Information Manipulation of Wishful Thinkers

Emel Filiz-Ozbay [*] Zhenxun Liu [†]

May 7, 2024

Abstract

Digital platforms may manipulate the customer reviews to impede the consumers' abilities of acquiring accurate information and to increase sales. In a rational inattention setting where consumers acquire information on the good's quality before making purchasing decisions, we show that the platform has incentives to restrict information acquisition in the presence of consumers with wishful thinking. These biased consumers are unaware of their bias and weigh any good news about the product quality more heavily than a Bayesian consumer. We identify the conditions for which the firm optimally constrains information acquisition, and characterize when competition can alleviate this type of exploitation. Our results extend to a class of other non-Bayesian rules beyond wishful thinking.

Keywords — Rational Inattention, Non-Bayesian Updating, Information Transparency

[*]University of Maryland, 3114 Tydings Hall, Department of Economics, College Park, MD 20742. Email: efozbay@umd.edu.

[†]University of Maryland, 3114 Tydings Hall, Department of Economics, College Park, MD 20742. Email: zliu1995@umd.edu.

1 Introduction

In the age of information, discovering product quality is still a challenging task for consumers. Many platforms (such as Amazon, Yelp etc.) provide independent product reviews for the products they host and online reviews are listed as the second most trusted source of information.[1] Nevertheless, anyone who has ever shopped online may know that reading those reviews is hard and they may not be fully informative even if one is willing to spend hours on them. Why does information remain so difficult to acquire and process? In this paper, we argue that the platforms may have an incentive not to provide transparent reviews in an easily accessible way and competing platforms can promote the accessibility of information, but only to a certain extent.

We model consumers (agents) who are rationally (in)attentive to information, and use the acquired information to guide their purchasing decision of a product from a platform (firm or principal). Moreover, a proportion of the consumers are biased when updating their initial beliefs - they exhibit wishful thinking in favor of the product being high quality.[2] We consider the case where the product in consideration is a third party product, and the platform has no control over its quality and price. Nevertheless, the platform still has the objective to maximize sales. While the platform cannot exactly choose the information on behalf of the consumers, it has the resources to control the quality of information by imposing an upper bound on the informativeness of signals.[3] We characterize the conditions under which the platform may benefit from limiting the accuracy of reviews to increase the sales. We also show that competition among platforms never hurts the consumers, and will sometimes strictly benefit consumers by eliminating information exploitation from the platforms. Our findings are consistent with the recent

[1] According to a Nielsen's report, online reviews have 70% approval rate among 25,000 survey respondents. See www.nielsen.com, 2009 Report "Global Advertising Consumers Trust Real Friends and Virtual Strangers the Most"

[2] Consider the scenario where the consumer has found the ideal product based on the product description, and she hopes that the quality is good as well.

[3] The platform may achieve this via deleting reviews or adding fake reviews to the customer review section. In practice, the sellers may also have the power to manipulate the reviews themselves, but we do not consider this case here.

federal register notice of the Federal Trade Commission[4] that is motivated by the lack of transparency in commercial reviews and that aims to promote honest businesses and prevent deception of consumers looking for real feedback on a product.

Our model is an application of the rational inattention approach introduced by Sims (2003) (and also Sims, 2006.)[5] This model is suitable for agents who hold an initial prior on the state of the world and would like to pick an information structure to reduce uncertainty before taking an action. Since the utility depends on the action and the state of the world, it is better to pay attention to a more informative signal structure to avoid mistakes but attention costs are proportional to the reduction of uncertainty in the agent's belief. This approach has been utilized in other principal/agent problems such as dynamic bargaining (Ravid, 2020) and sequential pricing (Martin, 2017; Boyacı and Akçay, 2018).[6] Rational inattention in these settings lead to non-standard results such as delayed attention in dynamic settings or uninformative prices in a quality discovery problem. Similar to these applications, the principal (the platform) benefits from sales in our setup. However, different from the other papers, the principal in our setting does not care about, nor can it influence prices.[7] Instead, the platform exerts influence on how informative the available information is. In applications, platforms design product review pages, control their accessibility (often with the help of algorithms), and implement certain review policies and features. These strategies or their absence might be interpreted as tools to improve or limit the informativeness of signals acquired by the consumers. To the best of our knowledge, ours is the first application of rational inattention where such constraints on consumers are analyzed.

---

[4] "Trade Regulation Rule on the Use of Consumer Reviews and Testimonials," 16 CFR Part 465-Matter Number: R311003, July 31, 2023.

[5] This framework has become a groundbreaking tool for information economics. It relaxed the assumption of agents being able to process all information freely. Instead, it models agents who understand the trade-off between using better information and the cost of acquiring that information.

[6] Several other papers apply rational inattention to pricing, where Mackowiak and Wiederholt (2009) and Woodford (2009) assume the price setters to be rationally inattentive, while Matějka and McKay (2012), Matějka (2015), and Steiner et al. (2017) maintain the assumption that the buyers are rationally inattentive. Caplin and Dean (2015), Matějka and McKay (2015), de Oliveira et al. (2017), and Steiner et al. (2017) also build on this model.

[7] As we are considering products sold by third parties, the seller (not modeled here) determines the price and the platform receives a fixed commission per unit of sale.

We first show that the platform cannot benefit from imposing an upper bound on the accuracy of information if consumers are Bayesian. Such consumers will choose their attention level optimally, update their beliefs in the Bayesian sense, and follow the signal they receive when deciding whether to purchase the product. There is, however, significant evidence against the Bayesian updating assumption and even those who intend to be Bayesian may fail to employ the rule when they are actually updating their priors (see e.g., Camerer, 1998; Benjamin, 2019; Liu, 2023b). Several non-Bayesian updating rules have been proposed and implemented in well-known decision environments in economics[8] and one of the most prevalent deviations from Bayesian updating is wishful thinking. This bias makes agents form beliefs more in line with their desire rather than rationality implied by Bayesianism. In our two-state environment, we interpret this as the possibility of some agents asymmetrically overweighting good news about the product quality relative to bad news, due to their desire to land a high quality product (see Mayraz, 2011; Kovach, 2020; Liu, 2023a). Such bias provides an incentive for the platform to limit the informativeness of the information. The intuition is simple. Let the consumption utilities of buying high and low-quality products be 1 and -1, respectively. Let the prior belief of quality being high be slightly less than being low. If the posterior belief on high quality exceeds 0.5, the consumer will buy the product (assuming that the utility of not buying is zero). A rationally inattentive, Bayesian person would follow the signal recommendation and in expectation buy exactly at the rate implied by the marginal probability of receiving good signals. Consider a wishful thinker instead. Even if she chooses her optimal information structure, her posterior will be more optimistic than the Bayesian after a good signal and less pessimistic after a bad signal. Hence, this consumer may under-interpret a bad signal or over-interpret a good signal and end up with a posterior of high-quality being more than 0.5. The existence of such non-Bayesians benefits the platform by increasing the volume of sales in expectation. We characterize the parameters of the model that will lead to such scenarios.

Next, we introduce competition into the model where the incumbent platform might

[8] Examples include Edwards (1982); Grether (1980); Rabin and Schrag (1999); Epstein (2006); Ortoleva (2012); Cripps (2018); Dominiak et al. (2023); Liu (2023a)

be challenged by an entrant platform. In this setup, consumers randomly choose their platform to collect information, and if they feel constrained then they switch to the other platform looking for more information. We characterize possible equilibrium outcomes of this competitive game. We show that for a certain range of parameters, the introduction of competition achieves unrestricted information structure and benefits consumers. However, there is a non-empty set of parameters where the platforms coordinate and both of them restrict the information. Given that some platforms have grown too big and competition against them by using standard tools, such as price and product variety, has become almost impossible, competition through information quality might be an effective tool against monopolies. Nevertheless, this result depends on the market power of the incumbent. For too large platforms, regulations for transparency are still justified.

Since the principal in our setup can control the information, our paper also relates to the persuasion literature (Gentzkow and Kamenica, 2014). de Clippel and Zhang (2022) extend the standard persuasion framework to games between a Bayesian sender and a non-Bayesian receiver. Wei (2021) analyzes optimal persuasion for rationally inattentive receivers. The game we study is related with these but ours is not a persuasion game since the platform does not know the state of the world in our setup, and it cannot choose a signal structure. It can only constrain the informativeness of the available signal structures of the consumer.

Finally, our paper relates to the literature on exploitative contracts. In these papers, at least a proportion of consumers are either unaware of their own bias or unaware of some product attribute, and the firms exploit this unawareness. DellaVigna and Malmendier (2004) discussed the implications of present bias for the optimal pricing strategy of subscription-based products such as gym memberships. Gabaix and Laibson (2006) showed that firms can also manipulate naive consumers by actively hiding relevant information about add-on purchases. The firm of our model exploits the consumer's unawareness of their biased updating, but distinct from this literature, it manipulates the consumer with blurry information rather than manipulated price.

In the following sections, we first start with introducing the game (Subsection 2.1) followed by characterizing the optimal attention level of the consumer whose information might be constrained (Subsection 2.2). Then we introduce non-Bayesian consumers to the model and study the purchasing decisions of Bayesian and biased consumers (Subsection 2.3.) In Subsection 2.4, we study the best response of a monopolist platform and characterize when it is optimal for it to constrain the information available to the consumers. In Subsection 2.5, we study the competition between platforms and characterize when competition on information can help consumers. Subsection 3 introduces extensions of our setup to other non-Bayesian rules besides wishful thinking. Subsection 4 discusses alternative models such as the case where the firm knows the product quality (the state of the world) before constraining information or where consumers have other types of heterogeneity rather than information costs. The concluding remarks finalizes Section 4.

2 Rational Inattention with Wishful Thinking

2.1 Setup

A platform (firm) sells a third-party product through its platform. A population of consumers use the platform to gather more information about the product quality (for instance, by reading the customer reviews) and decide whether to buy the product.[9]

Specifically, the product quality, ω, is unknown and it can be "High" or "Low", $\omega \in \{H, L\}$. The consumers (as well as the platform) have a common prior belief over the quality, $f(H) \leq 0.5$, and can choose one of two actions, "buy" or "reject" denoted by $a \in \{b, r\}$.[10] We assume a normalized utility:

$$u(b|H) = 1 > 0 = u(r|H) = u(r|L) > -1 = u(b|L)$$

Before making a purchasing decision, each consumer chooses how much information about the product quality she wants to gather.[11] We employ a standard rational inattention

[9]The price and quality are determined by a third party that is not modeled here.

[10]The exercise can be easily extended for the case of $f(H) > 0.5$ but the solution for that would be trivial, as we will discuss in Section 4.

[11]We assume that consumers always acquire signals, and in the case that acquiring any meaningful information is too costly, they will simply choose a totally uninformative signal

framework for information acquisition. Formally, a consumer chooses a signal structure $\pi(s|\omega)$ for two possible signals $s \in \{h, l\}$, which costs her $\lambda I(\pi)$.[12] λ represents the consumer's marginal cost of information acquisition. We assume that λ is private information and the consumers have heterogeneous cost parameters. The firm knows the distribution of this parameter, $\lambda \sim G[0, \infty)$. $I(\pi)$ measures informativeness of the signal by calculating the degree of correlation between the signal and the state.[13]

$$I(\pi) = \sum_s \sum_\omega \pi(s|\omega) f(\omega) ln(\frac{\pi(s|\omega)}{q(s)})$$

where $q(s)$ denotes the marginal probability of signals given the prior, f, and the signal structure, π, i.e. $q(s) = \sum_\omega \pi(s|\omega) f(\omega)$.

Consumer's problem of choosing π: We can write the consumer's information acquisition problem as

$$\begin{aligned} max_\pi \sum_\omega \sum_s V(s, \pi)\pi(s|\omega) f(\omega) - \lambda I(\pi) \\ \text{subject to } I(\pi) \leq \kappa \end{aligned} \tag{1}$$

Here, V denotes the expected utility of the consumer who believes that she will buy or reject optimally under information structure π. Formally,

$$V(s, \pi) = max_{a \in \{b,r\}} g^B(H|s)u(a, H) + g^B(L|s)u(a, L) \tag{2}$$

where $g^B(\omega|s)$ is the Bayesian posterior generated by π.

κ in the constraint of problem (1) is an upper bound on informativeness of the signal structure. Shortly, we will introduce the firm as the party which sets κ optimally. Before

structure.

[12] We have two possible signals since having more signals than the number of actions is redundant.

[13] This is called the "mutual information" between the signal and the state. It is also equivalent to the "relative entropy", or the Kullback-Leibler divergence, between the joint distribution of s and ω and the product of their marginal distributions. Entropy is a measure of uncertainty implied by a distribution, and for distribution $\gamma(x)$, Entropy is defined as $-\sum_x \gamma(x) ln(\gamma(x))$.

doing that, first we characterize the solution to problem (1).

