# Sparking curiosity or tipping the scales? Targeted advertising to rationally inattentive consumers*

Andrei Matveenko[1] and Egor Starkov [†1]

[1]University of Copenhagen, Department of Economics

September 15, 2021

**Abstract**

This paper argues, in the context of targeted advertising, that receivers' rational inattention and ability to independently acquire information have a non-trivial impact on the sender's optimal disclosure strategy. In our model, a monopolist has an opportunity to launch an advertising campaign and chooses a targeting strategy – which consumers to send its advertisement to. The consumers are uncertain about and heterogeneous in their valuations of the product, and are rationally inattentive in that they must incur a cost if they want to learn their true valuations. We discover that the firm generally prefers to target consumers who are either indifferent between ignoring and investigating the product, or between investigating and buying it unconditionally. If the firm is uncertain about the consumer appeal of its product, it targets these two distinct groups of consumers simultaneously but may ignore all consumers in between.

**JEL-codes**: D83, L15, M37.

**Keywords**: advertising, targeting, rational inattention, costly disclosure.

*We are grateful to Mogens Fosgerau, Germain Gaudin, Marit Hinnosaar, Alessandro Ispano, Igor Letina, Filip Matějka, Andrey Minaev, Xiaosheng Mu, Marco Schwarz, Adrien Vigier, and Alexander White, as well as participants of the 2020 Generalized Entropy Workshop, OLIGO 2020 and IIOC 2021 virtual conferences for helpful comments and valuable feedback. This project has received funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No 740369).

†The authors' contributions are equal and the order is alphabetical.
E-mail addresses: `anma@econ.ku.dk` (A. Matveenko), `egor.starkov@econ.ku.dk` (E. Starkov).

# 1 Introduction

The recent decade has witnessed a flood of website analytics systems and smartphone apps that collect detailed information on users and enable firms to track consumers' tastes and actions with increasing precision. The economic literature has largely focused on the implications of more extensive information collection for price discrimination allowed by the ever-growing amounts of information available to firms. However, much less attention has been devoted to what is arguably the more widespread use of personal information by firms, personalized advertising.

Knowledge of consumer tastes allows a firm to target its advertising towards audiences that can be most easily manipulated into purchasing the product, thus generating higher returns per dollar spent on advertising. The value added by targeting can be quite substantial, as illustrated (albeit in a political, rather than economic context) by the story of Cambridge Analytica, a now-infamous company that arguably played an important role in the outcomes of the 2016 US Presidential Election and the UK Brexit referendum by influencing voters via highly-targeted political advertising on Facebook.[1]

The folk wisdom implies that firms benefit most from advertising to consumers who pass on buying the product by a thin margin. Specifically, if consumers are uncertain about their valuation of the product, and advertising has a positive effect on their expectations, then a firm should target its advertising towards consumers whose beliefs are just a bit too pessimistic to justify a purchase. However, this reasoning ignores the fact that consumers can acquire more information after being subjected to advertising. Indeed, it seems implausible, especially with larger purchases such as automobiles or major appliances, that a consumer would make up her mind based solely on advertising. Instead, an ad would lead the consumer to investigate the product more or less closely, and affect her purchasing decision indirectly through this channel. We explore the trade-offs of targeted advertising in such a setting, where consumers strategically acquire information about a product. We show that this information acquisition layer overturns the folk wisdom outlined above, and the optimal advertising strategy changes dramatically when consumers can acquire information.[2]

To identify the optimal advertising strategy, we consider a theoretical model in

[1] See, for instance, `https://www.theguardian.com/uk-news/2018/mar/23/leaked-cambridge-analyticas-blueprint-for-trump-victory`.

[2] We focus on persuasive advertising, which aims to strengthen consumer willingness to pay for a product, rather than informative advertising meant to raise consumer awareness of a product.

which a population of rationally inattentive consumers is faced with an option to buy a product which yields an uncertain payoff. The consumers differ in their initial perceptions regarding the product's payoff, and they can acquire costly information about it. The monopolistic firm has an option to send a costly ad to a subset of consumers, improving their beliefs about the product, and the firm can freely select which consumers will receive the ad – i.e., target the ad towards the most susceptible consumers.

We discover that the optimal advertising strategy deviates quite significantly from the intuitive strategy of targeting consumers who are ex ante "almost indifferent" between buying and not buying the product. In particular, if the effect of the advertisement on consumer expectations is not too large, and the cost of advertising is significant, then it is optimal for the firm to target consumers in two distinct groups. The first group is the pessimistic consumers who ignore the product without investigating it. Sending an ad to these consumers will render them interested enough to acquire more information about the product, which translates into sales. The second group is the optimistic consumers who are close to buying the product but who want to acquire a little additional information "just to be sure". Advertising to these consumers will "tip the scales" and convince them to buy the product without further investigation.

The relative focus on the two groups depends on the firm's belief about its own product: an optimistic firm will push consumers towards acquiring more information, while a pessimistic one would rather prevent them from doing so. Curiously, if the firm is uncertain enough about the marketability of its product, then it will advertise to both groups simultaneously and at the same time ignore consumers with average beliefs – who are ex ante "closest to indifference" between buying and not buying the product. This is because such consumers choose to investigate the product regardless of advertising, meaning that advertising does not have much effect on their behavior. It is more valuable for the firm to affect the consumers' information acquisition decisions at the extensive margin (whether a consumer acquires any information) than at the intensive margin (how much and what kind of information a consumer acquires).

Our results hold regardless of whether consumers are sophisticated – i.e., are fully Bayesian, as is commonly assumed in Economics, – or "cursed" in the sense of Eyster and Rabin [2005]. A cursed consumer in our model does not draw any inferences from the fact that they did *not* receive an advertisement. This can happen because they are naive, or because they are unaware of the ad campaign unless it reaches them directly, or because they do not expect to be in the target group. While we mainly use the model with cursed consumers as a supplementary step in our analysis of the model with sophisticated consumers, cursedness is an empirically appealing modelling assumption on its own.[3]

Finally, the running example in our paper is that of a monopolist designing an advertising campaign. However, our results results have straightforward extensions to other similar environments, such as political advertising, and to non-monopolistic settings. In particular, in Section 6 we briefly explore an extension of our model to a political competition setting, in which two candidates are simultaneously choosing which voters to target with their promotional efforts. We show that our results continue to hold in that setting.

The remainder of this paper is organized as follows. Section 2 illustrates the main idea behind our result with a simple example. Section 3 reviews the relevant literature. Section 4 describes the model, and Section 5 derives the optimal advertising strategy and explores its properties. Section 6 explores an extension of our model to a political competition setting. Section 7 concludes. Most proofs are retained in the main text, since we believe they are concise and clear enough to convey valuable intuition. However, some of the longer proofs and supplementary lemmas are relegated to the Appendix.

# 2 Illustrative Example

The main driving force behind our result can be illustrated using the following highly stylized example. Suppose a firm sells an item of quality $s \in \{H, L\}$ at some fixed price normalized to 1. The firm values the item at zero. The consumer values a high-quality item at $v = w$ net of the price, and a low-quality item at $v = w - 1$. Both the firm and the consumer ex ante believe the product is of high quality with probability $p^0 = \mathbb{P}(H) = 1/3$. The firm receives a signal $y \in \{h, l\}$ about $s$ with precision $\rho = 0.8$ (thus with 80% probability the signal is correct, $\mathbb{P}(h|H) = \mathbb{P}(l|L) = 0.8$). The firm can reveal (advertise) this signal to the consumer at some cost $c < 1/3$. After receiving this ad, if any, the consumer can choose to investigate the quality (learn it perfectly) at cost $\lambda = 0.1$. Assume for this example

[3]Empirical evidence of the cursedness of consumers in various settings has been obtained by Li and Hitt [2008] and Brown, Camerer, and Lovallo [2012] in the field, and Jin, Luca, and Martin and Deversi, Ispano, and Schwardmann [2018] in the lab. Markets with cursed consumers have been explored by Matysková and Šípek [2017] and Ispano and Schwardmann [2018].

that the firm's signal is high, $y = h$, so observing $y$ increases a consumer's belief that $s = H$ to $\alpha(p^0) = \mathbb{P}(h|H)\mathbb{P}(H)/(\mathbb{P}(h|H)\mathbb{P}(H) + \mathbb{P}(h|L)\mathbb{P}(L)) = 2/3$. Assume also that if the consumer receives no ad, her belief remains at $p^0$.

The consumer effectively has three options: pass, buy, or investigate (and then buy if and only if $s = H$). Her expected payoffs from these options conditional on belief $p$ are given by:

$$\mathbb{E}U = \begin{cases} 0, & \text{if pass;} \\ w - (1-p) = pw + (1-p)(w-1), & \text{if buy;} \\ pw - \lambda = pw + (1-p)0 - \lambda, & \text{if investigate.} \end{cases}$$

The consumer then passes if $w \le \underline{w}$, investigates if $w \in (\underline{w}, \bar{w})$, and buys if $w \ge \bar{w}$, where $(\underline{w}, \bar{w}) = (0.3, 0.85)$ if the consumer does not receive the ad, and $(\underline{w}, \bar{w}) = (0.15, 0.7)$ if she does.

The firm's expected profit, depending on the consumer's decision, is given by: $\mathbb{E}\pi = 0$ if the consumer passes, $\mathbb{E}\pi = 1$ if she buys, and $\mathbb{E}\pi = \alpha(p) = 2/3$ if she investigates. Advertising to a consumer with $w \in [0.7, 0.85]$ pushes her from investigating the product to buying it immediately, increasing the firm's profit by 1/3. Advertising to a consumer with $w \in [0.15, 0.3]$ nudges her to at least investigate the product instead of ignoring it altogether, thereby increasing the firm's profit by 2/3. Advertising to a consumer with any other $w$ (including the ex ante indifferent consumer with $w = 2/3$) would not affect her behavior and is thus pointless.

Of note is the fact that the firm in this example profits more from advertising to pessimistic (low-$w$) consumers. This is driven by the firm's optimistic belief $p_f = 2/3$ in the quality of its product, conditional on $y = h$. Optimism leads the firm to expect that any investigating consumer is likely to learn that the product quality is high. Therefore, incentivizing otherwise unwilling consumers to investigate is valuable for the firm, while nudging the investigating consumers to buy unconditionally is less so. Conversely, if prior $p^0$ and signal strength $\rho$ were low, targeting optimistic (high-$w$) consumers would be better for the firm.

The observations above may appear as if they are due to the specific assumption that the consumers acquire information on an all-or-nothing basis (learn quality $s$ perfectly or learn nothing at all). The remainder of the paper uses a more general model to demonstrate that this is not the case – the intuition above holds under a more flexible information acquisition process, when the consumer can acquire any information she prefers, – and provide more insight into these and other issues.

# 3 Literature Review

Goldfarb and Tucker [2019] present an excellent survey of works on digital economics; see Chapter 6 for a survey of literature dealing with the implications of improved consumer tracking technologies. This literature is also discussed in Acquisti, Taylor, and Wagman [2016], a detailed survey on the economics of privacy.

Iyer, Soberman, and Villas-Boas [2005] is a seminal paper on targeted advertising. They model advertising as generating awareness of the product, and find that in the presence of competition, firms prefer to target consumers with a strong preference for their product, rather than those close to indifference. In the end, two firms in the market target their advertising to non-overlapping populations and do not fight for the "median consumer". We show that with informative advertising and a subsequent consumer information acquisition stage, a similar outcome materializes even in the absence of competition: a single monopolistic firm would target similar groups ("fans" and/or "haters" but not undecided consumers), although for very different reasons. Hefti and Liu [2020] extend the model of Iyer et al. [2005] to allow for inattentive consumers who may forget about one of the firms when making their choice. They show that this kind of inattention may lead firms to flood the whole market with ads even when they have the power to target their advertisements towards specific consumers.

