Strategic Obfuscation: Information Costs as a Tool to Influence Consumers

Monte, Daniel and Linhares, Luis Henrique

9 January 2023

Strategic Obfuscation: Information Costs as a Tool to Influence Consumers

Luis Henrique Linhares Daniel Monte*

March, 2025

Abstract

We develop a model of strategic obfuscation in which a firm chooses the cost for a rationally inattentive consumer to process product information, making learning costs endogenous. Our characterization of how obfuscation arises in equilibrium depends on the prior and the consumer's payoff structure, which determines when she is indifferent between actions. Surprisingly, the firm may encourage partial learning to sustain sales while inducing mistakes. The framework links persuasion and rational inattention, shifting attention from commitment in disclosure to information cost design. The model rationalizes hidden add-ons, fine print, and product complexity, and informs policies on disclosure and transparency.

Keywords: strategic obfuscation, rational inattention, costly information
JEL Codes: D11, D82, D83

*Linhares: Boston University. E-mail: lhrlinha@bu.edu. Monte: University of Turin, Collegio Carlo Alberto, and Sao Paulo School of Economics - FGV. E-mail: daniel.monte@unito.it. We are very thankful for conversations with Henrique de Oliveira, Lucas Maestri, Chiara Margaria, Laurent Mathevet, Jawwad Noor, Juan Ortner, Kareen Rozen, and Adrien Vigier.

1 Introduction

Obfuscation, by which we mean the conscious effort to hinder the collection and processing of information, is ubiquitous. When required to disclose evidence, lawyers may send excess material not directly linked to the case at hand, making it harder for other parties to discern what is relevant and what is not. Printers can be sold at a loss to attract consumers while hiding ink sales at a high margin. Terms and conditions are drawn out and written in legal jargon, so reading them before accepting is the exception, not the rule.

In this paper, we consider a closely related example from consumer product markets. A firm offers a product of uncertain value for a new consumer and wishes to induce her to buy it. Consumers can conduct their own costly investigation: reading reviews, looking at specifications, and comparing alternatives. To induce the consumer to buy, the firm can facilitate or hinder these efforts by controlling the availability of useful information.[1] For instance, the firm may or may not provide a free sample, or a test-period.

In a nutshell, we focus on a different method for a principal to induce a decision maker to take a certain action that does not rely on direct information transmission or on monetary transfers. Instead, the principal controls how difficult it is for the decision maker to learn about an exogenous and unknown (for both the decision maker and the principal) state of the world. Therefore, we provide a new framework in which we endogenize how costly it will be for the consumer to obtain new information.

Formally, we model a firm that chooses the marginal cost of obtaining information on its product to induce a rationally inattentive (Sims, 2003) consumer to buy it. When choosing how much to learn, the consumer will weigh the costs (e.g., time spent researching) and benefits of new information (lower probability of buying a bad product). We assume the firm can choose the value of a parameter $\lambda \in \mathbb{R}_+$ that linearly affects how much uncertainty,

[1] Richards et al. (2020) find evidence of firms offering products in different package sizes from their competition to create perceived product differentiation. The goal is not simply price discrimination, but also to hinder product comparison and hopefully induce the consumer to buy their products.

measured by Shannon Entropy (Shannon, 1948), is reduced by the acquired information.[2] Thus, by setting $\lambda = 0$, the firm effectively lets the consumer correctly identify the state of the world. By setting a large λ, think of $\lambda \to \infty$, the firm makes any information processing impossibly costly and the consumer must base her decisions on the common prior.

The firm's efforts to raise costs, or increase the value of λ, can be interpreted in different ways: restricting access to information, decreasing the time available to make a decision, or simply dumping irrelevant data.[3] Unlike other obfuscation papers, we do not focus on a single mechanism. Instead, we generalize how this behavior can appear if the consumer displays rational inattention. Moreover, we consider that the cost of information is chosen before the firm observes the state of the world and thus, the chosen λ is not a signal in itself. Think, for instance, of a family-owned online T-shirt seller and a prospective buyer. Both seller and buyer know that 30% of individuals like the shirts, but neither will know if the buyer falls in this category until the shirt is bought and tried on.