2.2 Characterizing the Optimal Signal Structure, π

Note that Matějka and McKay (2015) Lemma 1 applies to our setting, and hence, distinct signals should not lead to the same action if the information is selected optimally because it is inefficient to acquire information that will not be acted upon.[14] One can solve problem (1) using the standard Lagrangian method and the solution to it is summarized in Lemma 1 below. Note that the optimal information for this problem is a piece-wise function with a cutoff cost parameter, $\underline{\lambda}$. This cutoff is the cost parameter of the consumer who will optimally choose an information structure with exactly $I(\pi) = \kappa$.

In the proof of Lemma 1, we first observe that the upper bound on the informativeness of the information structure translates into an equivalent lower bound on the marginal cost of information, λ, pushing all those who have a lower cost act like they have the cost of $\underline{\lambda}$. The details are shown in the appendix. Intuitively, only those who have low enough λ would like to acquire a large amount of information and hence, they will be the ones who are constrained by the firm's choice of κ. The next best option for these constrained consumers is to mimic those at the threshold with $\lambda = \underline{\lambda}$. Those who do not acquire that much information to start with are unaffected by the constraint.

Lemma 1. *Given a constraint κ, the optimal information structure of a consumer with cost parameter $\lambda \in [0, \infty)$, is a piece-wise function defined by $\underline{\lambda}$ such that $I(\underline{\lambda}) = \kappa$ and*

$$\pi^*(h|H) = \begin{cases} 0 & q^*(h) = 0 \\ \frac{q^*(h)e^{1/max\{\underline{\lambda},\lambda\}}}{f(H)(1+e^{1/max\{\underline{\lambda},\lambda\}})} & q^*(h) \in (0,1) \\ 1 & q^*(h) = 1 \end{cases} \quad \pi^*(l|L) = \begin{cases} 0 & q^*(l) = 0 \\ \frac{q^*(l)e^{1/max\{\underline{\lambda},\lambda\}}}{f(L)(1+e^{1/max\{\underline{\lambda},\lambda\}})} & q^*(l) \in (0,1) \\ 1 & q^*(l) = 1 \end{cases} \tag{3}$$

with $\pi^*(h|\omega) = 1 - \pi^*(l|\omega)$, where

[14] This property is implied by the convexity of the entropy-based cost function that is used in the rational inattention model.

$$q^*(h) = \begin{cases} max\left\{0, min\left\{1, \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}}\right\}\right\} & \lambda \in [\underline{\lambda}, \infty) \\ max\left\{0, min\left\{1, \frac{f(L)-f(H)e^{1/\underline{\lambda}}}{1-e^{1/\underline{\lambda}}}\right\}\right\} & \lambda \in [0, \underline{\lambda}) \end{cases} \quad (4)$$

and $q^*(l) = 1 - q^*(h)$.

The marginal probability of observing signal h equals to the fractional term within the brackets in Equation (4) provided that it is in $(0, 1)$ interval. q^* is continuous in f and λ, and one may further show that π^* is also continuous in f and λ. One should also note that under our assumption of $f(H) \leq 0.5$, $q^*(h) \neq 1$. Nevertheless, the result of Lemma 1 applies without this assumption, hence, we state it in a more general way.

2.3 Bayesian and non-Bayesian consumers and their purchasing decisions

After picking the optimal information as described by Lemma 1, consumers observe their signals and update their beliefs. At this stage proportion $p \in (0, 1)$ of the consumers update their belief in a biased way, while the remaining $1 - p$ are Bayesian. Importantly, we let the biased consumers be naive, and not factor in their bias when choosing the signal structure.[15] In this section, we will use an updating rule that describes wishful thinking (Kovach, 2020). Wishful thinking is a natural explanation for why a person would hold optimistic beliefs about a product's quality. However, the results we will derive for wishful thinkers can be generalized to a much larger set of updating rules, which will be discussed in Section 3.

Let $g^\delta(\omega|s)$ denote the posterior belief of a consumer who distorts the Bayesian posterior with parameter vector $\delta = (\delta_H, \delta_L)$.

$$g^\delta(\omega|s) = \frac{f(\omega)\pi(s|\omega)\delta_\omega}{f(H)\pi(s|H)\delta_H + f(L)\pi(s|L)\delta_L} \quad (5)$$

[15]These consumers are unaware of their own bias, and thus expect themselves to interpret information objectively when they choose the optimal π while solving problem (1). Cognitive biases are largely subconscious processes. According to Korteling and Toet (2021): "The intuitive processes that precede biased judgments and decisions, and that are at the basis of the behavior that is ultimately shown, are largely implicit and unconscious".

Here, δ is the distortion parameter vector of a wishful thinker who uses a state-specific weight, as in Kovach (2020). g^{δ} is equal to the Bayesian posterior when $\delta_H = \delta_L$. We assume that biased consumers have $\delta_H \geq \delta_L > 0$, which represents the wishful thinking that the product is high quality.[16] Finally, we assume that being a biased consumer and the realization of λ are independent events, as it is intuitive to assume that having a high information processing cost of λ and being a wishful thinker are not correlated.

After obtaining a signal generated by the information structure, the consumers update their beliefs and decide to buy or reject by maximizing their (possibly biased) expected utility:

$$V^{\delta}(s,\pi) = max_{a\in\{b,r\}} g^{\delta}(H|s)u(a,H) + g^{\delta}(L|s)u(s,L)$$

The expected utility (V^{δ}) above depends on the posterior beliefs which might be biased as the bias only hits a portion of the consumers when they are updating their priors after collecting a signal. Hence, for a non-Bayesian consumer, $V^{\delta}(s,\pi)$ may be different from $V(s,\pi)$ calculated in Equation (2).

It is easy to observe that the optimal action a^* is to "buy" the product if $g^{\delta}(H|s) > g^{\delta}(L|s)$, and it is to "reject" it otherwise, due to the normalization of the utility function. Hence, we need to analyze which signals will lead to each of these for Bayesian and non-Bayesian consumers.

Before stating the overall optimal behavior of consumers in Lemma 2, we will provide the intuition behind the purchasing decision. First, one can obtain the actual posterior belief of each consumer using (5), and determine their expected utility of buying as a function of the signal received. The proof of Lemma 2 derives Equation (6) for $q(h) \in (0,1)$ and Equation (7) for $q(h) = 0$.

$$Eu^{\lambda}(b|s=l) = \frac{\delta_H - e^{\frac{1}{max\{\underline{\lambda},\lambda\}}}\delta_L}{\delta_H + e^{\frac{1}{max\{\underline{\lambda},\lambda\}}}\delta_L} \quad \text{and} \quad Eu^{\lambda}(b|s=h) = \frac{e^{\frac{1}{max\{\underline{\lambda},\lambda\}}}\delta_H - \delta_L}{e^{\frac{1}{max\{\underline{\lambda},\lambda\}}}\delta_H + \delta_L} \tag{6}$$

[16]For instance, a person who browses Amazon and sees a pair of shoes that she loves would also wish to obtain confirmation from the customer reviews that the product has good quality.

$$Eu^{\lambda}(b|l) = \frac{f(H)\delta_H - f(L)\delta_L}{f(H)\delta_H + f(L)\delta_L} \tag{7}$$

Recall that the utility of rejecting is always 0. Given the choice of π^*, the consumers' optimal action can be discussed under four cases, summarized in Table 1. The empty cells correspond to cases not applicable to the Bayesian consumers because they do not have a bias ($\delta_H = \delta_L$). Whether a consumer acquires informative or uninformative signal and her purchasing decision after collecting a signal varies with her cost, λ, and the relation between her prior and bias. We illustrate each case in Table (1) also on Figures (1) and (2) for different cost parameters. These figures will be useful when we analyze the firm's optimal strategy later.

	$q(h) \in (0,1)$		$q(h) \in \{0,1\}$	
	$\delta_H < e^{1/max\{\lambda,\underline{\lambda}\}}\delta_L$	$\delta_H \geq e^{1/max\{\lambda,\underline{\lambda}\}}\delta_L$	$f(H)\delta_H \geq f(L)\delta_L$	$f(H)\delta_H < f(L)\delta_L$
Bayesian ($\delta_H = \delta_L$)	Follow s	–	–	Always Reject
Biased ($\delta_H > \delta_L$)	Follow s	Always Buy	Always Buy	Always Reject

Table 1: Consumer's best response.

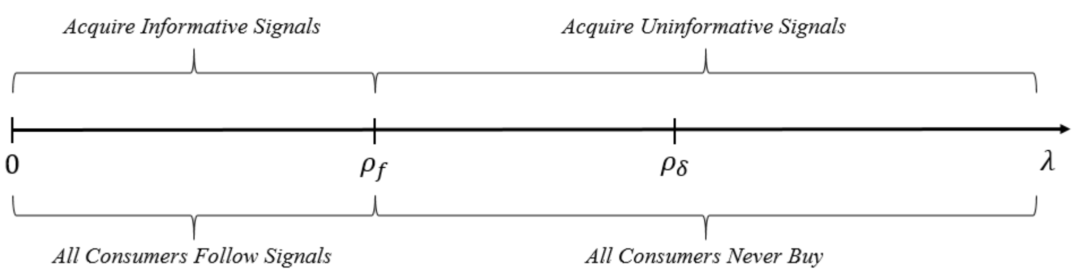

Figure 1: Information acquisition and purchasing decisions depending on λ, given $\rho_f < \rho_\delta$.

Below we explain the intuition behind the behavior described in each cell of Table 1.

Case 1: $0 < q(h) < 1$

Sub-Case 1.1: $\delta_H \geq e^{1/max\{\lambda,\underline{\lambda}\}}\delta_L$

This results in $Eu^{\lambda}(b|s = h) > Eu^{\lambda}(b|s = l) > 0$, and the biased consumer will always buy regardless of signal. In other words, the consumer's bias is large enough for them to ignore a bad signal. This case doesn't occur for a Bayesian since $\delta_H = \delta_L$.

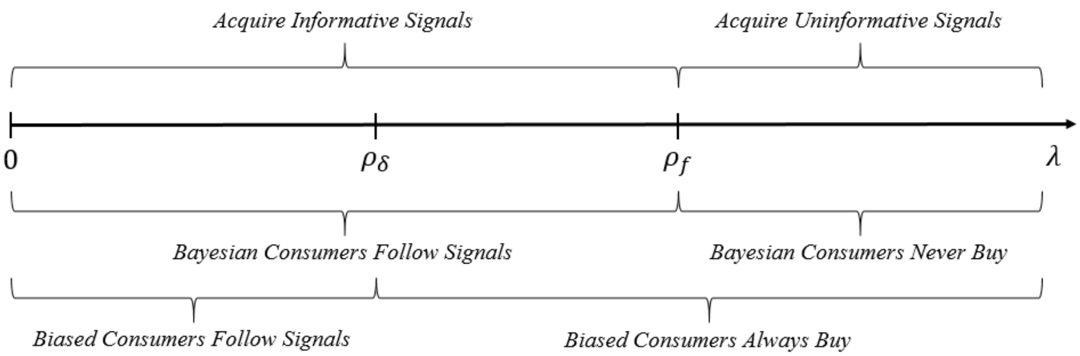

Figure 2: Information acquisition and purchasing decisions depending on λ, given $\rho_f \geq \rho_\delta$.

Sub-Case 1.2: $\delta_H < e^{1/max\{\lambda,\underline{\lambda}\}}\delta_L$

This results in $Eu^\lambda(b|s=h) > 0 > Eu^\lambda(b|s=l)$, and the biased consumer will follow the action recommended by their acquired signal, i.e. "buy" if $s=h$ and "reject" if $s=l$. For a Bayesian person, $\delta_H = \delta_L$, and this is always the case.

Next is the extreme case of $q(h) \in \{0,1\}$. Recall that, according to Lemma 1 the optimal signal structure is the uninformative one for this case.

Case 2: $q(h) \in \{0,1\}$

Observe that for this case Equation (7) applies and the utility does not depend on λ. For a Bayesian person, the expected utility is $2f(H)-1 \leq 0$, and thus will never buy.[17] For a biased person, there are again two cases to consider.

Sub-Case 2.1: $f(H)\delta_H \geq f(L)\delta_L$

The biased consumer always buys in this case, as the expected utility of buying is greater than or equal to 0.

Sub-Case 2.2: $f(H)\delta_H < f(L)\delta_L$

The biased consumer never buys in this case, as the expected utility of buying is less than 0.

[17]Note that $2f(H)-1<0$ when $f(H)<0.5$, and it is equal to 0 when $f(H)=0.5$. However, when $f(H)=0.5$, $q(h)=q(l)=0.5$, and the consumers will always acquire somewhat informative signals, and will never reach the case of $q(h)\in\{0,1\}$.

