f(\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R}) \cdot \left(1 - \frac{\partial \Delta(R)}{\partial \Delta(P^{\mathcal{R}})}\right) \cdot P_1 = 0, \tag{66}$$

$$\frac{\partial \pi_2}{\partial P_2^{\mathcal{R}}} = F(\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R}) - f(\Delta(P^{\mathcal{R}}) - \Delta \mathcal{R}) \cdot \left(1 - \frac{\partial \Delta(R)}{\partial \Delta(P^{\mathcal{R}})}\right) \cdot P_2^{\mathcal{R}} = 0. \tag{67}$$

where we use the fact that $\frac{\partial \Delta \mathcal{R}}{\partial P_1^{\mathcal{R}}} = \frac{\partial \Delta \mathcal{R}}{\partial \Delta(P^{\mathcal{R}})} = -\frac{\partial \Delta \mathcal{R}}{\partial P_2^{\mathcal{R}}}$. Then, in equilibrium,

$$1 - 2F(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R}) = f(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})\left(1 - \frac{\partial\Delta\mathcal{R}}{\partial\Delta(P^{\mathcal{R}})}\right) \cdot \Delta(P^{\mathcal{R}}) \qquad (68)$$

Suppose (by contradiction) that $\Delta(P^{\mathcal{R}}) \geq \Delta(Q)$. By Lemma 5, $\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R} \geq 0$. In this case, the LHS of Eq. (68) is weakly negative but the RHS is positive, a contradiction, so $\Delta(P^{\mathcal{R}}) < Q$.

Therefore, we have

$$0 < \Delta(P^{\mathcal{R}}) < \Delta\mathcal{R} = \tilde{\Delta}Q < \Delta(Q). \qquad (69)$$

Solving the FOCs to derive prices, we get

$$P_1^{\mathcal{R}} = \frac{1 - F(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}{f(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})\left(1 - \frac{\partial\Delta\mathcal{R}}{\partial\Delta(P^{\mathcal{R}})}\right)} \qquad (70)$$

and

$$P_2^{\mathcal{R}} = \frac{F(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}{f(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})\left(1 - \frac{\partial\Delta\mathcal{R}}{\partial\Delta(P^{\mathcal{R}})}\right)}. \qquad (71)$$

∎

Proof of Proposition 12

Taking derivatives w.r.t. $\Delta(P)$ on both sides of Eq. 39 and rearranging, we have

$$1 - \frac{\partial\Delta\mathcal{R}}{\partial\Delta(P)} = \left(1 + \frac{d\mathbb{E}(\theta_1|\Delta\theta > \Delta(P) - \Delta\mathcal{R})}{dk} + \frac{d\mathbb{E}(\theta_2|\Delta\theta < \Delta(P) - \Delta\mathcal{R})}{dk}\right)^{-1} \qquad (72)$$

When $\Delta(Q)$ is close to 0, the previous analysis tells us that in equilibrium $\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R}$ and $\Delta(P)^F - \Delta(Q)$ are close to 0. In contrast, $1 - \frac{\partial\Delta\mathcal{R}}{\partial\Delta(P)} < (1 + \epsilon)^{-1}$. Therefore, we have

$$P_1^{\mathcal{R}} > \frac{1 - F(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}{f(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}(1 + \epsilon) \approx \frac{1 - F(\Delta(P)^F - \Delta(Q))}{f(\Delta(P)^F - \Delta(Q))}(1 + \epsilon) > P_1^F > P_1^N$$

and

$$P_2^{\mathcal{R}} > \frac{F(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}{f(\Delta(P^{\mathcal{R}}) - \Delta\mathcal{R})}(1 + \epsilon) \approx \frac{F(0)}{f(0)}(1 + \epsilon) > P_2^N > P_2^F.$$

∎