A key contribution of our paper, relative to the Bayesian Persuasion framework,[4] is to shift the focus from commitment in information transmission to the design of which information can feasibly be learned.[5] As such, we have a story about a firm trying to persuade a consumer without requiring a strong commitment assumption.[6] In settings where this assumption does not hold, we interpret our results as what the firm could achieve if it could commit to an obfuscation level, and therefore provide an upper bound to the firm's payoffs when compared to no commitment.

[2]This parameter is standard in the Rational Inattention literature: $\lambda(H(\mu_0) - \mathbb{E}_\tau[H(\mu_a)])$.

[3]Hiding or providing information may be costly to the firm. In our model, we abstract from such costs. Appendix C briefly discusses an alternative model in which the firm incurs a cost when choosing λ. Adding these costs, however, does not alter our main conclusions.

[4]See Kamenica and Gentzkow (2011) but also Kamenica (2019) and Bergemann and Morris (2019) for comprehensive surveys. There, a Sender commits to an information disclosure policy, while in our paper, the firm (essentially) chooses how the consumer will learn, within the rational inattention framework.

[5]Indeed, Kamenica and Gentzkow (2011) use commitment in the information acquisition stage instead of commitment in the information transmission stage as a possible motivation for the commitment assumption in Bayesian Persuasion.

[6]We comment further on this assumption when presenting the model and describing practical examples. Martin (2017) and De Clippel and Rozen (2025) present models in which there is a strategic component to the fact that an agent is choosing to obfuscate.

To see our model in practice, consider an example in which a firm is attempting to sell a product of uncertain quality. The product may be *Good*, in which case a consumer would enjoy a positive benefit from purchasing it, or *Bad*, in which case consumers would regret the purchase. Both the firm and the consumer share a common prior belief that the product is Good with probability $\mu_0 \in (0,1)$. In contrast to the consumer whose payoff depends on the realized quality, the firm always benefits from a sale. The payoff structure is summarized as

	not buy	buy
Good	$0, 0$	$1, B$
Bad	$0, 0$	$1, -C.$

Here, the first entry in each cell represents the firm's payoff, while the second entry represents the consumer's payoff. The parameters satisfy $B > C > 0$, so that buying is desirable only when the product is Good. From the consumer's perspective, the decision can be summarized by a threshold belief. If the probability μ_0 that the product is Good exceeds a belief of indifference

$$\mu_{ind} = \frac{C}{B+C},$$

then the expected benefit of buying is positive, and the consumer will purchase if no new information is obtained. In this case, the firm's best strategy is to fully obfuscate, making additional information prohibitively costly, so that the consumer bases her decision solely on the favorable prior. This is the sales adage "don't talk yourself out of a sale".[7]

If the prior μ_0 falls below this threshold, the consumer would not buy without further investigation. The firm faces the following problem. If it provides full transparency, it will enable the consumer to perfectly observe whether the product is Good or Bad. This maximizes efficiency but yields sales only when the product is indeed Good, which may be unattractive to the firm when the prior is low. On the other hand, the firm may release

[7]Also quoted as "When you get the sale, stop talking." `https://www.inc.com/jim-schleckser/when-you-get-sale-stop-talking.html`

information in a way that is costly to process. By raising the marginal cost of learning, the firm induces the consumer to remain only partially informed. If this cost is appropriately chosen the consumer buys more frequently than under full transparency.