To simplify the notation for the statement of Lemma 2, define the following:

$$\rho_\delta = \frac{1}{ln(\delta_H) - ln(\delta_L)} \qquad \text{and} \qquad \rho_f = \frac{1}{ln(f(L)) - ln(f(H))} \tag{8}$$

By our assumptions on f and δ, these are positive (with the exception that the former is undefined for the Bayesian individual). We can potentially re-write the thresholds for each of the above cases with ρ_δ, ρ_f and λ. The following lemma summarizes our discussion above, while emphasizing on the important thresholds in relation to λ.

Lemma 2. *The individual acquires informative signals, i.e., $q(h) \in (0,1)$, if and only if $max\{\lambda, \underline{\lambda}\} < \rho_f$. The optimal purchasing behavior of Bayesian and non-Bayesian consumers are as follows:*

- *If the individual is Bayesian, they will follow the signal recommendation (i.e. buy the product if and only if the signal is high).*
- *If the individual is biased, they will ignore the bad signal of l, and always buy the product if and only if $max\{\underline{\lambda}, \lambda\} \geq \rho_\delta$ or $\rho_f \geq \rho_\delta$. They will follow the signal recommendation otherwise.*

The proof is in the appendix.

Since we have derived the optimal information choice and purchasing decisions of the consumers, next, we will examine the firm's problem of selecting the optimal constraint κ. The figures presented earlier are useful to visualize the information choice strategies and purchasing decisions of Bayesian and non-Bayesian consumers. We will refer to these figures in our derivation of the firm's optimal constraint, κ.

2.4 Firm's optimal constraint on information

The firm can limit information acquisition by setting κ. This may correspond to instances where platforms manipulate customer reviews to their advantage. This inevitably affects the views of the potential buyers, and the probability of purchasing the product. Recall that a consumer with cost λ purchases when her posterior satisfies $g^\delta(H|s) > g^\delta(L|s)$). Then the firm's problem is:

$$max_\kappa \int_\lambda E[\mathbb{1}(g^\delta(H|s) > g^\delta(L|s))|\lambda]dG(\lambda) \tag{9}$$

According to the previous discussions, for any κ that the firm chooses, there exists a unique $\underline{\lambda}$ such that $I(\pi^{\underline{\lambda}}) = \kappa$. Therefore, the ensuing analyses will assume that the firm is directly setting $\underline{\lambda}$. As observed in Lemma 1, having a constraint only affects those consumers whose constraints bind, i.e., $\lambda \leq \underline{\lambda}$. Proposition 1 identifies a threshold for the prior above which the optimal $\underline{\lambda}$ is non-zero, and hence, the firm wants to constrain the consumers. Otherwise, the firm cannot benefit from blurring the consumer's information. The proof is in the appendix and it analyzes the change in expected profit when the firm imposes a $\underline{\lambda}$ versus when it doesn't constrain at all.

Proposition 1. *Let $p > 0$, $\exists \hat{f} \in [\frac{\delta_L}{\delta_H+\delta_L}, 0.5]$ such that for all prior beliefs such that $f(H) \in [\hat{f}, 0.5]$, $\underline{\lambda}^* = \rho_\delta$. Otherwise, for $f(H) \notin [\hat{f}, 0.5]$, the firm is better off without imposing a constraint.*

In the proof of this result, we first notice that the firm cannot benefit from constraining the Bayesian consumers. This is because the Bayesian consumers either acquire an informative signal and follow it or acquire an uninformative signal and reject the product. Under a constraint either the behavior of such consumer does not change or it switches from following the signal to not buying at all (see Table 1 and Figures 1 and 2.) That cannot benefit the firm. Hence, it is crucial to have biased consumers for the firm to set non-trivial constraint optimally, i.e. we need $p > 0$.

The proof of Proposition 1 shows next that the firm can benefit from exploiting those biased consumers who would collect informative signal due to low marginal cost and follow the signal when unconstrained, and would have bought the product independent of the signal when constrained. Such consumers are the ones who have $\lambda \leq \rho_\delta$ when a constraint in the interval of (ρ_δ, ρ_f) is implemented. We note that it is sub-optimal to set $\underline{\lambda} \in (0, \rho_\delta)$, because the constrained consumers still have a low enough marginal cost to acquire informative signals and follow them, and decreasing $\underline{\lambda}$ will simply increase the probability that h signals are realized and dominate such $\underline{\lambda}$. Similarly, it is sub-optimal to set $\underline{\lambda} > \rho_\delta$, because further increasing $\underline{\lambda}$ only decreases the chances that Bayesian individuals would buy the product. In sum, the firm is deciding between $\underline{\lambda} \in \{0, \rho_\delta\}$, and it will find it more profitable to set $\underline{\lambda} = \rho_\delta$ when the prior belief of the good being high quality is not too low (so the case in Figure 2 rather than Figure 1 occurs.) Otherwise, the exploitable biased consumers may not have the intention to buy the product in the first place.

2.5 Entry Game

We have just discussed the conditions under which the monopolistic firm finds it profitable to constrain information acquisition, which was summarized in Proposition 1. In sum, this case would occur only for $\rho_f \geq \rho_\delta$. In what follows, we will analyze how the entry of a competitor firm affects the equilibrium strategies of the incumbent firm, which would constrain the consumers when it is a monopoly, i.e., $\rho_f \geq \rho_\delta$.

Consider a simple entry game with an entrant and an incumbent firm. Assume the incumbent firm has been facing the problem described in the previous section, with $p \in (0, 1)$, and the parameters are such that the optimal strategy for a monopolist is setting $\underline{\lambda}^* = \rho_\delta$. The entrant can choose to enter the market or not. The profit of not entering is set to zero. If the entrant chooses to enter, it also needs to make a choice of constraining information or not. To differentiate the two firms, denote the incumbent's choice as $\underline{\lambda}^I$ and the entrant's choice as $\underline{\lambda}^E$. Assume that initially, a proportion of $\eta \in [0, 1]$ of the consumers randomly choose the entrant's platform, while the rest use the incumbent's. However, each consumer will try out the other firm's sales platform if they feel constrained by the current one. Consumers stick with their initially picked platform when they are constrained by both platforms. As a result, the firm with the lower $\underline{\lambda}$ gets all the consumers with $min\{\underline{\lambda}^I, \underline{\lambda}^E\} \leq \lambda \leq max\{\underline{\lambda}^I, \underline{\lambda}^E\}$, and for the remaining consumers ($\lambda > max\{\underline{\lambda}^I, \underline{\lambda}^E\}$ or $\lambda < min\{\underline{\lambda}^I, \underline{\lambda}^E\}$), the entrant gets a share of $\eta \in [0, 1]$. We proceed to find conditions that characterize the equilibrium when the firms move sequentially and the incumbent moves first. The detailed proof of the following discussion is in the appendix (proof of Proposition 2), and we will outline the general logic here.

Firstly, note that neither firm will choose the strictly dominated strategy of $\underline{\lambda} > \rho_\delta$, which was established in the previous section. The addition of another firm only makes this choice less desirable due to the other firm's ability to undercut and attract all the constrained consumers with a less restrictive constraint. Similarly, any $0 < \underline{\lambda} < \rho_\delta$ is dominated by $\underline{\lambda} = 0$, due to the reasons stated earlier, as well as the undercutting by the competitor firm. Hence, it suffices to consider $\{0, \rho_\delta\}$ as the set of potential strategies for equilibrium. We first find the best responses of the entrant given the incumbent's choice of either ρ_δ (case 1) and 0 (case 2.) Afterward, we will characterize the optimal strategy of the incumbent given the best responses of the entrant.

Case 1: $\underline{\lambda}^I = \rho_\delta$

The entrant has two responses: 0 or ρ_δ. The expected profits of the entrant, $\Pi^E(\underline{\lambda}^I, \underline{\lambda}^E)$ for each strategy are:

$$\Pi^E(\rho_\delta, 0) = \int_0^{\rho_\delta} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) + \eta \left[p \int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p) \int_{\rho_\delta}^{\rho_f} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) \right]$$

$$\Pi^E(\rho_\delta, \rho_\delta) = \eta p \int_0^{\infty} dG(\lambda) + \eta(1-p) \left[\int_0^{\rho_\delta} \frac{1-f-fe^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} dG(\lambda) + \int_{\rho_\delta}^{\rho_f} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) \right]$$

The first equation is the expected profit of the entrant when it doesn't constrain the consumers, while the incumbent sets $\underline{\lambda}^I = \rho_\delta$. The first term is the expected probability of the entrant selling the product to the consumers who are constrained by the incumbent and, thus, come to acquire information and purchase the product from the entrant. The second term describes the consumers who are indifferent between the two firms because they are not constrained by either. A proportion of η of those consumers go to the entrant. p of them are biased (first term in the bracket) and always buy the product, and $1-p$ of them are Bayesian (second term in the bracket) and follow their signals. One may refer to Table 1 for these purchasing decisions.

The second equation is the expected profit of the entrant when it mimics the incumbent's strategy and sets $\underline{\lambda}^E = \rho_\delta$. The first term indicates that all the biased consumers will buy the product regardless of signal due to the firm's constraint. Again, only η share of them go to the entrant. The second term corresponds to the Bayesian consumers who go to the entrant and follow their signals. The first integral in the bracket is the profit generated by the constrained Bayesian consumers and hence, their λ ranges from zero to ρ_δ. The second integral in the bracket is the profit generated by the unconstrained Bayesian consumers and hence, their λ ranges from ρ_δ to ρ_f. Any Bayesian consumer with higher λ will pick an uninformative signal structure and reject the product. The entrant will find constraining information acquisition not profitable and hence, provide better information to the consumers as long as Inequality (10) is satisfied.

$$\Pi^E(\rho_\delta, 0) > \Pi^E(\rho_\delta, \rho_\delta) \iff \int_0^{\rho_\delta} \left[-\eta p + \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} - \eta(1-p)\frac{1-f-fe^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} \right] dG(\lambda) > 0 \quad (10)$$

Denote the set of values for η and p that violate condition (10) as $C = \{(\eta, p)|\Pi^E(\rho_\delta, 0) = \Pi^E(\rho_\delta, \rho_\delta)\}$. The dashed curve in Figure 3 is the boundary of C.[18] Since the left hand side of Inequality (10) is linear in η and p, for all (η, p) such that $(\eta, p) \geq (\hat{\eta}, \hat{p})$ for some $(\hat{\eta}, \hat{p}) \in C$, the entrant finds it optimal to also constrain information acquisition. For other (η, p) values, the entrant finds it optimal to not constrain information acquisition. Realistically, entrants would have a fairly low share η, and thus, it is possible that the inequality would hold for all positive p.

Case 2: $\underline{\lambda}^I = 0$

This case analyzes the entrant's strategy if the incumbent imposes no constraint to preempt a competition through information. First, we calculate the expected profits of the entrant, $\Pi^E(\underline{\lambda}^I, \underline{\lambda}^E)$ when it sets no constraint (i.e., $\underline{\lambda}^E = 0$) and when it sets $\underline{\lambda}^E = \rho_\delta$, given $\underline{\lambda}^I = 0$.

$$\Pi^E(0,0) = \eta \left[\int_0^{\rho_\delta} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) + p\int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p)\int_{\rho_\delta}^{\rho_f} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) \right]$$

[18] In the proof of Proposition 2, we will show that the boundary of set C can be expressed as a decreasing function $\hat{p}(\eta)$ and has the form illustrated in the figure.