Hoffmann, Inderst, and Ottaviani [2020] model targeted advertising as selective disclosure of information. They show that while the seller would always prefer to target his ads given the opportunity, his profits would be higher if targeting was not available. Consumers, on the other hand, always benefit from targeted information in their model. Villas-Boas and Yao [2021] tackle the problem of optimal ad targeting based on consumers' search behavior, which is informative of whether they are interested in the product. This is the only one of the above papers that recognizes the consumers' information acquisition layer, which is the focus of our paper and the main driving force behind our results, but the approach to information acquisition adopted by Villas-Boas and Yao [2021] is very different from ours.[4]

From the perspective of information economics literature, we develop a model of

[4] Other work on targeting includes papers by Athey and Gans [2010], Bergemann and Bonatti [2011], Farahat and Bailey [2012], Chen and Stallaert [2014], and Anderson, Baik, and Larson [2019]. These papers explore (theoretically and in the field) the effects of technologies that enable targeting and of legislation which limits it, such as GDPR in Europe. They find multiple surprising results pertaining to the price and amounts of advertising, as well as its social value. However, none of these papers account for the possibility that consumers acquire information independently, which, as we show, can drive predictions to a significant degree.

costly disclosure, in which receivers can strategically acquire information from other sources. The costly disclosure part effectively follows the classical model of Verrecchia [1983]. On the other hand, literature on strategic information acquisition in various settings is vast, going back at least to the sequential sampling model of Wald [1945]. More recently, it has also been closely intertwined with literature on rational inattention, pioneered by Sims (see Sims [2010] and Maćkowiak, Matějka, and Wiederholt [2018] for recent surveys). We adopt this rational inattention framework to model the information acquisition process. Hébert and Woodford [2017] and Morris and Strack [2019] show that under certain conditions, the Wald problem can be represented as a static problem with information costs given by mutual information, as in our model.

Ichihashi [2018] exemplifies the importance of accounting for information acquisition in economic analysis, making a point similar to what our paper suggests. He explores a setting in which a monopolist sets a price for the product, and consumers can acquire a costly signal about their valuation of the product. He shows that the threat of information acquisition incentivizes the monopolist to set lower prices. Our results are similar in spirit, but relate to the monopolist's communication choices instead.

Among the papers on communication with rationally inattentive receivers, the most closely related to this paper are works by Matysková [2018] and Montes [2020], who characterize an optimal Bayesian Persuasion mechanism subject to the constraint that the receiver may engage in costly information acquisition after hearing the sender's message. In related work, Jerath and Ren [2019] explore the firm's incentives to limit the amount of information consumers can potentially acquire about the product. Bloedel and Segal [2018] study a general Bayesian Persuasion problem with a rationally inattentive receiver who must exert a cost to understand the sender's message. Jain and Whitmeyer [2020] look at a similar model with competing senders; they find that competition encourages disclosure. Our model is different from all of these in that our sender (the firm) is constrained in which messages it can send – i.e., we explore a model of disclosure, rather than Bayesian Persuasion.

Related to the two latter papers is a study by Branco, Sun, and Villas-Boas [2015]. They focus on the optimal amount of information that the firm would disclose to a consumer who must exert effort to filter through this information. In their model, the seller may want to provide more information to consumers with more favorable prior expected valuations of the product, contrary to the "naive" intuition that might suggest the opposite. We reconcile the two results, showing that value of providing a fixed amount of positive information can peak at consumers with much or little ex ante interest in the product, depending on the firm's belief in its own product.

# 4 The Model

The market consists of a single firm and a continuum of consumers $\mathcal{I}$ with unit demand each. The firm offers for sale a product of unknown quality $s \in S \equiv \{H, L\}$ at a price exogenously fixed at 1 and with zero marginal costs of production.[5] All consumers and the firm share a common prior belief $p^0 \in \Delta(S)$ about product quality. We will often abuse the notation and use $p^0$ to denote $p^0(H)$, the probability assigned by the prior to state $s = H$; the same abuse is also applied for interim and posterior beliefs, as defined later.

The value that consumer $i \in \mathcal{I}$ extracts from consuming a product of quality $s$ is given by $w_i + v_s$. The quality term $v_s$ is given by $v_L = 0$ for a low-quality product and $v_H = 1$ for a high-quality product. The quality term is state-dependent and thus is not observed by the consumer prior to purchase. Term $w_i$ is known to the consumer and represents the idiosyncratic component of consumer $i$'s valuation – a composite of all factors affecting valuation that are orthogonal to quality.[6] The firm can use the information it has about consumer $i$ to estimate $w_i$. This information can include demographics, location data, search history, and other kinds of data commonly available to advertisers thanks to tracking technologies. In this paper we simply assume that the firm perfectly observes $w_i$.

The firm has an opportunity to advertise its product. Specifically, at the beginning of the game it receives a private signal $y \in Y \equiv \{h, l\}$ with precision $\rho \in (1/2, 1)$, meaning $\mathbb{P}(h|H) = \mathbb{P}(l|L) = \rho$. The firm can verifiably disclose this signal in an advertisement, and it chooses which consumers $\mathcal{T}(y) \subseteq \mathcal{I}$ will receive the ad. We see $y$ as a piece of hard but inconclusive evidence that the firm can disclose

[5]The exact price level is irrelevant to the analysis. In particular, the results still apply if the firm selects both the price and the advertising strategy simultaneously – as long as the firm can not price discriminate across consumers. To shut down price signaling in this case, it is important that either the firm prices uniformly across quality levels, or consumers ignore the information contained in prices (which relates to the case of cursed consumers). Another case of special interest is that of digital products with zero price (where the firm's revenue comes from other sources, such as selling customer data), in which advertising remains the only choice variable for the firm.

[6]The additive functional form we adopt for consumer valuations clarifies the exposition but is not necessary for the results. One could easily obtain similar results in models in which consumers differ instead in their values for quality, or in their prior beliefs about product quality.

to selected consumers; this may include awards, rankings standings, critics' reviews, results of external quality evaluations, etc. Hereinafter we will refer to $\mathcal{T} : Y \to 2^{\mathcal{I}}$ as the firm's (pure) advertising strategy. The cost of such an advertising campaign depends on its size and is given by $c \cdot |\mathcal{T}(y)|$ for some per-consumer cost $c > 0$, where $|\mathcal{T}(y)|$ is the Lebesgue measure of set $\mathcal{T}(y)$. The consumers do not explicitly observe the firm's choice of strategy $\mathcal{T}(y)$. More generally, we can allow for mixed strategies $\tau : Y \to \Delta(2^{\mathcal{I}})$, which lead the firm to randomize between different sets of consumers given $y$. We will use $\tau(\mathcal{T}|y)$ to denote the probability with which consumer set $\mathcal{T}$ is targeted given $y$ and $\tau(y)$ to denote the whole probability measure.

Each consumer $i$ chooses whether to purchase the product or not. Prior to making the decision (but after receiving the ad, if any), the consumer has an opportunity to acquire a noisy signal $x_i \in \mathbb{R}$ about her valuation for the product. The consumer can choose any distribution of quality-contingent signals $\mathcal{G}_i(x_i|s) : S \to \Delta(\mathbb{R})$. However, generating a signal is costly, as described further. Upon observing $x_i$, each consumer updates her belief using Bayes' rule, and then decides whether to purchase the product so as to maximize her expected payoff.

While we assume that consumers use Bayes' rule to update their belief upon receiving any ad or private signal, we allow for one possible deviation from Bayesianism when no ad is received. We analyze a specific case of the model in which consumers are *cursed* in the sense of Eyster and Rabin [2005]. In our setting, being cursed means that a consumer does not update her belief if she receives no ad. This version provides a valuable insight into the model with fully *sophisticated* consumers, who update their beliefs accurately using Bayes' rule after all histories.

The overall timing of the model is as follows:

1. state $s$ is drawn by Nature, not observed by anyone;
2. the firm observes $y$, updates its belief from the prior $p^0$ to $p_f$, and chooses an advertising strategy $\tau(y)$;
3. every targeted consumer $i \in \mathcal{T}$ receives the ad and every $i \in \mathcal{I}$ updates her belief from the prior $p^0$ to the interim belief $p_i^1$;
4. every consumer $i \in \mathcal{I}$ selects an information acquisition strategy $\mathcal{G}_i(x_i|s)$;
5. every consumer $i \in \mathcal{I}$ observes signal $x_i$ and updates her belief to posterior $p_i^2$;
6. every consumer $i \in \mathcal{I}$ decides whether to purchase the product given $p_i^2$;
7. payoffs are realized.

The consumers' purchasing decisions at stage 6 are mechanical: consumer $i$ buys the product if and only if the expected utility from doing so is positive:

$$w_i + \mathbb{E}\left[v_s | p_i^2\right] - 1 \geq 0 \iff p_i^2 \geq 1 - w_i$$

(for concreteness, we break ties in favor of the firm). We henceforth take this purchasing strategy as given and focus on other strategic layers of the game.

It is convenient to use the beginning of stage 4 as a point of reference. Let $\mathcal{D}(p_i^1, w_i, s)$ denote the probability with which consumer $i$ purchases the item (conditional on her optimal information acquisition strategy) when the realized quality is $s$. Let $\mathcal{D}(p_i^1, w_i)$ denote the respective unconditional probability from the consumer's perspective as of stage 4:[7]

$$\mathcal{D}(p_i^1, w_i) \equiv \sum_{s \in S} p_i^1(s) \mathcal{D}(p_i^1, w_i, s).$$

Let $\mathcal{D}_f(p_i^1, w_i)$ represent the probability that consumer $i$ buys the product *as perceived by the firm*, conditional on its respective signal $y$ and the advertising strategy implemented,

$$\mathcal{D}_f(p_i^1, w_i) \equiv \sum_{s \in S} p_f(s) \mathcal{D}(p_i^1, w_i, s).$$

Given these probabilities, we can define the players' payoffs. In particular, consumer $i$'s expected utility as of stage 4 is given by

$$\sum_{s \in S} p_i^1(s) \left[\mathcal{D}(p_i^1, w_i, s) \cdot (w_i + v_s - 1) + (1 - \mathcal{D}(p_i^1, w_i, s)) \cdot 0\right] - \lambda\kappa(\mathcal{G}_i; p_i^1), \tag{1}$$

where $\lambda\kappa(\mathcal{G}_i; p_i^1)$ is the cost of generating the signal structure $\mathcal{G}_i$, proportional to the expected reduction in entropy between the consumer's interim and posterior

[7] From this point onwards we incorporate a part of equilibrium reasoning into our definitions. In particular, we assume that consumers' behavior from stage 4 onwards only depends on her reservation utility $w_i$ and the induced interim belief $p_i^1$, but not on her identity $i$. This incorporates the purchasing strategy we fixed for stage 6, but also requires that all consumers use the same belief updating rules on and off the equilibrium path. This assumption ensures that $w_i$ and $p_i^1$ serve as sufficient characteristics of consumer $i$ as of stage 4 and we do not need to keep careful track of consumers' identities when defining sale probabilities and payoffs.

beliefs:[8,9]

$$\kappa(\mathcal{G}_i; p_i^1) \equiv -\sum_s p_i^1(s)\log p_i^1(s) + \sum_s \sum_{x_i} \mathcal{G}_i(x_i) p_i^2(s|x_i) \log p_i^2(s|x_i),$$

with $\lambda \in \mathbb{R}_{++}$ being the information cost factor. In the above, $\mathcal{G}_i(x_i) \equiv \sum_s p_i^1(s)\mathcal{G}_i(x_i|s)$ is the unconditional probability of observing signal $x_i$.

The firm's expected profit conditional on $y$ and strategy $\mathcal{T}(y)$ is given by

$$\int_{i\in\mathcal{I}} \mathcal{D}_f(p_i^1, w_i)di - c \cdot |\mathcal{T}(y)|. \tag{2}$$

This expression estimates the expected profit as of stage 4 conditional on the consumers' interim beliefs $\{p_i^1\}_{i\in\mathcal{I}}$. However, in equilibrium, the firm knows that all consumers' prior beliefs are $p^0$ and knows exactly how these beliefs will react to ads or lack thereof, hence (2) is also a valid representation of the firm's expected profit as of stage 2.

We look for a Perfect Bayesian Equilibrium of the game which consists of the firm's mixed advertising strategy $\tau(y)$ and the collection of consumers' information acquisition strategies $\{\mathcal{G}_i(x_i|s)\}_{i\in\mathcal{I}}$ such that:

1. information acquisition strategy $\mathcal{G}_i(x_i|s)$ maximizes expected payoff (1) for all consumers $i \in \mathcal{I}$;
2. mixed advertising strategy $\tau(y)$ assigns positive weight only to $\mathcal{T}$ which maximize the firm's expected profit (2) given signal $y \in \{h, l\}$;[10]
3. the firm's belief is updated using Bayes' rule.