Consider now a product replacement or recycling policy that cushions losses when the product is Bad, effectively lowering C. In that case, the consumer is more willing to buy under partial information, and the firm benefits from maintaining some degree of obfuscation. Eliminating such a policy increases C. This raises the threshold μ_{ind}, making the consumer more reluctant to buy without clear evidence. The higher stakes from bad outcomes can push the firm to choose full transparency, since obfuscation no longer yields profitable sales. Thus, policies that seem to reduce consumer payoffs in some states may in fact increase overall welfare by inducing firms to simplify disclosure and reveal more information. We discuss this example in greater detail in Appendix B.1.

So, interestingly enough, there are cases in which partial information can dominate both full transparency (too few sales) and full secrecy (no sales) for the firm. When the cost of buying a bad product is low, the consumer chooses more "optimistic signals" that recommend buying more often, so allowing them to learn a bit may lead to increased sales. In our Theorem 1, we characterize when partial information (or partial obfuscation) is the equilibrium outcome in a 2-state x 2-actions setting, when the firm has state-independent preferences.[8]

This example highlights the central role of strategic obfuscation. When beliefs are favorable, obfuscation sustains sales by preventing the consumer from discovering bad news. When beliefs are unfavorable, partial obfuscation allows some sales that complete transparency would eliminate. Our framework formalizes these tradeoffs and provides techniques that will allow us to (i) compute exactly how costly the firm will choose information processing to be; (ii) understand how the consumer chooses how much to learn given the chosen information processing costs; and (iii) perform some comparative statics given policies motivated at improving welfare, such as policies on consumer transparency.

[8]Appendix B.2 includes a case of partial obfuscation under state-dependent preferences, where the role of the consumer is replaced by that of another firm trying to enter a market.

Section 4 generalizes our model to multiple actions and states, and allows the firm to have state-dependent preferences. This connects our model to the general literature of decisions by rationally inattentive agents (Matějka and McKay, 2015; Caplin et al., 2019, 2022; Matysková and Montes, 2023), which we leverage to construct our results. In this setting, we characterize the beliefs of indifference between actions, illustrated in the example above, as the beliefs in which inducing full obfuscation is not possible. Intuitively, if the consumer is indifferent between any two actions at her prior, any signal, no matter how coarse, can break the tie. Therefore, even if information is extremely costly, the consumer will acquire an almost uninformative signal that she can still afford. We also reinterpret one of the examples of an agent with rational inattention from Caplin et al. (2019) to our setting. In the example, a consumer wants to buy the only good product among 9 bad ones. We can then predict the chosen level of obfuscation depending on which option the firm prefers to sell.

Obfuscation has clear policy implications. When the firm tries to obfuscate to increase its payoff, it comes at a cost for buyers. Understanding when obfuscation arises provides insights for policies that seek to increase the consumer's welfare. For instance, mandatory disclosure can be particularly ineffective (Ben-Shahar and Schneider, 2014). Indeed, in our setting, mandated disclosure does not imply a low λ, as the firm is free to include in the disclosure as much pointless information as it wants. In one of our examples, we see that an effective policy - that at first glance seems to harm consumers - can induce simplified disclosure, as difficulty in processing information might push customers not to purchase their products.

Related literature:

Information Design and the commitment assumption - We contribute to the large recent literature that seeks to understand how different information structures affect the problems at hand, and how incentives shape the information acquired in equilibrium. In particular, our paper does not require the commitment assumption in the information transmission stage of canonical models of Bayesian Persuasion and information design (Kamenica and Gentzkow (2011), Bergemann and Morris (2019), for example). By rethinking the usual commitment

assumption, our paper is related to the works by Lin and Liu (2024), Barros (2021), Lipnowski et al. (2022), Min (2021), Fréchette et al. (2022), Nguyen and Tan (2021), Perez-Richet and Skreta (2022) and Alonso and Camara (2024). Also related are Edmond and Lu (2021), who add confusion to the communication model by making the Sender capable of confusing the consumer at a cost, and Bizzotto et al. (2020), who study the case of persuasion when dealing with certifications and add a cost for obtaining information.

