Discussion Paper Series – CRC TR 224

Discussion Paper No. 425
Project B 03

Sparking Curiosity or Tipping the Scales? Targeted Advertising With Consumer Learning

Andrei Matveenko [1]
Egor Starkov [2]

May 2023

[1] University of Mannheim, Department of Economics, Email: matveenko@uni-mannheim.de
[2] University of Copenhagen, Department of Economics, Email: egor.starkov@econ.ku.dk

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Sparking curiosity or tipping the scales? Targeted advertising with consumer learning*

Andrei Matveenko[1] and Egor Starkov †[2]

[1]University of Mannheim, Department of Economics
[2]University of Copenhagen, Department of Economics

April 28, 2023

Abstract

This paper argues, in the context of targeted advertising, that receivers' ability to independently acquire information has 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 can engage in costly learning about 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.

*The paper was previously circulated under the title "Sparking curiosity or tipping the scales? Targeted advertising with rationally inattentive consumers". We are grateful to anonymous referees, Mogens Fosgerau, Germain Gaudin, Marit Hinnosaar, Alessandro Ispano, Igor Letina, Filip Matějka, Andrey Minaev, Xiaosheng Mu, Marco Schwarz, Nicolas Schutz, Adrien Vigier, and Alexander White, as well as participants of the 2020 Generalized Entropy Workshop, OLIGO 2020 and IIOC 2021 virtual conferences, workshop at Cergy and seminar participants at Aarhus, Bergen, Copenhagen and Mannheim 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). Funding by the German Research Foundation (DFG) through CRC TR 224 (project B03) is gratefully acknowledged.

†The authors' contributions are equal and the order is alphabetical.
E-mail addresses: `matveenko@uni-mannheim.de` (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 (see Acquisti, Taylor, and Wagman [2016] and Goldfarb and Tucker [2019] for excellent surveys). 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 [Lewis and Hilder, 2018].

The question that arises then is: which consumers should a firm advertise to? In other words, which consumers does advertising have the most effect on? The question becomes more complicated in the modern information-rich world, since advertising is not the only source of information, and consumers can choose to learn about the product further from other sources. We hence ask: which consumers are the most appealing targets for advertising in the presence of other learning channels?

To identify the optimal advertising strategy, we consider a theoretical model in which a population of heterogeneous consumers is faced with an option to buy a product which yields an uncertain payoff. The consumers differ in their initial predisposition towards the product and they can further acquire costly information about their actual valuation for it. A monopolistic seller 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, in terms of their predisposition.

It would appear natural to target consumers with the highest demand for information, since the possibly unfavorable information they may learn on their own is then replaced with the favorable information contained in the ad. However, we show that this is not the case, and it is instead best to target some of the consumers with no or low demand for information. In particular, if the effect of the advertisement on consumer expectations is not too large, and the cost of advertising is significant, 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 sparks their curiosity – renders 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 the consumers in the latter group choose to acquire information about 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 result extends the empirical findings of Blake, Nosko, and Tadelis [2015], who show that in case of eBay, the largest effect of sponsored search ads is on consumers who were either not using, or barely using eBay before that. These are the "pessimist" consumers in our model, who are nudged by the ad to acquire information. However, our results suggest that in general it is also worth targeting the optimists – those, who can be pushed to stop deliberating and make the purchase. This effect may be absent in the results of Blake et al. [2015] because their paper investigates sponsored search ads, which do not convey information about product or platform quality or fit with consumer's tastes, but rather serve to make consumers aware of the product. The effects of such "informative" advertising are different from those of "persuasive" advertising we explore in this paper (using the terminology of Bagwell [2007], see Section 7 for a discussion).

Our results hold regardless of whether consumers 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 she did *not* receive an advertisement, which is an empirically appealing assumption (see Section 7 for a discussion). This "cursedness" can manifest because consumers 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. Equivalently, the version of the model with cursed consumers can be seen as a non-informational model, in which advertising directly increases the consumers' willingness to pay. The bottom line is that our results are robust to a variety of extensions and alternative interpretations.

The connection between the models with cursed and sophisticated consumers deserves attention by itself. We show that the uniquely optimal advertising strategy with cursed consumers also constitutes an equilibrium with sophisticated consumers. The latter model features a multiplicity of self-fulfilling equilibria driven by the consumers' off-the-equilibrium path beliefs: a costly advertisement may be worth sending to a consumer who otherwise (in the absence of an ad) infers that the firm's product is bad, but not to a consumer who otherwise sticks to their prior belief. Such a multiplicity is inherent to any model of costly information disclosure. In the equilibrium with the smallest amount of advertising, the firm's advertising strategy is then exactly the same as if all consumers were cursed, i.e., retained their prior belief in the absence of an ad regardless of whether they expected to receive an ad. We further show that this least-advertising equilibrium is the least preferred by consumers, but the most preferred by the firm, out of all equilibria in the model with sophisticated consumers.

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 full model. In Section 5 we derive the optimal advertising strategy and explore its properties in the context of one salient equilibrium. Section 6 characterizes the set of equilibria and the players' preferences over them, as well as further justifies the salient equilibrium as being equivalent to the unique equilibrium with cursed consumers. Finally, Section 7 discusses the main assumptions behind the model, and Section 8 concludes. All proofs are relegated to the Appendix.

2 An Illustrative Example

The main driving force behind our result can be illustrated using the following highly stylized example. Suppose a firm sells a product of quality $s \in \{H, L\}$ at some fixed price normalized to 1. The firm's reservation utility for the product is zero. The consumer values a high-quality product at $v = w$ net of the price, and a low-quality product 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 $\rho = 0.8$ (i.e., the signal is correct with probability 0.8: $\mathbb{P}(h|H) = \mathbb{P}(l|L) = 0.8$). 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 that the firm's signal is high, $y = h$, so observing y increases a consumer's belief that $s = H$ to

$$p_i^1(h) = \frac{\mathbb{P}(h|H)\mathbb{P}(H)}{(\mathbb{P}(h|H)\mathbb{P}(H) + \mathbb{P}(h|L)\mathbb{P}(L))} = \frac{2}{3}.$$

Assume also that if the consumer receives no ad, her belief remains at p^0.[1]

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:

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

The consumer then passes if $w \leq \underline{w}$, investigates if $w \in (\underline{w}, \bar{w})$, and buys if $w \geq \bar{w}$, 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. Figure

[1] We refer to such consumers as "cursed" in the sense of Eyster and Rabin [2005]. This assumption is adopted to simplify the example as much as possible. In the full model, we consider both the version with fully Bayesian consumers and the version with cursed consumers (and justify this assumption properly).

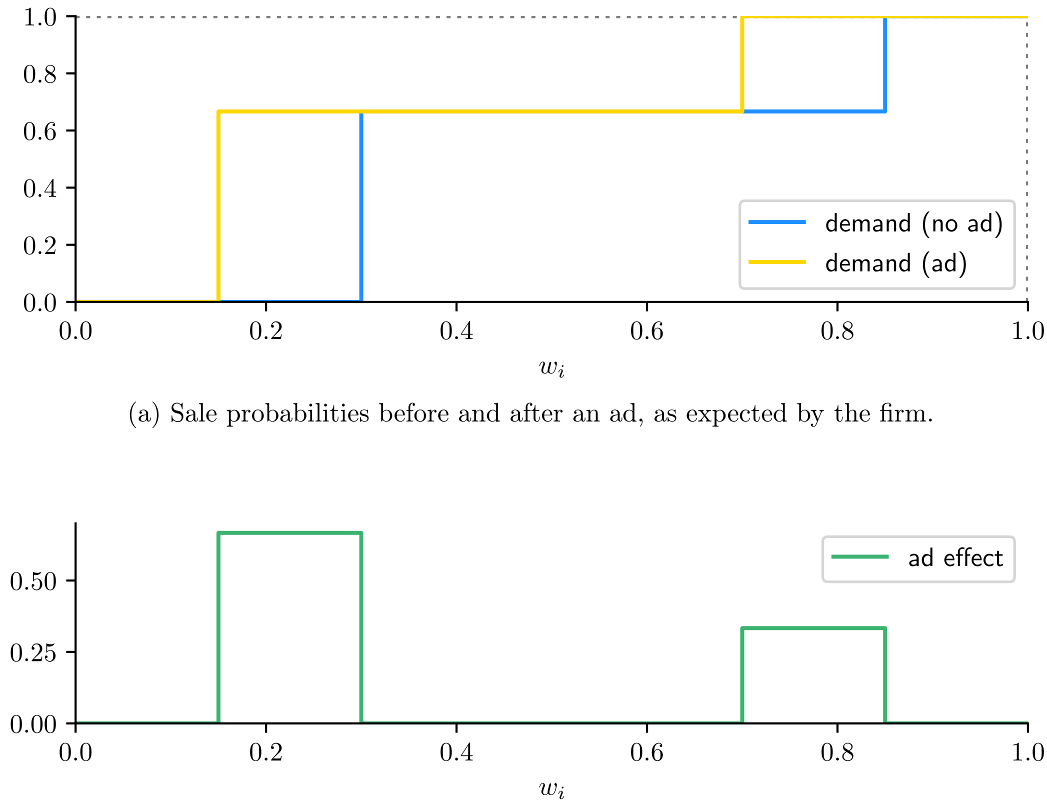

(a) Sale probabilities before and after an ad, as expected by the firm.

(b) Ad effect.

Figure 1: Advertising effect in the illustrative example.

1 plots the purchasing probabilities with and without an ad (panel (a)) and the effect of advertising on firm sales (panel (b)).

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 ρ 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 much more flexible information acquisition process, when the consumer can acquire any information she prefers, – and provides more insight into the issues and tradeoffs the firm faces when deciding on its advertising strategy. By interpolation, one could then hope that our intuition also applies to models "in between".

3 Literature Review

Bagwell [2007] and Renault [2015] present excellent surveys of literature on (non-targeted) advertising. For surveys of literature on consumer targeting and tracking, see Acquisti, Taylor, and Wagman [2016] and Chapter 6 of Goldfarb and Tucker [2019].

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". In our model advertising instead increases (the aware) consumers' perception of product quality, and consumers can acquire further information. We show that in such a model, 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.

Burguet and Petrikaite [2020] and Villas-Boas and Yao [2021] explore optimal ad targeting in search models. The former paper focuses on the issues of competition and pricing, while the latter asks how to optimally target ads based on consumers' search behavior, which is informative of whether they are interested in the product. These two are the only of the aforementioned papers that recognize the consumers' information acquisition layer, which is the focus of our paper and the main driving force behind our results. However, their approach to information acquisition and the questions that they focus on are different from ours.[2]

[2]Other work on targeting includes papers by Athey and Gans [2010], Bergemann and Bonatti [2011], Farahat and Bailey [2012], Chen and Stallaert [2014], Anderson, Baik, and Larson [2019], and Hoffmann, Inderst, and Ottaviani [2020]. 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 value to the firm and the consumers. However, none of these papers account for

The literature on strategic information acquisition in various settings is vast, going back at least to the sequential sampling model of Wald [1945]. Ichihashi [2018] exemplifies the importance of 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.

More recently, literature on strategic learning has become closely intertwined with literature on rational inattention (see Sims [2010] and Maćkowiak, Matějka, and Wiederholt [2018] for recent surveys). We adopt this rational inattention framework to model the information acquisition process and rely, in particular, on the results of Caplin and Dean [2013] and Matějka and McKay [2015]. The most closely related to our paper in this literature are the works of Cheng [2021], Matysková and Montes [2021], and Figueroa and Guadalupi [2021], who explore, respectively, Bayesian Persuasion mechanisms and the optimal costly signaling strategy, when the receiver is rationally inattentive in the sense of being able to engage in costly information acquisition after hearing the sender's message. We complement these two papers by looking at costly verifiable disclosure a lá Verrecchia [1983] at the communication stage. We show that the results can differ substantially across modes of communication: where Matysková and Montes [2021] show that the optimal persuasion mechanism disincentivizes the consumer from acquiring any further information,[3] we show that under costly disclosure the exact converse can be the case, with the sender in some cases using the disclosure opportunity to nudge the receiver to acquire *more* information that she otherwise would.

4 The Model

We now proceed to introduce the full model. An extensive discussion of the assumptions made is postponed until Section 7.

Our model builds upon a classic Hotelling market. A unit continuum $\mathcal{I}$ of consumers is distributed uniformly on an interval $[0, 1]$, with w_i denoting the location of a generic consumer $i \in \mathcal{I}$. A firm, located in position 1 on the interval, offers

the possibility that consumers acquire information independently, which, as we show, can drive predictions to a significant degree.