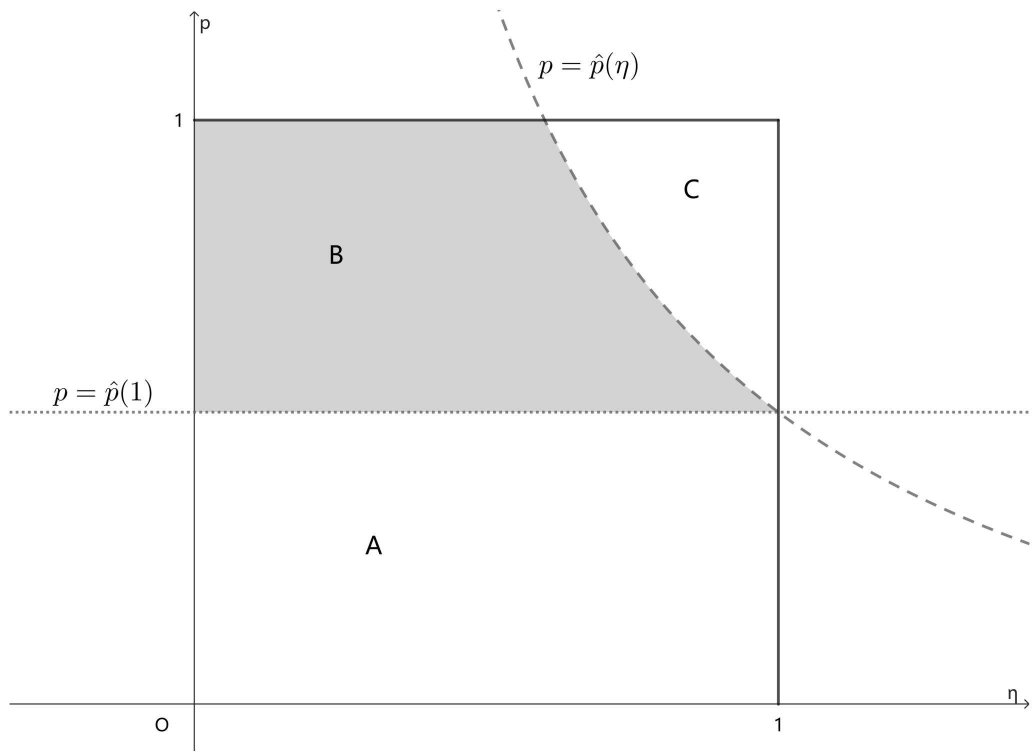

Figure 3: Region A: No constraint equilibrium under monopoly and competition. Region B: Equilibrium with constraint under monopoly but equilibrium without constraint under competition. Region C: Equilibrium with constraint in both monopoly and competition.

$$\Pi^E(0,\rho_\delta) = \eta\left[p\int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p)\int_{\rho_\delta}^{\rho_f} \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda)\right]$$

The first equation is the expected profit of the entrant when neither firm constrains information. The first term describes the consumers who follow their signals, whether they are biased or Bayesian. The second term describes the biased consumers who always buy (i.e., biased consumers with large λ, see Table 1.) The third term describes Bayesian consumers who follow their signal. Since the consumers are indifferent between the two firms, we multiply the whole paranthesis by η to calculate the entrant's share in the market. The second equation is the expected profit of the entrant when it constrains information acquisition by setting $\underline{\lambda}^E = \rho_\delta$ while the incumbent doesn't constrain. The first term represents the biased consumers who always buy and aren't constrained. The second term describes the Bayesian consumers who aren't constrained and follow their signals.

It is always the case that $\Pi^E(0,0) > \Pi^E(0,\rho_\delta)$. Hence, the entrant will not constrain information when the incumbent is not constraining. In sum, not constraining information is a dominant strategy for the entrant when neither η nor p are too high, with the thresholds as mentioned earlier.

Next, we will check the incumbent's expected profits to determine whether the incumbent will constrain information or not, anticipating the entrant's best responses above. By symmetry,

$$\Pi^I(0,0) = \frac{(1-\eta)}{\eta}\Pi^E(0,0) > \frac{(1-\eta)}{\eta}\Pi^E(0,\rho_\delta) = \Pi^I(\rho_\delta, 0)$$

Since $\Pi^I(0,0) > \Pi^I(\rho_\delta,0)$, the incumbent will not constrain information acquisition in anticipation of the entrant not constraining.

If the incumbent anticipates the entrant to constrain (recall that Case 1 characterized that strategy for large η and p), then the entrant's best response to $\underline{\lambda}^I = \rho_\delta$ is also $\underline{\lambda}^E = \rho_\delta$. Hence, we have either both not constraining, or both constraining at ρ_δ. The incumbent will determine the equilibrium by moving first.

The incumbent essentially chooses which symmetric strategy to be played by checking Inequality (11).

$$\Pi^I(0,0) > \Pi^I(\rho_\delta,\rho_\delta) \iff \int_0^{\rho_\delta}\left[\frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} - (1-p)\frac{1-f-fe^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} - p\right] dG(\lambda) > 0 \quad (11)$$

Note that this inequality is very similar to (10), but the left-hand side of (11) does not have η. If (11) holds, then (10) holds. Therefore, if $\underline{\lambda}^E = \rho_\delta$ is a best response to $\underline{\lambda}^I = \rho_\delta$ (which requires (10) to fail), $(\rho_\delta,\rho_\delta)$ is the equilibrium. Recall that this case occurred if there are enough exploitable biased consumers in the market. Proposition 2 summarizes this observation.

Let $p > 0$, and $\hat{p}$ and $\hat{\eta}$ be the thresholds described to make inequality (10) satisfied. These give us full characterization of the equilibrium of this sequential game.

Proposition 2. *Let $p > 0$ and $\eta \in (0,1)$. There exists $\hat{p}(\eta)$, decreasing in η such that the equilibrium of the sequential game is:*

$$(\underline{\lambda}^I, \underline{\lambda}^E) = \begin{cases} (0,0) & \text{if } p < \hat{p}(1) \\ (0,0) & \text{if } \hat{p}(1) \leq p < \hat{p}(\eta) \\ (\rho_\delta, \rho_\delta) & \text{if } \hat{p}(\eta) \leq p \end{cases}$$

Moreover, when $p < \hat{p}(1)$ a monopolist would not constrain information but when $p \geq \hat{p}(1)$, a monopolist would constrain information.

Proposition 2 identifies the set of parameters for which competition would or would not help consumers collect unconstrained information. It depends on whether there are enough biased people, and whether the entrant commands a large enough market share upon entry. The former guarantees that the benefit from exploiting the biased consumers is large, and the latter ensures that the entrant has no incentive to steal consumers by providing better information. For both conditions to be satisfied, the monopolist should have the incentive to constrain in the first place, i.e. $f \geq \hat{f}$. In all other cases, neither firm will constrain information acquisition. Figure 3 illustrates the equilibrium characterized by Proposition 2. With the two axes representing the values of η and p, the dotted line and the dashed curve divide the graph into three regions. Region A denotes the parameter values for which there will not be information constraint either by the monopolist or under the competition. Region B denotes the parameter values where competition will alleviate information exploitation. Region C denotes the region where there will be information constraints under both monopoly and competition. We can observe from the graph that competition can only be beneficial in alleviating information exploitation.

3 Other Non-Bayesian Updating Rules

In this section, we show that our earlier results generalize to some other non-Bayesian updating rules beyond wishful thinking described in (5). Denote a general updating rule:

$$g(\omega|s) = F_\omega(f, s)$$

where g is the posterior belief on state w after observing signal s given a prior distribution f. Properties 1-3 below are sufficient conditions for Proposition 1 to generalize and a monopolist to optimally constrain information acquisition for some parameters. Recall that the key idea behind Lemma 1 and Proposition 1 is that there needs to be a non-empty set of priors and λ values such that the biased posterior belief after receiving an l signal favors state H, while at the same time, the expected utility of buying after receiving an h signal is maintained positive. The properties below suffice to guarantee this idea.

Property 1. *$F_\omega(f, s)$ is continuous in both the prior probabilities $f(\omega')$ and the conditional probabilities $\pi(s|\omega')$ for all ω'.*

Property 2. *For any f and s, $F_H(f, s) \geq g^B(H|s)$. Furthermore, if the Bayesian posterior $g^B(\omega|s) \in (0, 1)$, then $F_H(f, s) > g^B(H|s)$.*

Property 3. *For any f and s, the posterior belief ratio $\frac{F_\omega(f,s)}{F_{\omega'}(f,s)}$ is non-decreasing in $\frac{\pi(s|\omega)}{\pi(s|\omega')}$.*

These are fairly loose restrictions. Property 1 is a continuity requirement for a well-behaved updating rule. The second property requires that unless the Bayesian individual is certain the product is of a certain quality, the biased individual will always be strictly more optimistic than a Bayesian person. The third property is also intuitive, as it means the biased individual understands the informational content of signals so that changing from "good news" to "better news" would not decrease the individual's optimism. For an updating function with these properties, we can show that, for $p > 0$, there exists a non-empty set of priors such that the monopolistic firm finds it profitable to constrain information acquisition. This is analogous to Proposition 1.

Proposition 3. *Let $p > 0$ and the updating rule of the biased consumers satisfy Properties 1-3, then there exists a non-empty set of model parameters and priors such that a monopolist optimally picks $\kappa > 0$, i.e. constrains consumers' information structure.*

It is straightforward to check that the wishful thinking model in Equation (5) satisfies these properties. Other updating rules that may satisfy the above properties under certain specifications include the following.

1. Affine Transformation: Define an updating rule that distorts the Bayesian posterior with an affine shift:

$$g^A(\omega|s) = \alpha g^B(\omega|s) + (1-\alpha)\bar{g}(\omega)$$

where $\alpha \in (0, 1)$ is the weight and $\bar{g}$ is an exogenous probability distribution. If, for instance, $\bar{g}(H) = 1$, the updating rule would satisfy all the properties.[19]

2. Distorted Prior: In this updating rule, the individual employs the Bayes rule on a distorted prior (as in Liu, 2023a and Dominiak et al., 2023).

$$g^P(\omega|s) = g^B(\omega|s, \chi(f)) = \frac{\chi(f(\omega))\pi(s|\omega)}{\sum \chi(f(\omega'))\pi(s|\omega')}$$

where χ is a distortion function. If the distortion function effectively shifts the prior towards state H, that is if $\frac{\chi(f(H))}{\chi(f(\omega))} > \frac{f(H)}{f(\omega)}$ for all $\omega \neq H$, the resulting Bayesian updating upon the distorted prior would satisfy the above properties.

3. Psychological Utility Maximization: If the individual directly gains utility from certain beliefs (as in Lipnowski and Mathevet, 2018) and chooses beliefs that maximize such utility, then this can distort beliefs in directions dependent on which beliefs are preferred. As discussed in de Clippel and Zhang (2022), this can accommodate a motivated belief updater who has a target belief $\bar{g}$ and suffers a cost from deviating away from both

[19] Any anchor would work as long as the anchoring belief $\bar{g}(H)$ is greater than the maximal $g^B(H|s)$. When we limit our discussion to a certain range of $f(H)$, it may be impossible to obtain a Bayesian posterior $g^B(H|s)$ close to 1 given any signal. In such cases, there would be a wide range of eligible values for $\bar{g}(H)$, generalizing this example to more applications.

$\bar{g}$ and the Bayesian posterior g^B.

$$g^U = argmax_g U(g, \bar{g}, g^B)$$

If $\bar{g}(H) = 1$, then similar to the first example, the updating rule would satisfy the above properties. In general, having $\bar{g}(H) > g^B(H)$ would suffice.

4 Discussions and Conclusion

In this section, we will revisit some of our modeling assumptions and discuss alternative ways of modeling. In our analysis, we implicitly assumed that firms do not have perfect knowledge of the products' quality. This makes sense for the main scenario we are considering, where the firm is a digital platform that hosts many products potentially from many different upstream sellers. Platforms like Amazon typically do not, and cannot, monitor the quality of every product sold on their site. However, it is still worthwhile to mention the possibility that the firm knows the product quality, as is usually the case when the seller is also the producer. This is a classic setting for signalling games and has been discussed extensively in the literature. The equilibrium typically would involve the firm's actions inevitably revealing information regarding their product quality. A similar signaling game is, for example, studied by Martin (2017). It is an open question to study rationally inattentive biased consumers in a signaling game where the firm can directly constrain the informativeness of the signals through κ.

Another variation that is worth mentioning is the possibility that each consumer may have a different state realization. Some may be unlucky and end up with a faulty product, while others are lucky enough to obtain a perfectly manufactured one. Assuming that the firms do not control the quality of each product the consumers buy and have to offer the same set of information structures to everyone, they can then only maximize expected profits over the unconditional probabilities. Then this exercise becomes the same as our model. Alternatively, if the firms could control, or have knowledge of, the quality of each product that each consumer gets and hence, personalize the information constraint, the problem coincides with the signalling game mentioned above.

One may also consider heterogeneity in other aspects of the model. In our setting, the only sources of heterogeneity are λ (the cost parameter) and being biased or not. When individuals all have different priors, we again have two possible scenarios. If the firm knows each consumer's prior and is capable of customizing its constraint for each individual, the firm's problem is reduced to the base problem we have discussed, but with one person. Of course, the firm has a much higher profit in this case, because it can choose to only apply the information constraints to biased consumers. However, if the firm is not capable of setting individual constraints, it will simply maximize expected profit over the distribution of priors, regardless of whether they know the consumers' idiosyncratic priors or not. Given this, the threshold for which the firm finds it profitable to constrain information acquisition will become a condition regarding the distribution of priors. It is easy to conceive that the more individuals holding higher priors, the more profitable such information exploitation is. Assuming that biased individuals have different magnitudes of bias (different δ_H and δ_L) will produce an analogous result, where the firm either faces many instances of the single consumer problem if they can apply individual constraints and knows each individual's bias, or it maximizes expected profit over the distribution of δ otherwise.

Alternatively, one may extend the discussion to continuous states (quality of the product). The optimal information structure (Lemma 1), given that the action set does not change, would retain the same form (the solution in the appendix also applies to continuous states, except with integrals instead of summations over the states). For continuous states, it may be difficult to obtain a closed form solution to both the marginal probabilities of signals q and the conditional probabilities π, unless one is willing to assume certain distributions for the prior belief. This is a standard complexity observed previously by the rational inattention literature. However, we expect that the intuition behind the results would not change, and there would still be a condition on the distribution of prior belief such that, for a prior distribution skewed enough towards higher product qualities, the firm would find it profitable to exploit information acquisition.

