

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

Tommaso Bondi

Cornell University

November 2, 2023

Abstract. We develop a dynamic model of naïve social learning from consumer reviews. In our model, consumers decide which of two products to buy based on both expected quality and idiosyncratic taste. 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, 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 of reviews; pricing in presence of reviews; the design of crowd-sourced exploration; 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. Throughout most of the paper, we assume that consumers learn from reviews naïvely, that is, by equating differences in mean reviews with differences in quality.

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. 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 is relatively lower (*“the curse of the best-seller”*). In other words, the review system is

biased against products with high quality: by attracting many consumers (not all of whom have a strong taste for the product), the product’s success is also its curse.

Many industry players (and scholars) have recognized that firms might face a quality-quantity trade-off in reviews, reaching contradicting conclusions on how to optimally solve it.¹ This paper is, to the best of our knowledge, the first to carefully formalize and quantify such trade-off in a fully dynamic setting.

The review dynamics observed in a variety of highly differentiated product markets display patterns that are strikingly 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 of reviews and their variance go up. By studying individual reviewer 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.

Conversely, we show that consumer reviews favor “polarizing” products, that is, products whose fit component has very high variance (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, we apply our model to characterize properties of the variance of reviews. We show that, counterintuitively, 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 doing so, polarizing products obtain both higher mean reviews and a lower variance of reviews: more homogeneous reviews means uniformly higher average product-consumer fit, and thus *better* reviews.

In other words: polarizing products usually split consumers into two very distinct

1. See for instance <https://qigo.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).

camps, one who loves them, 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.

The (admittedly anecdotal) patterns we observe on *Goodreads*, a popular consumer book reviews platform, provide strong support for our theory. For instance, bestsellers (books with more than 200 editions) average a variance of 1.03, compared to 0.89 overall; the results get *stronger* as we raise the threshold to 500, 750 and 1000 editions. The last, extremely selective, group includes all-time classics such as “*The Jungle Book*” and “*Alice in Wonderland*”, which have a variance of 0.95 and 1.05 respectively.

Meanwhile, “*Bullies: How the Left’s Culture of Fear and Intimidation Silences Americans*” by Ben Shapiro – which is likely extremely polarizing *ex-ante* – has a slightly below-average variance of 0.79, reflecting the arguably very homogeneous set of its buyers. It is highly implausible that these patterns are in line with *ex-ante* polarization², while they are well explained by the self-selection mechanism we propose.

Third, we discuss the robustness of our findings by looking at the full dynamics of ratings. Because consumers trade-off quality and fit, overrated products attract many poorly matched consumers. Therefore, shouldn’t the biases highlighted above be short-lived, and naturally self-correcting?

We show, surprisingly, that the aforementioned self-correction, while occurring, is only partial: long-run reviews display the same qualitative biases as short-run ones. In particular, 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 rating, 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.

Thus, the opposite nature of our results and those in Park et al. (2021) illustrates the opposite issues plaguing reviews-based learning in different type 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. Conversely, in markets without taste-based differentiation, the biases we highlight in this paper are not a concern, but robustness is lost, as initial conditions are disproportionately important, giving rise to “winner-take-all” dynamics similar to the ones in

2. That is, it is very difficult to imagine that if all US readers were to read, and rate, these three books, Ben Shapiro’s would have the lowest variance in reviews of the three, since all left-leaning readers would likely pan the book.

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. It is therefore natural to ask: would these biases disappear if consumers were to learn from reviews in a Bayesian fashion, that is, if they correctly internalized the informational content of reviews?

We show that, perversely, the presence of Bayesian consumers makes all the long-run biases we have highlighted *larger*. The reason for this is that, because Bayesians are able to identify, and choose, their subjectively preferred products, they leave higher reviews, on average, than their naïve 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.

These findings have important strategic implications for consumers, sellers and platforms alike. For consumers, the key takeaway is to be wary of simple aggregate statistics like 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 help.

Sellers should consider the non-trivial trade-offs of different designs, pricing and advertising strategies in contexts in which online word of mouth is key to long-term success, often opting for counterintuitive strategies.

When thinking about product design, a naïve marketer might assume that a mainstream design maximizes the chances to satisfy buyers, thus minimizing 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 This concern might be less relevant in markets – such as those for books or movies – in which prices are not as salient, or even fixed. Furthermore, in our model, high prices effectively function as a matching device: only consumers with a strong taste for the product will buy it. The opposite is true for low prices, which attract some consumers who do not have a strong taste for the product.

This prediction is in line with empirical evidence on the perverse reputational effects of deep discounts, see for instance Byers et al. (2012), Liu et al. (2019).⁵

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

4. E.g., see Luca and Reshef (2021) in the context of *Yelp* restaurant reviews.

5. For a much more detailed analysis of pricing incentives in the presence of consumer reviews, we refer to the recent work of Carnehl et al. (2021).

Similarly, when it comes to advertising, our findings can be seen as an additional justification for “demarketing” (Miklós-Thal and Zhang, 2013).

On the platform side, we 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” these products face.

The fact that reviews, while biased, are naturally self-correcting over time has a crucial implication for platform and sellers alike: long-run reviews do not depend on initial ones, or, in a similar vein, on time t reviews, for any $t \geq 1$. This suggests that the effects of reviews manipulation on the part of sellers might be short-lived: by inducing excessively high expectations about product quality, the product will sell more in the next period, but in doing so will attract a set of buyers for whom it will be, on average, a worse match. Thus, reviews will naturally go down 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 mechanisms and results. Section 4 introduces the model, and Section 5 presents the main results. Section 6 introduces new results and 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 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), fake reviews (Luca and Zervas (2016), He et al. (2022)), or – most relevant for 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. (2022)).

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, theoretically, 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, follow a decreasing pattern.

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. Using a similar sequential model, Besbes and Scarsini (2018) show that accurate learning can be achieved under weak assumptions on 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 study a competitive setting. This is key, because reviews are usually employed to alleviate choice overload problems, that is, to aid consumers deciding *what* – not *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 this paper we study biases in relative, not absolute, reviews. These are much harder for consumers to correct for than systematic platform-wide biases: for instance, it is well known that rating inflation (Filippas et al., 2022) is pervasive on many platforms, and that, for instance, an *Uber* driver with an average rating of 4.5 out of 5 is well below average.

Another important 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), Lee et al. (2022)).⁶ We show that studying the full dynamic evolution of reviews, while complex, is important, since it allows us to elucidate many phenomena that the study of two-period models does not. In doing so, our paper helps bridging the gap between the theoretical and empirical literature; the latter has investigated the dynamic evolution of reviews for over a decade (*e.g.*, Li and Hitt (2008), 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 rating*

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

*systems robustly aggregate information when agents use simple learning rules.”*⁷

3. Example

Before presenting a 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: $Q_1 = Q_2$.⁸ Without loss of 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”.⁹ Without loss of generality (given the symmetric role of prices and qualities – see Section 6 for a more thorough discussion of this point), we assume $P_1 = P_2 = 0$. Each consumer knows their fit with each product, but do not know products’ qualities, and employs reviews to learn them.