[3]Similar results – that the sender/firm would often find it optimal to limit consumer learning – were also reported by Wang [2017], Jain and Whitmeyer [2020], and Jerath and Ren [2021].

a product of unknown quality $s \in S \equiv \{H, L\}$. All consumers and the firm share a common prior belief about product quality that assigns probability $p^0 \in [0, 1]$ to state $s = H$. Price of the product to consumers is exogenously fixed at zero, whereas the firm yields benefit equal to 1 from every unit sold (from, e.g., selling consumer information).[4]

All consumers have unit demand for the product. The value that consumer i extracts from consuming a product of quality s is given by $v_s - (1 - w_i)$, and buying nothing yields utility zero. The first term in the consumption utility, v_s, represents the value added by product quality: $v_L = 0$ for a low-quality product and $v_H = 1$ for a high-quality product. The quality term is state-dependent and is thus not observed by the consumer prior to the purchase. The second term, $-(1 - w_i)$, stands for the transportation cost given by the distance between i and the firm on the line. This could be interpreted as either the physical transportation cost, or the fit of the firm's product to i's idiosyncratic tastes. In what follows, we stick to the latter interpretation, with higher w_i meaning better fit with consumer i's tastes. Consumer i's location w_i is commonly known: the consumer is aware of her own tastes, and the firm (or the advertising agency it employs) 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 owing to tracking technologies.

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 interpret y as a piece of hard but inconclusive evidence that the firm can disclose to selected consumers; this may include certificates, 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 (equivalently, per-impression) cost $c > 0$, where $|\mathcal{T}(y)|$ is the Lebesgue measure of set $\mathcal{T}(y)$. Consumer i observes her "ad signal" $a_i \in \{\emptyset, h, l\}$, with $a_i = y$ if $i \in \mathcal{T}(y)$ and $a_i = \emptyset$ if $i \notin \mathcal{T}(y)$, but does not explicitly observe the firm's choice of strategy $\mathcal{T}(y)$.

Each consumer i chooses whether to purchase the product or not. Prior to making the decision (but after observing a_i), the consumer has an opportunity to

[4] As discussed in Section 7, neither the fact that the price is zero, nor the firm's inability to set the price is important to our results, so long as the firm cannot price discriminate.

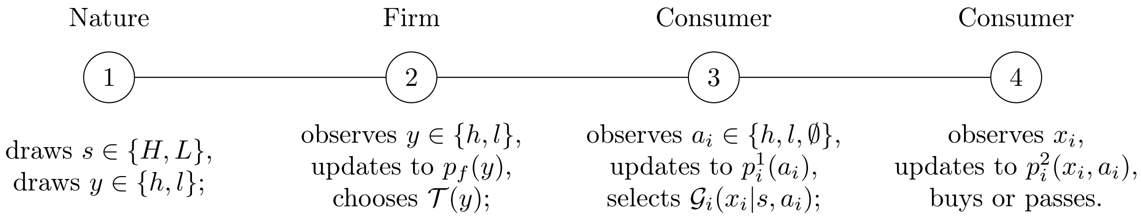

Figure 2: The timing of the model.

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 $\mathcal{G}_i(x_i|s, a_i) : S \to \Delta(\mathbb{R})$. 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.

The overall timing of the model is as follows (also depicted in Figure 2):

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

We now turn to formulating the firm's and consumers' objective functions maximized in stages 2-4, required to define an equilibrium of the game. Proceeding by backwards induction, at stage 4 the consumers' purchasing decisions are mechanical: consumer i buys the product if and only if the expected utility from doing so is positive:

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

(w.l.o.g. we break ties in firm's favor).

Moving on to stage 3, let $\mathcal{D}(p_i^1, w_i, s)$ denote the probability with which consumer i purchases the product (conditional on her belief p_i^1 and her optimal information acquisition strategy) when the realized quality is s. Further, with some abuse of notation let $\mathcal{D}(p_i^1, w_i)$ denote the respective buying probability unconditional on s from the consumer's perspective, and $\mathcal{D}_f(p_i^1, w_i)$ represent i's expected demand (the probability that consumer i buys the product) as perceived by the firm, conditional on its belief $p_f(y)$. The three probabilities are summarized as follows:

$$\begin{aligned}
\mathcal{D}(p_i^1, w_i, s) &\equiv \mathbb{P}\left(p_i^2 \geq 1 - w_i \mid p_i^1, w_i, s\right), \\
\mathcal{D}(p_i^1, w_i) &\equiv \sum_{s \in S} p_i^1(s|a_i)\mathcal{D}(p_i^1, w_i, s), \\
\mathcal{D}_f(p_i^1, w_i) &\equiv \sum_{s \in S} p_f(s|y)\mathcal{D}(p_i^1, w_i, s).
\end{aligned}$$

Given these probabilities, we can define the players' payoffs. In particular, consumer i's expected utility as of stage 3 is given by

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

where $\lambda\kappa(\mathcal{G}_i; p_i^1)$ is the cost of generating signal structure $\mathcal{G}_i$.

To model the cost of information, we follow a large literature on "rational inattention" and assume that it is proportional to the expected reduction in entropy between the consumer's interim and posterior beliefs. This functional form has been used extensively in the literature to model costs of both information processing and information acquisition and has been justified on many grounds, see Section 7 for an extended discussion and Maćkowiak et al. [2018] for a survey of the literature. The entropy cost is given by[5]

$$\begin{aligned}
\kappa(\mathcal{G}_i; p_i^1) \equiv &- \sum_s p_i^1(s|a_i) \log p_i^1(s|a_i) \\
&+ \sum_s \sum_{x_i} \left(\sum_s p_i^1(s|a_i)\mathcal{G}_i(x_i|s, a_i) \right) p_i^2(s|x_i, a_i) \log p_i^2(s|x_i, a_i),
\end{aligned}$$

with $\lambda \in \mathbb{R}_{++}$ being the information cost factor.

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

$$\begin{aligned}
&\int_{i \in \mathcal{I}} \mathcal{D}_f(p_i^1, w_i)di - c \cdot |\mathcal{T}(y)| = \\
&\qquad = \int_{i \in \mathcal{T}(y)} \left[\mathcal{D}_f(p_i^1(y), w_i) - c\right] di + \int_{i \in \mathcal{I} \setminus \mathcal{T}(y)} \mathcal{D}_f(p_i^1(\emptyset), w_i)di. \qquad (2)
\end{aligned}$$

This expression estimates the expected profit as of stage 3 conditional on the consumers' interim beliefs $\{p_i^1(a_i)\}_{i \in \mathcal{I}}$. In equilibrium, the firm knows that all con-

[5] 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.

sumers' prior beliefs are p^0 and knows exactly how these beliefs react to ads or lack thereof, hence (2) is also a valid representation of the firm's expected profit as of stage 2.

Having defined the game and the players' objective functions, we can now introduce the equilibrium concept. We look for sequential equilibria (SE) of the game, hereinafter referred to simply as equilibria.[6]

Definition (Sequential Equilibrium)**.** *A sequential equilibrium of the game consists of the firm's advertising strategy $\mathcal{T}(y)$, the collection of the consumers' information acquisition strategies $\{\mathcal{G}_i(x_i|s,a_i)\}_{i\in\mathcal{I}}$, and the updating rules $p_f(y)$, $p_i^1(a_i)$, and $p_i^2(x_i,a_i)$, for the firm's and the consumers' beliefs respectively, such that:*

1. *consumer i buys the product if and only if $p_i^2(x_i,a_i) \geq 1-w_i$;*
2. *information acquisition strategy $\mathcal{G}_i(x_i|s,a_i)$ maximizes expected payoff (1) for all consumers $i \in \mathcal{I}$, all $a_i \in \{\emptyset, l, h\}$;*
3. *advertising strategy $\mathcal{T}(y)$ maximizes the firm's expected profit* (2) *given any signal $y \in \{h,l\}$;*[7]
4. *the firm's belief $p_f(y)$ is updated using Bayes' rule after any $y \in \{h,l\}$;*
5. *there exists a sequence of fully mixed advertising strategies $\{\mathcal{T}_n\}_{n\in\mathcal{N}}$ (i.e., $\mathcal{T}_n : Y \to \Delta(2^{\mathcal{I}})$) that converges to $\mathcal{T}$ in probability for each y, such that for every consumer i, beliefs derived from $\mathcal{T}_n$ using Bayes' rule converge to $p_i^1(a_i)$ and $p_i^2(x_i,a_i)$ for all respective a_i and x_i.*

It is not necessary to impose the consistency requirements on the firm's belief $p_f(y)$, since both realizations of y are on equilibrium path for the firm, and hence the belief can always be derived using Bayes' rule.

5 Analysis

5.1 The Consumers' Problem

We solve the model via backwards induction. As mentioned previously, the optimal strategy in stage 4 is trivial: consumer i buys the product if and only if $p_i^2(x_i,a_i) \geq 1-w_i$. This section is hence devoted to the consumers' information acquisition

[6]For a general definition of sequential equilibria, see, e.g., Osborne and Rubinstein [1994], Ch.12. We use the SE concept instead of a weaker Perfect Bayesian Equilibrium (PBE) concept in order to avoid consumers having unreasonable beliefs after observing ads off the equilibrium path – which could happen in a PBE despite the ads containing only hard information in our model.

[7]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.

problem in stage 3. Since 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 in stage 3 is realizing that, given some interim belief p_i^1, every signal x_i produced by the consumer's optimal strategy must induce a different action. In particular, Caplin and Dean [2013] show that if two signals x_i induce the same action, the consumer can save on information costs by pooling them into one signal. In our setting, this means that the consumer's optimal signal structure $\mathcal{G}_i$ will generate at most two distinct signals x_i – a recommendation to buy and a recommendation to pass. If the consumer's location w_i makes the decision obvious enough, then she will choose a trivial (uninformative) signal $\mathcal{G}_i$, either the one that always recommends to buy, or the one suggesting to pass. Using the terminology of Lindbeck and Weibull [2020], we say that consumer i is in *sour* conditions if she passes on the product without investigation, *sweet* conditions if she buys the product without investigation, and *normal* conditions if she decides to acquire additional information.

Suppose the consumer is in normal conditions. Then Matějka and McKay [2015] show (see their Theorem 1) that the consumer's resulting choice probabilities (evaluated at stage 3) should satisfy the generalized logit rule conditional on the true product quality s:

$$\mathcal{D}(p_i^1, w_i, H) = \frac{\mathcal{D}(p_i^1, w_i)e^{\frac{w_i}{\lambda}}}{\mathcal{D}(p_i^1, w_i)e^{\frac{w_i}{\lambda}} + (1 - \mathcal{D}(p_i^1, w_i))}, \tag{3}$$

$$\mathcal{D}(p_i^1, w_i, L) = \frac{\mathcal{D}(p_i^1, w_i)e^{\frac{w_i - 1}{\lambda}}}{\mathcal{D}(p_i^1, w_i)e^{\frac{w_i - 1}{\lambda}} + (1 - \mathcal{D}(p_i^1, w_i))}, \tag{4}$$

where the exponent terms capture the terminal utilities associated with different actions, adjusted for information cost λ (note that $e^{\frac{0}{\lambda}} = 1$ is omitted), and different actions (whether to buy or not) are additionally weighed by their unconditional choice probabilities $\mathcal{D}(p_i^1, w_i)$ and $1 - \mathcal{D}(p_i^1, w_i)$ respectively. The law of total probability then implies that

$$\mathcal{D}(p_i^1, w_i) = p_i^1 \mathcal{D}(p_i^1, w_i, H) + (1 - p_i^1)\mathcal{D}(p_i^1, w_i, L). \tag{5}$$

Solving the system of equations (3)–(5) w.r.t. the three choice probabilities yields the closed-form solution

$$\mathcal{D}(p_i^1, w_i) = \begin{cases} 0 & \text{if } p_i^1 \in [0, \underline{p}_i], \\ \hat{\mathcal{D}}(p_i^1, w_i) & \text{if } p_i^1 \in (\underline{p}_i, \hat{p}_i), \\ 1 & \text{if } p_i^1 \in [\hat{p}_i, 1]; \end{cases} \tag{6}$$

where

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

In the above, the cutoffs $\hat{p}_i$ and $\underline{p}_i$ are determined from the clipping conditions $\hat{\mathcal{D}}(\hat{p}_i, w_i) = 1$ and $\hat{\mathcal{D}}(\underline{p}_i, w_i) = 0$, respectively (which, by construction, implies that $\mathcal{D}(p_i^1, w_i)$ is continuous in p_i^1). One can solve for the two cutoffs to obtain

$$\underline{p}_i \equiv \frac{e^{\frac{1-w_i}{\lambda}} - 1}{e^{\frac{1}{\lambda}} - 1} \qquad \text{and} \qquad \hat{p}_i \equiv \frac{1 - e^{-\frac{1-w_i}{\lambda}}}{1 - e^{-\frac{1}{\lambda}}} = e^{\frac{w_i}{\lambda}} \underline{p}_i. \tag{8}$$

Then consumer i is in sweet conditions whenever $p_i^1 \geq \hat{p}_i$ and sour conditions 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 signal $x_i \in \{x^{buy}, x^{pass}\}$, where the two signal realizations bring her posterior belief to either $p_i^2(x^{buy}) = \hat{p}_i$, or $p_i^2(x^{pass}) = \underline{p}_i$.