Finally, while we assumed $f(H) \leq 0.5$ in our discussion, it is certainly feasible to assume otherwise. However, having $f(H) > 0.5$ would ultimately lead to a trivial discussion, because the consumers lean towards buying the product by default. In this case, the firm can simply set $\kappa = 0$ (or equivalently $\underline{\lambda} \to \infty$), that is completely shutting down the customer review section and not allowing anyone to acquire information before purchasing. Assuming this in itself does not convey malicious intentions, all consumers would buy the product. In fact, it is beneficial for the firm to do this even when all consumers are Bayesian, since information serves to warn the consumers of potential low quality products now (as opposed to signal high quality products), and stripping away the consumers' ability to acquire meaningful information is always desirable for the firm.

In sum, we provide a model where the firms may find it profitable to constrain the quality of information they provide about the product they sell. We characterize when competition among firms cannot mitigate the vagueness of information. Our results rely on the existence of biased consumers in the market and they generalize to some other non-Bayesian updating rules beyond wishful thinking. We also list applicable open questions for more complex or alternative modeling assumptions. To the best of our knowledge, this is the first model introducing competition through quality of information which has fruitful applications for modern business models.

A Proofs

Proof of Lemma 1: Solution to the Constrained Problem

We first note that Lemma 1 and Corollary 1 of Matějka and McKay (2015) still holds with the information constraint present. Their Lemma 1 shows that having more signals than actions while keeping the consumer's expected payoff the same will only increase information cost. Their Corollary 1 follows naturally by stating that, since the signals and actions are one-to-one, then the consumer is essentially choosing the joint probability distribution of states and actions. The optimality of the simplest signal structures, is not affected by the firm's extra constraint that we have in our model. Without loss of generality, assume that consumers would expect themselves to buy if they receive an h signals, and reject if they receive an l signal. Denote this correspondence as a_s where $a_h = b$ and $a_l = r$. The following claim provides a general solution for optimal signal structure for the constrained optimization problem. After proving the claim we will proceed with the closed form formula of π^* and q^* specific to our setup and derive Equations (12) and (4).

Claim 1: A consumer with cost parameter λ chooses optimal signal structure

$$\pi^*(s|\omega) = \begin{cases} \pi^{\lambda}(s|\omega) & \lambda \in [\underline{\lambda}, \infty) \\ \pi^{\underline{\lambda}}(s|\omega) & \lambda \in [0, \underline{\lambda}) \end{cases} \tag{12}$$

where

$$\pi^{\lambda}(s|\omega) = \frac{q(s)e^{u(a_s|\omega)/\lambda}}{\sum_s q(s')e^{u(a_{s'}|\omega)/\lambda}}$$

Proof of Claim 1: The following is the ex-ante expected utility of choosing information structure π when the consumer expects to follow the signal either because she is a Bayesian or a naive non-Bayesian:

$$\sum_{\omega}\sum_{s} V(s,\pi)\pi(s|\omega)f(\omega) = \sum_{s} V(s,\pi)\sum_{\omega}\pi(s|\omega)f(\omega) = \sum_{s} V(s,\pi)q(s)$$
$$= \sum_{s}\sum_{\omega} g^B(\omega|s)u(a_s|\omega)q(s) = \sum_{s}\sum_{\omega} u(a_s|\omega)\pi(s|\omega)f(\omega)$$

The first equality holds because $V(s,\pi)$ is based on the expectation of the state of the world. The second equation holds by the definition of $q(s)$. The third one follows from the definitions of v and a_s. Finally, the last equation is an implication of the Bayes rule. We first consider the case where $\lambda > 0$. Adding in the firm's constraint, as well as the default constraints that $\pi(s|\omega)$ needs to be non-negative and that they add to 1 across signals, we have the Lagrangian for the consumer's problem:

$$L = \sum_{\omega}\sum_{s} u(a_s|\omega)\pi(s|\omega)f(\omega) - \lambda I(\pi) + \zeta(\kappa - I(\pi)) + \sum_{s}\sum_{\omega}\gamma_{s,\omega}\pi(s|\omega)$$
$$+ \sum_{\omega}\mu_{\omega}[1 - \sum_{s}\pi(s|\omega)]$$

where ζ, $\gamma_{s,\omega}$ and μ_ω are the Lagrangian multipliers. The first order condition with respect to $\pi(s|\omega)$ is:

$$u(a_s|\omega)f(\omega) - (\lambda+\zeta)[ln(\pi(s|\omega) - ln(q(s))]f(\omega) + \gamma_{s,\omega} - \mu_\omega = 0$$

If $q(s) > 0$, and since $u(a_s|\omega) > -\infty$, we must have $\pi(s|\omega) > 0$ for all ω. Otherwise, if $\pi(s|\omega') = 0$ for some ω', $ln(\pi(s|\omega') = -\infty$, causing $\mu_{\omega'} = \infty$. This implies that, for all $s' \neq s$, $\pi(s'|\omega') = 0$. But this contradicts with $\sum_s \pi(s|\omega) = 1$ for all ω.

Therefore, $\gamma_{s,\omega} = 0$. By solving the first order condition above, we get:

$$\begin{aligned}
\frac{\mu_\omega}{f(\omega)} &= u(a_s|\omega) - (\lambda+\zeta)[ln(\pi(s|\omega) - ln(q(s))] \\
&\Rightarrow e^{[u(a_s|\omega) - \frac{\mu_\omega}{f(\omega)}]/(\lambda+\zeta)} = \frac{\pi(s|\omega)}{q(s)} \\
&\Rightarrow \sum_s e^{[u(a_s|\omega) - \frac{\mu_\omega}{f(\omega)}]/(\lambda+\zeta)} q(s) = 1 \\
&\Rightarrow \sum_s e^{u(a_s|\omega)/(\lambda+\zeta)} q(s) = e^{[\frac{\mu_\omega}{f(\omega)}]/(\lambda+\zeta)} \\
&\Rightarrow \pi(s|\omega) = \frac{q(s)e^{u(a_s|\omega)/(\lambda+\zeta)}}{\sum_s q(s')e^{u(a_{s'}|\omega)/(\lambda+\zeta)}}
\end{aligned}$$

The last equation is derived from plugging the previous observation into the second line.

If $q(s) = 0$, then we simply have $\pi(s|\omega) = 0$ for all ω. If $\lambda = 0$, the problem is trivial, where the consumers want to gather as much information as possible. Since they are constrained by the firm, they acquire information up to the limit $I(\pi) = \kappa$.

To obtain the value of ζ, we need to note that, since it is the Lagrangian multiplier for the firm's constraint, it is only positive for consumers when it is a binding constraint. Furthermore, it acts purely as an additive amount to the marginal cost of information λ, and the consumers behave as if they have a higher marginal cost. Let π^λ denote the optimal π given λ. The constraint binds when consumers' unconstrained optimization would lead to an information structure such that $I(\pi^\lambda) > \kappa$, which happens when consumers have naturally low λ. This is because $I(\pi^\lambda)$ is decreasing in λ. To see this, assume that the optimal signal structures (without the firm's constraint) are such that, for some $\lambda > \lambda'$, $I(\pi^\lambda) > I(\pi^{\lambda'})$. Denoting the optimal expected payoffs as $EV(\pi^\lambda)$ and $EV(\pi^{\lambda'})$ respectively, we have the following due to their optimality:

$$EV(\pi^\lambda) - \lambda I(\pi^\lambda) \geq EV(\pi^{\lambda'}) - \lambda I(\pi^{\lambda'})$$

$$EV(\pi^{\lambda'}) - \lambda' I(\pi^{\lambda'}) \geq EV(\pi^\lambda) - \lambda' I(\pi^\lambda)$$

And thus

$$\lambda'[I(\pi^\lambda) - I(\pi^{\lambda'})] \geq EV(\pi^\lambda) - EV(\pi^{\lambda'}) \geq \lambda[I(\pi^\lambda) - I(\pi^{\lambda'})]$$

which is impossible as $\lambda > \lambda'$.

Define $\underline{\lambda}$ such that, without the firm's constraint, the optimal signal structure is such that $I(\pi^{\underline{\lambda}}) = \kappa$. Then, for the constrained consumers, it must be that $I(\pi^{\lambda+\zeta}) = \kappa$, which would give $\lambda + \zeta = \underline{\lambda}$ due to the monotonicity of $I(\pi)$ in λ.

Lastly, the case of $\lambda = 0$ becomes trivial, as the above solution would still hold.

In sum, the consumer who has marginal cost λ acquires the following optimal signal structure:

$$\pi^*(s|\omega) = \begin{cases} \pi^\lambda(s|\omega) & \lambda \in [\underline{\lambda}, \infty) \\ \pi^{\underline{\lambda}}(s|\omega) & \lambda \in [0, \underline{\lambda}) \end{cases}$$

where

$$\pi^{\lambda}(s|\omega) = \frac{q(s)e^{u(a_s|\omega)/\lambda}}{\sum_s q(s')e^{u(a_{s'}|\omega)/\lambda}}$$

This proves Claim 1.

Next, we solve for the closed form solutions of q and π within the setting of our model. Note that the following discussion will only include those with $\lambda \in [\underline{\lambda}, \infty)$ for simplicity, but the case where $\lambda \in [0, \underline{\lambda})$ carries over by replacing λ with $\underline{\lambda}$. Plugging in the signals $s \in \{h, l\}$, and the correspondence $a_h = b$, $a_l = r$ (recall that $u(r|\omega) = 0$), we obtain

$$\pi^*(s = l|\omega) = \frac{q(l)}{q(l) + q(h)e^{u(b|\omega)/\lambda}} \quad \text{and} \quad \pi^*(s = h|\omega) = \frac{q(h)e^{u(b|\omega)/\lambda}}{q(l) + q(h)e^{u(b|\omega)/\lambda}}$$

We plug this into $q(s) = \sum_{\omega} \pi(s|\omega)f(\omega)$ and obtain the following:

$$\frac{q(h)e^{1/\lambda}}{q(l) + q(h)e^{1/\lambda}} f(H) + \frac{q(h)e^{-1/\lambda}}{q(l) + q(h)e^{-1/\lambda}} f(L) = q(h)$$

Denoting for simplicity $q = q(h) = 1 - q(l)$, $f = f(H) = 1 - f(L)$, and assuming $1 > q > 0$, we get

$$\frac{e^{1/\lambda} f}{1 - q + qe^{1/\lambda}} + \frac{f}{1 - qe^{1/\lambda} + q} = 1$$

$$\Rightarrow e^{2/\lambda} f + e^{1/\lambda} qf - e^{2/\lambda} qf + 1 - f - q + qf + e^{1/\lambda} q - e^{1/\lambda} qf$$
$$= e^{1/\lambda} - 2e^{1/\lambda} q + e^{1/\lambda} q^2 + e^{2/\lambda} q - e^{2/\lambda} q^2 + q - q^2 + e^{1/\lambda} q^2$$

$$\Rightarrow [1 - f - e^{1/\lambda} + e^{2/\lambda} f] - q[1 - f - e^{1/\lambda} + e^{2/\lambda} f + 1 + e^{2/\lambda} - 2e^{1/\lambda}] + q^2[1 + e^{2/\lambda} - 2e^{1/\lambda}] = 0$$

$$\Rightarrow (1 - e^{1/\lambda})(1 - f - e^{1/\lambda} f)(1 - q) + (q^2 - q)(1 - e^{1/\lambda})^2 = 0$$

$$\Rightarrow q = \frac{1 - f - e^{1/\lambda} f}{1 - e^{1/\lambda}}$$

Hence,

$$q(h) = \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}} \quad \text{and} \quad q(l) = \frac{f(H) - f(L)e^{1/\lambda}}{1 - e^{1/\lambda}}$$

The optimal $q(h)$ derived from this can actually be greater than 1 or less than 0, which can happen when $f(H)$ is too large or too small, respectively. (Under the assumption that $f(L) \geq f(H)$, having larger λ will lower $q(h)$ as well.) Since it is a probability, it must satisfy $1 > q(h) > 0$. Otherwise, we have a corner solution.