Rational Inattention - Particularly useful to our analysis is the branch of this literature that developed mathematical foundations for solving the problems of agents that exhibit Rational Inattention and expand on the possible information cost structures the agent faces (Matějka and McKay, 2015; De Oliveira et al., 2017; Caplin et al., 2019, 2022; Dean and Neligh, 2023). Ellis (2018) shows an axiomatic model of choice behavior for an agent with limited attention and gives the conditions under which the behavior can be considered the result of a choice of an optimal level of inattention. Similarly to our applications, Cusumano et al. (2024) applies Rational Inattention to an IO context where two competing firms try to sell to a boundedly rational consumer. Our model provides an alternative microfoundation for the commonly exogenous cost of information in this literature. For instance, Asriyan et al. (2023) demonstrates the impact of a buyer's imperfect observation of a product's quality on the complexity chosen by the product designer. Mensch and Ravid (2022) discuss the optimal menu that a monopolist should offer to an agent who costly learns about their type.

Information Design and Rational Inattention - A few authors have recently combined the study of persuasion with Rational Inattention. Gentzkow and Kamenica (2014) is a direct extension of the classic Bayesian Persuasion model in which the sender has to buy the signal at a cost given by Shannon Entropy. Matysková and Montes (2023) develop a model in which the Sender sends information at no cost. However, upon seeing a signal realization the consumer can decide if she wants to buy further information at a cost. They derive how the sender will choose the optimal signal and an algorithm that simplifies the Sender's problem. Bloedel and Segal (2021) detail how to optimally persuade a rationally inattentive receiver,

as the usual concavication results are not attainable. Lipnowski et al. (2020) show where a benevolent Sender may choose how to restrict the information that a Rational Inattention consumer sees. Le Treust and Tomala (2019) have the Sender persuading the consumer by sending a message through a limited capacity channel. Some papers have shown that rational inattention may be beneficial for the receiver. Indeed, Fabbri (2024) shows that it can serve as a "commitment" that diminishes the holdup problem and Wei (2020) also shows it serves as a commitment that they will not pay attention if the Sender does not send enough useful information. Our approach differs from all these papers in two main ways: we abstract from messages and signals, and our cost of information is endogenous.

Strategic Obfuscation - A set of papers studying obfuscation, usually focused on IO environments, motivated the economic problem that we detail here. Some of the seminal papers in this literature are Gabaix and Laibson (2006); Ellison (2005) and Ellison and Wolitzky (2012). Kalayci and Potters (2011) and Jin et al. (2022) obtain outcomes of agents actively engaging in obfuscation, but in a laboratory environment. Ben-Shahar and Schneider (2014) extensively describe the possible problems with disclosure policies. More recently, De Clippel and Rozen (2025) created and empirically tested a model of communication in which obfuscation intensity is endogenous and may appear in settings with mandatory disclosure. They add strategic inference from the agent, i.e., the agent sees obfuscating behavior as a signal and draws inference from that. Nocke and Rey (2024) shows that firms can benefit from limiting the information available to consumers. We see our model as a bridge between persuasion and strategic obfuscation. Our results show when full obfuscation is feasible, and describe the possible outcomes when the firm does not have strategic concerns about obfuscation.[9]

[9]In the sense that the cost of information is chosen before the firm knows the state.

2 The 2x2 Model

In this section, we focus our attention on a simplified class of problems in which there are two states, two actions, and the firm has state-independent preferences. This allows us to quickly explain how the model works and more directly introduce our main results. We generalize our model in Section 4.

There are two players, the firm (*it*) and the consumer (*she*). The consumer must choose an action a_i from the set $A = \{a_H, a_L\}$. The set of payoff relevant states of nature is $\Omega = \{\omega_H, \omega_L\}$ and the realization of a state $\omega_i \in \Omega$ is drawn from a distribution $\mu_0 \in int(\triangle(\Omega))$, the prior, known by both individuals.[10] In this simplified case with two states, we abuse notation and identify a belief $\mu \in \triangle(\Omega)$ by its corresponding probability that the state is ω_H, i.e., with the number $\mu = \mu(\omega_H) \in [0, 1]$.