In particular, the solution to the consumer's problem in the normal conditions is such 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. At the same time, the cutoffs $\underline{p}_i$ and $\hat{p}_i$ *do* depend on w_i – different consumers acquire different signals, which vary in both absolute precision and the balance between type-I and type-II errors. This also means that we can invert the result and characterize a consumer's learning behavior based on her idiosyncratic preference: given some fixed belief p_i, there exist cutoffs $\underline{w}_i$ and $\hat{w}_i$ such that consumer i with interim belief p_i is in normal conditions if $w_i \in (\underline{w}_i, \hat{w}_i)$, sour if $w_i \leq \underline{w}_i$, and sweet if $w_i \geq \hat{w}_i$. These cutoffs can be obtained by inverting the expressions in (8) and will be useful later.

5.2 The Firm's Problem

From this point onwards we explore the firm's decision in stage 2 of the game – namely, its choice of advertising strategy $\mathcal{T}$. The previous subsection implies that for any firm's strategy, there exists a solution to the consumer's problem that it generates in all contingencies – a "continuation equilibrium". Therefore, if a firm's advertising strategy is optimal given the consumers' response, all of them together constitute an equilibrium.

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

$$\mathcal{D}_f(p_i^1, w_i) = p_f\mathcal{D}(p_i^1, w_i, H) + (1-p_f)\mathcal{D}(p_i^1, w_i, L). \tag{9}$$

Note that this "demand" function for consumer i depends both on the firm's belief p_f about the state of the world, for obvious reasons, and on the consumer's belief p_i^1, since the latter determines both the consumer's ex ante valuations and how much and what kind of extra information the consumer acquires.

The question now is how the consumer's interim belief $p_i^1(a_i)$ responds to the firm's advertisement or a lack thereof. Lemma 1 below demonstrates that as expected, the belief increases relative to the prior p^0 after a good ad $a_i = h$ and decreases after a bad ad $a_i = l$, even after taking the strategic nature of the firm's disclosure decision into account. Since advertising with signal l is costly and depresses the consumer's belief, it is never optimal for the firm to advertise with $y = l$, as long as consumers' beliefs do not react too adversely to the lack of advertising.[8]

Lemma 1. *In any sequential equilibrium, the following are true for any consumer i:*

1. $p_i^1(h) = \frac{p^0\rho}{p^0\rho+(1-p^0)(1-\rho)} > p^0$,
2. $p_i^1(l) = \frac{p^0(1-\rho)}{p^0(1-\rho)+(1-p^0)\rho} < p^0$,
3. $p_i^1(\emptyset) \in [p_i^1(l), p^0]$,
4. $i \notin \mathcal{T}(l)$.

We can thus focus on the firm's decision when it has a high signal $y = h$ in hand. The probability that consumer i purchases the product is given by $\mathcal{D}_f(p_i^1(h), w_i)$ after ad h and $\mathcal{D}_f(p_i^1(\emptyset), w_i)$ in the absence of any ad. It is then immediate from the firm's profit function (2) that it is optimal to target consumer $i \in \mathcal{I}$ if and only if

$$\mathcal{A}(w_i) \equiv \mathcal{D}_f(p_i^1(h), w_i) - \mathcal{D}_f(p_i^1(\emptyset), w_i) > c. \tag{10}$$

Hereinafter, we refer to $\mathcal{A}(w_i)$ as the "ad effect" on consumer i. This ad effect depends, among other things, on the consumer's interim belief $p_i^1(\emptyset)$ in the absence of an ad, which is endogenous in equilibrium. However, since $D_f(p, w_i)$ is weakly

[8]Disclosing adverse information may be optimal in richer settings; see Smirnov and Starkov [2022] for one example and a survey of related papers.

increasing in p (see Lemma 4 in the Appendix), part 3 of Lemma 1 allows us to impose bounds on the magnitude of the ad effect:

$$\begin{aligned} \mathcal{A}(w_i) &\in \left[\underline{\mathcal{A}}(w_i), \bar{\mathcal{A}}(w_i)\right], \\ \text{where } \underline{\mathcal{A}}(w_i) &\equiv \mathcal{D}_f(p_i^1(h), w_i) - \mathcal{D}_f(p^0, w_i), \\ \bar{\mathcal{A}}(w_i) &\equiv \mathcal{D}_f(p_i^1(h), w_i) - \mathcal{D}_f(p_i^1(l), w_i). \end{aligned} \tag{11}$$

Expressions (10) and (11) together directly imply the following property of the firm's optimal targeting strategy:

Lemma 2. *In any equilibrium, $i \in \mathcal{T}(h)$ if $\underline{\mathcal{A}}(w_i) > c$ and only if $\bar{\mathcal{A}}(w_i) > c$.*

Lemma 2 provides a partial characterization of the game's equilibria, which allows us to guess how an equilibrium might look like. In particular, note that the non-targeted consumers do not react to a lack of advertisement in equilibrium: $p_i^1(\emptyset) = p^0$ for $i \notin \mathcal{T}(h)$, hence $\mathcal{A}(w_i) = \underline{\mathcal{A}}(w_i) \le c$ for these consumers. This implies that there exists an equilibrium in which $i \in \mathcal{T}(h)$ if *and only if* $\underline{\mathcal{A}}(w_i) > c$. We will refer to it as the *reticent equilibrium*, since it features the least amount of advertising, as shown by Proposition 1 below.

Definition. *The* reticent equilibrium *is given by the firm's targeting strategy $\mathcal{T}_R(l) = \emptyset$ and $\mathcal{T}_R(h) = \{i | \underline{\mathcal{A}}(w_i) > c\}$ and the consumers' corresponding beliefs and strategies.*

Proposition 1. *The reticent equilibrium exists. Further, for any targeting strategy $\mathcal{T}$ played in any other sequential equilibrium, $\mathcal{T}_R(y) \subseteq \mathcal{T}(y)$ for any $y \in \{h, l\}$.*

The idea behind the proposition is immediate: since $\mathcal{A}(w_i) \ge \underline{\mathcal{A}}(w_i)$ in all equilibria for all w_i, the actual ad effect can never be lower than $\underline{\mathcal{A}}(w_i)$ – and so if the firm sends an ad to consumer i in the reticent equilibrium, then it is optimal to do so in any other equilibrium.

In the following subsections, we analyze the reticent equilibrium and verify that it has the twin-peak structure we are looking for. We show later in Section 6 that there are many other reasons to focus on this equilibrium in particular, including it being the seller-preferred equilibrium and it being robust to the precise way we model the ads' effects on consumers. Further, in Section 6, we also look at other equilibria and confirm that they look qualitatively similar.

5.3 The Reticent Equilibrium Characterization

In this section we characterize $\mathcal{T}_R$, the firm's advertising strategy in the reticent equilibrium, and state our main result, which relates to this characterization. We begin by reminding that $\mathcal{T}_R(h)$ is constructed using the lower bound on the ad effect, $\underline{A}(w_i)$, which compares consumer i's demand when she receives a good ad to her demand in the absence of one if she remains at her prior belief: $p_i^1(\emptyset) = p^0$ (which is the highest belief she can have in this case). It will prove useful to introduce notation for the "critical" consumers, for whom one of these demands is at the boundary:

$$\begin{aligned} \underline{w} &\equiv \max\{w \mid \mathcal{D}_f(p^0, w) = 0\}, \\ \bar{w} &\equiv \min\{w \mid \mathcal{D}_f(p_i^1(h), w) = 1\}. \end{aligned} \tag{12}$$

Here $\underline{w}$ denotes the type of consumer, who is moderately pessimistic about the product – when her interim belief equals p^0, she 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 is moderately optimistic – after receiving a good ad, she 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 (see Lemma 4 in the Appendix), and $\mathcal{D}_f(p, 1) = 1$, $\mathcal{D}_f(p, 0) = 0$ for any $p \in (0, 1)$. We consider two cases depending on the relation between these two thresholds, as defined further.

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

Lemma 3. *The following are equivalent:*

(a) ad is strong;

(b) $\underline{w} \geq \bar{w}$;

(c) $\rho \geq e^{\frac{1}{\lambda}}/(1 + e^{\frac{1}{\lambda}})$.

That is, an ad is strong if and only if there exists some consumer type $w \in [\bar{w}, \underline{w}]$ – one that faces sour conditions without an ad, and would enter sweet conditions after an ad – equivalently, if and only if $\underline{w} \geq \bar{w}$. Another way to frame this division is in terms of a bound on ρ: to be considered strong, an ad should be *informative enough* to persuade a consumer who was so pessimistic as to not even investigate the product without an ad to buy the product without investigation after an ad.

If an ad is strong, then the ad effect $\underline{A}(w)$ is trivially maximized at $w \in [\bar{w}, \underline{w}]$: advertising to these consumers converts them from sour to sweet conditions, i.e., converts them from non-sales into guaranteed sales. This is the largest effect one could feasibly achieve, and it would gradually fade away for $w < \bar{w}$ or $w > \underline{w}$. It is more interesting, however, how the ad effect behaves when the ad is weak.

If an ad is weak, it is immediate that the ad effect $\underline{\mathcal{A}}(w)$ is, all else equal, increasing in w for $w \leq \underline{w}$: all of these consumers are confidently ignoring the product in the absence of an ad ($\mathcal{D}_f(p^0, w) = 0$), but if the ad moves them to normal conditions, their interim demand $\mathcal{D}_f(p_i^1(h), w)$ is increasing in w. Similarly, $\underline{\mathcal{A}}(w)$ is decreasing in w for $w \geq \bar{w}$: all of those consumers buy the product after an ad ($\mathcal{D}_f(p_i^1(h), w) = 1$), but if they are in the normal conditions in the absence of an ad, their demand $\mathcal{D}_f(p^0, w)$ is increasing in w, hence the difference between the two demands is decreasing in w.

Yet it is not obvious how the ad effect behaves when the ad is weak for consumers with $w \in [\underline{w}, \bar{w}]$, who are in the normal conditions with and without the ad. Lemma 5 in the appendix demonstrates that the ad effect is convex in w in that region, implying that either $\underline{w}$, or $\bar{w}$, or both are the local maxima of the $\underline{\mathcal{A}}(w)$. The convexity intuitively means that the firm would rather randomize between sending an ad to one of two consumers with relatively extreme idiosyncratic valuations (so long as $w \in [\underline{w}, \bar{w}]$) than send an ad to a consumer with an average w for sure. In other words, ads are relatively more impactful closer to the edges of this interval, where a consumer is relatively certain about which decision is optimal and hence acquires little information. Small changes in consumer's beliefs then have large repercussions for how much information she acquires and, hence, for her demand.

This leads us to our main result: the claim that the optimal advertising strategy in the reticent equilibrium has a bipartite structure, targeting groups of consumers centered at $\underline{w}$ and $\bar{w}$ (though one of two groups might be empty if the firm's belief is sufficiently close to 0 or 1, as discussed below, and if the advertising cost c is low then the two groups can merge into one).