After consumption, 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).¹⁰ 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 the buyers of product i , by $\mathcal{R}_i = \{\mathcal{R}_{ij}\}_{j \in \mathcal{J}_i}$ the entire set of reviews for product i , and by $E(\mathcal{R}_{ij})$ and $Var(\mathcal{R}_{ij})$ the expected value and variance of product i ’s reviews respectively.

Because consumers use reviews to decide *what*, not *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 $\theta_{1j} \sim N(0, 1)$, $\theta_{2j} \sim N(0, 4)$).

At each $t = 0, 1, \dots$, 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. 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. Moreover, denote by $\tilde{\Delta}_t(Q) := E_t(\mathcal{R}_1) - E_t(\mathcal{R}_2)$,

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

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

9. Notice that mainstream and niche do not necessarily have a market share interpretation, but rather simply refer to the spread in consumers’ taste for each product.

10. 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.

that is, the difference in average reviews at time t .

Now assume $\tilde{\Delta}_t(Q) = 0$. Then, in the following period, we have

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

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

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

$$E_{t+1}(\mathcal{R}_1) - E_{t+1}(\mathcal{R}_2) \approx -1.42 + 0.35 = -1.07.$$

Therefore, naïve period $t + 2$ consumers will believe $\tilde{\Delta}_{t+1}(Q) = E_{t+1}(\mathcal{R}_1) - E_{t+2}(\mathcal{R}_2) = -1.07$. 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.¹¹

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 products look less polarizing than they are – 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 polarizing product 2:

$$\begin{aligned} \frac{Var(\theta_{2j} | \theta_{2j} > \theta_{1j}) - Var(\theta_{2j})}{Var(\theta_{2j})} &= \frac{4 - 1.96}{4} = 0.51 \\ &> 0.13 = \frac{1 - 0.87}{1} = \frac{Var(\theta_{1j} | \theta_{1j} > \theta_{2j}) - Var(\theta_{1j})}{Var(\theta_{1j})}. \end{aligned}$$

11. 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) $E(\theta_1 | \theta_1 > \theta_2) < E(\theta_2 | \theta_2 > \theta_1)$, *ii*) $P(\theta_1 > E(\theta_1 | \theta_1 > \theta_2)) > P(\theta_2 > E(\theta_2 | \theta_2 > \theta_1))$, *iii*)

$\frac{\partial E(\theta_i | \theta_i > \theta_{-i})}{\partial \sigma_{-i}} < 0$, $\frac{\partial E(\theta_i | \theta_i > \theta_{-i})}{\partial \sigma_i} > 0$ for $i = 1, 2$ and *iv*) $\frac{\partial P(\theta_i > E(\theta_i | \theta_i > \theta_{-i}))}{\partial \sigma_{-i}} > 0$,

$\frac{\partial P(\theta_i > E(\theta_i | \theta_i > \theta_{-i}))}{\partial \sigma_i} < 0$ for $i = 1, 2$.

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 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 mislearning from 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

$$\begin{aligned} E_{t+2}(\mathcal{R}_1) &= \underbrace{E(E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \tilde{\Delta}_{t+1}(Q)))}_{\text{Average conditional taste for product 1, given bias } \tilde{\Delta}_{t+1}(Q) \text{ in reviews}} \approx 0.49 \\ E_{t+2}(\mathcal{R}_2) &= \underbrace{E(E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \tilde{\Delta}_{t+1}(Q)))}_{\text{Average conditional taste for product 2, given bias } \tilde{\Delta}_{t+1}(Q) \text{ in reviews}} \approx 0.95 \end{aligned}$$

While the reviews of product 2 remain higher, the difference in rating 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 $E_{t+2}(\mathcal{R}_1) - E_{t+2}(\mathcal{R}_2)$ 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

$$E(E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \tilde{\Delta}_\infty(Q))) - E(E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \tilde{\Delta}_\infty(Q))) = -\tilde{\Delta}_\infty(Q)$$

In other words, choices and reviews are self-confirming: 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, $\tilde{\Delta}_\infty(Q) = -0.66$. 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 ($E_\infty(\mathcal{R}_2) = 1.1 > 0.44 = E_\infty(\mathcal{R}_1)$). 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.




Figure 2

The bias in reviews starts at $B_0 \approx -1.07$ and then converges to its long-run average, $B_\infty \approx -0.66$.




Figure 3

The market share of product 2 starts at $M_0 = 0.50$ and then converges to its steady state, $M_\infty \approx 0.62$.

Assume now that the two products do not just differ in how polarizing they are, but also in their qualities. Following Johnson and Myatt (2006), assume that the

higher quality product is the less polarizing one: $Q := Q_1 - Q_2 > 0$.¹² Now, our fixed point equation becomes

$$E(E(\theta_{2j} | \theta_{2j} > \theta_{1j} + -\tilde{\Delta}_\infty^{1/2}(Q))) - \alpha \cdot E(E(\theta_{1j} | \theta_{1j} > \theta_{2j} + \tilde{\Delta}_\infty^{1/2}(Q))) = \tilde{\Delta}_\infty^{1/2}(Q) - Q$$

Let $Q = 1$. Numerically, we obtain $\tilde{\Delta}_\infty(Q) = 0.16$. That is, the bias in ratings, *relative to qualities*, got worse:

$$\mathcal{B}_\infty := \tilde{\Delta}_\infty(Q) - Q = 0.16 - 1 = -0.84 < -0.66.$$

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), thus decreasing its market share and increasing its relative reviews.




Figure 4

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.

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




Figure 5

The market share of product 2 starts at $M_0 = 0.50$ (period 0 consumers are unaware that $Q > 0$). Then, as period $t = 1, 2, \dots$ consumers become aware of quality differences, it goes down to its steady state, $M_\infty \approx 0.30$. This is still higher than if reviews were unbiased: $P(\theta_{2j} > \theta_{1j} + 1) \approx 0.18$.

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 differentiation in terms of product design, s . Design measures how polarizing (or, following Johnson and Myatt (2006), “niche”¹³) 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

$$U_{ij} = Q_i + \theta_{ij} - P_i.$$

13. Note that “niche” here should not be interpreted as having small market shares: in fact, in our main example (Section 3) the niche product obtains higher market shares in equilibrium, in light of its inflated reviews.

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.¹⁴ 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 $\theta_{s'_i}^\dagger$ such that*

$$F_{s_i}(\theta) < (>) F_{s'_i}(\theta) \iff \theta < (>) \theta_{s'_i}^\dagger.$$

Intuitively, $F_{s'_i}(\cdot)$ concentrates more mass around $\theta_{s'_i}^\dagger$ than $F_{s_i}(\cdot)$ does. In economic terms, this measures the difference between more mainstream designs, which are 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 fit are always observable to consumers before purchasing¹⁵, 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.¹⁶ Emphasizing the self-selection in choices and reviews, we have:

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

14. Given the two distributions both have 0 mean, this notion simplifies to that of a mean preserving spread.

15. 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 fit 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.