Recall that we have fixed the consumers' purchasing behavior at stage 6 as "$i$ buys iff $p_i^2 \geq 1 - w_i$", hence we do not include this in the equilibrium definition. Furthermore, we do not include the consumers' belief updating rules in the equilibrium definition, since, as mentioned previously, we explore different versions of the model in which consumers follow different updating rules. Finally, whenever the equilibrium advertising strategy $\tau(y)$ is degenerate, i.e., assigns probability one to some pure strategy $\mathcal{T}(y)$ for all $y$, we use $\mathcal{T}(y)$ to refer to the equilibrium strategy.

[8]It can be shown that the optimal $\mathcal{G}_i$ has finite support in our problem (see Matějka and McKay [2015]), hence we assume this from the start.

[9]For a detailed treatment of entropy and information theory, see Cover and Thomas [2012].

[10]Hereinafter we assume that if the firm is indifferent between advertising or not to consumer $i$, it does *not* advertise. Our results are not dependent on this restriction, up to indifference-specific formulations.

# 5 Analysis

## 5.1 The Consumers' Problem

We solve the model via backwards induction. As noted in the setup, in stage 6, consumer $i$ buys the product if and only if $p_i^2 \geq 1 - w_i$. Moving up in the timing, we now solve the consumers' information acquisition problem in stage 4. Because this problem is analogous to a problem explored by Matějka and McKay [2015] (see Problem 1 in Section F1 of their Online Appendix), we use their results to characterize the solution to the information acquisition problem in our model. This solution is described below; an interested reader is welcome to refer to their paper for more details.

The first step to solving the consumer's problem is realizing that, given some interim belief, every signal $x_i$ produced by the consumer's optimal strategy must induce a different action. If two signals $x_i$ induce the same action, the consumer can save on information costs by pooling the two signals into one. This insight was explored in detail by Caplin and Dean [2013]. In our setting, this means that the consumer's optimal signal structure $\mathcal{G}_i$ will generate at most two distinct signals – a recommendation to buy and a recommendation to pass.

One of the findings of Caplin and Dean [2013] is that the consumer's posterior $p_i^2$ after a given recommendation generated by the optimal signal structure $\mathcal{G}_i$ is independent of her interim belief $p_i^1$ (as long as she decides to acquire any information at all given $p_i^1$). In particular, if we define

$$\underline{p}_i \equiv \frac{e^{\frac{1-w_i}{\lambda}} - 1}{e^{\frac{1}{\lambda}} - 1} \qquad \text{and} \qquad \hat{p}_i \equiv \frac{1 - e^{-\frac{1-w_i}{\lambda}}}{1 - e^{-\frac{1}{\lambda}}} = e^{\frac{w_i}{\lambda}} \underline{p}_i, \tag{3}$$

then consumer $i$ buys the product without investigation whenever $p_i^1 \geq \hat{p}_i$, and passes on the product without investigation whenever $p_i^1 \leq \underline{p}_i$. If, however, $p_i^1 \in (\underline{p}_i, \hat{p}_i)$, then consumer $i$ investigates the product in such a way as to generate a binary recommendation, with a recommendation to buy bringing her belief up to $p_i^2 = \hat{p}_i$, and a recommendation to pass pushing her belief down to $p_i^2 = \underline{p}_i$. We can use the terminology introduced in Lindbeck and Weibull [2020]: when consumer $i$ passes on the product without investigation we say that she is in *sour* conditions, if she buys the product without investigation we say that she is in *sweet* conditions and when she decides to acquire some extra information we say that she is in *normal* conditions.

Note that cutoffs $\underline{p}_i$ and $\hat{p}_i$ vary across consumers, since they depend on idiosyncratic valuation $w_i$. This also means that we can invert the result: given some fixed belief $p_i$, there exist cutoffs $\hat{w}_i$ and $\underline{w}_i$ such that consumer $i$ with interim belief $p_i$ investigates the product if $w_i \in (\underline{w}_i, \hat{w}_i)$, passes if $w_i \leq \underline{w}_i$, and buys if $w_i \geq \hat{w}_i$. These cutoffs can be obtained by inverting the expressions in (3).

The most relevant for our purposes are probabilities $\mathcal{D}(p_i, w_i)$ generated by consumer $i$'s optimal information acquisition strategy. These represent the probability with which consumer $i$ expects to eventually buy the product, as evaluated at stage 4. These probabilities are given by

$$\mathcal{D}(p_i^1, w_i) = \begin{cases} 0, & \text{if } \hat{\mathcal{D}}(p_i^1, w_i) \leq 0; \\ \hat{\mathcal{D}}(p_i^1, w_i), & \text{if } 0 < \hat{\mathcal{D}}(p_i^1, w_i) < 1; \\ 1, & \text{if } 1 \leq \hat{\mathcal{D}}(p_i^1, w_i); \end{cases} \tag{4}$$

where
$$\hat{\mathcal{D}}(p_i^1, w_i) = \frac{p_i^1 \left(e^{\frac{1}{\lambda}} - 1\right) - \left(e^{\frac{1-w_i}{\lambda}} - 1\right)}{\left(e^{\frac{1}{\lambda}} - e^{\frac{1-w_i}{\lambda}}\right)\left(1 - e^{-\frac{1-w_i}{\lambda}}\right)}. \tag{5}$$

Now we can use Theorem 1 in Matějka and McKay [2015] to calculate the respective choice probabilities conditional on the true product quality $s$:

$$\mathcal{D}(p_i^1, w_i, H) = \frac{\mathcal{D}(p_i^1, w_i) e^{\frac{w_i}{\lambda}}}{\mathcal{D}(p_i^1, w_i) e^{\frac{w_i}{\lambda}} + (1 - \mathcal{D}(p_i^1, w_i))}, \tag{6}$$

$$\mathcal{D}(p_i^1, w_i, L) = \frac{\mathcal{D}(p_i^1, w_i) e^{-\frac{1-w_i}{\lambda}}}{\mathcal{D}(p_i^1, w_i) e^{-\frac{1-w_i}{\lambda}} + (1 - \mathcal{D}(p_i^1, w_i))}. \tag{7}$$

## 5.2 The Firm's Problem: Preliminaries

From this point onwards we explore the firm's decision in stage 2 of the game – namely, its choice of advertising strategy $\mathcal{T}$.

Given the consumers' information acquisition strategy, the firm can calculate its expected sales. Specifically, the firm can compute $\mathcal{D}_f(p_i^1, w_i)$, the expected probability that consumer $i$ will buy the product conditional on her interim belief $p_i^1$. This probability is given by

$$\mathcal{D}_f(p_i^1, w_i) = p_f \mathcal{D}(p_i^1, w_i, H) + (1 - p_f) \mathcal{D}(p_i^1, w_i, L), \tag{8}$$

where $p_f$ is the firm's belief which incorporates both the prior belief $p^0$ and the good private signal, i.e., $p_f = p^0\rho/\left(p^0\rho + (1-p^0)(1-\rho)\right)$. The question now is how the consumer's interim belief $p_i^1$ responds to the firm's advertisement or lack thereof.

The firm can only advertise using hard information in our model, meaning that given private signal $y$, the firm can disclose $y$ to consumers in an ad or stay silent, but cannot modify the ad or send any other messages. Therefore, consumer $i$'s interim belief upon receiving an ad with signal $h$ is

$$p_i^1(h) = \alpha(p^0) \equiv \frac{p^0\rho}{p^0\rho + (1-p^0)(1-\rho)} > p^0, \tag{9}$$

which coincides with $p_f$, while upon hearing an ad with signal $l$, the belief is

$$p_i^1(l) = \beta(p^0) \equiv \frac{p^0(1-\rho)}{p^0(1-\rho) + (1-p^0)\rho} < p^0. \tag{10}$$

Since advertising with signal $l$ is costly and depresses the consumer's belief, it is never optimal for the firm to advertise with $l$.[11] This is true as long as consumers' beliefs react at least somewhat rationally to the lack of advertising, in the sense of beliefs not dropping further than they would have after observing a low signal.

**Lemma 1.** *In any equilibrium, if in the absence of an ad consumer $i$'s interim belief is $p_i^1 \geq \beta(p^0)$, then $i \notin \mathcal{T}(l)$ for all $\mathcal{T}(l)$ in the support of the equilibrium strategy $\tau(l)$.*

*Proof.* Sale probability $\mathcal{D}_f(p_i^1, w_i)$ is weakly increasing in $p_i^1$, which follows immediately from (4)-(8). Hence, sending an ad with $y = l$ to such a consumer $i$ would decrease the sale probability compared to not advertising, while costing $c > 0$ to the firm. Therefore, not advertising to $i$ is optimal. □

The condition in the lemma holds for all consumers in all cases that we consider. Hence from this point onwards we focus on the problem of a firm with a high private signal $y = h$, implicitly assuming that the firm with signal $l$ never advertises ($\mathcal{T}(l) = \emptyset$).

## 5.3 The Firm's Problem with Cursed Consumers

We begin by solving a version of the problem with *cursed* consumers who react to advertisements, but not to the lack thereof. In other words, if consumer $i$ is cursed

[11] Although disclosing adverse information may, more generally, be optimal in richer settings; see Smirnov and Starkov [2020] for one example and a survey of related papers.

and she receives an ad with signal $h$ or $l$ then her belief changes to $\alpha(p^0)$ or $\beta(p^0)$ respectively, but if she hears nothing then her belief remains at $p^0$, regardless of the firm's equilibrium strategy (even though a sophisticated consumer would infer from the firm's strategy that silence is suggestive of bad news). The fact that consumers are cursed is common knowledge. The next subsection demonstrates the connection of this problem to the problem with fully sophisticated consumers.

The condition in Lemma 1 holds trivially for all $i \in \mathcal{I}$ since $p^0 > \beta(p^0)$, thus $\mathcal{T}(l) = \emptyset$ and we focus on the high signal $y = h$. The probability that (cursed) consumer $i$ purchases the product is given by $\mathcal{D}_f(p_i^1, w_i) = \mathcal{D}_f(\alpha(p^0), w_i)$ after ad $h$ and $\mathcal{D}_f(p_i^1, w_i) = \mathcal{D}_f(p^0, w_i)$ in the absence of any ad. The firm's expected profit from targeting set $\mathcal{T}(h)$ of consumers with ad $h$ is given by (2), which can then be rewritten as

$$\left[\int_{i\in\mathcal{T}(h)} \mathcal{D}_f(\alpha(p^0), w_i)di + \int_{i\in\mathcal{I}\setminus\mathcal{T}(h)} \mathcal{D}_f(p^0, w_i)di\right] - c \cdot |\mathcal{T}(h)|.$$

Therefore, if we let

$$\mathcal{A}(w_i) \equiv \mathcal{D}_f(\alpha(p^0), w_i) - \mathcal{D}_f(p^0, w_i)$$

denote the effect of the ad on demand as a function of $w_i$, then the firm will find it optimal to target consumer $i \in \mathcal{I}$ if and only if $\mathcal{A}(w_i) > c$. Notably, the resulting strategy $\mathcal{T}_{cursed}(h) \equiv \{i|\mathcal{A}(w_i) > c\}$ is uniquely optimal given consumers' beliefs, hence the firm will never want to use mixed strategies. Furthermore, beliefs of cursed consumers are independent of the firm's strategy (conditional on a given observation), hence the equilibrium is also unique.[12]

The linearity of ad campaign cost w.r.t. $|\mathcal{T}|$ is the driver of separability obtained above: i.e., of the fact that the decision to advertise to consumer $i$ only depends on $w_i$, but not on the whole distribution of $w$ in the population or on other advertising decisions. This also implies that instead of choosing which consumers $\mathcal{T}(y) \subseteq \mathcal{I}$ to target with signal $y$, the firm can equivalently choose which consumer *types* $\mathcal{W}(y) \equiv \{w_i|i \in \mathcal{T}(y)\}$ to target given its private signal $y$. Hereinafter, we refer to $\mathcal{W}$ as the type-representation of a given pure targeting strategy $\mathcal{T}$.