Theorem 1. *If $\mathcal{T}_R(h) \neq \emptyset$, then it satisfies the following for some $\underline{w}_d, \underline{w}_u, \bar{w}_d, \bar{w}_u$:*

1. *if the ad is weak: $\mathcal{T}_R(h) = \{i | w_i \in (\underline{w}_d, \underline{w}_u) \cup (\bar{w}_d, \bar{w}_u)\}$, 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;*
2. *if the ad is strong: $\mathcal{T}_R(h) = \{i | w_i \in (\bar{w}_d, \underline{w}_u)\}$, with $\bar{w}_d \leq \bar{w} \leq \underline{w} \leq \underline{w}_u$.*

Figure 3 illustrates the result: it plots the reticent equilibrium for $\lambda = 0.35$, $p^0 = 1/3$, $\rho = 2/3$, and $c = 0.22$. Panel (a) depicts the expected sale probabilities after an ad, $\mathcal{D}_f(p_i^1(h), w_i)$, as a function of w_i, and the best-case demand (best for the firm) in the absence of an ad, $\mathcal{D}_f(p_i^0, w_i)$. Panel (b) shows the ad effect lower bound, $\underline{\mathcal{A}}(w_i)$, as a function of w_i, as well as the equilibrium advertising strategy

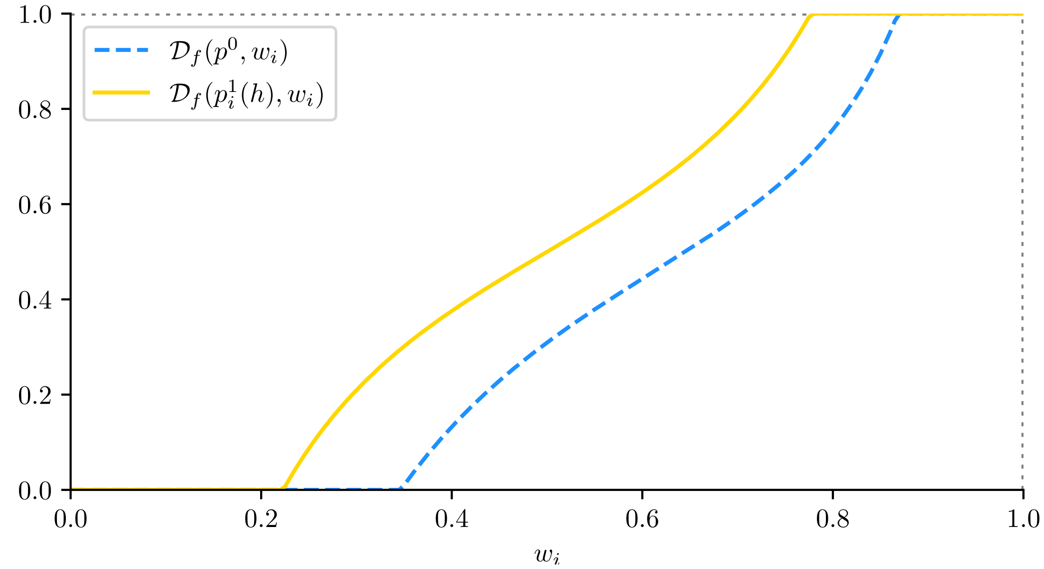

(a) Sale probabilities before and after an ad, as expected by the firm.

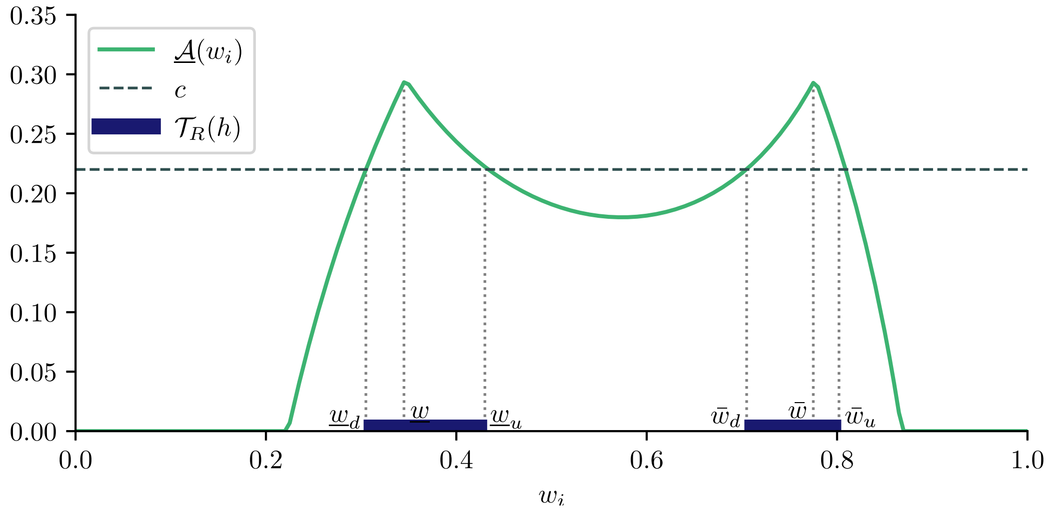

(b) Ad effect lower bound $\underline{\mathcal{A}}(w_i)$ and the optimal advertising strategy $\mathcal{T}_R(h)$.

Figure 3: The reticent 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 demand from consumers with different w_i after receiving an ad and the best-case estimate without an ad. The green solid line in panel (b) shows the difference between the two probabilities – the lower bound for the ad effect, $\underline{\mathcal{A}}(w_i)$. Presented in dark blue in panel (b) is $\mathcal{T}_R(h)$: the set of consumers, in terms of their types w_i, that the firm targets in equilibrium.

$\mathcal{T}_R(h) = \{i|\underline{\mathcal{A}}(w_i) > c\}$. 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 1 claims that this twin-peak structure holds regardless of the parameter values.

The right peak in Figure 3b 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. A 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.

The nice twin-peak nature of Figure 3 is partly due to the firm being unsure of the quality of its own product (the parameters in Figure 3 are such that $p_f = 1/2$). 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 $\mathcal{T}_R$ focuses more on one of the two groups. A confident firm (with a high p_f) prefers to target either sour consumers (low w_i), or normal consumers who would not put much effort into investigating its quality. Such a firm realizes that any consumer that investigates its product will realize that it is worth buying. Hence the firm tries to induce the consumers to pay attention to its product. A firm that is skeptic (has low p_f) has the opposite incentives: it would prefer to discourage information acquisition as much as possible, since it knows that by looking closely at its product the consumers will mostly get discouraged 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. We do not frame this statement as a formal result, but instead illustrate the intuition behind it using Figure 4.[9]

[9] Figure 4 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 $\underline{\mathcal{A}}(w)$ 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.

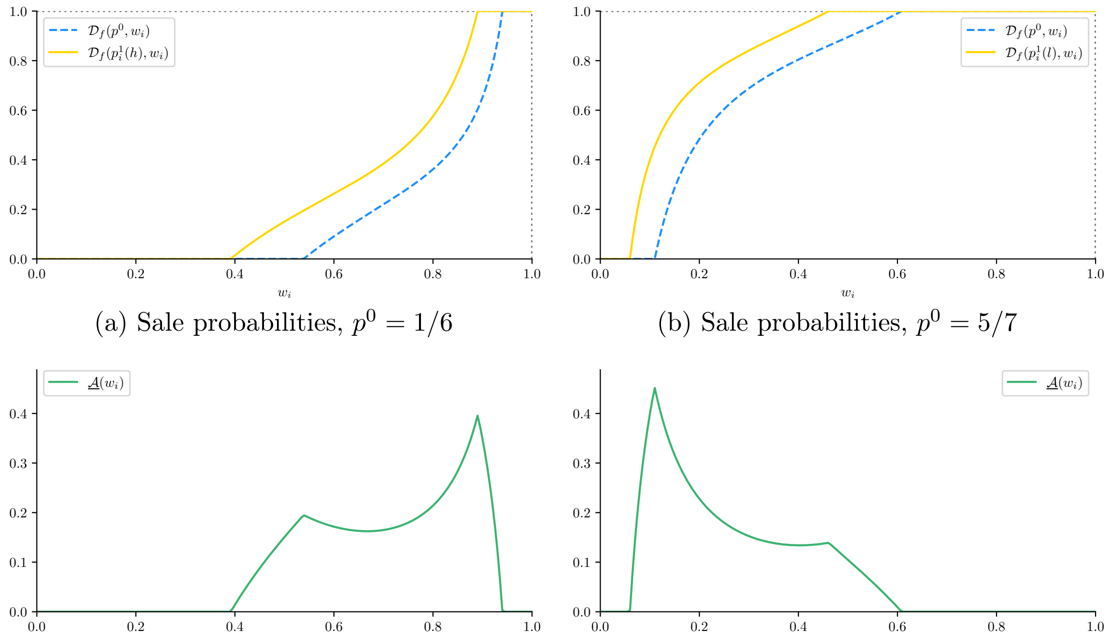

(a) Sale probabilities, $p^0 = 1/6$

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

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

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

Figure 4: Effects of the prior belief p^0 ($\lambda = 0.35$, $\rho = 2/3$).

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

6 Analysis: Other Equilibria

6.1 Equilibrium Multiplicity

The reticent equilibrium strategy $\mathcal{T}_R$ is not the only advertising strategy that could be a part of a sequential equilibrium. 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 $p_i^1(\emptyset) = p_i^1(l)$ to $p_i^1(h)$ – since by not receiving the ad she infers that $y = l$, – while in the latter case the increase is from $p^0 > p_i^1(l)$ to $p_i^1(h)$.

Proposition 1 demonstrated that the logic above implies that the reticent equilibrium features the least amount of advertising among all equilibria. However, the main idea of our result applies equally well to other equilibria. For example, the following proposition demonstrates that the equilibrium with the most advertising (exists and) has the exact same structure as the reticent equilibrium.

Proposition 2. *There exists a sequential equilibrium with the firm's advertising strategy* $\bar{\mathcal{T}}(l) = \emptyset$ *and* $\bar{\mathcal{T}}(h) = \{i|\bar{\mathcal{A}}(w_i) > c\}$*. For any* $\mathcal{T}$ *played in any other*

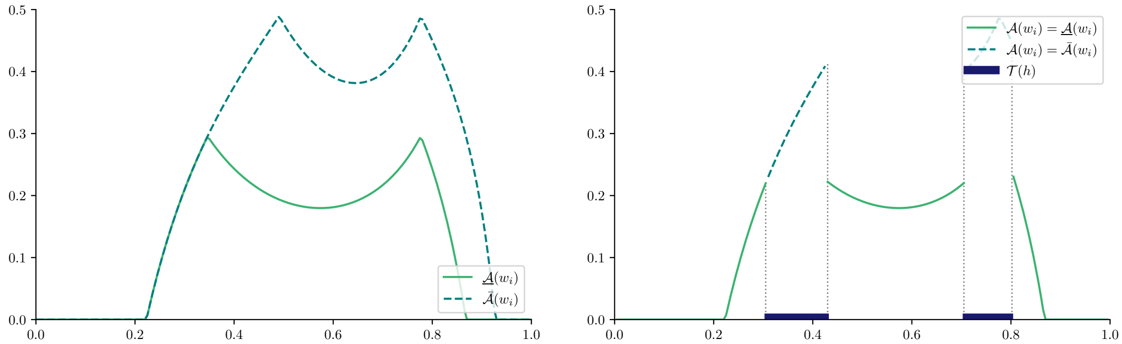

Figure 5: Ad effect bounds $\underline{\mathcal{A}}(w_i)$ and $\bar{\mathcal{A}}(w_i)$ (left panel); the actual ad effect in the reticent equilibrium (right panel).

sequential equilibrium, $\mathcal{T}(y) \subseteq \bar{\mathcal{T}}(y)$ *for any* $y \in \{h, l\}$*. Furthermore, strategy* $\bar{\mathcal{T}}$ *satisfies properties 1-2 in Theorem 1.*[10]

Recall from Lemma 2 that it is optimal to target consumer $i \in \mathcal{I}$ if $\underline{\mathcal{A}}(w_i) > c$ and only if $\bar{\mathcal{A}}(w_i) > c$. Similarly to how we constructed the reticent equilibrium by only looking at the smallest possible ad effect, we can construct an equilibrium using the upper bound $\bar{\mathcal{A}}(w_i)$. Consider in particular the targeting strategy $\bar{\mathcal{T}}(h) = \{i | \bar{\mathcal{A}}(w_i) > c\}$ from Proposition 2. It is immediate that non-targeted consumers are not worth advertising to, since even in the best case the ad effect does not cover the cost. Hence to verify that $\bar{\mathcal{T}}$ constitutes an equilibrium, we need to show that the targeted consumers are worth targeting – which is true because in the absence of an ad they infer that $y = l$ and update their belief accordingly, $p_i^1(\emptyset) = p_i^1(l)$, and this, in turn, implies that $\mathcal{A}(w_i) = \bar{\mathcal{A}}(w_i)$. Proposition 2 and Figure 5 then show that $\bar{\mathcal{A}}(w)$ satisfies all the same properties that $\underline{\mathcal{A}}(w)$ does, and yields a targeting strategy that looks qualitatively the same as in the reticent equilibrium.

The reasoning above highlights that due to the fact that $\bar{\mathcal{A}}(w) \geq \underline{\mathcal{A}}(w)$ for all w, a large variety of targeting sets can be self-sustaining in equilibrium. If a consumer expects to receive an ad, it is more profitable to target this consumer than if she was unexpecting, since her belief in the absence of an ad is lower in the former case. While we cannot claim that the optimal advertising strategy in *any* such equilibrium has the neat bipartite structure outlined in Theorem 1, 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.

[10] The criteria for strong and weak ads would be different in this case: an ad would be strong if there exists w such that $\mathcal{D}_f(p_i^1(l), w) = 0$ and $\mathcal{D}_f(p_i^1(h), w) = 1$, and weak otherwise.