Plugging the interior solutions for q into the conditional probabilities above, the denominators of $\pi^*(s|H)$ and $\pi^*(s|L)$ are

$$q(l)+q(h)e^{1/\lambda} = 1-q(h)+q(h)e^{1/\lambda} = 1+q(h)(e^{1/\lambda}-1) = 1+f(H)e^{1/\lambda}-f(L) = f(H)(1+e^{1/\lambda})$$

$$q(l) + q(h)e^{-1/\lambda} = f(L)(1 + e^{1/\lambda}),$$

respectively. And we get:

$$\pi^*(h|H) = \frac{q(h)e^{1/\lambda}}{f(H)(1+e^{1/\lambda})} = 1 - \pi^*(l|H) = 1 - \frac{q(l)}{f(H)(1+e^{1/\lambda})}$$

$$\pi^*(l|L) = \frac{q(l)e^{1/\lambda}}{f(L)(1+e^{1/\lambda})} = 1 - \pi^*(h|L) = 1 - \frac{q(h)}{f(L)(1+e^{1/\lambda})}$$

for $q(h) \in (0, 1)$.

Additionally, when $q(h) = 0$, $\pi^*(h|\omega) = 0 = 1 - \pi^*(l|\omega)$

And when $q(h) = 1$, $\pi^*(h|\omega) = 1 = 1 - \pi^*(l|\omega)$

In sum, the optimal marginal signal probabilities are (and including the case for constrained consumers this time)

$$q^*(h) = \begin{cases} max\left\{0, min\left\{1, \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}}\right\}\right\} & \lambda \in [\underline{\lambda}, \infty) \\ max\left\{0, min\left\{1, \frac{f(L)-f(H)e^{1/\underline{\lambda}}}{1-e^{1/\underline{\lambda}}}\right\}\right\} & \lambda \in [0, \underline{\lambda}) \end{cases}$$

and $q^*(l) = 1 - q^*(h)$.

The optimal information structure is

$$\pi^*(h|H) = \begin{cases} 0 & q^*(h) = 0 \\ \frac{q^*(h)e^{1/\lambda}}{f(H)(1+e^{1/\lambda})} & \lambda \in [\underline{\lambda}, \infty), q^*(h) \in (0,1) \\ \frac{q^*(h)e^{1/\underline{\lambda}}}{f(H)(1+e^{1/\underline{\lambda}})} & \lambda \in [0, \underline{\lambda}), q^*(h) \in (0,1) \\ 1 & q^*(h) = 1 \end{cases}$$

$$\pi^*(l|L) = \begin{cases} 0 & q^*(l) = 0 \\ \frac{q^*(l)e^{1/\lambda}}{f(L)(1+e^{1/\lambda})} & \lambda \in [\underline{\lambda}, \infty), q^*(l) \in (0,1) \\ \frac{q^*(l)e^{1/\underline{\lambda}}}{f(L)(1+e^{1/\underline{\lambda}})} & \lambda \in [0, \underline{\lambda}), q^*(l) \in (0,1) \\ 1 & q^*(l) = 1 \end{cases}$$

with $\pi^*(h|\omega) = 1 - \pi^*(l|\omega)$.

□

Proof of Lemma 2: Optimal Purchasing Decision

Again, we mainly show the calculations for the unconstrained consumers. The calculations for the constrained consumers are analogous with $\underline{\lambda}$ instead of λ. First, we consider the case where $q(h) \in (0,1)$.

By the updating rule (5),

$$g^\delta(H|h) = \frac{f(H)\pi^*(h|H)\delta_H}{f(H)\pi^*(h|H)\delta_H + f(L)\pi^*(h|L)\delta_L} = \frac{e^{1/\lambda}\delta_H}{e^{1/\lambda}\delta_H + \delta_L} = 1 - g^\delta(L|h)$$

$$g^\delta(L|l) = \frac{f(L)\pi^*(l|L)\delta_L}{f(H)\pi^*(l|H)\delta_H + f(L)\pi^*(l|L)\delta_L} = \frac{e^{1/\lambda}\delta_L}{\delta_H + e^{1/\lambda}\delta_L} = 1 - g^\delta(H|l)$$

The expected utilities for buying the product given an h signal or an l signal are:

$$Eu^\lambda(s = h) = g^\delta(H|h) - g^\delta(L|h) = \frac{e^{1/\lambda}\delta_H - \delta_L}{e^{1/\lambda}\delta_H + \delta_L}$$

$$Eu^\lambda(s = l) = g^\delta(H|l) - g^\delta(L|l) = \frac{\delta_H - e^{1/\lambda}\delta_L}{e^{1/\lambda}\delta_L + \delta_H}$$

Note that since $\delta_H \geq \delta_L$, we always have $Eu^\lambda(s = h) > 0$. Hence observing an h signal always leads to buying the product, for both biased and Bayesian consumers. When observing an l signal, it would depend on whether

$$\delta_H - e^{1/\lambda}\delta_L \geq 0$$

For Bayesian consumers ($\delta_H = \delta_L$), this will always be negative, and they always reject if they see an l signal. However, it is possible for the biased consumers to ignore the bad signal and buy regardless, as long as δ_H is large enough compared to δ_L, or when λ is large enough.

To transform this condition as a threshold for λ, let

$$\rho_\delta = \frac{1}{ln(\delta_H) - ln(\delta_L)}$$

And the condition becomes

$$\lambda \geq \rho_\delta$$

Adding in the possibility of the firm's constraint, effectively setting a lower bound for the consumers' marginal cost to $\underline{\lambda}$, we can state that the biased consumer with marginal cost will ignore bad signals and buy regardless if

$$max\{\lambda, \underline{\lambda}\} \geq \rho_\delta$$

Contrarily, the Bayesian consumer will always be consistent and follow their signals.

Now we discuss the case where the optimal $q(h)$ is a corner solution. It is important to note that a corner solution is not possible when $f(H) = f(L) = 0.5$, as that would put $q(h) = q(l) = 0.5$ as well. Also note that given our assumption that $f(L) \geq f(H)$, it is actually impossible for $q(h) = 1$ and $q(l) = 0$, as this would require

$$f(L)e^{1/\lambda} < f(H)$$

And $q(h) = 1$ is only possible when $f(L) < f(H)$. However, regardless of whether $q(h) = 1$

or $q(h) = 0$, the following results don't change.

When $q(h) = 0$, that is when $f(L) - f(H)e^{1/\lambda} \leq 0$ the consumer chooses an uninformative signal structure with only l signals. First, we define

$$\rho_f = \frac{1}{ln(f(L)) - ln(f(H))}$$

And consumers will acquire uninformative signals when

$$f(L) - f(H)e^{1/\lambda} \leq 0 \iff \lambda \geq \rho_f$$

While it doesn't make a difference for Bayesian individuals if they actually acquire a signal or not, as it would be uninformative, it does matter for the biased individual, as the posteriors are different from the prior due to the bias:

$$g^\delta(L|l) = \frac{f(L)\pi^*(l|L)\delta_L}{f(H)\pi^*(l|H)\delta_H + f(L)\pi^*(l|L)\delta_L} = \frac{f(L)\delta_L}{f(H)\delta_H + f(L)\delta_L} = 1 - g^\delta(H|l)$$

The expected utility for buying is then:

$$Eu^\lambda(s = l) = g^\delta(H|l) - g^\delta(L|l) = \frac{f(H)\delta_H - f(L)\delta_L}{f(H)\delta_H + f(L)\delta_L}$$

This is positive if and only if

$$f(H)\delta_H - f(L)\delta_L \geq 0 \iff \rho_f \geq \rho_\delta$$

Note that for the Bayesian individual who has $\delta_H = \delta_L$, this would only be possible for the specific scenario where $f(H) = f(L) = 0.5$. However, that possibility has already been ruled out.

Hence, consumers will acquire uninformative signals when $max\{\lambda, \underline{\lambda}\} \geq \rho_f$. In that case, Bayesian individuals will not buy the product, and biased individuals will buy if and only if $\rho_f \geq \rho_\delta$. □

Proof of Proposition 1: Optimal constraint of the firm.

Choosing optimal κ is equivalent to choosing optimal $\underline{\lambda}$ since $I(\pi^{\underline{\lambda}}) = \kappa$ has a unique solution for each κ. To analyze the firm's optimal choice of $\underline{\lambda}$, we will divide the discussion into two cases depending on the relation between ρ_f and ρ_δ. The firm's profit is calculated as the expected sales, which is the unconditional probability of selling the item summed over all consumers.

Case 1: $\rho_f < \rho_\delta$

Under this assumption, we have $f(H)\delta_H < f(L)\delta_L$. Figure 1 illustrates this case. Furthermore, for any individual that acquires informative signals, that is with $\lambda < \rho_f$, we automatically have $\lambda < \rho_\delta$. And whenever $\lambda \geq \rho_f$, the consumer will acquire uninformative signals. Combined with $\rho_f < \rho_\delta$, the consumer will never buy. Therefore, the firm has two potential strategies:

Strategy 1.1: $0 < \underline{\lambda} < \rho_f$

By setting $0 < \underline{\lambda} < \rho_f$, the constrained consumers, with $\lambda < \underline{\lambda}$, will act as if they have a marginal cost of $\underline{\lambda}$. By the results from Lemmas 1 and 2, and since $\underline{\lambda} < \rho_f < \rho_\delta$, these constrained consumers, whether biased or not, still acquire informative signals and buy the product only when they see an h signal.

We take the difference in expected sales between the firm setting $0 < \underline{\lambda} < \rho_f$ and imposing no constraint:

$$\begin{aligned} E[\Delta\Pi](\underline{\lambda} < \rho_f | \rho_\delta > \rho_f) &= \int_{\lambda < \underline{\lambda}} (q^{\underline{\lambda}}(h) - q^{\lambda}(h)) dG(\lambda) \\ &= \int_{\lambda < \underline{\lambda}} \left(\frac{f(L) - f(H)e^{1/\underline{\lambda}}}{1 - e^{1/\underline{\lambda}}} - \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}} \right) dG(\lambda) < 0 \end{aligned}$$

The change only comes from people who are constrained by the $\underline{\lambda}$, hence the integral is only calculated for those $\lambda \leq \underline{\lambda}$. The first term is the marginal probability of receiving $s = h$ with the constraint, which all of the consumers will follow. The second term is the marginal probability of $s = h$ without the constraint. The difference is negative as $q(h) = \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}}$ decreases in λ.

Strategy 1.2: $\underline{\lambda} \geq \rho_f$

As illustrated in Figure 1, pushing consumers' marginal cost higher than ρ_f will cause them to obtain uninformative signals and not buy the product. This is reflected in a negative change in profit:

$$E[\Delta\Pi](\underline{\lambda} \geq \rho_f | \rho_\delta > \rho_f) = \int_{\lambda<\rho_f} (q^{\underline{\lambda}}(h) - q^{\lambda}(h))dG(\lambda)$$
$$= \int_{\lambda<\rho_f} (0 - \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}})dG(\lambda) < 0$$

The upper limit for the integral is ρ_f instead, because any consumers with $\underline{\lambda} \geq \lambda \geq \rho_f$ do not buy the product with or without the firm imposing the constraint.

Case 2: $\rho_f \geq \rho_\delta$

Under this assumption, we have $f(H)\delta_H \geq f(L)\delta_L$. This case is illustrated in Figure 2. Furthermore, by Lemmas 1 and 2, any biased individual with $\rho_\delta \leq \lambda < \rho_f$, will always buy the product. Any individual with $\rho_\delta \leq \rho_f \leq \lambda$, will acquire uninformative signals and buy the product regardless if she is biased, and reject the product if she is Bayesian. The firm's strategy can be discussed under three cases.

Strategy 2.1: $0 < \underline{\lambda} < \rho_\delta$

We can refer to Figure 2 to see that the constrained consumers are those who follow their signals without the constraint, and will continue to follow their signals with the constraint, as $\underline{\lambda} < \rho_\delta < \rho_f$. Similar to the previous case, the change in profit is negative.

$$E[\Delta\Pi](\underline{\lambda} < \rho_\delta | \rho_\delta \leq \rho_f) = \int_{\lambda<\underline{\lambda}} (q^{\underline{\lambda}}(h) - q^{\lambda}(h))dG(\lambda)$$
$$= \int_{\lambda<\underline{\lambda}} (\frac{f(L) - f(H)e^{1/\underline{\lambda}}}{1 - e^{1/\underline{\lambda}}} - \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}})dG(\lambda) < 0$$

Strategy 2.2: $\rho_\delta \leq \underline{\lambda} < \rho_f$

For this case, any biased constrained consumer will ignore bad signals. If the consumer originally had $\lambda < \rho_\delta$, then she went from buying the product only when seeing an h signal to always buying the product regardless of signal. If the biased consumer originally had $\lambda \geq \rho_\delta$, their behavior doesn't change. On the other hand, constrained Bayesian consumers will continue to follow their informative signals consistently. With proportion

p of the population being biased, the change in profit is as follows.

$$E[\Delta\Pi](\rho_\delta \le \underline{\lambda} < \rho_f | \rho_\delta \le \rho_f) =$$

$$p\int_{\lambda<\rho_\delta}(1-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)+(1-p)\int_{\lambda<\underline{\lambda}}(\frac{f(L)-f(H)e^{1/\underline{\lambda}}}{1-e^{1/\underline{\lambda}}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)$$

This can be positive as the first integral is positive.