This subsection is summarized by (and proves) the following Proposition.

**Proposition 1.** *If consumers are cursed, the firm's optimal advertising strategy in the unique equilibrium is pure and given by* $\mathcal{W}(l) = \emptyset$ *and* $\mathcal{W}(h) = \{w_i|\mathcal{A}(w_i) > c\}$.

[12] Unique up to indifference – i.e., other equilibria may exist in which the firm or some consumer chooses a different action when indifferent between some actions available to them (advertise or not to consumer $i$; buy the product or not).

In what follows, we use $\mathcal{T}_{cursed}$ to refer to the firm's advertising strategy in the equilibrium with cursed consumers, and $\mathcal{W}_{cursed}$ for its type-representation. The following subsection demonstrates that the same advertising strategy is optimal when consumers are sophisticated. Section 5.5 then demonstrates that this optimal strategy has the bipartite structure described in the Introduction.

## 5.4 The Firm's Problem with Sophisticated Consumers

We now expand the analysis above to allow for sophisticated (fully Bayesian) consumers. Sophisticated consumer $i$ would still update her belief to $\alpha(p^0)$ or $\beta(p^0)$ after an ad with signal $h$ or $l$ respectively. However, in case she does not receive an ad, she would make an inference from that fact based on the firm's equilibrium advertising strategy, potentially implying $p_i^1 \neq p^0$. Notably, $p_i^1 \in [\beta(p^0), \alpha(p^0)]$ in the absence of an ad for all $i$, so Lemma 1 applies, and in all equilibria with sophisticated consumers $\mathcal{T}(l) = \emptyset$ is uniquely optimal.

We now show that the optimal advertising strategy in this case is the same as with cursed consumers (in one equilibrium). The intuition behind this lies in the exact way cursed and sophisticated consumers differ in our problem. Specifically, if consumer $i$ does not expect to receive an ad, then it does not matter to the firm whether she is cursed or sophisticated, since in the absence of an ad her belief remains at $p_i^1 = p^0$ in either case. Therefore, if advertising to $i$ is not optimal if $i$ is cursed, it is also not optimal when she is sophisticated, since both the cost and the benefit of an ad are the same. Conversely, if consumer $i$ does expect to receive an ad and it was optimal to advertise to her when she was cursed, then it is surely optimal to advertise if she is sophisticated. This is because an ad improves her belief to $\alpha(p^0)$ in either case, but her reaction to the absence of an ad is weakly worse when she is sophisticated ($p_i^1 \leq p^0$).

Therefore, when a firm is developing its advertising strategy, it is a safe bet for it to assume that all consumers are cursed and unresponsive to a lack of advertising – regardless of whether this is actually true, or consumers are sophisticated but correctly anticipate the firm's strategy. This idea is formalized in our first theorem below.

**Theorem 1.** *There exists an equilibrium of the game with sophisticated consumers in which the firm plays a pure strategy $\mathcal{T} = \mathcal{T}_{cursed}$.*

*Proof.* According to Proposition 1, $\mathcal{T}_{cursed}$ is such that $\mathcal{W}_{cursed}(l) = \emptyset$ and $\mathcal{W}_{cursed}(h) = \{w_i | \mathcal{A}(w_i) > c\}$. Optimality of $\mathcal{W}(l) = \emptyset$ is given by Lemma 1.

Now consider $y = h$. If consumer $i$ is sophisticated and $w_i \in \mathcal{W}(h)$, then upon not receiving an ad she realizes that this is due to the firm's signal being $y = l$, so she updates her belief to $p_i^1 = \beta(p^0)$. Then the actual effect of an ad with $y = h$ in terms of firm-expected increase in the probability of a sale is given by

$$\hat{\mathcal{A}}(w_i) \equiv \mathcal{D}_f(\alpha(p^0), w_i) - \mathcal{D}_f(\beta(p^0), w_i), \tag{11}$$

which is weakly greater than $\mathcal{A}(w_i)$ for all $w_i$, since $\mathcal{D}_f(\beta(p^0), w_i) \leq \mathcal{D}_f(p^0, w_i)$ (see Lemma 2 in the Appendix). Hence if advertising to consumer $i$ was optimal when she was cursed, it is still optimal to advertise if she is sophisticated. Conversely, if $w_i \notin \mathcal{W}(h)$ then the ad effect is still given by $\mathcal{A}(w_i) \leq c$ (the consumer does not expect to receive an ad, hence does not update her belief after not receiving one), advertising to this consumer is not optimal. □

Note that this result is specific to costly disclosure settings, but does not in any way rely on the consumers' information acquisition layer of the problem. In other words, it applies equally well if the cost of information is $\lambda = +\infty$, but loses its bite (while remaining formally true) when the cost of advertising is $c = 0$, since in the latter case the firm has no reason to ever withhold a positive signal from any consumer.

Hereinafter we refer to the equilibrium in Theorem 1 as the *cursed equilibrium* of the game with sophisticated consumers. The following subsection illustrates and explores this equilibrium in greater detail. Section 5.6 then studies other equilibria of the game.

## 5.5 The Cursed Equilibrium Characterization

In this section we analyze the cursed equilibrium of the game with sophisticated consumers (on-path equivalent to the unique equilibrium with cursed consumers) and state our main result. However, before we can do that, additional notation and definitions are required. First we introduce two threshold consumer types:

$$\bar{w} \equiv \min\{w \mid \mathcal{D}_f(\alpha(p^0), w) = 1\},$$
$$\underline{w} \equiv \max\{w \mid \mathcal{D}_f(p^0, w) = 0\}.$$

Here $\underline{w}$ denotes the type of consumer, who, when her interim belief equals $p^0$, is indifferent between passing on the product and investigating it (i.e., she is at the border between sour and normal conditions; $\underline{w}$ here is analogous to $\underline{w}_i$ in Section 5.1). On the other hand, $\bar{w}$ denotes the type of consumer, who, after receiving a good ad, is indifferent between buying the product immediately and investigating it (i.e., she is at the border between normal and sweet conditions). Both thresholds are well defined since $\mathcal{D}_f$ is weakly increasing and continuous in $w_i$ (see Lemma 2 in the Appendix), and $\mathcal{D}_f(p,1) = 0$, $\mathcal{D}_f(p,0) = 1$ for any $p \in (0,1)$. We will consider two cases depending on the relation between these two thresholds.

**Definition.** *Advertisement $\rho$ is* strong *if there exists $w$ such that $\mathcal{D}_f(p^0, w) = 0$ and $\mathcal{D}_f(\alpha(p^0), w) = 1$, and* weak *otherwise.*

Equivalently, an ad is strong if and only if $\underline{w} \geq \bar{w}$, meaning if and only if there exists some consumer type $w \in [\bar{w}, \underline{w}]$ that faces sour conditions without an ad, and would enter sweet conditions after an ad. Another way to frame this division is in terms of a bound on $\rho$: an ad is strong if and only if $\rho \geq \bar{\rho}$ for some cutoff value $\bar{\rho} > 1/2$, which depends on model parameters.[13]

We now state the main result: our claim that the optimal advertising strategy in the cursed equilibrium has a bipartite structure, targeting groups of consumers centered at $\underline{w}$ and $\bar{w}$.

**Theorem 2.** *If the cursed equilibrium prescribes a non-zero level of advertising ($\mathcal{T}_{cursed} \neq \emptyset$), then the optimal targeting strategy satisfies the following for some $\underline{w}_d, \underline{w}_u, \bar{w}_d, \bar{w}_u$:*

1. *$\mathcal{W}_{cursed}(l) = \emptyset$;*
2. *if the ad is weak: $\mathcal{W}_{cursed}(h) = (\underline{w}_d, \underline{w}_u) \cup (\bar{w}_d, \bar{w}_u)$, with $\underline{w}_d \leq \underline{w} \leq \underline{w}_u, \bar{w}_d \leq \bar{w} \leq \bar{w}_u$, and with one of the two intervals possibly empty;*
3. *if the ad is strong: $\mathcal{W}_{cursed}(h) = (\bar{w}_d, \underline{w}_u)$, with $\bar{w}_d \leq \bar{w} \leq \underline{w} \leq \underline{w}_u$.*

*Proof.* See Appendix. □

In the case of a weak ad, Theorem 2 says that in the cursed equilibrium the firm targets some group of consumers close to the border between sour and normal conditions ($w_i \approx \underline{w}$), and/or some other group close to the border between normal and sweet conditions ($w_i \approx \bar{w}$). If the ad is weak but the advertising cost is low, then the two groups may merge into one – then $\mathcal{W}_{cursed} = (\underline{w}_d, \bar{w}_u)$, and it is not possible to pin down cutoffs $\underline{w}_u = \bar{w}_d$.

[13] The equivalence of this formulation is easy to see given that $\underline{w}$ does not depend on $\rho$, while $\bar{w}$ is decreasing in $\rho$, and for $\rho = 1/2$ we have $\bar{w} > \underline{w}$.

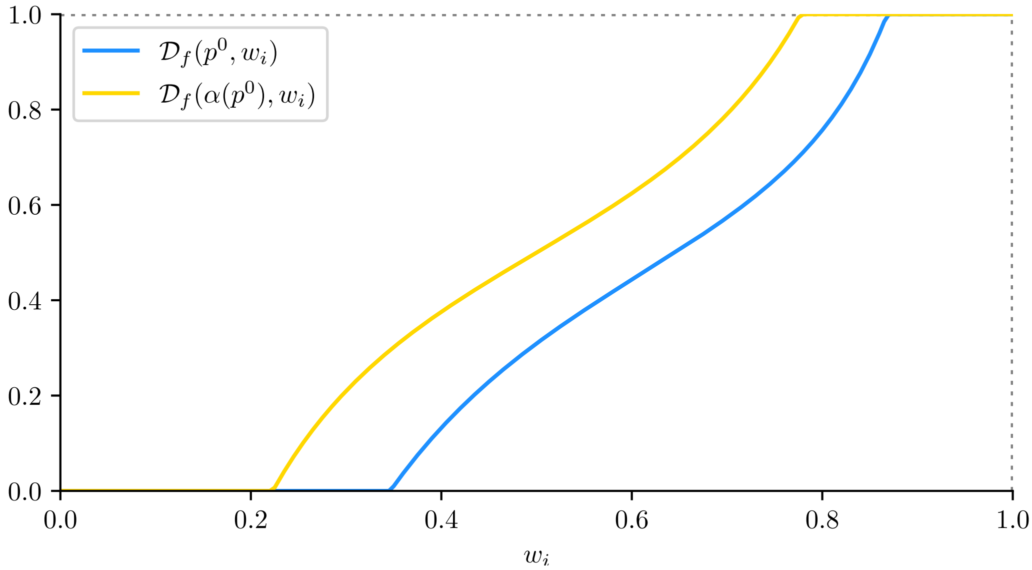


(a) Sale probabilities before and after an ad, as expected by the firm.

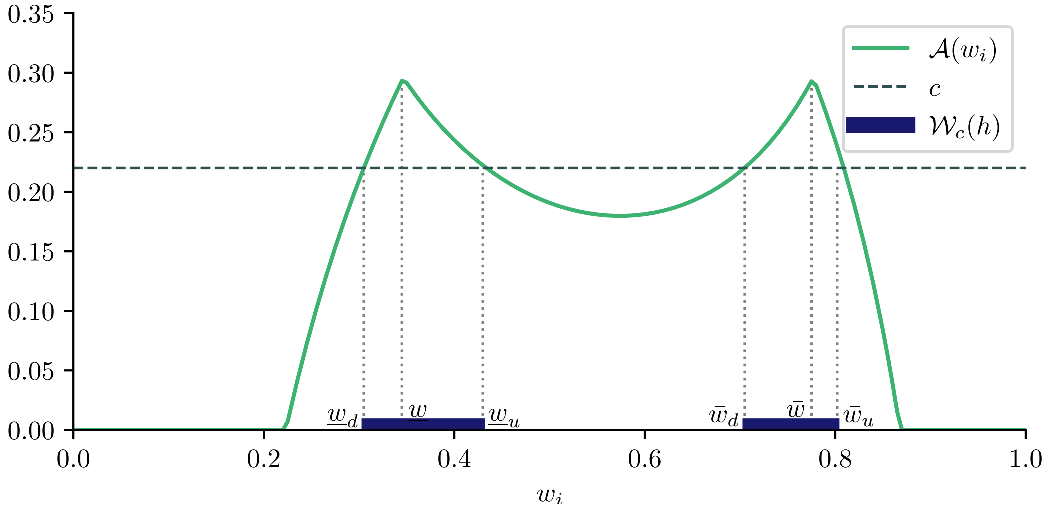


(b) Ad effect $\mathcal{A}(w_i)$ and the optimal advertising strategy $\mathcal{W}(h)$.