6.2 Firm Profit Across Equilibria

While the previous section characterizes the full set of equilibria, it offers no insights into which equilibria are more or less preferred by the firm and the consumers, which is what this and the following sections focus on. We begin by looking at the firm's preferences over equilibria. The equilibria differ in terms of the targeting sets – i.e., any given consumer may receive a good ad and make poor inferences in its absence in one equilibrium, and receive no ads at all and stay with her prior belief in another equilibrium. Therefore, we need to calculate the difference in firm's profit between these two cases, as a function of the consumer's location w_i.

In particular, let $\mathcal{A}_E(w_i)$ denote the ex ante expected effect of advertising to consumer i, which accounts for the negative inference that the consumer makes in the absence of an ad:

$$\mathcal{A}_E(w_i) \equiv \left[p^* \mathcal{D}_f(p_i^1(h), w_i) + (1-p^*)\mathcal{D}_f(p_i^1(l), w_i)\right] - \mathcal{D}_f(p^0, w_i),$$
$$\text{where } p^* \equiv p^0\rho + (1-p^0)(1-\rho).$$

I.e., $\mathcal{A}_E(w_i)$ compares the expected demand from a consumer i who receives the ad when $y = h$ *and* updates her beliefs correctly upon not receiving one when $y = l$ to the expected demand when i never receives any information and sticks with her prior belief p^0. Here, p^* is the unconditional probability of generating a high signal $y = h$. For it to be (ex ante) profitable for the firm to advertise to consumer i, it must be that $\mathcal{A}_E(w_i) > p^* c$: the expected ad effect should exceed the expected advertising cost, where the latter is given by $p^* c$, since the actual ad is only sent if $y = h$. We can write this optimality condition down as

$$\begin{aligned} &\mathcal{A}_E(w_i) - p^* c > 0 \\ \iff\ & p^* \left[\underline{\mathcal{A}}(w_i) - c\right] - (1-p^*)\bar{\mathcal{A}}(w_i) > 0 \\ \iff\ & \underline{\mathcal{A}}(w_i) > c + \frac{1-p^*}{p^*}\bar{\mathcal{A}}(w_i). \end{aligned}$$

The expression above implies that the firm's *ex ante expected* net profit from advertising to i is lower than the ex post ad effect when i does not expect an ad in equilibrium (and the latter ad effect is weakly lower than the ad effect for any reasonable belief that i may have, by Lemmas 1 and 4). In other words, the firm would prefer to advertise even less than it does in the reticent equilibrium.

The firm thus favors equilibria with the least advertising. By Proposition 1, it then prefers the reticent equilibrium to any other equilibrium. The reason is that the ex post benefit of advertising to i is larger than the ex ante expected benefit, since the former does not account for the downside – the negative inference in the absence of an ad that arises in equilibrium. Figure 6 plots the ex ante expected ad effect $\mathcal{A}_E(w_i)$ and the interim lowed bound for the ad effect after $y = h$, $\underline{\mathcal{A}}(w_i)$, together with the marginal cost of advertising and the expected cost of advertising. The figure demonstrates that if the firm had commitment power ex ante, it would ideally advertise to some low-w consumers (yet less than it does in the reticent equilibrium), but it actually dislikes advertising to high-w consumers, with the expected ad effect being negative even ignoring the advertising costs.

We can also think of the firm's interim preferences: which equilibrium would the firm prefer to be in once it has observed its private signal y? If $y = l$, it is immediate that since the firm never advertises with $y = l$, it prefers to be in the equilibrium with the least advertising (or, ideally, none at all) to minimize the number of consumers expecting an ad and becoming disenchanted in the product if none arrives. If, on the other hand, $y = h$, then the firm prefers to be in the reticent equilibrium, since the trade-off we are thinking of – sending $a_i = h$ at cost c to increase i's belief to $p_i^1(h)$ or leaving i with her prior p^0 in no-ad-for-i equilibria – is exactly the trade-off that the firm faces in the reticent equilibrium. Hence the reticent equilibrium represents the bliss point for the firm with $y = h$, but features too much advertising for the firm with $y = l$, which in expectation amounts to the conclusion above regarding the firm's ex ante preferences. All of the arguments above are summarized by and prove the following proposition.

Proposition 3. *Of all the equilibria of the game, the firm's ex ante and interim expected profit after any $y \in \{h, l\}$ is maximized in the reticent equilibrium.*

The shape of $\mathcal{A}_E(w)$ deserves a separate investigation, since it can be interpreted as the firm's expected profit from the ad if it could *commit* ex ante to its targeting strategy. This variation can be interpreted as a model of *informative* (as opposed to persuasive) advertising, in which the firm decides whether to reveal a piece of information about the product that would inform the consumer about her fit with the product, but the firm itself does not know ex ante the consumer's fit y (but knows the consumer's general affinity towards the product, w_i). If the shape of $\mathcal{A}_E(w)$ presented in Figure 6 is general (which it is, see Proposition 4 below), then in such a setting, the firm strictly prefers to target only a small set of the skeptics and none of the believers – tipping the scales from normal to sweet conditions is never worth it for the firm, since the downside risk always outweighs the upside benefit. The question is why exactly $\mathcal{A}_E(w)$ looks the way it does.

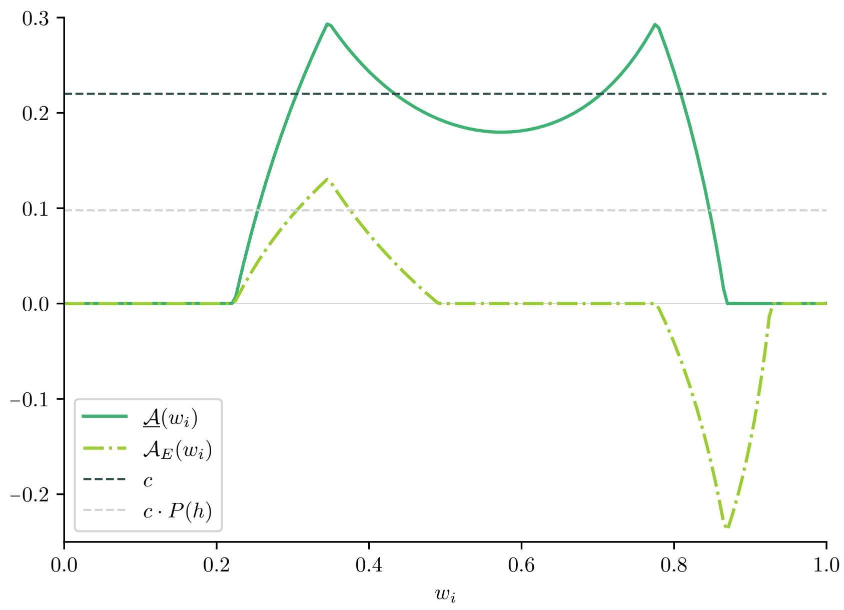

Figure 6: The expected ad effect.

Notes: the parameter values are the same as in Figure 3. The light-grey dashed line plots the expected cost of advertising (realized by sending an ad iff $y = h$).

The idea is that if the consumer is in sour conditions in the absence of an ad ($w \leq \underline{w}$, where $\underline{w}$ is defined by (12)), then $\mathcal{A}_E(w) \geq 0$, since the firm gets the positive effect of good news and does not suffer from revealing bad news: $\mathcal{D}_f(p_i^1(l), w) = \mathcal{D}_f(p^0, w) = 0$ in that case. Symmetrically, $\mathcal{A}_E(w) \leq 0$ for $w \geq \bar{w}$, since the opposite happens: $\mathcal{D}_f(p_i^1(h), w) = \mathcal{D}_f(p^0, w) = 1$, and some of that demand may be lost as a result of bad news.[11] For $w \in (\underline{w}, \bar{w})$, if w is close to either edge of the interval, the conclusions above hold since $\mathcal{A}_E(w)$ is continuous. Proposition 4 below formalizes this intuition. An interested reader is welcome to refer to the proof in the appendix for the intuition behind the flat middle interval.

Proposition 4. *The expected ad effect $\mathcal{A}_E(w)$ is such that, for some $\underline{w}_d, \underline{w}_u, \bar{w}_d, \bar{w}_u$:*

$$\begin{aligned}
\mathcal{A}_E(w) &= 0 \quad && \text{if } w \geq \bar{w}_u, \\
\mathcal{A}_E(w) &< 0 \quad && \text{if } w \in (\bar{w}_d, \bar{w}_u), \\
\mathcal{A}_E(w) &= 0 \quad && \text{if } w \in [\underline{w}_u, \bar{w}_d], \\
\mathcal{A}_E(w) &> 0 \quad && \text{if } w \in (\underline{w}_d, \underline{w}_u), \\
\mathcal{A}_E(w) &= 0 \quad && \text{if } w \leq \underline{w}_d.
\end{aligned}$$

Further, $\underline{w}_d < \underline{w} < \underline{w}_u, \bar{w}_d < \bar{w} < \bar{w}_u$, and if the ad is sufficiently strong, then

[11] Note that $\mathcal{A}_E(w) = 0$ for very high and very low w, since in those cases a contrarian signal can not overpower the consumer's predisposition, so $\mathcal{D}_f(p_i^1(h), w) = \mathcal{D}_f(p^0, w) = \mathcal{D}_f(p_i^1(l), w) \in \{0, 1\}$.

$\underline{w}_u = \bar{w}_d$.[12]

6.3 Consumer Welfare Across Equilibria

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 main question to ask is: do ads benefit consumers? The answer in our model is "yes". Being in the target group is equivalent to receiving a free signal: if $y = h$ a consumer receives an ad, and if $y = l$ she can perfectly infer this from the lack of an ad. Thus the consumers in the target group can effectively observe the realization of the firm's signal y. The Blackwell principle states that free information is always beneficial for decision-makers, hence consumers benefit from being targeted with advertisements. This implies the following proposition.

Proposition 5. *Suppose $\mathcal{T}(l) = \emptyset$. Consumer i's ex ante expected utility is larger if $i \in \mathcal{T}(h)$ than if $i \notin \mathcal{T}(h)$.*

This proposition suggests that the equilibrium with the largest amount of advertising described in Proposition 2 is the one most preferred by consumers. Proposition 5 implies that we can proxy consumer welfare in our model with $|\mathcal{T}(y)|$ – the number of consumers who receive ads. The more advertising there is in equilibrium, the happier consumers are in aggregate. Therefore, the answer to the broader question of welfare implications of targeting technologies depends on whether these technologies lead to more or less total advertising.

The result above, however, is at odds with the popularity of ad-blockers among Internet users in the real world – while online ads are often perceived as nuisance by the consumers, they should still be perceived as a welcome nuisance if the information therein was useful to the consumers.[13] A possible explanation for this discrepancy is the partial equilibrium nature of our model, which ignores the firm's pricing decisions (or other consumer utility-relevant decisions in case of zero-price products) and the effect that advertising may have on them. The conclusion presented in Proposition 5 should thus be viewed critically.

[12] To clarify, thresholds $\underline{w}_d, \underline{w}_u, \bar{w}_d, \bar{w}_u$ are different from those in Theorem 1. "Sufficiently strong" means strong in the sense of Proposition 2 (see Footnote 10).

[13] According to Perrin [2019], in 2019 about a quarter of all Internet users in the US and Europe were using some kind of ad-blocking software.

6.4 Cursed Consumers

In this section we consider the firm's problem when consumers are *cursed* in the sense of Eyster and Rabin [2005]. They define cursedness as a failure on the player's behalf to fully comprehend the correlation between other players' actions and their information. In our setting, a cursed consumer reacts to advertisements like Bayesian consumers would, but does not update her beliefs when she does not receive an advertisement. We argue that the assumption that consumers could be cursed is particularly fitting in our setting (like in most disclosure settings). Indeed, Li and Hitt [2008] and Brown, Camerer, and Lovallo [2012] provide empirical evidence from various markets that consumers typically infer too little from the lack of signal but infer correctly after observing a signal. Jin, Luca, and Martin [2021] and Deversi, Ispano, and Schwardmann [2021] have obtained further evidence in the lab suggesting very strongly that in games of disclosure of verifiable information (such as ours), a significant share of receivers can be classified as cursed. I.e., they infer too little from the lack of signal – even after playing a number of rounds as the sender and playing fully rationally (i.e., concealing adverse information) in those rounds. Markets with cursed consumers have also been explored theoretically by Matysková and Šípek [2017] and Ispano and Schwardmann [2018].

Formally, in addition to the sequential equilibria of the game, we also seek to characterize cursed equilibria as defined below. As shown further, the game turns out to have a unique cursed equilibrium, which coincides in terms of strategies (but not beliefs) with the reticent equilibrium.