16. 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.

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,

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

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 , $i = 1, 2$:

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

We start with a simple Lemma, which guarantees that studying the properties of reviews is equivalent to studying patterns of taste-based self-selection.

Lemma 1. *Denoting by $G_{\mathcal{R}_i}(\cdot)$ the CDF of product i reviews, we have that $G_{\mathcal{R}_i}(\cdot)$ satisfies, for every $k \geq 0$,*

$$G_{\mathcal{R}_i}(k) = F_{s_i}^{\mathcal{J}_i}(k - Q_i + P_i).$$

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

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

$$\mathbb{E}(\mathcal{R}_i) = \int \mathcal{R}_{ij} d\mathcal{J}_i = Q_i + \mathbb{E}_{F_{s_i}^{\mathcal{J}_i}}(\theta_i) - P_i, \quad \text{Var}(\mathcal{R}_i) = \text{Var}_{F_{s_i}^{\mathcal{J}_i}}(\theta_i).$$

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 $\text{Var}_{F_{s_i}^{\mathcal{J}_i}}(\theta_i)$. 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 $P_1 = P_2 = 0$), with the understanding that quality can be thought of as net of price, and that each result speaking to the impact of quality differences can be immediately rephrased into (an opposite) one on the impact of price differences. We discuss the role of pricing in greater detail in Section 6.

Third, due to self-selection, consumers usually buy products for which they have positive taste:

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

Thus, we generally have $E(\mathcal{R}_i) > Q$ 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* the two products¹⁷ – that is, choices of *what*, not *if* to buy. We start with the following:

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

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

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: $E(\mathcal{R}_1) - Q_1 > E(\mathcal{R}_2) - Q_2$.

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.¹⁸

17. 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 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, $\Delta(Q)$.

18. 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.

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 $t = 0, 1, 2, \dots$, 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 consumers¹⁹, and then discuss the consequence of relaxing this assumption in Section 5.4. Naïve consumers simply take reviews at face value:

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

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

Before moving on to 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 2. *Assume consumers are uninformed about their fit with each product: $E(\theta_{1j}) = E(\theta_2) = 0 \quad \forall j \in \mathcal{J}$. Then, $E(\mathcal{R}_i) = Q_i$ and $Var(\mathcal{R}_i) = Var_{F_{s_i}}(\theta_i)$, for all Q_i and s_i , $i = 1, 2$.*

19. 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.

Table 1Main 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 rating for product i |
| \mathcal{R}_i | set of all reviews $\{\mathcal{R}_{ij}\}_{j \in \mathcal{J}_i}$ for product i |
| $E(\mathcal{R}_i)$ | mean of reviews for product i |
| $var(\mathcal{R}_i)$ | variance of reviews for product i |
| M_i | market share of product i |
| \mathcal{B} | relative bias in reviews in favor of product 1 |

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. 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.²⁰

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: $E(\theta_{ij} + \epsilon_{ij} \mid \text{Choice}) = E(\theta_{ij} \mid \text{Choice})$. 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.²¹
- **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 of generation $t - 2, t - 3, \dots, 1$. We chose this approach for analytical tractability and ease of exposition. One can think of our model as one in which the platform quickly discount past opinions as new ones arrive, or one in which consumers only read the most recent reviews. While the exposition becomes considerably more cumbersome, none of our major results would be affected by assuming, instead, that more (all) reviews are accessible to consumers.

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 $E_0(\mathcal{R}_i) = Q_i + E(E(\theta_{ij} \mid \theta_{ij} \geq \theta_{-ij}))$, $i = 1, 2$. More generally, given period t reviews $E_t(\mathcal{R}_1)$ and $E_t(\mathcal{R}_2)$, we have that a naïve consumer in period $t + 1$ chooses

20. See Lee et al. (2022) for a decomposition of the variance of reviews into a taste and a quality-variability component, and its implications for consumer learning.

21. 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.

product 1 if and only if $E_t(\mathcal{R}_1) + \theta_{1j} > E_t(\mathcal{R}_2) + \theta_{2j}$. Thus, product 1's period $t + 1$ reviews from naïve consumers are given by

$$E_{t+1}(\mathcal{R}_1) = Q_1 + E(E(\theta_{1j} \mid E_t(\mathcal{R}_1) + \theta_{1j} \geq E_t(\mathcal{R}_2) + \theta_{2j})),$$

and similarly for product 2.

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

$$\begin{cases} E_\infty(\mathcal{R}_1) = Q_1 + E(E(\theta_{1j} \mid E_\infty(\mathcal{R}_1) + \theta_{1j} \geq E_\infty(\mathcal{R}_2) + \theta_{2j})) \\ E_\infty(\mathcal{R}_2) = Q_2 + E(E(\theta_{2j} \mid E_\infty(\mathcal{R}_2) + \theta_{2j} \geq E_\infty(\mathcal{R}_1) + \theta_{1j})) \end{cases} \quad (1)$$

Clearly, this is independent on initial conditions. Furthermore, notice how that the left hand side of both equations is increasing in $E(\mathcal{R}_i)$, while the right hand side is decreasing in $E(\mathcal{R}_i)$ and increasing in $E(\mathcal{R}_{-i})$. This follows straightforwardly from properties of conditional expectations. Two important facts follow: first, for a fixed $E(\mathcal{R}_{-i})$, there is at most one $E(\mathcal{R}_i)$ 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}_∞ 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 averages²², 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: $Q_1 = Q_2$, $s_1 = H$, $s_2 = L$. Furthermore, assume that $M_H \geq M_L$ and that the function $\theta F_{s_L}(\theta)$ is weakly convex.²³ Then, in equilibrium ($t = \infty$):*

22. See e.g., https://www.researchgate.net/figure/Histogram-of-venues-rating-for-restaurants-fast-foods-and-bars_fig4_314160022

23. This assumption is satisfied by most well known distributions, including the uniform and the normal. Moreover, it represents a sufficient, but not necessary, condition.

- The more polarizing product 2 is relatively overrated: $\mathcal{B}_\infty < 0$,
- and thus captures a higher market share than it otherwise would: $M_{2,\infty} > M_L$.
- Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_0 < \mathcal{B}_\infty < 0$, $M_{2,1} > M_{2,\infty} > M_L$.

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 1's reviews will not be particularly upward-biased.

On the flip side, when observing a consumer choosing product 2, it will be relatively more likely this is due to a strong taste for it than to a (again, statistically rare) strong distaste for the mainstream alternative.

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 stronger product-consumer fit – and thus relatively more upward-biased reviews – for niche products.

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, in equilibrium, product 2 often captures a *larger* market share ($M_{2,\infty} > 1/2$) *and* obtain higher reviews, even though $M_L \leq 1/2$. We have seen a concrete instance of this in the Example of Section 3: there, $M_L = 0.50$, $M_{2,1} \approx 0.67$ and $M_{2,\infty} \approx 0.62$.