Figure 1: The cursed equilibrium ($\lambda = 0.35$, $p^0 = 1/3$, $\rho = 2/3$, and $c = 0.22$).

*Notes*: Panel (a) depicts the firm's equilibrium estimates $\mathcal{D}_f$ of the probabilities with which cursed consumers will buy the product without receiving an ad and after receiving an ad, depending on their types $w_i$. The green line in panel (b) shows the difference between the two probabilities – the net effect of an ad, $\mathcal{A}(w_i)$. Presented in dark blue in panel (b) are $\mathcal{W}_c(h)$: the sets of consumers, as characterized by $w_i$, whom the firm targets with its ad in equilibrium.

Figure 1 illustrates the result: it plots the cursed equilibrium for $\lambda = 0.35$, $p^0 = 1/3$, $\rho = 2/3$, and $c = 0.22$. Panel (a) depicts the expected sale probabilities with and without an ad, $\mathcal{D}_f(\alpha(p_i^0), w_i)$ and $\mathcal{D}_f(p_i^0, w_i)$ respectively, as functions of consumers' reservation utility $w_i$. Panel (b) shows the effect of ads on expected sales as a function of $w_i$, as well as the equilibrium advertising strategy.[14] As panel (b) shows, the ad effect is double-peaked in $w_i$, which results in a bipartite optimal advertising strategy, depending on $c$. Theorem 2 claims that this twin-peak structure holds in general.

The right peak in Figure 1b is at $w_i = \bar{w}$, i.e., the consumer who is at the border between normal and sweet conditions after receiving the ad. By advertising to this consumer and those with slightly larger $w_i$, the firm eliminates any risk that these consumers will investigate the product on their own and arrive at the (possibly incorrect) conclusion that the product is bad. The possible sale is converted into a sure sale – or an almost sure sale in the case of consumers with $w_i$ slightly below $\bar{w}$. On the other hand, the left peak is at $w_i = \underline{w}$ – the consumer who is at the border between sour and normal conditions. By advertising to this consumer, the firm does not automatically generate a sale, but rather sparks the consumer's interest and leads her to investigate the product, as opposed to simply walking past it.

It is quite surprising that the firm prefers to primarily target two disjointed groups of consumers, while consumers who investigate the product regardless are less profitable to advertise to. This is partly due to the firm being unsure of the quality of its own product (the parameters in Figure 1 are such that $p_f = \frac{1}{2}$). We do not frame this statement as a formal result, but instead illustrate the intuition behind it using Figure 2.

Specifically, if the firm's belief (equal to the common prior $p^0$ adjusted for good news $y = h$) dictates that product is likely good or likely bad, then its targeting strategy $\mathcal{W}_{cursed}$ focuses more on one of the two groups. The confident firm (with a high $p_f$) prefers to target consumers who would otherwise either not look at the product (low $w_i$), or those who would not put much effort into investigating its quality. Such a firm realizes that any consumer that investigates its product carefully enough will realize that it is worth buying. Hence the firm tries to induce the consumers to investigate its product carefully. A firm that is not confident (has low $p_f$) has the opposite incentives: it would prefer to discourage information acquisi-

[14]Recall that $\mathcal{A}(w_i)$ measures the ad effect for cursed consumers and those sophisticated consumers who do not expect an ad. The actual effect of the ad on sophisticated consumers who expect it is greater: see Sections 5.4 and 5.6.

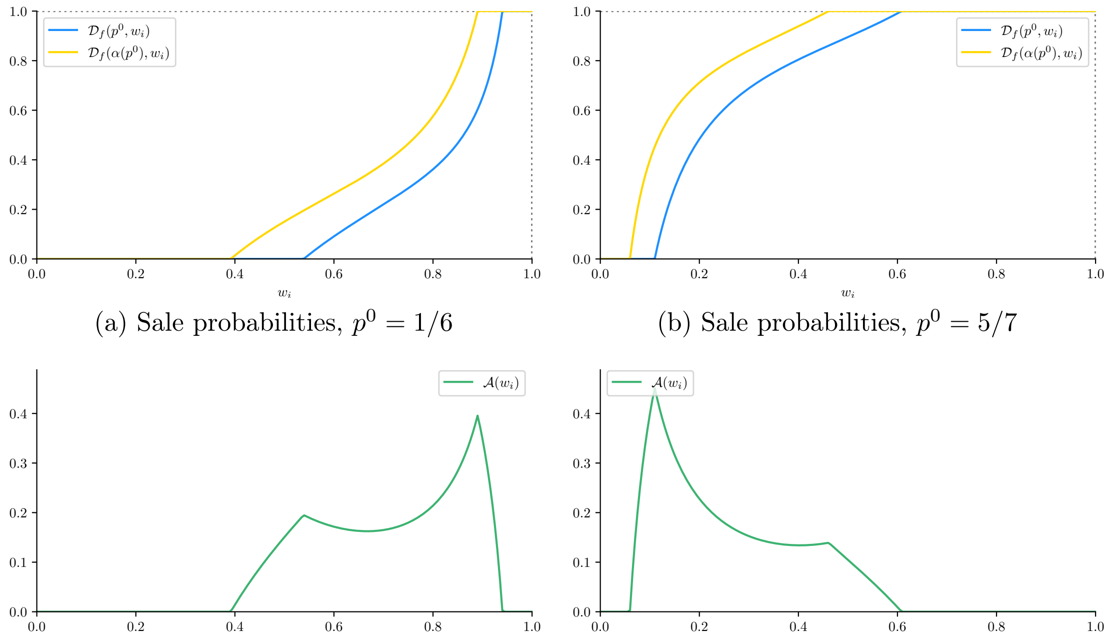


(a) Sale probabilities, $p^0 = 1/6$

(b) Sale probabilities, $p^0 = 5/7$

(c) Ad effect, $p^0 = 1/6$

(d) Ad effect, $p^0 = 5/7$

Figure 2: Effects of prior $p^0$ ($\lambda = 0.35$, $\rho = 2/3$).

*Notes*: The figure presents the expected demands and the ad effects when the common prior is pessimistic (panels (a) and (c)) and optimistic (panels (b) and (d)).

tion as much as possible, since it knows that by looking closely at its product the consumers will mostly be encouraged from buying it. Therefore, it primarily targets the consumers for whom the ad can tip the scales in favor of buying the product without further investigation (high $w_i$). A firm that is uncertain about the quality of its own product thus prefers to hedge and diversify its targeting across the two options above.[15]

Finally, even for a firm which is uncertain regarding the quality of its product, the double-peaked strategy only arises if the amount $\rho$ of information contained in its ad is low enough relative to the information acquisition/contemplation cost factor $\lambda$ – i.e., if the ad is weak, which is the case explored so far. In contrast, if the ad is strong ($\rho$ is high) then there exists a group of consumers whom the firm can fully convert – turn from sour to sweet conditions, completely avoiding the normal phase. This group would trivially be the most profitable for the firm to advertise

[15]Figure 2 presents the effects of changing $p_0$, which is both the firm's and the consumers' prior belief. However, it can be easily verified using a variation of our model with a non-common prior that the shape of $\mathcal{A}(w)$ is indeed determined by the firm's belief $p_f$, while a change in the consumers' prior would simply shift the demand functions and the ad effect function along the $w$ axis.

to, as compared to consumers who investigate the product – and thus buy it with probability strictly between zero and one, – in one of the two cases (with or without the ad). In other words, the two consumer groups described above merge into one group in this scenario.

## 5.6 Other Equilibria with Sophisticated Consumers

The cursed equilibrium is not a unique equilibrium in our model when consumers are sophisticated. The multiplicity stems from the fact that equilibria are, to some extent, self-reinforcing: the benefit from advertising to some consumer $i$ is larger when she expects to receive an ad (from a firm with signal $h$) than when she does not. This is because in the former case sending the ad increases her belief from $\beta(p)$ to $\alpha(p)$ – since by not receiving the ad she infers that $y = l$, – while in the latter case the increase is from $p > \beta(p)$ to $\alpha(p)$.

However, this also implies that the cursed equilibrium features the least amount of advertising among all equilibria, as formalized by the following proposition.

**Proposition 2.** *In any equilibrium with sophisticated consumers, for any $y \in \{h, l\}$: any $\mathcal{T}$ in the support of the firm's equilibrium strategy $\tau(y)$ is such that $\mathcal{T}_{cursed}(y) \subseteq \mathcal{T}$.*

*Proof.* By Lemma 1 and the consumers' belief updating rule, $\tau(\emptyset|l) = 1$, hence for the remainder of the proof, consider $y = h$. Consider an arbitrary equilibrium of the game. Advertising to sophisticated consumer $i$ increases the probability of a sale by an amount between $\hat{\mathcal{A}}(w_i)$ and $\mathcal{A}(w_i)$, since her belief $p_i^1$ after no ad is $p_i^1 \in [\beta(p^0), p^0]$. The consumer is targeted if and only if the ad benefit is strictly larger than $c$. It is true that $\hat{\mathcal{A}}(w_i) \geq \mathcal{A}(w_i)$ (see proof of Theorem 1), and in the cursed equilibrium, a consumer with prior $p$ is targeted if and only if $\mathcal{A}(p) > c$, hence if the consumer is targeted in the cursed equilibrium, she is also targeted in the equilibrium under consideration. □

The proposition above validates the cursed equilibrium as the one least tainted by advertising. This is on top of its appeal as being the unique equilibrium – up to the firm's indifference – when the consumers are cursed, which is an assumption that is empirically appealing by itself in this setting (see footnote 3 for references to works providing empirical and experimental evidence supporting receivers' cursedness in disclosure games). However, the main idea of our result applies equally well to other equilibria with sophisticated consumers. For example, the following proposition

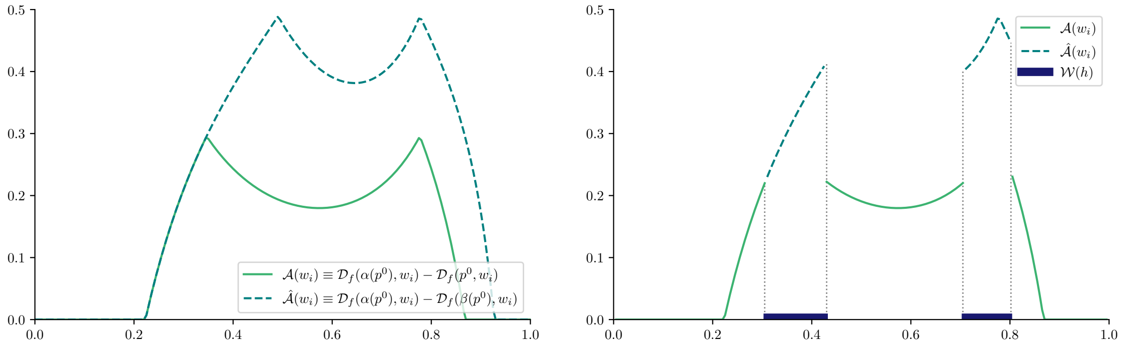


Figure 3: Ad effects $\mathcal{A}(w_i)$ and $\hat{\mathcal{A}}(w_i)$ for cursed and unexpecting sophisticated vs expecting sophisticated consumers, respectively (left panel); the actual ad effect in the cursed equilibrium (right panel).

demonstrates that the equilibrium with the most advertising (exists and) has the exact same structure as the cursed equilibrium.