Definition (Cursed Equilibrium)**.** *A Cursed Equilibrium of the game consists of the firm's advertising strategy $\mathcal{T}(y)$, the collection of the consumers' information acquisition strategies $\{\mathcal{G}_i(x_i|s, a_i)\}_{i\in\mathcal{I}}$, and the updating rules $p_f(y)$, $p_i^1(a_i)$, and $p_i^2(x_i, a_i)$, for the firm's and the consumers' beliefs respectively, such that:*

1. *$p_i^1(\emptyset) = p^0$ for all i;*
2. *all other objects satisfy the requirements imposed by sequential equilibria.*

It is easiest to interpret a "cursed equilibrium" as an "equilibrium when all consumers are cursed". If consumer i is cursed and she receives an ad with signal h or l then her belief changes to $p_i^1(h)$ or $p_i^1(l)$ 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 in a cursed equilibrium.

Lemma 1 continues to hold in cursed equilibria (the expressions for the beliefs $p_i^1(a_i)$ for $a_i \in \{l, h\}$ hold true regardless of the firm's strategy, and the proof of $\mathcal{T}(l) = \emptyset$ relies on $p^0 > p_i^1(l)$, which holds by assumption in Cursed Equilibria). We can thus focus on the firm's decision when it has a high signal $y = h$ in hand. The probability that (a cursed) consumer i purchases the product is given by $\mathcal{D}_f(p_i^1(h), w_i)$ after ad h and $\mathcal{D}_f(p_i^1(\emptyset), w_i) = \mathcal{D}_f(p^0, w_i)$ 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

$$\left[\int_{i\in\mathcal{T}(h)} \mathcal{D}_f(p_i^1(h), w_i)di + \int_{i\in\mathcal{I}\setminus\mathcal{T}(h)} \mathcal{D}_f(p^0, w_i)di\right] - c \cdot |\mathcal{T}(h)| =$$
$$= \int_{i\in\mathcal{T}(h)} \big(\mathcal{D}_f(p_i^1(h), w_i) - \mathcal{D}_f(p^0, w_i) - c\big)\, di + \int_{i\in\mathcal{I}} \mathcal{D}_f(p^0, w_i)di.$$

Therefore, it is optimal for firm to include i in the target group if and only if the expression in the first integral is positive for this consumer, which is equivalent to $\underline{\mathcal{A}}(w_i) > c$, where $\underline{\mathcal{A}}(w_i)$ is given by (11). This means that the advertising strategy in the cursed equilibrium coincides with that in the reticent (sequential) equilibrium. Notably, the resulting strategy $\mathcal{T}(h) = \{i|\mathcal{A}(w_i) > c\}$ is uniquely optimal given the consumers' beliefs, and the beliefs of cursed consumers are independent of the firm's strategy (conditional on a_i), hence the equilibrium is also unique.[14]

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

Proposition 6. *There is a unique cursed equilibrium, and the firm's optimal advertising strategy in this equilibrium is the same as in the reticent equilibrium:* $\mathcal{T}_R(l) = \emptyset$ *and* $\mathcal{T}_R(h) = \{i|\mathcal{A}(w_i) > c\}$.

Finally, the conclusion of Proposition 5, stating that consumers benefit from advertisements, continues to hold for cursed consumers. If $y = h$ then targeted cursed consumers receive a piece of accurate information and make a decision that is better on average than if they observed no ad/signal. If, on the other hand, $y = l$, then cursed consumers are left with their prior belief p^0, same as if they were not targeted. Thus cursed consumers win from being targeted if $y = h$ and lose nothing if $y = l$, hence they win on average.

[14] Unique up to indifference – i.e., other equilibria may exist in which the firm or some consumers choose a different action when indifferent between some actions available to them (advertise or not to consumer i; buy the product or not). It is, however, straightforward that all such equilibria are payoff-equivalent for all parties.

7 Discussion

We now discuss the assumptions made by the model and relate them to the literature.

Hotelling model and vertical differentiation. Our model features a combination of both horizontal and vertical differentiation. Vertical differentiation via the product quality s is the main focus of the model, since both the firm's advertising decision and the consumers' information acquisition decisions concern information about s. It can, however, be replaced by idiosyncratic variance in consumers' tastes (so that $s_i \in S$ is i.i.d. across consumers), with the firm receiving a signal $y_i \in Y$ about each consumer's match value. All calculations and results would be the same in this case.

Horizontal differentiation via w_i, on the other hand, is effectively used for comparative statics: we assume that both the firm and consumer i know w_i and the advertising costs are linear, so the interaction between the firm and consumer i is not affected by the firm's interaction with another consumer j. The Hotelling model is used due to it being one of the standard frameworks for horizontal differentiation. Yet the exact way in which horizontal differentiation is captured is not important, as demonstrated by the earlier version of the paper using heterogeneity in the consumers' beliefs instead of preferences. Similarly, the exact distribution of w_i on the Hotelling line is irrelevant, but uniform is specified for concreteness.

Finally, we put the firm at the end of the interval purely for notational convenience. If a monopolist could choose the location, it would prefer to be located at the center of the interval, but the two cases are qualitatively equivalent. We can extend the market to $[0, 2]$ so that the monopolist is in the center; this would not affect our analysis of the (sub)market $[0, 1]$, and the analysis of $[1, 2]$ would be equivalent up to symmetry.

No price/quality choice by the firm. The zero-price assumption is made to motivate the setting in which advertising is the only margin of control that the firm has. The results continue to hold if the firm controls both the price and the advertising strategy, since we can decompose such a problem into the firm choosing the price first and making the advertising decisions later, with our model capturing the second stage of this decision problem given any fixed price level. This interpretation does, however, require the assumptions that the firm is unable to price discriminate based on w_i (which does indeed appear to be more a challenging task than targeted advertising, since under targeted pricing the firm needs to shut down arbitrage opportunities among consumers). Price signaling some information about y is another concern that would need to be addressed in a full pricing model.

The same applies to quality choice: while some models of vertical differentiation incorporate the firm's quality choice, we treat it as given since it is made on a very different time scale – by the time the advertising decisions are made, the production process is well underway, and quality is difficult to affect. Quality choice can be incorporated into the model, but this decision would not interact with the rest of the model in any meaningful way and would not affect the results.

Disclosure/Persuasive advertising. In the classification of Bagwell [2007] and Renault [2015], advertising (and models thereof) can be roughly split into *persuasive* and *informative*. The latter refers to advertising that either creates awareness of the product, or conveys product information that may be interpreted favorably by some consumers and negatively by others. While we do briefly consider (the latter form of) informative advertising in Section 6.2, our main focus is on persuasive advertising, which improves the consumers' opinion about the product. To reiterate, we assume the ad contents y to be a piece of verifiable evidence about the firm's product (or the firm itself), such as expert certifications ("Certified organic!"), results of external audits and evaluations ("Nine out of ten dentists recommend!"), critics' opinions ("«Magnificent!» –New York Times"), or consumer reviews ("Rated 90 on Rotten Tomatoes!"). The realization of such a signal can persuade consumers that s is high (or low), and the firm can decide whether to disclose any given realization. Alternatively, an earlier version of this paper considered advertising that did not affect the consumers' *beliefs*, but rather increased consumers' *valuations* for the good, be it due to directly affecting consumer tastes [Dixit and Norman, 1978] or creating some kind of "hype" [Johnson and Myatt, 2006]. We did show that Theorem 1 holds in that setting as well.

Linear cost of advertising. Our model assumes that the cost of advertising is linear in campaign size. The literature commonly assumes convex cost of reach (e.g., Butters [1977], Stahl [1994]). Yet in our setting, such a convexity would more reasonably correspond to a deeper penetration within a cohort of consumers in a given "location" w. In turn, covering broader locations should, if anything, be cheaper per recipient than sending highly targeted ads. It is therefore up for debate which cost function is more appropriate in this setting. Regardless, our main focus is on the benefits from advertising, hence we adopt the simplest functional form for the costs of advertising.

Entropy cost of information for consumers. Following the vast literature on rational inattention (see Maćkowiak et al. [2018] for one survey), we adopt the entropy cost of information. This could capture both cognitive costs and the time and effort costs of collecting information. To clarify, "inattention" in our setting relates not to the consumer being inattentive w.r.t. the ad that the firm sends (the ad is instead salient enough to be perceived and comprehended fully and at no cost). We rather interpret it as the consumer being inattentive to *other* information about the product that is freely available, but requires a costly effort to sift through – e.g., detailed product description, customer reviews, critics' opinions, etc.

One does not need to interpret the entropy costs literally: Hébert and Woodford [2017] and Morris and Strack [2019] show that under certain conditions, the classic Wald problem (that features a decision-maker sequentially sampling costly i.i.d. signals) can be represented as a static problem with information costs given by mutual information, as in our model. Further, its use has been justified both axiomatically and through links to optimal coding in information theory [Denti, Marinacci, and Montrucchio, 2019]. Finally, Matějka and McKay [2015] show that entropy cost generates logit-like choice rules, which is especially fitting for our model of consumer behavior, seeing how logit is the workhorse of the empirical IO literature. Due to all of the above, we elected to stay close to the literature and assume entropy cost of information, as opposed to adopting an ad hoc functional form. This is despite the entropy cost function admittedly also possessing some questionable properties (such as the cost of a given signal structure $\mathcal{G}_i(x_i|s, a_i)$ depending on the consumer's belief p_i^1, see Mensch [2018], Denti, Marinacci, and Rustichini [2022]). That said, the example explored in Section 2 demonstrates that our results are not specific to entropy cost of information.

Monopoly. For simplicity, we work with a monopolistic market, which raises the question of whether our results extend to competitive markets. This question is made more relevant by the fact that one of the primary goals of advertising is to give the firm an advantage over its competitors. One response to this is that our model can be interpreted as a partial equilibrium of this richer setting, with the consumers' outside options incorporating everything related to competing products. While this seemingly ignores the possible strategic interactions in advertising decisions between the competing firms, we showed in an earlier version of this paper that if the consumers' choice problem remains binary (whether to buy one firm's product or the other firm's product), such strategic interaction is minimal. The optimal advertising strategies in such a duopoly are then qualitatively the same as those we obtain for monopoly.

8 Conclusion

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. We further show that the model features multiple equilibria (among which the structure above applies to the two with the least and the most advertising). Perhaps surprisingly, among those, the firm prefers the equilibrium with the least advertising (but not necessarily a total ban on advertising), while the consumers prefer the one with the most advertising.

In our analysis we make a connection to the behavioral economics 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 optimal strategy for the sender.

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. Further, to focus on advertising, our model abstracts from some relevant aspects of the setting, such as flexible pricing (including personalized pricing/price discrimination), quality choice, and competition among firms, as discussed in Section 7. Exploring these and other extensions in greater detail would be important for understanding the scope of the applicability of our results and constitute prospective avenues for further work.

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.

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.

K. Bagwell. Chapter 28 The Economic Analysis of Advertising. In M. Armstrong and R. Porter, editors, *Handbook of Industrial Organization*, volume 3, pages 1701–1844. Elsevier, January 2007. doi: 10.1016/S1573-448X(06)03028-7.

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.

T. Blake, C. Nosko, and S. Tadelis. Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica*, 83(1):155–174, 2015. ISSN 1468-0262. doi: 10.3982/ECTA12423.

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.

R. Burguet and V. Petrikaite. Targeted advertising and costly consumer search. Working paper, 2020.

G. R. Butters. Equilibrium Distributions of Sales and Advertising Prices. *Review of Economic Studies*, 44(3), 1977.

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.

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

X. Cheng. Deterrence, Diversion, and Encouragement: How Freely Provided Information May Distort Learning. working paper, 2021.

T. Denti, M. Marinacci, and L. Montrucchio. A note on rational inattention and rate distortion theory. *Decisions in Economics and Finance*, pages 1–15, April 2019. ISSN 1593-8883, 1129-6569. doi: 10.1007/s10203-019-00243-0.

T. Denti, M. Marinacci, and A. Rustichini. Experimental cost of information. *American Economic Review*, 112(9):3106–23, September 2022. ISSN 0002-8282. doi: 10.1257/aer.20210879.

M. Deversi, A. Ispano, and P. Schwardmann. Spin doctors: An experiment on vague disclosure. *European Economic Review*, 139:103872, 2021. ISSN 0014-2921. doi: 10.1016/j.euroecorev.2021.103872.

A. Dixit and V. Norman. Advertising and welfare. *The Bell Journal of Economics*, pages 1–17, 1978. Publisher: JSTOR.

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.

N. Figueroa and C. Guadalupi. Testing the sender: When signaling is not enough. *Journal of Economic Theory*, 197:105348, October 2021. ISSN 0022-0531. doi: 10.1016/j.jet.2021.105348. Publisher: Academic Press.