If anything, the fact that in equilibrium the more polarizing product obtains higher reviews *while* capturing a higher market share is surprising, especially in light of Proposition 2 and Corollary 2 in the next Section. Again, market share and 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.

Before moving on to our second main result, we emphasize that, while our main focus is on the impact of reviews on competitive markets, the result does not strictly depend on competition, so that it is applicable to the much more studied setting of online reviews in monopoly.

Corollary 1. *Let each consumer choose between a product of quality Q and design s and an outside option, c . Then, for c sufficiently high,*

- *The presence of reviews lead consumers to excessive consumption: compared to the outside option, the product is always overrated.*
- *More polarizing products are more overrated, and thus obtain higher equilibrium market shares:*

$$E_\infty(\mathcal{R}(s_L)) > E_\infty(\mathcal{R}(s_H)), \quad M_\infty(s_L) > M_\infty(s_H).$$

Corollary 1 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. The requirement that c is high enough is a sufficient condition that can easily be removed if we assume some more regularity properties of $F_s(\theta)$.²⁴

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:

24. Proof available upon request. The example in Section 3, for instance, shows that the condition is not necessary: in fact, for $X \sim N(0, \sigma^2)$, we have that $E(X | X > c) = \sigma^2 \cdot \frac{\phi(c)}{1 - \Phi(c)}$, which is increasing in σ independently of $c \in \mathbb{R}$.

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 equilibrium ($t = \infty$):*

- *The higher quality product 1 has higher reviews: $E_\infty(\mathcal{R}_1) > E_\infty(\mathcal{R}_2)$,*
- *Despite being relatively underrated: $\mathcal{B}_\infty < 0$.*
- *It thus captures a higher market share: $M_{1,\infty} > 1/2$, 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: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $M_{1,\infty} > M_{1,2}$.*
- *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:

$$E(E(\theta_1 \mid Q_1 + \theta_1 > Q_2 + \theta_2)) < E(E(\theta_2 \mid Q_2 + \theta_2 > Q_1 + \theta_1)),$$

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 $\mathcal{J}_1 \subset \mathcal{J}_1$ such that for every $j \in \mathcal{J}_1$ 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}_∞ , here we use \mathcal{B}_1 and \mathcal{B}_∞ instead.²⁵ The reason is that, since qualities are initially unknown, and the bias in Proposition 2 comes from the trade-off between qualities and taste, here we have $\mathcal{B}_0 = 0$. The bias begins appearing 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 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 rating, 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

25. The same is true for market shares, which always lag one period compared to \mathcal{B} .

we see that the average review of a product with fewer than 100 reviews is almost as that of the average product with over 10000 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) books 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.²⁶

Proposition 2 suggests that – for a fixed mean of reviews – products with a high number of reviews should be of higher quality. Thus, rational consumers should have a preference for them. This offers a 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 2. 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.

26. For a “dual” interpretation of Proposition 2, one that centers around prices as opposed to qualities, see Section 6.4.

The following corollary follows straightforwardly from Proposition 2.

Corollary 2 (High reviews Are Self-Defeating.) $E_{t+1}(\mathcal{R}_i)$ is decreasing in $E_t(\mathcal{R}_i)$ and increasing in $E_t(\mathcal{R}_{-i})$. Jointly, these imply that the same is true, a fortiori, for relative reviews: \mathcal{B}_{t+1} is decreasing in \mathcal{B}_t .

Corollary 2 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 market shares, and thus lower average matches, and worse 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 3 (Fake Reviews Backfire) Let product 1’s relative reviews be artificially inflated due to seller 1’s manipulation in period t : $E_t(\mathcal{R}^F) > E_t(\mathcal{R})$, where \mathcal{R}^F indicates the presence of fake reviews. Then:

- Fake reviews are effective in the short-run: product 1’s market shares will be higher in period $t + 1$: $M_{1,t+1}^F > M_{1,t+1}$.
- But absent additional manipulation, they will backfire in period: product 1’s reviews in $t + 1$, and thus market shares in $t + 2$, will be lower: $E_{t+1}(\mathcal{R}^F) < E_{t+1}(\mathcal{R})$, $M_{1,t+2}^F < M_{1,t+2}$.
- Long-run ($t = \infty$) reviews and market shares will be unaffected, no matter how large the manipulation: $E_\infty(\mathcal{R}^F) = E_\infty(\mathcal{R})$, $M_{2,\infty}^F = M_{2,\infty}$.

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 $E_t(\mathcal{R}_i)$ is conceptually equivalent to a decrease in $E_t(\mathcal{R}_{-i})$, as formalized in Corollary 2: future relative reviews are decreasing in today’s relative reviews.

Second, and importantly, this result does not just imply that reviews will 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 3 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 $\mathcal{B}_\infty = 0$.

Before moving on to the next Section, we state an analogous to Corollary 1 for the case of quality asymmetries, showing that like Proposition 1, Proposition 2 also has a natural monopolistic interpretation.

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 higher equilibrium market shares:*

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

First of all, notice that in a monopolistic setting $\mathcal{B}_\infty = E_\infty(\mathcal{R}) - Q$. 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 high market shares but a weak 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 low market shares and very upward-biased reviews.

5.3. 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?

Johnson and Myatt (2006), Bar-Isaac et al. (2012), Sun (2012) and more recently Menzio (2023) 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:

Proposition 3. *(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.*

5.4. The Impact of 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

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

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

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

Notice, however, the inherent complexity of such an inversion: besides observing current reviews $E_{t+1}(\mathcal{R}_i)$ 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 positive measure of 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 Equation (2) and back up quality differences.

This result has an interesting implication for reviews design and platform strategy, in that it suggests that the platform should not incentivize reviews (early on in the product lifecycle, or otherwise) from Bayesian consumers in an attempt to “seed” the correct information.

5.5. 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 5 (The Variance of Reviews Needs Not Proxy Polarization) *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 5 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.²⁷ 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 and Additional Results

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 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 $M_1^2 + M_2^2$ and thus maximized when either $M_1 = 1$ or $M_2 = 1$), 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

27. 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: https://www.goodreads.com/list/show/6199.The_Most_Polarizing_Books_Of_All_Time.

the model with Bayesian consumers only make matters worse, exacerbating these effects.

6.2. Exploration vs Exploitation (and Barriers to Entry)

An influential recent stream of research (*e.g.*, Kremer et al. (2014), Papanastasiou and Savva (2016), Che and Hörner (2017), Vellodi (2022)) analyzes the role of platform’s reviews design in incentivizing experimentation by consumers, as opposed to simply rewarding myopic exploitation, as well as participation by firms.

In particular, Vellodi (2022) shows that, even abstracting from reviews biases, 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. Similarly, 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.

There are several important differences between our models 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 the platform’s explicit intervention.