**Proposition 3.** *If consumers are sophisticated, there exists an equilibrium with the firm's advertising strategy $\bar{\mathcal{T}}$ such that for any $y \in \{h, l\}$ and any $\mathcal{T}$ in the support of the firm's strategy $\tau(y)$ played in any other equilibrium, $\mathcal{T} \subseteq \bar{\mathcal{T}}(y)$.*

*Furthermore, strategy $\bar{\mathcal{T}}$ is such that its type-representation $\bar{\mathcal{W}}(y) = \{w_i | i \in \bar{\mathcal{T}}(y)\}$ satisfies properties 1-3 in Theorem 2.*

*Proof.* See Appendix. □

Both Propositions 2 and 3 can be illustrated using Figure 3. The left panel plots the two values of ad effects, depending on the type of consumer $i$: $\mathcal{A}(w_i)$ applies when $i$ is either cursed, or is sophisticated but does not expect to receive any ad; $\hat{\mathcal{A}}(w_i)$ as defined in (11) is the ad effect for consumer $i$ who expects to receive an ad in equilibrium. In other words, in equilibrium with sophisticated consumers, the ad effect is given by $\hat{\mathcal{A}}(w_i)$ for all $i \in \mathcal{T}(h)$ and by $\mathcal{A}(w_i)$ for all $i \in \mathcal{I} \backslash \mathcal{T}(h)$. So targeting strategy $\mathcal{T}(h)$ is an equilibrium with sophisticated consumers if and only if $\hat{\mathcal{A}}(w_i) \geq c$ for all $i \in \mathcal{T}(h)$ and $\mathcal{A}(w_i) < c$ for all $i \notin \mathcal{T}(h)$.

The equilibrium multiplicity is produced by the fact that $\hat{\mathcal{A}}(w) \geq \mathcal{A}(w)$ for all $w$, which can lead to quite arbitrary targeting sets being self-sustaining in equilibrium. However, the point of Proposition 3 and Figure 3 is that $\hat{\mathcal{A}}(w)$ satisfies all the same properties that $\mathcal{A}(w)$ does, and yields (in the equilibrium with the largest $\mathcal{T}(h)$) qualitatively the same targeting strategy as in the cursed equilibrium.

In the end, while we cannot guarantee that the optimal advertising strategy in any equilibrium has the neat bipartite structure outlined in Theorem 2, the optimal target groups must belong to one of the two broad pools defined there (one group close to the border between normal and sweet conditions and the other group close to the border between sour and normal conditions). The only caveat is that the benefit of advertising to some of these consumers is only outweighed by the implicit threat of that they may grow skeptical of the product quality in case they do not see it advertised. The existence of such a threat in the context of advertising is debatable, but may be more justified in other settings.

## 5.7 Consumer Welfare and the Amount of Advertising

Our analysis so far has been focused around the implications of targeting strategies for the firm's profit. This section discusses instead the effects on consumers.

The first question to ask is: do ads benefit consumers? The answer our model suggests is "yes". For sophisticated consumers, being in the target group is equivalent to receiving a free signal: if $y = h$ they receive an ad, and if $y = l$ they can perfectly infer this from the lack of an ad. Thus the sophisticated consumers in the target group can effectively observe the realization of the firm's signal $y$. Free information is always beneficial for decision-makers, hence sophisticated consumers benefit from being targeted with advertisements.

For cursed consumers the argument above does not work. Nevertheless, we can show that they also benefit from advertisements via the following simple coupling argument. If $y = h$ then targeted cursed consumers receive a piece of accurate information and make a decision that is better on average, the same as sophisticated consumers would make. If, on the other hand, $y = l$, then cursed consumers are left with their prior belief $p^0$ and, hence, make the same decision as if they had not been targeted. Thus cursed consumers win from being targeted if $y = h$ and lose nothing if $y = l$, hence they win on average.

**Proposition 4.** *Suppose $\mathcal{T}(l) = \emptyset$. Consumer $i$'s ex ante expected utility is larger if $i \in \mathcal{T}(h)$ than if $i \notin \mathcal{T}(h)$, regardless of whether she is cursed or sophisticated.*

*Proof.* See Appendix. □

This proposition suggests that the equilibrium with the largest amount of advertising described in Proposition 3 is the one most preferred by sophisticated consumers. This result is at odds with the popularity of ad-blockers among Internet users in the real world.[16] A possible explanation for this discrepancy is the partial equilibrium nature of our model, which ignores the firm's pricing decisions and the effect that advertising may have on them. The results presented in this subsection should thus be viewed critically.

Proposition 4 implies that we can measure average consumer welfare in our model via $|\mathcal{T}(y)|$ – the number of consumers who receive ads. The more advertising there is in equilibrium, the happier consumers are in aggregate. We can use this observation to compute the welfare implications of targeting technologies.

Consider a benchmark scenario in which the firm has no power to target its ads and the consumers are cursed, for simplicity. In other words, the firm can only choose its campaign size (breadth) $b \in [0, 1]$, which defines the share of consumers who will receive an ad, but the identities of the recipients are pulled uniformly from $\mathcal{I}$ (without replacement). Let $F(w_i)$ denote the c.d.f. of the distribution of consumers' preferences $w_i$ in population $\mathcal{I}$. The firm's problem can then be written as

$$\max_{b \in [0,1]} \left\{ b \int_0^1 \mathcal{A}(w_i) dF(w_i) - bc \right\}. \tag{12}$$

Since the campaign cost is assumed to be linear in $b$, the firm's optimal choice is $b \in \{0, 1\}$ – either advertise to all consumers if the average ad benefit is large enough, or to no one. In the former case, targeting would lead to less advertising overall (the firm will serve targeted ads to some but not all consumers), which harms consumers who miss out on advertisements. Conversely, if in the absence of targeting, the firm chooses not to advertise at all, then targeting opportunities will increase the amount of advertising and benefit the consumers. It is clear from (12) that the latter case applies whenever $c$ is high enough, meaning that the firm's targeting capabilities are beneficial for consumers whenever non-targeted advertising is too expensive.[17]

# 6 Political Advertising and Competition

Our model can also be used to describe targeted political advertising, which appears to be a growing industry (see footnote 1). In this framework, our firm would become a candidate who is using costly communication channels to win voters' al-

[16] According to eMarketer, in 2019 about a quarter of all Internet users in the US and Europe were using some kind of ad-blocking software. (`https://www.emarketer.com/content/consumer-attitudes-on-marketing-2019`)

[17] This argument assumes that the cost of advertising $c$ is the same for targeted and non-targeted ads; the assumption can easily be relaxed.

legiance, and is selecting the optimal voter group to target. One example of such behavior would be the U.S. presidential candidates committing to different levels of promotional efforts in states with different levels of ex ante support.

There are two features missing from our model that are more prominent in such a political setting than in firm-consumer interactions. Firstly, while a firm may be a monopolist in its given market, elections are inherently competitive in that a candidate can advertise him/herself, but a competitor can do the same. Secondly, in the presence of competition, advertising can not only promote one's self (or product) but can also focus on tanking the opponent. In this section we show that our results are robust to these aspects of the environment.

Suppose there are two candidates running for office, $j \in \{1, 2\}$, referred to as C1 and C2. The state of the world is still binary: $s \in \{H, L\}$, and there is some common prior belief $p^0$. Voter $i$'s utility from electing C1 is given by $u_1 \equiv w_i + v_{1,s}$, while that for C2 is $u_2 \equiv 1 - w_i + v_{2,s}$, where $v_{1,H} = 1$, $v_{1,L} = 0$, and $v_{2,s} = 1 - v_{1,s}$. Here $v_j$ represent candidates' "objective fitness" for the office, and these types are perfectly negatively correlated. In other words, in state $s = H$ C1 is more fit for the office, while in state $s = L$ C2 is more fit. On the other hand, $w_i$ is the degree to which voter $i$ is originally leaning towards C1: if $w_i \geq 1$ then $i$ always votes for C1, while if $w_i \leq 0$ then $i$ always votes for C2. Both candidates observe a realization of some common signal $y \in \{h, l\}$ about $s$ and simultaneously select advertising strategies $\mathcal{T}_j(y)$. Voter $i$ then gets to observe the realization of $y$ if $i \in \mathcal{T}_1(y) \cup \mathcal{T}_2(y)$. Advertising costs per voter, $c$, are the same for both candidates. The remainder of the model is as in Section 4.

The first thing to note in this version is that there is no real distinction between "disclosing good news about self" and "disclosing bad news about the opponent". All that matters is that signal $y$ shifts public belief in the direction favorable for the candidate, but the interpretation of the contents of that signal is not important for us.[18] The second thing to note is that, given some interim belief $p_i^1$, the voter's problem in this model is fully equivalent to the consumer's problem explored in Section 5.1. This is because the net utility of choosing C1 over C2 is $u_1 - u_2 = 2(w_i - 1 + v_{1,s})$, exactly the same as the net utility of buying the item in Section 4.

Using the argument in Lemma 1 we can claim that C1 will only advertise if $y = h$, while C2 will only advertise if $y = l$. Then when $y = h$, the targeting problem of

[18]This equivalence relies on the uncertainty being one-dimensional. Self-promotion and tanking the competition could be meaningfully different in a model where, e.g., each candidate has an independent type $s_j$, or when there are more than two candidates.

C1 is exactly the same as the firm's problem we explored in Section 5, and our results still apply. It is trivial to verify that the results apply to C2's strategy as well, since the model is symmetrical. Therefore, competition does not have any real effect on the conclusions of our model: the optimal strategy for each candidate will still, in general, involve targeting two groups of voters, as per Theorem 2, and political advertising is beneficial for the voters as per Proposition 4. All other results continue to hold as well.[19]

# 7 Discussion

We explore the optimal ad targeting strategy of a monopolistic firm which is facing a population of rationally inattentive consumers with heterogeneous outside options. We show that this strategy is generally bimodal, focusing on two distinct groups of consumers: (i) those who are relatively optimistic about the product and are close to buying it without information acquisition and (ii) those who are relatively pessimistic about the product and are close to acquiring some information about the product. The relative focus on these two groups depends on the firm's expectation regarding its own product.

In our analysis we make a connection to behavioral literature studying "cursed" consumers, who deviate from being fully Bayesian by making no inferences from observing a lack of signal. Though it is an empirically relevant phenomenon, this "cursedness" contradicts the paradigm of a fully rational, fully Bayesian economic agent as described in Economic literature. We show that in costly disclosure games, this distinction may sometimes be ignored, as it does not affect the equilibrium strategies.

While we see this paper as primarily normative (prescribing how firms ought to advertise to maximize impact on sales, but not in the sense of prescribing a particular regulation), it may be interesting to see whether similar strategies are employed in the real world – or, if not, whether using our strategy would improve upon those currently used in real-life scenarios.

The model abstracts from some relevant aspects of the setting, most notably flexible pricing (including personalized pricing/price discrimination). Exploring this

[19] It may appear that this lack of effect is due to competition not being explicit, with the candidates never advertising simultaneously. This is not the case, and can be verified by using, e.g., a version of the model with two signals $y_1, y_2$ observed by the candidates – if $y_1 = h$ and $y_2 = l$ then both candidates would advertise using the respective signal, and their targeting strategies will be exactly the same as argued above (and will conform to Theorem 2).

and other extensions in greater detail would be important for understanding the scope of the applicability of our results. As we discuss in footnote 5, if a firm must set the price uniformly for all consumers before observing its private signals, then our main results (Theorems 1 and 2) hold. This is because our model can then be seen as the second stage of the firm's problem. In other cases, however, pricing has the power to alter the nature of the problem quite significantly, both via price signaling of the firm's private information as a competing communication channel, and by enabling personalized pricing as an alternative way to influence consumers' information acquisition and purchasing behavior.

# References

A. Acquisti, C. Taylor, and L. Wagman. The Economics of Privacy. *Journal of Economic Literature*, 54(2):442–492, June 2016. ISSN 0022-0515. doi: 10.1257/jel.54.2.442.

S. P. Anderson, A. Baik, and N. Larson. Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts. SSRN Scholarly Paper ID 3428313, Social Science Research Network, Rochester, NY, June 2019. URL `https://papers.ssrn.com/abstract=3428313`.

S. Athey and J. S. Gans. The Impact of Targeting Technology on Advertising Markets and Media Competition. *American Economic Review*, 100(2):608–613, May 2010. ISSN 0002-8282. doi: 10.1257/aer.100.2.608.

D. Bergemann and A. Bonatti. Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443, 2011. ISSN 1756-2171. doi: 10.1111/j.1756-2171.2011.00143.x.