Strategy 2.3: $\rho_\delta \le \rho_f \le \underline{\lambda}$

Any biased constrained consumer will acquire uninformative signals and always buy the product. However, any biased consumer with $\lambda \ge \rho_\delta$ always buys the product without the firm's constraint anyways. Meanwhile, constrained Bayesian consumers will not buy the product. This gives the following change in profit.

$$E[\Delta\Pi](\rho_f \le \underline{\lambda} | \rho_\delta \le \rho_f) =$$

$$p\int_{\lambda<\rho_\delta}(1-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)+(1-p)\int_{\lambda<\underline{\lambda}}(0-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)$$

This can be positive as the first integral is positive. However, this strategy is dominated by the best $\underline{\lambda}$ analyzed in Case 2.2. This is because the domain over which the second integral is taken is larger, and additionally,

$$0-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} < \frac{f(L)-f(H)e^{1/\underline{\lambda}}}{1-e^{1/\underline{\lambda}}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} < 0$$

Focusing on Strategy 2.2, we observe that $E[\Delta\Pi]$ is decreasing in $\underline{\lambda}$ due to the second term:

$$(1-p)\int_{\lambda<\underline{\lambda}}(\frac{f(L)-f(H)e^{1/\underline{\lambda}}}{1-e^{1/\underline{\lambda}}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})$$

As λ increases, the domain of the integral is larger, and the expression inside is smaller (more negative). Therefore, under Strategy 2.2, the highest change in expected profit is achieved by $\underline{\lambda}^* = \rho_\delta$.

For this $\underline{\lambda}^*$ to be optimal, we need $E[\Delta\Pi](\underline{\lambda}^* = \rho_\delta | \rho_\delta \leq \rho_f) \geq 0$.

$$p\int_{\lambda<\rho_\delta}(1-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)+(1-p)\int_{\lambda<\rho_\delta}(\frac{f(L)-f(H)e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}})dG(\lambda)\geq 0$$

Which gives us

$$\int_{\lambda<\rho_\delta}\left[p+(1-p)\frac{f(L)-f(H)e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}}\right]dG(\lambda)\geq 0 \tag{13}$$

Note that as $f(H)$ approaches 0.5, the integrand becomes:

$$\lim_{f(H)\to 0.5} p+(1-p)\frac{f(L)-f(H)e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}}-\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}}=0.5p$$

This is positive whenever we have a positive p.

When $p=0$, the term is unambiguously negative, indicating that it is only profitable to constrain information in the presence of biased consumers.

Note that the condition $\rho_\delta \leq \rho_f$ is equivalent to a lower bound on $f(H)$:

$$\delta_H f(H) \geq \delta_L f(L) \iff f(H) \geq \frac{\delta_L}{\delta_H+\delta_L}$$

When $f(H) \to \frac{\delta_L}{\delta_H+\delta_L}$, we have

$$f(L)-f(H)e^{1/\rho_\delta} \to \frac{\delta_H}{\delta_H+\delta_L}-\frac{\delta_L}{\delta_H+\delta_L}\frac{\delta_H}{\delta_L}=0$$

while for $\lambda < \rho_\delta$,

$$\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} \to \frac{\delta_L e^{1/\lambda}-\delta_H}{(e^{1/\lambda}-1)(\delta_H+\delta_L)}>0$$

The integrand becomes

$$p-\frac{\delta_L e^{1/\lambda}-\delta_H}{(e^{1/\lambda}-1)(\delta_H+\delta_L)}$$

Therefore, the integral can be positive or negative at the limit. However, notice that the integrand is monotonically decreasing in $f(H)$, as $\frac{e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} < \frac{e^{1/\lambda}}{1-e^{1/\lambda}}$. Hence, given $p>0$, there exists a threshold for $f(H)$ such that the expected change in profit is positive for all

$f(H)$ values above the threshold.

For any $p > 0$, if $E[\Delta\Pi](\underline{\lambda}^* = \rho_\delta | \rho_\delta \leq \rho_f) \geq 0$ when $f(H) \to \frac{\delta_L}{\delta_H + \delta_L}$, then the firm finds it profitable to constrain information, setting $\underline{\lambda}^* = \rho_\delta$ whenever $p > 0$ and $\rho_\delta \leq \rho_f$. For any $p > 0$, if $E[\Delta\Pi](\underline{\lambda}^* = \rho_\delta | \rho_\delta \leq \rho_f) < 0$ when $f(H) \to \frac{\delta_L}{\delta_H + \delta_L}$, then there exists an $\hat{f} \in [\frac{\delta_L}{\delta_H + \delta_L}, 0.5)$ such that for $p > 0$ and $f(H) \geq \hat{f}$, the firm finds it profitable to constrain information by setting $\underline{\lambda}^* = \rho_\delta$. $\square$

Proof of Proposition 2: Optimal constraints under competition.

The premise of the following derivations is the assumption that $\rho_f \geq \rho_\delta$ and $p > 0$, as these are necessary for the monopolist to have an incentive to constrain information in the first place.

Denote the incumbent's and the entrant's strategy as $\underline{\lambda}^I$ and $\underline{\lambda}^E$ respectively, and let their profits be $\Pi^I(\underline{\lambda}^I, \underline{\lambda}^E)$ and $\Pi^E(\underline{\lambda}^I, \underline{\lambda}^E)$ respectively. First, we will rule out some strategies. Note that given any $\underline{\lambda}^I$, setting $\underline{\lambda}^E = k \in (0, \rho_\delta)$ is strictly dominated by setting $\underline{\lambda}^E = 0$ for the entrant. To see this, consider the change in expected profit from switching strategies. If $k > \underline{\lambda}^I$,

$$\Pi^E(\underline{\lambda}^I, 0) - \Pi^E(\underline{\lambda}^I, k) = \int_0^{\underline{\lambda}^I} [g^\lambda(h) - \eta g^k(h)] dG(\lambda) + \eta \int_{\underline{\lambda}^I}^k g^\lambda(h) dG(\lambda)$$

The first integral represents the entrant getting all the consumers with $\lambda \in [0, \underline{\lambda}^I]$, instead of just η of them and with lower purchase probability. The second integral represents attracting η of the consumers with $\lambda \in (\underline{\lambda}^I, k]$, instead of none of them. The change is clearly positive. Alternatively, if $k \leq \underline{\lambda}^I$, then the change becomes

$$\Pi^E(\underline{\lambda}^I, 0) - \Pi^E(\underline{\lambda}^I, k) = \int_0^k [g^\lambda(h) - \eta g^k(h)] dG(\lambda)$$

This represents the entrant attracting all the customers with $\lambda \in [0, k]$, instead of just η of them and with lower purchase probability. Again, this is positive. Furthermore, since the incumbent's profits are symmetric to the entrant's except with $1 - \eta$ instead of η, this argument holds for the incumbent as well. As such, we do not need to consider any strategy within $(0, \rho_\delta)$.

Next, we note that given any strategy by the incumbent, setting $\underline{\lambda}^E = k \in (\rho_\delta, \infty)$ is strictly dominated by setting $\underline{\lambda}^E = \rho_\delta$ for the entrant. To see this, again take the difference from switching strategies. If $k > \rho_\delta > \underline{\lambda}^I$,

$$\Pi^E(\underline{\lambda}^I, \rho_\delta) - \Pi^E(\underline{\lambda}^I, k) = \eta[p \int_{\rho_\delta}^{k} dG(\lambda) + (1-p) \int_{\rho_\delta}^{min\{k,\rho_f\}} g^\lambda(h) dG(\lambda)]$$

The first integral represents the entrant attracting η of the biased consumers with $\lambda \in [\rho_\delta, k]$, who will always buy. The second integral represents the entrant attracting η of the Bayesian consumers with $\lambda \in [\rho_\delta, k]$, who will follow their signal recommendations. Alternatively, if $k > \underline{\lambda}^I > \rho_\delta$,

$$\Pi^E(\underline{\lambda}^I, \rho_\delta) - \Pi^E(\underline{\lambda}^I, k) = \eta[p \int_{\underline{\lambda}^I}^{k} dG(\lambda) + (1-p) \int_{\underline{\lambda}^I}^{min\{k,\rho_f\}} g^\lambda(h) dG(\lambda)] + \int_{\rho_\delta}^{\underline{\lambda}^I} p + (1-p) g^\lambda(h) dG(\lambda)$$

The first two integrals represent the types of customers the entrant now shares with the incumbent, while the second integral represents the types of customers the entrant has full market share over due to looser constraints. Lastly, if $\underline{\lambda}^I > k > \rho_\delta$, then

$$\Pi^E(\underline{\lambda}^I, \rho_\delta) - \Pi^E(\underline{\lambda}^I, k) = \int_{\rho_\delta}^{min\{k,\rho_f\}} p + (1-p) g^\lambda(h) dG(\lambda)$$

The entrant is fully attracting all consumers with $\lambda \in [\rho_\delta, k]$. All three of the above scenarios yield positive gains, hence we need not consider any strategy setting $\underline{\lambda}^E > \rho_\delta$. The same goes for the incumbent.

We find the equilibrium where both firms choose from one of two strategies $\{0, \rho_\delta\}$. We begin by examining the entrant's optimal response to an entrant that sets $\underline{\lambda}^I = \rho_\delta$. The profit for the entrant given $\underline{\lambda}^E \in \{0, \rho_\delta\}$ is:

$$\Pi^E(\rho_\delta, 0) = \int_0^{\rho_\delta} \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}} dG(\lambda) + \eta \left[p \int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p) \int_{\rho_\delta}^{\rho_f} \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}} dG(\lambda) \right]$$

$$\Pi^E(\rho_\delta, \rho_\delta) = \eta p \int_0^{\infty} dG(\lambda) + \eta(1-p) \left[\int_0^{\rho_\delta} \frac{f(L) - f(H)e^{1/\rho_\delta}}{1 - e^{1/\rho_\delta}} dG(\lambda) + \int_{\rho_\delta}^{\rho_f} \frac{f(L) - f(H)e^{1/\lambda}}{1 - e^{1/\lambda}} dG(\lambda) \right]$$

We find the conditions for which it is optimal for the entrant to not constrain information acquisition (Inequality (10) from the main text.)

$$\Pi^E(\rho_\delta,0) > \Pi^E(\rho_\delta,\rho_\delta) \iff$$
$$\int_0^{\rho_\delta}\left[-\eta p + \frac{1-f-fe^{1/\lambda}}{1-e^{1/\lambda}} - \eta(1-p)\frac{1-f-fe^{1/\rho_\delta}}{1-e^{1/\rho_\delta}}\right] dG(\lambda) > 0$$

Note that when either η or p is 0, the integrand is positive. However, we are not considering $p = 0$, as it would then never be profitable to impose any constraints. When both η and p are 1, the integrand is negative. The integral value is linear and decreasing in both η and p. Denote

$$\beta_0 = \int_0^{\rho_\delta} \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda)$$
$$\beta_1 = \int_0^{\rho_\delta} dG(\lambda)$$
$$\beta_2 = \int_0^{\rho_\delta} \frac{f(L)-f(H)e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} dG(\lambda)$$

Observe that $\beta_1 > \beta_0 > \beta_2 > 0$. Then

$$\Pi^E(\rho_\delta,0) > \Pi^E(\rho_\delta,\rho_\delta) \iff \beta_0 - \beta_1\eta p - \beta_2\eta(1-p) > 0$$

The entrant's best response is not constraining information if $\eta < \frac{\beta_0}{(\beta_1-\beta_2)p+\beta_2}$. That is, for any point (η, p) to the bottom left of the curve $\beta_0 - \beta_1\eta p - \beta_2\eta(1-p) = 0$.

Now consider the entrant's best response to $\underline{\lambda}^I = 0$.

$$\Pi^E(0,0) = \eta\left[\int_0^{\rho_\delta} \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda) + p\int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p)\int_{\rho_\delta}^{\rho_f} \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda)\right]$$

$$\Pi^E(0,\rho_\delta) = \eta\left[p\int_{\rho_\delta}^{\infty} dG(\lambda) + (1-p)\int_{\rho_\delta}^{\rho_f} \frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} dG(\lambda)\right]$$

It is quite obvious that $\Pi^E(0,0) > \Pi^E(0,\rho_\delta)$. Therefore, the entrant would best respond with no constraint as well.

Lastly, we check what strategy the incumbent would employ in anticipation of the entrant's best response. When $\underline{\lambda}^E = 0$, by symmetry,

$$\Pi^I(0,0) = \frac{1-\eta}{\eta}\Pi^E(0,0)$$

$$\Pi^I(\rho_\delta,0) = \frac{1-\eta}{\eta}\Pi^E(0,\rho_\delta)$$

Immediately, we see $\Pi^I(0,0) > \Pi^I(\rho_\delta,0)$.