6.3. A “Love for Large Numbers” is Rational

Powell et al. (2017) show experimentally that consumers prefer products with many reviews. This is evidence, they argue, for poor statistical reasoning: fewer reviews means higher variance, thus greater upside. That is, a product with an average rating of, say, 4.1 out of 5 and 20 reviews in total might be much better than the 4.1 indicates due to the intrinsic variability in reviews, but the same can not be said for a product with the same average over 2000 reviews.

By endogenizing the interplay between the average and the number of reviews, our findings challenge this view: more reviews effectively correspond to a higher burden of proof (Proposition 2), and thus to lower reviews for given quality. In fact, a heuristic that rewards products for both the average and the number of reviews can outperform one solely based on the average (and the variance) of reviews.

In fact, Equation (2), which describes how Bayesian consumers learn in this context, makes the optimality of a “love for large numbers” explicit. (The Bayesian estimate of) Q_i in Equation (2) is increasing in market shares, for a given $E_{t+1}(\mathcal{R}_i)$: an increase in $E_t(\mathcal{R}_i) - E_t(\mathcal{R}_{-i})$ increases market shares in t , and thus the right hand side in Equation (2). See Appendix C for a more in-depth analysis of this point.

This suggests an important complementarity between observational learning and learning from reviews: a hybrid form of learning, one that rewards products that are simultaneously popular and well reviewed, performs best. Interestingly, the fact that

both positive and negative opinions are present prevents cascades from happening, while the observational component of learning mitigates the biases we highlight in this paper.

6.4. Price as a Matching Device

It is easy to see that an increase in Q_i in our theoretical framework maps directly into a decrease in P_i . By virtue of this symmetry, Proposition 2 can be straightforwardly restated in terms of prices.²⁸

Proposition 6 (Cheaper Products Are Relatively Underrated) *Let the two products differ only in their prices: $P_1 < P_2$, $Q_1 = Q_2$, $s_1 = s_2$. Then, in equilibrium ($t = \infty$):*

- *The cheaper product 1 has higher reviews: $E_\infty(\mathcal{R}_1) > E_\infty(\mathcal{R}_2)$,*²⁹
- *Despite being relatively underrated: $\mathcal{B}_\infty < 0$.*
- *It thus captures a higher market share: $M_{1,\infty} > 1/2 > M_{2,\infty}$, but less than it would if consumers internalized the role of prices on self-selection.*
- *Nevertheless, some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $M_{1,\infty} > M_{1,1} > 1/2$.*
- *Reviews unambiguously increase consumer welfare.*

Much like higher quality, lower prices persuade consumers to buy the product even when it is not a great fit for them. Byers et al. (2012) document exactly this effect in the context of restaurant’s Groupon promotions, showing that the average Groupon consumer is more likely to be experimenting with less liked cuisines, and thus ends up reviewing less favorably, in line with our model. Similarly, in the context of mobile apps Liu et al. (2019) demonstrate the perverse reputational effects of a particularly salient low price: 0. Higher prices, on the contrary, will deter everyone but die-hard fans, thus leading to a more favorable consumer self-selection.

To get some additional intuition, define by $P_i^{*,t}$ the optimal price for firm i in period t absent reviews, and note that in the presence of reviews (and a positive measure of naïve consumers) firm i ’s period $t + 1$ market shares contain an additional term proportional to

$$\frac{\partial E_{t+1}(\mathcal{R}_i)}{\partial P_i^t} - \frac{\partial E_{t+1}(\mathcal{R}_{-i})}{\partial P_i^t},$$

This term is positive for every P_i^t : the first term indicates the increasing average match between product i and its period t consumers, while the second indicates

28. Once more, we emphasize that the lack of prices in our main model does not imply that reviews represent quality instead of “quality per dollar”: rather, it is simply a normalization. Our conclusions hold independent of this feature of the model.

29. The fact that $E_\infty(\mathcal{R}_1) > E_\infty(\mathcal{R}_2)$ need not hold if reviews do not internalize prices, namely if $\mathcal{R}_{ij} = Q_i + \theta_{ij}$ (even when $P_1 \neq P_2$). Nevertheless, the more general point of cheaper products being underrated is unaffected by this modelling choice.

the decreasing average match between product $-i$ and its period t consumers. The opposite is true for the reviews of product $-i$, doubling the benefits of high prices in a competitive setting compared to monopolistic settings. Thus, the optimal price in the presence of social learning from reviews will be higher than it would be otherwise. For a much more in depth analysis of reviews-induced pricing incentives in a similar framework, see Carnehl et al. (2021).³⁰

6.5. Tendency towards Personalization

All else equal, Proposition 1 shows that mainstream products have a hard(er) time creating uniformly positive word of mouth. Therefore, for any quality level, sellers are incentivized to move towards more niche designs. Furthermore, notice that personalization is associated with an increase in local monopoly power, providing an additional channel for price increases.

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

By attracting precisely the type of consumers the firm believes to have a higher taste for its products, 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 adds to theirs by showing that, due to the self-correcting nature of reviews, these effects might be short-lived.

7. Conclusion

We theoretically study social learning from consumer reviews in a horizontally differentiated market. In particular, we model the dynamic feedback loop between reviews, beliefs and choices: reviews today influence consumer beliefs and choices – and thus product market shares – tomorrow, but these in turn influence tomorrow’s reviews.

We first build a parsimonious duopoly 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 prod-

30. Also see Sayedi (2018) for a model of duopoly pricing with observational learning.

ucts 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 have an incentive to market more polarizing products, and to target their advertising (or even “demarket”, Miklós-Thal and Zhang (2013)).

Clearly, our study is not without limitations, and in seeking a tractable model, we had to make 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.

For instance, 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’ opinion about this product?*”) and extremity bias (“*Do I feel strongly enough to take the time to leave a rating?*”), as well as their interaction. While we believe that the results in this paper are largely robust to such extensions (Appendix B), 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 matchmakers)? 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 inquiry.

Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

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 Louis Stevens. *The Good, The Bad and the Picky*: Reference dependence and the reversal of product ratings. 2022.
- Leif Brandes, David Godes, and Dina Mayzlin. Extremity bias in online reviews: The role of attrition. *Journal of Marketing Research*, 59(4):675–695, 2022.
- 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. *Available at SSRN 3305217*, 2021.
- 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):112–49, 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.
- 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 (JMR)*, 2022.

- 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.
- Guido Menzio. Optimal product design: Implications for competition and growth under declining search frictions. *Econometrica (forthcoming)*, 2023.
- Jeanine Miklós-Thal and Juanjuan Zhang. (de) marketing to manage consumer quality inferences. *Journal of Marketing Research*, 50(1):55–69, 2013.
- 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.
- Sungsik Park, Woochoel Shin, and Jinhong Xie. The fateful first consumer review. *Marketing Science*, 40(3):481–507, 2021.
- John W Payne, John William 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):1944–71, 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*, 2022.
- 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: The more polarizing product 2 is relatively overrated: $\mathcal{B}_\infty < 0$.