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. 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.

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.

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 rational (in)attention to favorable and unfavorable product information, and firm information design. *Journal of Marketing Research*, 58(2):343–362, 2021. doi: 10.1177/0022243720977830.

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*, 13(2):141–73, 2021. doi: 10.1257/mic.20180217.

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

P. Lewis and P. Hilder. Leaked: Cambridge analytica's blueprint for trump victory. *The Guardian*, March 2018. URL `https://www.theguardian.com/uk-news/2018/mar/23/leaked-cambridge-analyticas-blueprint-for-trump-victory`. Retrieved Apr 18, 2023.

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 and A. Montes. Bayesian Persuasion with Costly Information Acquisition. Discussion paper no. 296, University of Bonn, 2021.

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.

J. Mensch. Cardinal Representations of Information. *SSRN Electronic Journal*, 2018. ISSN 1556-5068. doi: 10.2139/ssrn.3148954.

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.

M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT press, 1994.

N. Perrin. Consumer attitudes on marketing 2019. *eMarketer*, Aug 2019. URL `https://www.insiderintelligence.com/content/consumer-attitudes-on-marketing-2019`. Retrieved Apr 18, 2023.

R. Renault. Advertising in markets. In *Handbook of Media Economics*, volume 1, pages 121–204. Elsevier, 2015.

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.

A. Smirnov and E. Starkov. Bad news turned good: reversal under censorship. *American Economic Journal: Microeconomics*, 14(2):506–60, May 2022. doi: 10.1257/mic.20190379.

D. O. Stahl. Oligopolistic pricing and advertising. *Journal of economic theory*, 64 (1):162–177, 1994. Publisher: Elsevier.

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.

C. Wang. Advertising as a search deterrent. *The RAND Journal of Economics*, 48 (4):949–971, 2017. ISSN 1756-2171. doi: 10.1111/1756-2171.12209.

Appendix

Proof of Lemma 1. The first two parts follow trivially from Bayes' rule regardless of the firm's strategy – hence they apply both on and off the equilibrium path.

Let us now show that $p_i^1(\emptyset) \geq p_i^1(l)$. For some advertising strategy $\mathcal{T}_n$ of the firm, let $\tau_n(i, y) \in (0, 1)$ denote the probability that i is sent an ad y according to $\mathcal{T}_n$: i.e., $\tau_n(i, y) \equiv \mathbb{P}\{i \in \mathcal{T}_n(y)\}$. Then i's belief $p_i^1(\emptyset)$ after not receiving an ad, calculated from $\mathcal{T}_n$ using Bayes' rule, satisfies

$$\frac{p_i^1(\emptyset)}{1 - p_i^1(\emptyset)} = \frac{p^0}{1 - p^0} \cdot \frac{1 - \rho\tau_n(i, h) - (1 - \rho)\tau_n(i, l)}{1 - (1 - \rho)\tau_n(i, h) - \rho\tau_n(i, l)}. \tag{13}$$

Routine calculations confirm that the RHS of (13) is weakly decreasing in $\tau_n(i, h)$ and increasing in $\tau_n(i, l)$. Then plugging in the extreme values $(\tau_n(i, h), \tau_n(i, l)) = (1, 0)$ and $(\tau_n(i, h), \tau_n(i, l)) = (0, 1)$ yields that $p_i^1(\emptyset)$ is bounded from below by $p_i^1(l)$ and from above by $p_i^1(h)$. If $a_i = \emptyset$ is on path for i in equilibrium (meaning $(\tau(i, h), \tau(i, l)) \neq (1, 1)$ in equilibrium), then this directly yields the first part of the statement. If $a_i = \emptyset$ is off path for i, then there should exist a sequence $\{\mathcal{T}_n\}_{n \in \mathbb{N}}$ of fully mixed strategies that converges to the equilibrium strategy $\mathcal{T}$, and such that (13) holds in the limit as $n \to \infty$ for the equilibrium belief $p_i^1(\emptyset)$, which implies that $p_i^1(\emptyset) \in [p_i^1(l), p_i^1(h)]$.

The final statement of the lemma then follows trivially: from (6)-(9), sale probability $\mathcal{D}_f(p_i^1, w_i)$ is weakly increasing in p_i^1. 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. Since i was arbitrary, we also get that $\mathcal{T}(l) = \emptyset$.

Finally, $\mathcal{T}(l) = \emptyset$ then implies by (13) that $p_i^1(\emptyset) \leq p^0$, which allows us to conclude that in equilibrium, $p_i^1(\emptyset) \in [p_i^1(l), p^0]$. This concludes the proof. $\square$

Lemma 4. *Function $\mathcal{D}_f(p, w)$ is weakly increasing in w and p.*

Proof. It is immediate from (6)-(9) that $\mathcal{D}_f(p, w)$ is continuous and differentiable in both arguments almost everywhere, with the exceptions being the edge cases in (6). 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:

$$\frac{\partial \mathcal{D}_f(p, w)}{\partial w} = \frac{\partial \mathcal{D}_f(p, w)}{\partial \mathcal{D}(p, w)} \left(\frac{\partial \mathcal{D}(p, w)}{\partial w} + (1 - \mathcal{D}(p, w)) \frac{\mathcal{D}(p, w)}{\lambda} \right), \tag{14}$$

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. In the latter term of (14), $(1 - \mathcal{D}(p, w)) \frac{\mathcal{D}(p,w)}{\lambda} > 0$, 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

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

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

$\square$

Proof of Proposition 1. By Lemma 1, $\mathcal{T}(l) = \emptyset$ in all sequential equilibria, hence it only remains to show the claim for $y = h$.

To show existence, we start from the observation established in Section 5.2 that in any equilibrium, $i \in \mathbb{T}(h)$ if and only if (10) holds. With the construction for the reticent equilibrium, we have that in equilibrium, for all $i \in \mathcal{T}_R$: $p_i^1(\emptyset) = p_i^1(l)$ by the Bayes' rule, thus $\mathcal{A}(w_i) = \bar{\mathcal{A}}(w_i) \geq \underline{\mathcal{A}}(w_i) > c$, hence (10) is satisfied. For all $i \notin \mathcal{T}_R$: $p_i^1(\emptyset) = p^0$, so $\mathcal{A}(w_i) = \underline{\mathcal{A}}(w_i) < c$, hence (10) is satisfied as well. We conclude that targeting strategy $\mathcal{T}_R(y)$ does indeed generate an equilibrium.

The second part of the claim follows from the definition $\mathcal{T}_R(h) = \{i | \underline{\mathcal{A}}(w_i) > c\}$ and Lemma 2 saying that if $\underline{\mathcal{A}}(w_i) > c$ then $i \in \mathcal{T}(h)$ in any equilibrium. □

Proof of Lemma 3. The equivalence of statements (a) and (b) follows immediately from the definitions and Lemma 4, which establishes that $\mathcal{D}_f(p, w)$ is weakly increasing in w. From here onwards we focus on the monotonicity of (b) and (c).

Due to the aforementioned monotonicity, (b) is equivalent to $\mathcal{D}_f(p_i^1(h), \underline{w}) = 1$. Since $\mathcal{D}_f(p, w)$ is also weakly increasing in p, (b) is further equivalent to

$$p_i^1(h) \geq \hat{p}(\underline{w}) = e^{\frac{\underline{w}}{\lambda}} \underline{p}(\underline{w}) = e^{\frac{\underline{w}}{\lambda}} p^0, \tag{15}$$

where the cutoffs $\hat{p}$ and $\underline{p}$ are defined by (8), and the latter equality follows from the definition of $\underline{w}$. Using definition (8), specifically that $\underline{p}(w) = \frac{e^{\frac{1-w}{\lambda}} - 1}{e^{\frac{1}{\lambda}} - 1}$, we can obtain that $e^{\frac{\underline{w}}{\lambda}} = \frac{e^{\frac{1}{\lambda}}}{1 - p^0 + p^0 e^{\frac{1}{\lambda}}}$.

Further, using the Bayes' rule, we get that $p_i^1(h) = \frac{p^0 \rho}{p^0 \rho + (1 - p^0)(1 - \rho)}$, hence the inequality in (15) is equivalent to

$$\begin{aligned}
& \frac{p^0 \rho}{p^0 \rho + (1 - p^0)(1 - \rho)} \geq e^{\frac{\underline{w}}{\lambda}} p^0 \\
\iff \quad & \frac{p^0 \rho}{p^0 \rho + (1 - p^0)(1 - \rho)} \geq \frac{e^{\frac{1}{\lambda}}}{1 - p^0 + p^0 e^{\frac{1}{\lambda}}} p^0 \\
\iff \quad & \rho \left((1 - p^0) e^{-\frac{1}{\lambda}} + p^0 \right) \geq p^0 \rho + (1 - p^0)(1 - \rho) \\
\iff \quad & \rho \geq \frac{1 - p^0}{(1 - p^0) e^{-\frac{1}{\lambda}} + p^0 - 2p^0 + 1} = \frac{1}{1 + e^{-\frac{1}{\lambda}}},
\end{aligned}$$

which is equivalent to the condition presented in the statement. □

Lemma 5. *Function $\underline{\mathcal{A}}(w_i)$ 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 weakly increasing in w and p, and $p_i^1(h) > p^0$, there exist

$\bar{\bar{w}}, \underline{\underline{w}} \in [0,1]$ such that:

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

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

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

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

Proof of Theorem 1. By definition of the reticent equilibrium, $i \in \mathcal{T}_R(h) \iff \underline{\mathcal{A}}(w_i) > c$. From the assumption that $\mathcal{T}_R \neq \emptyset$ it follows that $\max_{w \in [0,1]} \underline{\mathcal{A}}(w) > c$. From the continuity of $\underline{\mathcal{A}}(w)$ together with $\underline{\mathcal{A}}(0) = \underline{\mathcal{A}}(1) = 0$, it follows that the upper contour set $\{w | \underline{\mathcal{A}}(w) > c\}$ is open for any $c > 0$. If the ad is strong then by Lemma 5 the maximum of $\underline{\mathcal{A}}(w)$ is attained by all $w \in [\bar{w}, \underline{w}]$. By the monotonicity of $\underline{\mathcal{A}}(w)$ for $w \leq \bar{w}$ and for $w \geq \underline{w}$ described in Lemma 5, part 2 of the statement follows.

If the ad is weak then by Lemma 5 $\underline{\mathcal{A}}(w)$ is strictly convex for all $w \in [\underline{w}, \bar{w}]$. Together with the monotonicity of $\underline{\mathcal{A}}(w)$ in the remaining regions, this implies that $\arg\max_{w \in [0,1]} \underline{\mathcal{A}}(w) \in \{\underline{w}, \bar{w}\}$. Consider two cases depending on whether there exists a $\tilde{w} \equiv \arg\min_{w \in (\underline{w}, \bar{w})} \underline{\mathcal{A}}(w)$. If it does, then consider $\underline{\mathcal{A}}(w)$ separately on $[0, \tilde{w}]$ and $[\tilde{w}, 1]$. On both intervals $\underline{\mathcal{A}}(w)$ is single-peaked, hence quasi-concave, meaning that its upper contour sets $\{w | \underline{\mathcal{A}}(w) > c\}$ are convex within each interval and include the respective peaks $\underline{w}$ and $\bar{w}$. If no such $\tilde{w}$ exists then $\underline{\mathcal{A}}(w)$ is strictly monotone on $[\underline{w}, \bar{w}]$. Then $\underline{\mathcal{A}}(w)$ is single-peaked on $[0,1]$, so again its upper contour set $\{w | \underline{\mathcal{A}}(w) > c\}$ is convex and includes the global maximum.[15] This proves part 1 of the theorem. $\square$

[15] The global maximum is then one of the points in $\{\underline{w}, \bar{w}\}$, but the UCS may or may not include the other point.

Proof of Proposition 2. The proof is constructive. Let $\bar{\mathcal{A}}(w)$ be defined as in (11). Let $\bar{\mathcal{T}}(l) \equiv \emptyset$ and $\bar{\mathcal{T}}(h) \equiv \{i|\bar{\mathcal{A}}(w_i) > c\}$. In what follows, we show that such a strategy $\bar{\mathcal{T}}$ satisfies the proposition requirements.

Firstly, we show that there exists a sequential equilibrium with $\bar{\mathcal{T}}$ as the firm's advertising strategy. By Lemma 1, $\bar{\mathcal{T}}(l) = \emptyset$. As for $y = h$, follow the same argument as for Proposition 1. For all $i \in \bar{\mathcal{T}}$: $p_i^1(\emptyset) = p_i^1(l)$, thus $\mathcal{A}(w_i) = \bar{\mathcal{A}}(w_i) > c$, hence (10) is satisfied. For all $i \notin \mathcal{T}_R$: $p_i^1(\emptyset) = p^0$, so $\mathcal{A}(w_i) = \underline{\mathcal{A}}(w_i) \leq \bar{\mathcal{A}}(w_i) < c$, hence (10) is satisfied as well. We conclude that targeting strategy $\bar{\mathcal{T}}(y)$ does indeed generate an equilibrium.