Therefore, if $\eta < \frac{\beta_0}{(\beta_1-\beta_2)p+\beta_2}$, the entrant has dominant strategy $\underline{\lambda}^E = 0$, and the incumbent would preemptively respond with $\underline{\lambda}^I = 0$. If $\eta \geq \frac{\beta_0}{(\beta_1-\beta_2)p+\beta_2} > 0$, the entrant's best response to $\underline{\lambda}^I = \rho_\delta$ would be to match it $\underline{\lambda}^E = \rho_\delta$. Here the incumbent gets to choose the desirable equilibrium between $(\rho_\delta,\rho_\delta)$ and $(0,0)$. By symmetry,

$$\Pi^I(\rho_\delta,\rho_\delta) = \frac{1-\eta}{\eta}\Pi^E(\rho_\delta,\rho_\delta)$$

We check the conditions for the following inequality (which was Inequality 11 in the main text.)

$$\Pi^I(0,0) > \Pi^I(\rho_\delta,\rho_\delta) \iff \int_0^{\rho_\delta}\left[\frac{f(L)-f(H)e^{1/\lambda}}{1-e^{1/\lambda}} - (1-p)\frac{f(L)-f(H)e^{1/\rho_\delta}}{1-e^{1/\rho_\delta}} - p\right]dG(\lambda) > 0$$

Expressing this in terms of the β's as defined above,

$$\beta_0 - \beta_1 p - \beta_2(1-p) > 0$$

The incumbent prefers $(\rho_\delta,\rho_\delta)$ if this inequality holds, and prefers $(0,0)$ otherwise. This is the same inequality as (13), meaning that the incumbent will prefer $(\rho_\delta,\rho_\delta)$ over $(0,0)$ if and only if it prefers to constrain information without the entrant present.

This makes it seem like competition has no impact, but that is not the case. For $(\rho_\delta,\rho_\delta)$ (and for the information constraint under monopoly) to be preferable, we need

$$p \geq \frac{\beta_0-\beta_2}{\beta_1-\beta_2} > 0$$

However, (10) states that for the entrant to best respond to constraining information by mirroring the incumbent's move, we would need

$$p \geq \frac{\beta_0 - \beta_2 \eta}{(\beta_1 - \beta_2)\eta} > \frac{\beta_0 - \beta_2}{\beta_1 - \beta_2}$$

Therefore, competition shrinks the set of eligible parameter values for which information manipulation is profitable. Specifically, for $p \in [\frac{\beta_0 - \beta_2}{\beta_1 - \beta_2}, \frac{\beta_0 - \beta_2 \eta}{(\beta_1 - \beta_2)\eta})$, the only equilibrium is $(0, 0)$. While the incumbent would prefer otherwise, the entrant would always choose not to constrain information.

In sum, given $p > 0$ and $\rho_f > \rho_\delta$, it is profitable to constrain information, with or without competition, if $p \geq \frac{\beta_0 - \beta_2}{\beta_1 - \beta_2}$. With competition from an entrant, the equilibrium is $(0, 0)$ when $p < \frac{\beta_0 - \beta_2 \eta}{(\beta_1 - \beta_2)\eta}$, and $(\rho_\delta, \rho_\delta)$ otherwise. This means that for $p \in [\frac{\beta_0 - \beta_2}{\beta_1 - \beta_2}, \frac{\beta_0 - \beta_2 \eta}{(\beta_1 - \beta_2)\eta})$, the incumbent is forced to switch from information constraint to no constraint.

□

Proof of Proposition 3: Generalized Updating Rule

Unless $f(H) = 0$ or $\lambda = 0$, the individual will not obtain perfectly accurate signals such that they have posteriors at 0 or 1. Due to the second assumption, we always have

$$Eu(b|s = h) = F_H(f, s = h) - F_L(f, s = h) > g^B(H|s = h) - g^B(L|s = h) \geq 0$$

We just need to check the behavior of biased individuals when they receive l signals.

The third assumption is a sufficient condition to guarantee that expected utility of buying after receiving an l signal is non-decreasing in λ. To see this, note that with two states

$$Eu(b|s = l) = F_H(f, s = l) - F_L(f, s = l) = 2F_H(f, s = l) - 1$$

And

$$\frac{F_H(f, s = l)}{F_L(f, s = l)} = \frac{F_H(f, s = l)}{1 - F_H(f, s = l)}$$

Both increasing in $F_H(f, s = l)$. Hence $Eu(b|s = l)$ would be non-decreasing in $\frac{\pi(l|H)}{\pi(l|L)}$.

Furthermore, the optimal signal structure satisfies

$$\frac{\pi^*(l|H)}{\pi^*(l|L)} = \frac{\frac{q^*(l)}{f(H)(1+e^{1/\lambda})}}{\frac{q^*(l)e^{1/\lambda}}{f(L)(1+e^{1/\lambda})}} = \frac{f(L)}{f(H)e^{1/\lambda}}$$

which is increasing in λ. For high enough λ, the individual acquires pure noise, with $q(h) = 0$. In this region, $\frac{\pi^*(l|H)}{\pi^*(l|L)} = 1$, and the posterior is a constant function of λ. Therefore, $Eu(b|s = l)$ is non-decreasing in λ. Furthermore, since the posterior is continuous in signal probabilities, $Eu(b|s = l)$ is also continuous in λ.

Note that when we increase λ such that $\lambda \to \infty$, signals become completely uninformative at the limit. When $f(H) = 0.5$ and signals are completely uninformative, we would have $g^B(H|s) = 0.5$. Thus, $F_H(f(H) = 0.5, s) > 0.5$, and consequently, $Eu(b|s = l) > 0$ in this scenario. While the optimal signal structure is never completely uninformative when $f(H) = 0.5$ and $\lambda < \infty$, it still allows there to exist $f(H) < 0.5$ and finite λ such that $Eu(b|s = l) > 0$ by the continuity of the updating rule. However, it is crucial to point out that, with optimal information acquisition, the exogenous prior may be within a region such that the expected utility may be negative for all λ. Hence, the following discussion.

If the updating rule is such that $Eu(b|s = l) > 0$, or equivalently $F_H(f, s) > 0.5$, given any f and λ, then the problem becomes trivial, and the firm need not do anything. We focus on the scenario where the posterior over state H could be below 0.5.

Given this, we can solve the following for λ:

$$Eu(b|s = l) = 0$$

Due to continuity, and the expression being monotone in λ, we have the set of solutions $\{\lambda | Eu(b|s = l) = 0\}$ being a closed interval (a single point if the expected utility is strictly increasing in λ). Take the minimum of the set and denote it as $\rho^*(F)$.

Recall that the threshold (for λ) for the individual to seek uninformative signals is ρ_f. If $\rho_f < \rho^*(F)$, then the biased individuals with $\lambda < \rho_f$ are acquiring informative signals, while following the action recommended by the signals as well. That is, they don't buy if they see l signals. For $\lambda \geq \rho_f$, as λ increases, the posterior does not change, as the individual is starting to acquire pure noise. Therefore, the condition $\rho_f < \rho^*(F)$ exactly indicates the region for f where it is impossible to obtain $Eu(b|s = l) \geq 0$. If $\rho_f \geq \rho^*(F)$, then for individuals with $\lambda \in [\rho^*(F), \rho_f]$, we have $Eu(b|s = l) \geq 0$ as the expected utility is non-decreasing in λ. Furthermore, since posteriors stay constant when λ increases beyond ρ_f (due to individuals acquiring pure noise from that point onward), we have $Eu(b|s = l) \geq 0$ for all $\lambda \geq \rho^*(F)$. (Note that the case where the posterior $F_H(f, s = l)$ is always higher that 0.5 can be included as the special case of $\rho^*(F) = 0$.)

Since the set $\mathcal{F} = \{f(H) \in [0, 0.5] | \rho_f \geq \rho^*(F)\}$ must include the neighborhood around $f(H) = 0.5$, the firm's strategy carries out the same as in Proposition 1. We know that the firm either chooses not to constrain information setting $\underline{\lambda} = 0$, or constrain information by setting $\underline{\lambda} = \rho^*(F)$. There exists $\underline{f} \in \mathcal{F}$ such that for all $f(H) \in \mathcal{F}$ satisfying $f(H) \geq \underline{f}$, the firm finds it profitable to set $\underline{\lambda} = \rho^*(F)$. This eligible set of priors must be non-empty, since $f(H) = 0.5$ must be in it.

□

References

Benjamin, D. J. (2019): "Chapter 2 - Errors in probabilistic reasoning and judgment biases," in *Handbook of Behavioral Economics - Foundations and Applications 2*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, North-Holland, vol. 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, 69–186.

Boyaci, T. and Y. Akçay (2018): "Pricing When Customers Have Limited Attention," *Management Science*, 64, 2995–3014.

Camerer, C. (1998): "Bounded Rationality in Individual Decision Making," *Experimental Economics*, 1, 163–183.

Caplin, A. and M. Dean (2015): "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105, 2183–2203.

Cripps, M. W. (2018): "Divisible updating," *Manuscript, UCL.*

de Clippel, G. and X. Zhang (2022): "Non-Bayesian Persuasion," *Journal of Political Economy*, 130, 2594–2642.

de Oliveira, H., T. Denti, M. Mihm, and K. Ozbek (2017): "Rationally inattentive preferences and hidden information costs," *Theoretical Economics*, 12, 621–654.

DellaVigna, S. and U. Malmendier (2004): "Contract Design and Self-Control: Theory and Evidence*," *The Quarterly Journal of Economics*, 119, 353–402.

Dominiak, A., M. Kovach, and G. Tserenjigmid (2023): "Inertial Updating," *working paper.*

Edwards, W. (1982): *Conservatism in human information processing*, Cambridge University Press, 359–369.

Epstein, L. G. (2006): "An axiomatic model of non-Bayesian updating," *The Review of Economic Studies*, 73, 413–436.

Gabaix, X. and D. Laibson (2006): "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*," *The Quarterly Journal of Economics*, 121, 505–540.

Gentzkow, M. and E. Kamenica (2014): "Costly Persuasion," *American Economic Review*, 104, 457–62.

Grether, D. M. (1980): "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *The Quarterly Journal of Economics*, 95, 537–557.

Korteling, J. E. and A. Toet (2021): "Cognitive biases," in *Encyclopedia of behavioral neuroscience*, ed. by S. D. Sala, Elsevier Science Amsterdam-Edinburgh, vol. 3, 610–619.

Kovach, M. (2020): "Twisting the truth: Foundations of wishful thinking," *Theoretical Economics*, 15, 989–1022.

Lipnowski, E. and L. Mathevet (2018): "Disclosure to a Psychological Audience," *American Economic Journal: Microeconomics*, 10, 67–93.

Liu, Z. (2023a): "A Model of Confirmation Bias," *working paper, University of Maryland.*

——— (2023b): "Dual Elicitation of Motivated and Unmotivated Confirmation Bias," *working paper, University of Maryland.*

Mackowiak, B. and M. Wiederholt (2009): "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769–803.

Martin, D. (2017): "Strategic pricing with rational inattention to quality," *Games and Economic Behavior*, 104, 131–145.

Matějka, F. (2015): "Rationally Inattentive Seller: Sales and Discrete Pricing," *The Review of Economic Studies*, 83, 1125–1155.

Matějka, F. and A. McKay (2012): "Simple Market Equilibria with Rationally Inattentive Consumers," *American Economic Review*, 102, 24–29.

——— (2015): "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105, 272–98.

Mayraz, G. (2011): "Priors and Desires," CEP Discussion Papers dp1047, Centre for Economic Performance, LSE.

Ortoleva, P. (2012): "Modeling the change of paradigm: Non-Bayesian reactions to unexpected news," *American Economic Review*, 102, 2410–2436.

Rabin, M. and J. L. Schrag (1999): "First Impressions Matter: A Model of Confirmatory Bias," *The Quarterly Journal of Economics*, 114, 37–82.

Ravid, D. (2020): "Ultimatum Bargaining with Rational Inattention," *American Economic Review*, 110, 2948–63.

Sims, C. A. (2003): "Implications of rational inattention," *Journal of Monetary Economics*, 50, 665–690, swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

——— (2006): "Rational Inattention: Beyond the Linear-Quadratic Case," *American Economic Review*, 96, 158–163.

Steiner, J., C. Stewart, and F. Matějka (2017): "Rational Inattention Dynamics: Inertia and Delay in Decision-Making," *Econometrica*, 85, 521–553.

Wei, D. (2021): "Persuasion under costly learning," *Journal of Mathematical Economics*, 94.

Woodford, M. (2009): "Information-constrained state-dependent pricing," *Journal of Monetary Economics*, 56, S100–S124, supplement issue: December 12-13, 2008 Research Conference on 'Monetary Policy under Imperfect Information' Sponsored by the Swiss National Bank (http://www.snb.ch) and Study Center Gerzensee (www.szgerzensee.ch).