Proof of Claim 1:

Subtracting the second line in Equation 1 from the first we get:

$$\begin{aligned} E_\infty(\mathcal{R}_1) - E_\infty(\mathcal{R}_2) &= (Q_1 - Q_2) + E(E(\theta_{1j} \mid E_\infty(\mathcal{R}_1) + \theta_{1j} > E_\infty(\mathcal{R}_2) + \theta_{2j})) \\ &\quad - E(E(\theta_{2j} \mid E_\infty(\mathcal{R}_2) + \theta_{2j} > E_\infty(\mathcal{R}_1) + \theta_{1j})). \end{aligned} \quad (4)$$

Using the definition of

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

we can simplify this expression to

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

Now note that $Q_1 = Q_2$ by assumption, and thus $\mathcal{B}_\infty = (E_\infty(\mathcal{R}_1) - E_\infty(\mathcal{R}_2))$.

Therefore, the above equation can be rewritten as

$$\mathcal{B}_\infty = E(E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_\infty)) - E(E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_\infty)). \quad (5)$$

We now have one equation in one variable, \mathcal{B}_∞ . To show that a solution exists and that it is unique, first notice that the LHS of Equation 5 is (trivially) increasing in \mathcal{B}_∞ . The RHS, on the other hand, is decreasing in \mathcal{B}_∞ : this follows from the fact that $E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \mathcal{B}_\infty)$ is decreasing in \mathcal{B}_∞ , due to basic properties of conditional expectations, while the opposite is true for $E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \mathcal{B}_\infty)$.

To show that the (only) solution \mathcal{B}_∞ is negative, therefore, we have to show that *i*) if $\mathcal{B}_\infty = 0$, the RHS is negative and *ii*) if \mathcal{B}_∞ 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

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

Therefore, we have to show the following:

Lemma 3. *Let the assumptions of Proposition 1 hold. Then*

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

Proof: To prove the desired inequality, we start by defining the two conditional density functions,

$$f_L^{max}(\theta) := f_L(\theta \mid \theta_L > \theta_H)$$

and symmetrically for $f_H^{max}(\theta)$.

By independence, we have that

$$f_L^{max}(\theta) = \frac{f_L(\theta) \cdot \text{Prob}(\theta_H \leq \theta)}{P(\theta_L > \theta_H)} = \frac{f_L(\theta) \cdot F_H(\theta)}{M_L}.$$

Therefore,

$$\mathbb{E}(E(\theta_L \mid \theta_L > \theta_H)) = \frac{1}{M_L} \int_{\underline{\theta}}^{\bar{\theta}} \theta f_L(\theta) F_H(\theta) d\theta = \frac{1}{M_L} \int_{\underline{\theta}}^{\bar{\theta}} (\theta F_H(\theta)) f_L(\theta) d\theta = \frac{1}{M_L} \cdot \mathbb{E}_L(\theta F_H(\theta)),$$

where \mathbb{E}_L indicates that the expectation is taken over the $F_L(\cdot)$ distribution. Symmetrically, we have that $E(\theta_H \mid \theta_H > \theta_L) = \frac{1}{M_H} \cdot \mathbb{E}_H(\theta F_L(\theta))$.

Given that $M_H \geq M_L$, we are left with having to prove that

$$\mathbb{E}_L(\theta F_H(\theta)) > \mathbb{E}_H(\theta F_L(\theta)).$$

Proving the inequality directly is hard given that we are taking expectations of two different functions under two different distributions. Therefore, instead of comparing the two directly, we compare each of them to $\mathbb{E}_L(\theta F_L(\theta))$, to show that

$$\mathbb{E}_L(\theta F_H(\theta)) > \mathbb{E}_L(\theta F_L(\theta)) \geq \mathbb{E}_H(\theta F_L(\theta)).$$

Under the convexity assumption in Proposition 1, the second inequality follows from the fact $F_H(\cdot)$ SOSD $F_L(\cdot)$. This is because for any two random variables such that the first SOSD the second, expectations of convex functions are larger under the second. To prove the first, appeal to the definition of rotation directly: $F_L(\theta) > (<)F_H(\theta)$ for every $\theta < (>)\theta^\dagger$. In other words, both $\mathbb{E}_L(\theta F_H(\theta))$ and $\mathbb{E}_L(\theta F_L(\theta))$ are weighted averages of θ under $F_L(\cdot)$, but the second puts more weight on lower values of θ and less on higher ones. Therefore, $\mathbb{E}_L(\theta F_H(\theta)) > \mathbb{E}_L(\theta F_L(\theta))$, which concludes the proof. ■

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

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

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 captures a higher market share than it otherwise would: $M_{2,\infty} > M_L$.

Proof of Claim 2:

We have that

$$M_{2,\infty} = \text{Prob}(\theta_{2j} > \theta_{1j} + \mathcal{B}_\infty).$$

Since $\mathcal{B}_\infty < 0$, we have $M_{2,\infty} > M_L$. ■

Claim 3: Some self-correction occurs, and both biases are less severe than in the short-run: $\mathcal{B}_0 < \mathcal{B}_\infty < 0$, $M_{2,1} > M_{2,\infty} > M_L$.

Proof of Claim 3:

We have that

$$E_0(\mathcal{R}_1) = Q_1 + E(E(\theta_{1j} | \theta_{1j} > \theta_{2j})) \tag{6}$$

and similarly for product 2. Thus, $\mathcal{B}_0 = E_1(\mathcal{R}_1) - E_1(\mathcal{R}_2)$ implies

$$\mathcal{B}_0 = E(E(\theta_{1j} | \theta_{1j} > \theta_{2j})) - E(E(\theta_{2j} | \theta_{2j} > \theta_{1j})), \tag{7}$$

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

$$\mathcal{B}_0 = E(E(\theta_{1j} | \theta_{1j} > \theta_{2j} - \mathcal{B}_0)) - E(E(\theta_{2j} | \theta_{2j} > \theta_{1j} + \mathcal{B}_0)). \tag{8}$$

Therefore,

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

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 $M_{2,1} > M_{2,\infty}$ follows straightforwardly from $\mathcal{B}_0 < \mathcal{B}_\infty$ using the same argument as in the Proof of **Claim 2**. ■

Proof of Corollary 1

Assume without loss of generality the $Q = 0$. Thus, in equilibrium, the ratings for a product of design s are given by

$$E_\infty(\mathcal{R}(s)) = E(\theta_j | E_\infty(\mathcal{R}(s)) + \theta_j > c).$$

Assume by contradiction that $E_\infty(\mathcal{R}(s_H)) \geq E_\infty(\mathcal{R}(s_L))$. This implies

$$E(\theta_H \mid E_\infty(\mathcal{R}(s_H)) + \theta_H > c) \geq E(\theta_L \mid E_\infty(\mathcal{R}(s_L)) + \theta_L > c).$$

But then, *a fortiori*, we have

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

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

$$E(\theta_L \mid \theta_L > c) > E(\theta_H \mid \theta_H > c)$$

We have that

$$E(\theta_L \mid \theta_L > c) = \frac{\int_c^{\bar{\theta}} (1 - F_{s_L}(\theta)) d\theta}{1 - F_{s_L}(c)} \quad (9)$$

and similarly

$$E(\theta_H \mid \theta_H > c) = \frac{\int_c^{\bar{\theta}} (1 - F_{s_H}(\theta)) d\theta}{1 - F_{s_H}(c)}. \quad (10)$$

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

As a result,

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

and thus, substituting $c = \theta^\dagger$ in Equations (9) and (10), we have

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

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, $E(\mathcal{R}(s_L)) > E(\mathcal{R}(s_H))$ as desired.