A. W. Bloedel and I. Segal. Persuasion with Rational Inattention. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3164033.

F. Branco, M. Sun, and J. M. Villas-Boas. Too Much Information? Information Provision and Search Costs. *Marketing Science*, 35(4):605–618, December 2015. ISSN 0732-2399. doi: 10.1287/mksc.2015.0959.

A. L. Brown, C. F. Camerer, and D. Lovallo. To Review or Not to Review? Limited Strategic Thinking at the Movie Box Office. *American Economic Journal: Microeconomics*, 4(2):1–26, 2012. ISSN 1945-7669. doi: 10.2307/23249431.

A. Caplin and M. Dean. Behavioral Implications of Rational Inattention with Shannon Entropy. Technical Report w19318, National Bureau of Economic Research, Cambridge, MA, August 2013. URL `http://www.nber.org/papers/w19318.pdf`.

J. Chen and J. Stallaert. An economic analysis of online advertising using behavioral targeting. *MIS Quarterly*, 38(2):429–450, 2014.

T. M. Cover and J. A. Thomas. *Elements of information theory.* John Wiley & Sons, 2012.

M. Deversi, A. Ispano, and P. Schwardmann. Spin Doctors: A model and an Experimental Investigation of Vague Disclosure. SSRN Scholarly Paper ID 3275418, Social Science Research Network, Rochester, NY, 2018. URL `https://papers.ssrn.com/abstract=3275418`.

E. Eyster and M. Rabin. Cursed Equilibrium. *Econometrica*, 73(5):1623–1672, September 2005. ISSN 0012-9682, 1468-0262. doi: 10.1111/j.1468-0262.2005.00631.x.

A. Farahat and M. C. Bailey. How effective is targeted advertising? In *Proceedings of the 21st international conference on World Wide Web*, WWW 2012, pages 111–120, Lyon, France, April 2012. Association for Computing Machinery. ISBN 978-1-4503-1229-5. doi: 10.1145/2187836.2187852. URL `https://doi.org/10.1145/2187836.2187852`.

A. Goldfarb and C. Tucker. Digital Economics. *Journal of Economic Literature*, 57 (1):3–43, March 2019. ISSN 0022-0515. doi: 10.1257/jel.20171452.

B. Hébert and M. Woodford. Rational Inattention and Sequential Information Sampling. Working Paper 23787, National Bureau of Economic Research, September 2017. URL `http://www.nber.org/papers/w23787`. Series: Working Paper Series.

A. Hefti and S. Liu. Targeted information and limited attention. *The RAND Journal of Economics*, 51(2):402–420, 2020. ISSN 1756-2171. doi: 10.1111/1756-2171.12319. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1756-2171.12319.

F. Hoffmann, R. Inderst, and M. Ottaviani. Persuasion Through Selective Disclosure: Implications for Marketing, Campaigning, and Privacy Regulation. *Management Science*, March 2020. ISSN 0025-1909. doi: 10.1287/mnsc.2019.3455. Publisher: INFORMS.

S. Ichihashi. Consumer-Optimal Disclosure With Costly Information Acquisition. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3253110.

A. Ispano and P. Schwardmann. Competition Over Cursed Consumers. SSRN Scholarly Paper ID 3208067, Social Science Research Network, Rochester, NY, May 2018. URL `https://papers.ssrn.com/abstract=3208067`.

G. Iyer, D. Soberman, and J. M. Villas-Boas. The Targeting of Advertising. *Marketing Science*, 24(3):461–476, August 2005. ISSN 0732-2399, 1526-548X. doi: 10.1287/mksc.1050.0117.

V. Jain and M. Whitmeyer. Competing to Persuade a Rationally Inattentive Agent. *arXiv:1907.09255 [econ]*, February 2020.

K. Jerath and Q. Ren. Consumer Attention Allocation and Firm Information Design. working paper, 2019.

G. Z. Jin, M. Luca, and D. Martin. Is No News (Perceived As) Bad News? An Experimental Investigation of Information Disclosure. *American Economic Journal: Microeconomics*. ISSN 1945-7669. doi: 10.1257/mic.20180217.

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

A. Lindbeck and J. Weibull. Delegation of investment decisions, and optimal remuneration of agents. *European Economic Review*, page 103559, August 2020. ISSN 00142921. doi: 10.1016/j.euroecorev.2020.103559.

B. Maćkowiak, F. Matějka, and M. Wiederholt. Rational inattention: A disciplined behavioral model. *Goethe University Frankfurt mimeo*, 2018.

F. Matějka and A. McKay. Rational Inattention to Discrete Choices: A new Foundation for the Multinomial Logit Model. *American Economic Review*, 105(1): 272–298, 2015. doi: 10.1257/aer.20130047.

L. Matysková. Bayesian Persuasion with Costly Information Acquisition. SSRN Scholarly Paper ID 3161174, Social Science Research Network, Rochester, NY, April 2018. URL https://papers.ssrn.com/abstract=3161174.

L. Matysková and J. Šípek. Manipulation of Cursed Beliefs in Online Reviews. SSRN Scholarly Paper ID 2960745, Social Science Research Network, Rochester, NY, April 2017. URL https://papers.ssrn.com/abstract=2960745.

A. Montes. Bayesian Persuasion with Rational Inattentive Receiver. Technical report, University of Bonn, 2020.

S. Morris and P. Strack. The Wald Problem and the Relation of Sequential Sampling and Ex-Ante Information Costs. SSRN Scholarly Paper ID 2991567, Social Science Research Network, Rochester, NY, February 2019. URL https://papers.ssrn.com/abstract=2991567.

C. A. Sims. Rational Inattention and Monetary Economics. In B. M. Friedman and M. Woodford, editors, *Handbook of Monetary Economics*, volume 3, pages 155–181. Elsevier, January 2010. doi: 10.1016/B978-0-444-53238-1.00004-1. URL http://www.sciencedirect.com/science/article/pii/B9780444532381000041.

A. Smirnov and E. Starkov. Bad news turned good: reversal under censorship. *American Economic Journal: Microeconomics*, forthcoming, 2020. doi: 10.1257/mic.20190379.

R. E. Verrecchia. Discretionary disclosure. *Journal of Accounting and Economics*, 5:179–194, January 1983. ISSN 0165-4101. doi: 10.1016/0165-4101(83)90011-3.

J. M. Villas-Boas and Y. J. Yao. A dynamic Model of Optimal Retargeting. *Marketing Science*, February 2021. ISSN 0732-2399. doi: 10.1287/mksc.2020.1267. Publisher: INFORMS.

A. Wald. Sequential tests of statistical hypotheses. *The annals of mathematical statistics*, 16(2):117–186, 1945.

# Appendix

**Lemma 2.** *Function $\mathcal{D}_f(p, w)$ is weakly increasing in $w$ and $p$.*

*Proof.* It is immediate from (4)-(8) that $\mathcal{D}_f(p, w)$ is continuous and differentiable in both arguments almost everywhere, with the exceptions being the edge cases in (4). If $(p, w)$ are such that $\mathcal{D}(p, w) \in \{0, 1\}$ then $\mathcal{D}_f(p, w) = \mathcal{D}(p, w) \in \{0, 1\}$, hence monotonicity holds. Otherwise $\mathcal{D}(p, w) \in (0, 1)$ – then take the partial derivative of $\mathcal{D}_f(p, w)$ w.r.t. $w$:

$$\frac{\partial \mathcal{D}_f(p, w)}{\partial w} = \frac{\partial \mathcal{D}_f(p, w)}{\partial \mathcal{D}(p, w)} \left( \frac{\partial \mathcal{D}(p, w)}{\partial w} + (1 - \mathcal{D}(p, w)) \frac{\mathcal{D}(p, w)}{\lambda} \right), \tag{13}$$

where $\dfrac{\partial \mathcal{D}_f(p, w)}{\partial \mathcal{D}(p, w)} = \dfrac{p_f e^{\frac{w}{\lambda}}}{\left(\mathcal{D}(p, w) e^{\frac{w}{\lambda}} + (1 - \mathcal{D}(p, w))\right)^2} + \dfrac{(1 - p_f) e^{-\frac{1-w}{\lambda}}}{\left(\mathcal{D}(p, w) e^{-\frac{1-w}{\lambda}} + (1 - \mathcal{D}(p, w))\right)^2}.$

The $\frac{\partial \mathcal{D}_f(p,w)}{\partial \mathcal{D}(p,w)}$ term is nonnegative, since all numerators and denominators are strictly positive if $w \in (0, 1)$ (which is necessary for $\mathcal{D}(p, w) \in (0, 1)$). In the latter term of (13), $(1 - \mathcal{D}(p, w)) \frac{\mathcal{D}(p,w)}{\lambda} > 0$, and

$$\frac{\partial \mathcal{D}(p, w)}{\partial w} = \frac{\partial \hat{\mathcal{D}}(p, w)}{\partial w} = \frac{p e^{\frac{1-w}{\lambda}} \left(e^{\frac{w}{\lambda}} - 1\right)^2 + (1 - p) e^{\frac{w}{\lambda}} \left(e^{\frac{1-w}{\lambda}} - 1\right)^2}{\lambda \left[\left(e^{\frac{w}{\lambda}} - 1\right)\left(e^{\frac{1-w}{\lambda}} - 1\right)\right]^2} \geq 0.$$

Therefore, $\frac{\partial \mathcal{D}_f(p,w)}{\partial w} \geq 0$. On the other hand, the partial derivative of $\mathcal{D}_f(p, w)$ w.r.t. $p$ when $\mathcal{D}(p, w) \in (0, 1)$ is given by

$$\frac{\partial \mathcal{D}_f(p, w)}{\partial p} = \frac{\partial \mathcal{D}_f(p, w)}{\partial \mathcal{D}(p, w)} \cdot \frac{\partial \mathcal{D}(p, w)}{\partial p} \geq 0,$$

$$\text{since } \frac{\partial \mathcal{D}(p, w)}{\partial p} = \frac{\partial \hat{\mathcal{D}}(p, w)}{\partial p} = \frac{e^{\frac{1}{\lambda}} - 1}{(e^{\frac{1}{\lambda}} - e^{\frac{1-w}{\lambda}})(1 - e^{-\frac{1-w}{\lambda}})} \geq 0.$$

□

**Lemma 3.** *Function $\mathcal{A}(w_i)$ is continuous, weakly increasing for $w_i \leq \min\{\underline{w}, \bar{w}\}$, weakly decreasing for $w_i \geq \max\{\underline{w}, \bar{w}\}$. It is also constant for $w_i \in [\bar{w}, \underline{w}]$ if the ad is strong, and strictly convex for $w_i \in [\underline{w}, \bar{w}]$ if the ad is weak.*

*Proof.* Since $\mathcal{D}_f(p, w)$ is a strictly decreasing function of $w$ truncated to $[0, 1]$, is weakly increasing in $p$, and $\alpha(p^0) > p^0$, there exist $\bar{\bar{w}}, \underline{\underline{w}} \in [0,1]$ such that:

$$\mathcal{A}(w_i) = \begin{cases} 0, & \text{if } w_i \geq \bar{\bar{w}}, \\ 1 - \mathcal{D}_f(p^0, w_i), & \text{if } w_i \in [\max\{\underline{w}, \bar{w}\}, \bar{\bar{w}}], \\ \mathcal{D}_f(\alpha(p^0), w_i) - \mathcal{D}_f(p^0, w_i), & \text{if } w_i \in [\min\{\underline{w}, \bar{w}\}, \max\{\underline{w}, \bar{w}\}], \\ \mathcal{D}_f(\alpha(p^0), w_i), & \text{if } w_i \in [\underline{\underline{w}}, \min\{\underline{w}, \bar{w}\}], \\ 0, & \text{if } w_i \leq \underline{\underline{w}}. \end{cases}$$

The continuity and monotonicities follow immediately from this representation. If the ad is strong, then $\mathcal{A}(w_i) = 1$ for $w_i \in (\bar{w}, \underline{w})$. If the ad is weak, then the second derivative of $\mathcal{A}(w_i)$ for $w_i \in (\min\{\underline{w}, \bar{w}\}, \max\{\underline{w}, \bar{w}\})$ is given by

$$\frac{d^2\mathcal{A}(w)}{dw^2} = \frac{p_f - p^0}{\lambda^2 p^0(1-p^0)} \cdot \left[ p^0 e^{\frac{1-w}{\lambda}} \frac{e^{\frac{1-w}{\lambda}} + 1}{\left(e^{\frac{1-w}{\lambda}} - 1\right)^3} + (1-p^0) e^{\frac{w}{\lambda}} \frac{e^{\frac{w}{\lambda}} + 1}{\left(e^{\frac{w}{\lambda}} - 1\right)^3} \right],$$

which is strictly positive since $p_f > p^0$. □

*Proof of Theorem 2.* Part 1 holds by Proposition 1 (or, equivalently, Lemma 1).