Secondly, to show that in any other sequential equilibrium, the firm's ad strategy $\mathcal{T}$ is such that $\mathcal{T}(y) \subseteq \bar{\mathcal{T}}(y)$ for , invoke Lemma 2, which implies that if $\bar{\mathcal{A}}(w_i)$ then $i \notin \mathcal{T}(h)$, hence $i \notin \bar{\mathcal{T}}(h) \Rightarrow i \notin \mathcal{T}(h)$.

Finally, we need to show that $\bar{\mathcal{T}}$ satisfies the properties listed in Theorem 1. Recall that $\bar{\mathcal{T}}(h) = \{i|\bar{\mathcal{A}}(w_i) > c\}$, while for the reticent equilibrium we have $\mathcal{T}_R(h) = \{i|\underline{\mathcal{A}}(w_i) > c\}$. Consider a fictitious ("prime") world in which all players share a common prior $p' \equiv p_i^1(l)$ and the precision of the firm's signal y' is $\rho' \equiv \frac{\rho^2}{\rho^2+(1-\rho)^2}$. Then $\alpha'(p') = p_i^1(h)$. But then for any i, $\bar{\mathcal{A}}(w_i) = \underline{\mathcal{A}}'(w_i)$, where $\underline{\mathcal{A}}'(w)$ is calculated in the same way as $\underline{\mathcal{A}}(w)$, except with ρ' instead of ρ. Consequently, Theorem 1, as applied to this fictitious world, implies that $\mathcal{T}_R'(h) \equiv \{i|\underline{\mathcal{A}}'(w_i) > c\}$ satisfies all the required properties. But we have $i \in \mathcal{T}_R'(h) \iff i \in \bar{\mathcal{T}}(h)$, meaning that $\bar{\mathcal{T}}(h)$ satisfies all required properties as well. $\square$

Proof of Proposition 4. Start with the case when $w < \underline{w}$. For such w, $\mathcal{D}_f(p_i^1(l), w) = \mathcal{D}_f(p^0, w) = 0$. Let $\underline{w}_d \equiv \max\{w : \mathcal{D}_f(p_i^1(h), w) = 0\}$. By (12) and the monotonicity of $\mathcal{D}_f$ (see Lemma 4), $\underline{w}_d \leq \underline{w}$. Then for $w < \underline{w}_d$: $\mathcal{D}_f(p_i^1(h), w) = 0$, and hence $\mathcal{A}_E(w) = 0$, which proves part 5 of the proposition.

Symmetrically, when $w > \bar{w}$: $\mathcal{D}_f(p_i^1(h), w) = \mathcal{D}_f(p^0, w) = 1$. Let us denote $\bar{w}_u \equiv \min\{w : \mathcal{D}_f(p_i^1(l), w) = 1\}$. By (12) and Lemma 4, $\bar{w}_u \geq \bar{w}$. For $w > \bar{w}_u$: $\mathcal{D}_f(p_i^1(l), w) = 1$, and hence $\mathcal{A}_E(w) = 0$, which proves part 1 of the proposition.

Let $\hat{w}_1 \equiv \max\{w : \mathcal{D}_f(p_i^1(l), w) = 0\}$ and $\hat{w}_2 \equiv \min\{w : \mathcal{D}_f(p_i^1(h), w) = 1\}$. To prove the remaining parts of the proposition, we need to consider two cases.

Case 1: $\hat{w}_1 \leq \hat{w}_2$ (ad is relatively weak).[16] Then denote $\underline{w}_u \equiv \hat{w}_1$ and $\bar{w}_d \equiv \hat{w}_2$. All consumers i with $w_i \in (\underline{w}_u, \bar{w}_d)$ choose to acquire information after any $a_i \in \{\emptyset, h, l\}$. As shown in Section 5.1, this means that all such i choose a binary signal structure

[16] This corresponds to the notion of weakness that arises in Proposition 2, but not the one defined in Section 5.3.

$x \in \{x^{buy}, x^{pass}\}$, such that $p_i^2(x^{buy}, a_i) = \hat{p}_i$, or $p_i^2(x^{pass}, a_i) = \underline{p}_i$ for any a_i. Then by belief consistency, if i receives no ad, then

$$\begin{aligned} p^0 = p_i^1(\emptyset) =& \underline{p}_i \cdot \left[p^0 \cdot \mathcal{G}_i(x^{pass}|H, \emptyset) + (1-p^0) \cdot \mathcal{G}_i(x^{pass}|L, \emptyset)\right] \\ &+ \hat{p}_i \cdot \left[p^0 \cdot \mathcal{G}_i(x^{buy}|H, \emptyset) + (1-p^0) \cdot \mathcal{G}_i(x^{buy}|L, \emptyset)\right] \\ =& \underline{p}_i \cdot \big(1 - \mathcal{D}(p^0, w_i)\big) + \hat{p}_i \cdot \mathcal{D}(p^0, w_i), \end{aligned} \tag{16}$$

where the latter equality follows from the definition of $\mathcal{D}(p_i^1, w_i)$ and the consumer's strategy (buy after x^{buy}, pass after x^{pass}). Further, $\mathcal{D}_f(p^0, w_i) = \mathcal{D}(p^0, w_i)$, since $p^0 = p_f(h) \cdot p^* + p_f(l) \cdot (1 - p^*)$.

In turn, if the firm advertises to consumer i in equilibrium, then we obtain in a similar way:

$$\begin{aligned} p_i^1(h) =& \underline{p}_i \cdot \big(1 - \mathcal{D}(p_i^1(h), w_i)\big) + \hat{p}_i \cdot \mathcal{D}(p_i^1(h), w_i) \\ p_i^1(l) =& \underline{p}_i \cdot \big(1 - \mathcal{D}(p_i^1(l), w_i)\big) + \hat{p}_i \cdot \mathcal{D}(p_i^1(l), w_i), \\ \text{so since } p^0 =& p_i^1(h) \cdot p^* + p_i^1(l) \cdot (1 - p^*), \\ \text{it follows that } p^0 =& \hat{p}_i \cdot \left[p^* \mathcal{D}(p_i^1(h), w_i) + (1-p^*) \mathcal{D}(p_i^1(l), w_i)\right] \\ &+ \underline{p}_i \cdot \left[1 - p^* \mathcal{D}(p_i^1(h), w_i) - (1-p^*) \mathcal{D}(p_i^1(l), w_i)\right]. \end{aligned} \tag{17}$$

Further, $\mathcal{D}_f(p_i^1, w_i) = \mathcal{D}(p_i^1, w_i)$ for either p_i^1, since $p_i^1(y) = p_f(y)$ for any $y \in \{h, l\}$.

Since $p^0 = z_1 \hat{p}_i + (1 - z_1)\underline{p}_i = z_2 \hat{p}_i + (1 - z_2)\underline{p}_i$ implies that $z_1 = z_2$, (16) and (17) together imply that $\mathcal{D}_f(p^0, w_i) = p^* \mathcal{D}_f(p_i^1(h), w_i) + (1 - p^*) \mathcal{D}_f(p_i^1(l), w_i)$, meaning that $\mathcal{A}_E(w_i) = 0$ for all $w_i \in (\underline{w}_u, \bar{w}_d)$. This proves part 3 of the proposition.

For $w \in [\underline{w}_d, \underline{w}_u]$, we have that $p_i^1(l) < \underline{p}_i$, and so $\mathcal{D}_f(p_i^1(l), w_i) = \mathcal{D}(p_i^1(l), w_i) = 0$. Then the latter equation in (17) can be rewritten as

$$p^0 = p^* \cdot \left[\underline{p}_i \cdot \big(1 - \mathcal{D}(p_i^1(h), w_i)\big) + \hat{p}_i \cdot \mathcal{D}(p_i^1(h), w_i)\right] + (1 - p^*) p_i^1(l),$$

meaning that if i is advertised to, we can write p^0 as a weighted average of $\hat{p}_i, \underline{p}_i$, and $p_i^1(l)$, with the ex ante expected demand $p^* \mathcal{D}(p_i^1(h), w_i)$ being equal to the weight on $\hat{p}_i$. If i is not advertised to, then (16) still applies, and p^0 is a weighted average of $\hat{p}_i$ and $\underline{p}_i$, with the demand again equal to the weight on $\hat{p}_i$. Since $p_i^1(l) < \underline{p}_i$, the weight on $\hat{p}_i$ in the latter case (no ad) must be smaller than in the former (ads to i). Hence $\mathcal{A}_E(w_i) > 0$ in this case, which proves part 2 of the proposition. A mirror argument can then be used to prove part 4.

Case 2: $\hat{w}_1 > \hat{w}_2$ (as is relatively strong). For $w \in [\underline{w}_d, \hat{w}_2]$, the same logic as in part 2 above can be used to show that $\mathcal{A}_E(w) > 0$ in that case. Further, a symmetric argument could again be used to show that $\mathcal{A}_E(w) < 0$ if $w \in [\hat{w}_1, \bar{w}_u]$. Finally, if $w \in [\hat{w}_1, \hat{w}_2]$, then $\mathcal{D}_f(p_i^1(h), w) = 1$ and $\mathcal{D}_f(p_i^1(l), w) = 0$. Since $\mathcal{D}_f(p^0, w)$ is increasing (strictly so when $\mathcal{D}_f(p^0, w) \in (0, 1)$) and all $\mathcal{D}_f$ are continuous, it follows that $\mathcal{A}_E(w)$ is continuous and decreasing. As shown above, $\mathcal{A}_E(\hat{w}_2) \geq 0 \geq \mathcal{A}_E(\hat{w}_1)$, which, by the intermediate value theorem, implies that there exists $\hat{w}_3$ such that $\mathcal{A}_E(\hat{w}_3) = 0$, and the strict part of monotonicity implies that such $\hat{w}_3$ is unique. Letting $\underline{w}_u \equiv \bar{w}_d \equiv \hat{w}_3$, we then obtain parts 2-4 of the proposition in this case.

This concludes the proof. $\square$

Proof of Proposition 5. Let

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

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

$$\begin{aligned} U(p, p) &= \max_{\mathcal{G}} \left\{ \sum_{s \in S} p(s)\mathcal{D}(p(s), w_i, s) \cdot (v_s - (1 - w_i)) - \lambda\kappa(\mathcal{G}; p) \right\} \\ \Rightarrow U(p, p) &= \max_{q} U(p, q). \end{aligned} \tag{18}$$

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 prove the claim for the baseline model. This is equivalent to showing that having more information at the interim stage is always weakly beneficial for the consumer. Note that the interim belief of a Bayesian observer who directly observes the firm's private signal y coincides with the firm's belief $p_f^1(y)$. Consumer i's belief p_i^1 coincides with $p_f^1(y)$ for both y if $i \in \mathcal{T}(h)$. Her ex ante expected utility in this case is

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

where $P(y)$ are the ex ante probabilities of different realizations of y. If, on the other hand, $i \notin \mathcal{T}(h)$, then such a consumer's belief remains at $p_i^1(\emptyset) = p^0$ regardless of y. Her expected utility in this case is $U(p^0, p^0) = \mathbb{E}_y[U(p_f^1(y), p^0)]$, since $\mathbb{E}_y[p_f^1(y)] = p^0$ and $U(p, q)$ is linear in p. From (18) it follows that $U(p_f^1(y), p_f^1(y)) \geq U(p_f^1(y), p^0)$, hence we conclude that $\mathbb{E}_y[U(p_f^1(y), p_f^1(y))] \geq U(p^0, p^0)$ – being targeted by ad h (which is equivalent to observing y) is indeed beneficial for the consumer. This proves the statement of the proposition for the baseline model.

We now proceed to also prove the statement of the proposition for the cursed consumers as introduced in Section 6.4, following 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

$$P(h)U(p_f^1(h), p_f^1(h)) + P(l)U(p_f^1(l), p^0), \tag{19}$$

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 $p_f^1(h)$. The ex ante expected utility of a cursed consumer if $i \notin \mathcal{T}(h)$ (as evaluated by an objective observer) is

$$P(h)U(p_f^1(h), p^0) + P(l)U(p_f^1(l), p^0), \tag{20}$$

since her belief then remains at p^0 regardless of y. By (18) we have $U(p_f^1(h), p_f^1(h)) \geq U(p_h^1, p^0)$, which implies that (19) is weakly greater than (20). This proves the statement of the proposition for cursed consumers and concludes the proof. $\square$