To show that $M_\infty(s_L) > M_\infty(s_H)$, notice that, for every $c > \theta^\dagger$, we have

$$M_\infty(s_H) = 1 - F_{s_H}(c) < 1 - F_{s_L}(c) = M_\infty(s_L).$$

This concludes the proof. \blacksquare

Proof of Proposition 2

Claim 1: The higher quality product 1 has higher reviews: $E_\infty(\mathcal{R}_1) > E_\infty(\mathcal{R}_2)$

Proof of Claim 1:

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

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

But then, we have

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

which, using the definition of

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

simplifies to

$$\begin{aligned} \mathcal{B}_\infty &= E(E(\theta_{1j} \mid E_\infty(\mathcal{R}_1) + \theta_{1j} > E_\infty(\mathcal{R}_2) + \theta_{2j})) \\ &\quad - E(E(\theta_{2j} \mid E_\infty(\mathcal{R}_2) + \theta_{2j} > E_\infty(\mathcal{R}_1) + \theta_{1j})). \end{aligned} \tag{11}$$

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

$$\begin{aligned} &E(E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} + E_\infty(\mathcal{R}_2) - E_\infty(\mathcal{R}_1))) \\ &\quad - E(E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + E_\infty(\mathcal{R}_1) - E_\infty(\mathcal{R}_2))) \\ &= E(E(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(\mathcal{R}))) \\ &\quad - E(E(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(\mathcal{R}))) \\ &> 0 \end{aligned} \tag{12}$$

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 $E_\infty(\mathcal{R}_1) \geq E_\infty(\mathcal{R}_2)$.

To rule out equality, notice that if $E_\infty(\mathcal{R}_1) = E_\infty(\mathcal{R}_2)$ the RHS is Equation (11) 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 $E_\infty(\mathcal{R}_1) > E_\infty(\mathcal{R}_2)$. ■

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

Proof of Claim 2:

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

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

Using the definition of

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

we can simplify this expression to

$$\begin{aligned} \mathcal{B}_\infty &= \mathbb{E}(\theta_{1j} \mid \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j} > \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j}) \\ &\quad - \mathbb{E}(\theta_{2j} \mid \mathbb{E}_\infty(\mathcal{R}_2) + \theta_{2j} > \mathbb{E}_\infty(\mathcal{R}_1) + \theta_{1j}). \end{aligned} \quad (13)$$

Now assume $\mathcal{B}_\infty = 0$. Then, Equation (13) becomes

$$\mathbb{E}(\mathbb{E}(\theta_{1j} \mid Q_1 + \theta_{1j} > Q_2 + \theta_{2j})) = \mathbb{E}(\mathbb{E}(\theta_{2j} \mid Q_2 + \theta_{2j} > Q_1 + \theta_{1j})). \quad (14)$$

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

$$\mathbb{E}(\mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(Q))) = \mathbb{E}(\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(Q))). \quad (15)$$

But this is a contradiction, since:

$$\begin{aligned} \mathbb{E}(\mathbb{E}(\theta_{1j} \mid \theta_{1j} > \theta_{2j} - \Delta(Q))) &= \mathbb{E}(\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} - \Delta(Q))) \\ &< \mathbb{E}(\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j})) \\ &< \mathbb{E}(\mathbb{E}(\theta_{2j} \mid \theta_{2j} > \theta_{1j} + \Delta(Q))). \end{aligned} \quad (16)$$

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 captures a higher market share: $M_{1,\infty} > 1/2$, but less than it would if consumers were fully informed.

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

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

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: $\mathcal{B}_1 < \mathcal{B}_\infty < 0$, $M_{1,\infty} > M_{1,2}$.

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

$$E_1(\mathcal{R}_1) = Q_1 + E(E(\theta_1 \mid \theta_1 + E_0(\mathcal{R}_1)) > \theta_2 + E_0(\mathcal{R}_2)) = Q_1 + E(E(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q)))$$

and similarly

$$E_1(\mathcal{R}_2) = Q_2 + E(E(\theta_2 \mid \theta_2 + E_0(\mathcal{R}_2) > \theta_1 + E_0(\mathcal{R}_1)) = Q_2 + E(E(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q))).$$

Jointly, these two imply that

$$\begin{aligned} \mathcal{B}_1 &= (Q_1 + E(E(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q)))) - (Q_2 + E(E(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q)))) - Q_1 + Q_2 \\ &= E(E(\theta_1 \mid \theta_1 > \theta_2 - \Delta(Q))) + E(E(\theta_2 \mid \theta_2 > \theta_1 + \Delta(Q))). \end{aligned} \tag{17}$$

But then, comparing the expressions in Equation (13) and (17), we obtain $\mathcal{B}_1 < \mathcal{B}_\infty$ if and only if $\Delta(Q) > E_\infty(\mathcal{R}_1) - E_\infty(\mathcal{R}_2)$. 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 $M_{1,2} < M_{1,\infty}$ follows straightforwardly from $B_1 < \mathcal{B}_\infty$. ■

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. Notice that $\Delta_\infty(R) < \Delta(Q)$ – or, equivalently, \mathcal{B}_∞ , 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

$$1 - Prob(j \in \mathcal{J}_{NR}) < 1 - Prob(j \in \mathcal{J}_{BR}) < 1, \quad (18)$$

which concludes the proof. \blacksquare

Proof of Corollary 2

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

$$\mathcal{J}_i^{t+1} = \{j \mid \theta_{ij} + E_t(R_i) > \theta_{-ij} + E_t(R_{-i})\}$$

Now let $E_t(R_i)$ increase to $E_t(R_i) + \epsilon$. 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:

$$E(E(\theta_{ij} \mid \theta_{ij} + E_t(R_i) + \epsilon > \theta_{-ij} + E_t(R_{-i}))) < E(E(\theta_{ij} \mid \theta_{ij} + E_t(R_i) > \theta_{-ij} + E_t(R_{-i}))),$$

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

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

Proof of Corollary 3

The proof follows directly from our previous results. \blacksquare

Proof of Corollary 4

Omitting the depend of \mathcal{R} on Q , we have that, in equilibrium, ratings satisfy

$$E_\infty(\mathcal{R}) = Q + E(\theta \mid \theta + E_\infty(\mathcal{R}) > c). \quad (19)$$

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

To show this in a straightforward fashion, notice that as Q increases by ϵ , so does the RHS of Equation (19).