As argued in the text, the firm targets consumer $i \in \mathcal{I}$ with $y = h$ if and only if $\mathcal{A}(w_i) > c$. From the assumption that $\mathcal{T}_{cursed} \neq \emptyset$ it follows that $\max_{w \in [0,1]} \mathcal{A}(w) > c$. From the continuity of $\mathcal{A}(w)$ together with $\mathcal{A}(0) = \mathcal{A}(1) = 0$, it follows that the upper contour set $\{w | \mathcal{A}(w) > c\}$ is open for any $c > 0$. If the ad is strong then by Lemma 3 the maximum of $\mathcal{A}(w)$ is attained by all $w \in [\bar{w}, \underline{w}]$. By the monotonicity of $\mathcal{A}(w)$ for $w \leq \bar{w}$ and for $w \geq \underline{w}$ described in Lemma 3, part 3 of the statement follows.

If the ad is weak then by Lemma 3 $\mathcal{A}(w)$ is strictly convex for all $w \in [\underline{w}, \bar{w}]$. Together with the monotonicity of $\mathcal{A}(w)$ in the remaining regions, this implies that $\arg\max_{w \in [0,1]} \mathcal{A}(w) \in \{\underline{w}, \bar{w}\}$. Consider two cases depending on whether there exists a $\tilde{w} \equiv \arg\min_{w \in (\underline{w}, \bar{w})} \mathcal{A}(w)$. If it does, then consider $\mathcal{A}(w)$ separately on $[0, \tilde{w}]$ and $[\tilde{w}, 1]$. On both intervals $\mathcal{A}(w)$ is single-peaked, hence quasi-concave, meaning that its upper contour sets $\{w | \mathcal{A}(w) > c\}$ are convex within each interval and include the respective peaks $\underline{w}$ and $\bar{w}$. If $\tilde{w}$ does not exist then $\mathcal{A}(w)$ is strictly monotone on $[\underline{w}, \bar{w}]$. Then $\mathcal{A}(w)$ is single-peaked on $[0,1]$, so again its upper contour set $\{w | \mathcal{A}(w) > c\}$ is convex and includes the global maximum.[20] This proves part 2 of

[20] Note that the UCS may or may not include the second point $\{\underline{w}, \bar{w}\}$ in this case.

the theorem. $\square$

*Proof of Proposition 3.* The proof is constructive. Let $\hat{\mathcal{A}}(w_i)$ be defined as in (11). Let $\bar{\mathcal{W}}(l) \equiv \emptyset$ and $\bar{\mathcal{W}}(h) \equiv \{w_i | \hat{\mathcal{A}}(w_i) > c\}$ and, correspondingly, $\bar{\mathcal{T}}(l) \equiv \emptyset$ and $\bar{\mathcal{T}}(h) \equiv \{i \in \mathcal{I} | w_i \in \bar{\mathcal{W}}(h)\}$.

Firstly, we show that there exists an equilibrium with $\bar{\mathcal{T}}$ as the firm's advertising strategy. Any consumer $i \in \bar{\mathcal{T}}(h)$ updates her belief to $p_i^1 = \alpha(p^0)$ if she receives ad $h$ and to $p_i^1 = \beta(p^0)$ if she receives no ad or ad $l$ – since she expects to not receive an ad if and only if $y = l$. The ad effect is given by $\hat{\mathcal{A}}(w_i) > c$, hence it is strictly optimal to advertise to $i$. On the other hand, any consumer $i \notin \bar{\mathcal{T}}(h)$ updates to $p_i^1 = \alpha(p^0)$ after ad $h$, $p_i^1 = \beta(p^0)$ after ad $l$ (both are zero-probability events from her perspective), and $p_i^1 = p^0$ in the absence of any ad. The ad effect is then given by $\mathcal{A}(w_i) < \hat{\mathcal{A}}(w_i) \leq c$, hence advertising to $i$ is strictly suboptimal.

Secondly, we show that in any other equilibrium, the firm's ad strategy $\mathcal{W}$ is such that $\mathcal{W} \subseteq \bar{\mathcal{W}}$ (which implies $\mathcal{T} \subseteq \bar{\mathcal{T}}$). Fix any such equilibrium and the respective firm's strategy. By Lemma 1, $\mathcal{T}(l) = \emptyset = \bar{\mathcal{T}}(l)$, hence we only need to show the claim for $y = h$. The interim belief of sophisticated consumer $i$ who does not receive any ad must be $p_i^1 \in [\beta(p^0), p^0]$ in this equilibrium. It cannot be lower because the worst private information the firm can have is $y = l$, and observing such a signal directly would lead the consumer to have belief $\beta(p^0)$, hence no action of the firm can drop consumer $i$'s belief below $\beta(p^0)$. On the other hand, $p_i^1$ cannot be higher than $p^0$ because the firm never advertises signal $y = l$ (see Lemma 1) – so $p_i^1$ after no ad is given by

$$p_i^1 = \frac{p^0\left(\rho(1-\sigma_i) + 1 - \rho\right)}{p^0\left(\rho(1-\sigma_i) + 1 - \rho\right) + (1-p^0)\left(\rho + (1-\rho)(1-\sigma_i)\right)} < p^0,$$

$$\text{where } \sigma_i \equiv \int_{\mathcal{T} \in 2^{\mathcal{I}}} \mathbb{I}\{i \in \mathcal{T}\} d\tau(\mathcal{T}|h) \in [0,1]$$

is the probability that signal $y = h$ is transmitted to $i$, and $\mathbb{I}\{\cdot\}$ is the indicator function. It follows that the ad effect $A_i$ on $i$ is such that $A_i \in [\mathcal{A}(w_i), \hat{\mathcal{A}}(w_i)]$. Advertising to $i$ is profitable if and only if $A_i > c$, hence only if $\hat{\mathcal{A}}(w_i) > c$, hence only if $i \in \bar{\mathcal{T}}(h)$.

Finally, we need to show that $\bar{\mathcal{W}}$ satisfies the properties listed in the Theorem. Recall that $\bar{\mathcal{W}}(h) = \{w_i | \hat{\mathcal{A}}(w_i) > c\}$, while for the cursed equilibrium we have $\mathcal{W}_{cursed}(h) = \{w_i | \mathcal{A}(w_i) > c\}$. The two are defined in terms of $\hat{\mathcal{A}}(w_i) = \mathcal{D}_f(\alpha(p^0), w_i) - \mathcal{D}_f(\beta(p^0), w_i)$ and $\mathcal{A}(w_i) = \mathcal{D}_f(\alpha(p^0), w_i) - \mathcal{D}_f(p^0, w_i)$ respectively.

Consider a fictitious ("prime") world in which all players share a common prior $p' \equiv \beta(p^0)$ and the precision of signal $y'$ is $\rho' \equiv \frac{\rho^2}{\rho^2+(1-\rho)^2}$. Then $\alpha'(p') = \alpha(p^0)$. But then for any $i$, $\hat{\mathcal{A}}(w_i) = \mathcal{A}'(w_i)$, where $\mathcal{A}'(w)$ is calculated same as $\mathcal{A}(w)$, except with $\rho'$ instead of $\rho$. Consequently, Theorem 2, as applied to this fictitious world, implies that $\mathcal{W}'_{cursed}(h) \equiv \{w_i | \mathcal{A}'(w_i) > c\}$ satisfies all the required properties. But for a given $i$, we have $w_i \in \mathcal{W}'_{cursed}(h) \iff w_i \in \bar{\mathcal{W}}(h)$, meaning that $\bar{\mathcal{W}}(h)$ satisfies all required properties as well. $\square$

*Proof of Proposition 4.* Let

$$U(p,q) \equiv \sum_{s \in S} p(s)\mathcal{D}(q(s), w_i, s) \cdot (w_i + v_s - 1) - \lambda\kappa(\mathcal{G}_i(q); q)$$

denote the expected utility received by consumer $i$ with idiosyncratic valuation $w_i$ and interim belief $q$, as estimated by an outside observer with belief $p$. Here $\mathcal{G}_i(q)$ is the signal structure chosen optimally by such a consumer, and $\mathcal{D}(\cdot)$ are the purchase probabilities generated by $\mathcal{G}_i(q)$ and the optimal purchasing behavior. In particular, $\mathcal{G}_i(q)$ is chosen so as to maximize the consumer's expected utility as evaluated by her, meaning that

$$\begin{aligned} U(p,p) &= \max_{\mathcal{G}} \left\{ \sum_{s \in S} p(s)\mathcal{D}(p(s), w_i, s) \cdot (w_i + v_s - 1) - \lambda\kappa(\mathcal{G}; p) \right\} \\ \Rightarrow U(p,p) &= \max_{q} U(p,q). \end{aligned} \tag{14}$$

It is optimal for the consumer to have the correct belief – otherwise she acquires information suboptimally and may not attain the utility maximum.

We now show the claim for sophisticated consumers. This is equivalent to showing that having more information at the interim stage is always weakly beneficial for the consumer. Fix the prior belief $p^0$ and let $p^1_y$ denote the interim belief of a Bayesian observer who directly observes the firm's private signal $y$. Note that $p^1_y$ coincides with the beliefs of a sophisticated consumer after all $y$ if $i \in \mathcal{T}(h)$, while if $i \notin \mathcal{T}(h)$ then such a consumer's belief remains at $p^0$. Her expected utility in the latter case is $U(p^0, p^0)$, and we want to show that this is weakly smaller than her ex ante expected utility in the former case:

$$\mathbb{E}_y[U(p^1_y, p^1_y)] = \sum_{y \in \{h,l\}} P(y) U(p^1_y, p^1_y),$$

where $P(y)$ are the ex ante probabilities of different realizations of $y$.

Consider now a Bayesian observer with the same prior $p^0$ who also sees $y$. Consider this observer's estimates of the expected utility of an uninformed consumer, as calculated at two stages: ex ante and interim stages (i.e., before and after the observer sees $y$). The former coincides with the consumer's own estimate, $U(p^0, p^0)$. The latter, conditional on the realization of $y$, is given by $U(p^1_y, p^0)$. The interim estimates for different $y$ must average out to the ex ante estimate – otherwise one of the estimates is inaccurate, – so

$$U(p^0, p^0) = \sum_{y \in \{h,l\}} P(y) U(p^1_y, p^0).$$

From (14) it follows that $U(p^1_y, p^1_y) \geq U(p^1_y, p^0)$, hence we get that $\mathbb{E}_y[U(p^1_y, p^1_y)] \geq U(p^0, p^0)$ – observing $y$ is indeed beneficial for the consumer. This proves the statement of the proposition for sophisticated consumers.

For cursed consumers, follow the same argument. The ex ante expected utility of a cursed consumer if $i \in \mathcal{T}(h)$ (as evaluated by an objective observer) is

$$P(h)U(p^1_h, p^1_h) + P(l)U(p^1_l, p^0), \tag{15}$$

since her interim belief in case $y = l$ is same as the prior $p^0$, because the firm is not advertising in that case, while in case $y = h$ the firm advertises and the consumer updates her belief to $p^1_h$. The ex ante expected utility of a cursed consumer if $i \notin \mathcal{T}(h)$ (as evaluated by an objective observer) is

$$P(h)U(p^1_h, p^0) + P(l)U(p^1_l, p^0), \tag{16}$$

since her belief then remains at $p^0$ regardless of $y$. We can now see that (15) is weakly greater than (16), since $U(p^1_h, p^1_h) \geq U(p^1_h, p^0)$ by (14). This proves the statement of the proposition for cursed consumers and concludes the proof. $\square$