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

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

Thus, we have

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

Market shares are given by

$$M_\infty(Q) = \text{Prob}(\theta + E_\infty(\mathcal{R}) > c),$$

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

Proof of Proposition 3

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

$$\mathcal{J}_1^{NR} = \emptyset \tag{20}$$

whereas in the presence of biased reviews we have

$$\mathcal{J}_1^{BR} = \{j \mid \theta_{1j} < \theta_{2j} - \Delta_\infty(\mathcal{R}) \ \& \ \theta_{2j} < \theta_{1j}\}. \tag{21}$$

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 $1 - \text{Prob}(j \in \mathcal{J}_1^{BR})$ 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

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

and

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

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

But then,

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

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

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

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

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

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

The denominator is given by

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

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 equation from the first in (25), we get:

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

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

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

and

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

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

But then,

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

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

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

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

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

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

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

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 3 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 5

We want to show that

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

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

$$\text{Var}(\theta_H \mid \theta_H > \theta_H - \Delta(Q)) \rightarrow \text{Var}(\theta_H).$$

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

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

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

B. Alternative Rating 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. (2022) 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.³¹ 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 $\mathbb{E}(\theta_{ij} \mid \theta_{ij} \in \mathcal{J}_i^{90+})$.

31. 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.

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.

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), Sun (2012) and more recently Menzio (2023) – the conclusions are ambiguous. Again, spelling these out in greater detail seems like a promising avenue for future research.

Rating to Persuade

In our model, consumers are not strategic in their rating 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. 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 rating 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

$$L^S(\mathcal{R}_{ij} \mid Q_i, \theta_{ij}, \mathcal{R}_i) := -(Q_i + \theta_{ij} - \mathbb{E}_{t+1}(\mathcal{R}_i \mid \mathcal{R}_{ij}))^2$$

instead of the purely individual (I) one

$$L^I(\mathcal{R}_{ij} \mid Q_i, \theta_{ij}, \mathcal{R}_i) := -(Q_i + \theta_{ij} - \mathcal{R}_{ij})^2.$$

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 $E_{t+1}(\mathcal{R}_i | \mathcal{R}_{ij}) \approx E_{t+1}(\mathcal{R}_i)$ whenever the number of reviews 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, $L^S(Q_i - \theta_{ij} - P_i | Q_i, \theta_{ij}, P_i, \mathcal{R}_i) \approx \max_{\mathcal{R}_{ij}} L^S(\mathcal{R}_{ij} | Q_i, \theta_{ij}, \mathcal{R}_i)$. 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, $\mathcal{R}_{ij}^I - \mathcal{R}_{ij}^S$:

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

In other words, consumer j^* strategic rating 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 2, 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,

$$\frac{\partial \mathcal{R}_{ij}^{t+1}(Q_i, \theta_{ij}, E_t(\mathcal{R}_i))}{\partial E_t(\mathcal{R}_i)} > 0.^{32}$$

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 B. There, consumers effectively look as contrarians (despite their lack of 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 reputation, in line with what described in our Proposition 2 and Corollaries 2 and 3.

Understanding when social influence is the dominant force, and when, on the contrary, taste-based self-selection leads to fragmented market outcomes, seems like a promising research question moving forward.

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

C. “Observational” Learning: the Number of Reviews

Anecdotal evidence suggests that consumers have a preference for goods which are purchased by many of their peers. One possible explanation is that there are network effects of some sort. However, even absent these network effects there is a premium for high-sales products. Caminal and Vives (1996) provide an explanation: consumers use market shares as a signal of product quality. Powell et al. (2017) suggest that the quantity premium results from a psychological bias (“love of large numbers”).

When it comes to consumer reviews, a similar empirical regularity is observed: consumers seem to have greater regard for products with a greater number of reviews (Powell et al., 2017). In this case, the puzzle is deeper since the explanation in Caminal and Vives (1996) seems to fail: if consumers have direct information regarding quality (average quality reviews) what additional information can the number of reviews contain?

Our model – slightly adapted and expanded in here – provides one possible explanation based on two observations/assumptions: (a) consumers rate products according to their subjective experience, which in turn reflects both objective product quality and subjective fit between product characteristics and consumer preferences; (b) higher-quality products attract more consumers, including in particular consumer for whom the fit component is lower. A combination of (a) and (b) implies that rational consumers should use both average review and number of reviews as a measure of quality.

To see how this may be the case, consider the following stylized model. There are n sellers, each of which sells one product. The product can be one of two types, type a and type b . In each of two periods, there is a measure 1 of consumers who are equally divided in terms of preferences for product type (that is, a measure $\frac{1}{2}$ has a preference for type a products. Let τ be the disutility from consuming a product of type different from the preferred type.

In addition to this element of horizontal product differentiation, we also assume that each product is characterized by vertical quality q , where quality units are the same as utility units.

Since the focus of the analysis is on learning about quality and match value, we assume that prices are exogenously given; and with no additional loss of generality, we assume prices are zero.³³

There are two periods. First-period consumers do not have access to reviews. Second-period consumers, by contrast, have access to the reviews issued by first-period consumers.

Consider the case of first-period consumers. We assume they are randomly presented with one product and learn the product’s quality q as well as its type t . The decision problem is then easy: If the product’s type is identical to the consumer’s type, then the consumer makes a purchase if and only if $q > 0$. If, by contrast, the product and the consumer are of different type, then the consumer makes a purchase

33. Alternatively, one may think of q as quality net of price.

if and only if $q > \tau$. Finally, we assume that consumers leave an honest review after purchase, that is, the review score equals utility level q if there is a match of types and $q - \tau$ if there is no such match.




Figure 9

Relation between product quality q and first periods sales x (left panel) and between quality q and average consumer review r (right panel)

Figure 9 shows the relation between product quality q and first periods sales x (left panel) and between quality q and average consumer review r (right panel). As the left panel shows, if quality q is greater than the τ threshold, then sales double, as the product attracts from both consumers with good fit and consumers with bad fit. As the right-hand panel shows, the relation between q and r is non-monotonic. This is because, as quality increases, the product attracts buyers for whom the fit component is lower, resulting in lower reviews.

Second period consumers do not observe the value of q (or don't need to make the investment into finding the value of q). Rather, they observe two summary statistics regarding first period purchase decisions: x and q . For simplicity, suppose that all buyers rate the product, so that x is both sales and number of reviews.

Similarly to first-period consumers, second-period consumers make a purchase if and only if net utility is positive. If there is a product fit, then the rule is $E(q) > 0$. If there isn't a product fit, then the rule is $E(q) > \tau$. Rational consumers "invert" the mapping on the right-hand panel of Figure , and so

$$E(q) = \begin{cases} r & \text{if } x \leq \frac{1}{2n} \\ r + \frac{1}{2}\tau & \text{if } x > \frac{1}{2n} \end{cases}$$

The main point is that rational consumers use both r and x as measures of quality. This is not a bounded rationality issue or a network effects issue. It simply results from correcting for the subjective element in reviews and the "curse of success" in consumer reviews highlighted in Proposition 2 and Corollaries 2 and 3.