The consumer and firm have material utility functions $u : A \times \Omega \to \mathbb{R}$ and $v : A \to \mathbb{R}$, respectively. We let $a^*(\mu) \equiv \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)]$ represent the set of actions that the consumer may take given a belief $\mu \in \triangle(\Omega)$. We break ties in favor of the firm as this helps ensure that a solution to the firm's problem exists.[11] We further assume that $u(a_i, \omega_i) > u(a_{-i}, \omega_i)$ for $i \in \{H, L\}$, so no action is optimal for the consumer in both states. Finally, to ensure that full obfuscation is possible (as we discuss in Section 4), we assume that the consumer is not indifferent between the two actions at the prior, so $a^*(\mu_0)$ is a singleton. In the 2×2 case, there is only one such belief of indifference, which is given by

$$\mu_{ind} = \frac{u(a_L, \omega_L) - u(a_H, \omega_L)}{u(a_H, \omega_H) - u(a_H, \omega_L) + u(a_L, \omega_L) - u(a_L, \omega_H)}. \tag{1}$$

The consumer may purchase a signal structure (or Blackwell experiment) π that consists of the set A of signals - that can be seen as recommended actions, - and a family of distributions

[10] We denote by $\triangle(\Omega)$ the space of all probability distributions on Ω, and by $int(X)$ the interior of set X. We endow $\mathbb{R}^M$ with the Lebesgue measure.

[11] This is a common assumption in the literature (Kamenica and Gentzkow, 2011). However, the implication of this assumption in our model is two-fold: when the consumer is indifferent between different distributions of posteriors, she chooses the one favored by the firm; in addition, if at an induced posterior the consumer is indifferent between any two actions, she chooses again the one favored by the firm.

$\{\pi(\cdot|\omega)\}_{\omega\in\Omega}$ over A.[12] Instead of working with the experiment π, it will be convenient to work directly with the corresponding induced distribution of posteriors τ (Kamenica and Gentzkow, 2011), which we will do moving forward.

The consumer is rationally inattentive and incurs an information cost given by

$$c_\lambda(\tau) \equiv \lambda(H(\mu_0) - \mathbb{E}_\tau[H(\mu^a)]), \tag{2}$$

where $H : \triangle(\Omega) \to \mathbb{R}_+$ is the Shannon Entropy (Shannon, 1948) given by

$$H(\mu) = -\mu \ln(\mu) - (1-\mu)\ln(1-\mu).$$

and $\lambda \in \mathbb{R}_+$ is the marginal cost of information. We assume that the firm chooses the value of λ; it is the only parameter the firm has control over, and it is the only way it can persuade the consumer to take a certain action. As we mentioned in the introduction, this can be thought of as the firm hindering or easing the process for the consumer to gather new information.

Entropy can be interpreted as a measure of uncertainty, so the cost function charges for how much uncertainty the signal reduces.[13] There are two ways a signal can have zero cost: if the firm sets $\lambda = 0$ and "reveals the state", or if the consumer chooses an uninformative signal so that the posterior equals the prior and there is no change in entropy, i.e., the uncertainty is not reduced.

For a given cost λ, the consumer chooses τ, and a signal $a_i \in A$ is realized, leading to a posterior $\mu^{a_i} \in \triangle(\Omega)$. Since the consumer knows the conditional probability of seeing the signal realization in each of the possible states of the world, she updates her beliefs to the posterior μ^{a_i} according to Bayes' rule.[14] She then takes an action a_i, and utilities $u(a_i, \omega)$ and $v(a_i)$ are realized (given that the state is ω).

[12] With rational inattention, it is without loss of generality to consider the signal realization set to be A and treat the signal realization as a recommended action. If more than one signal induces the same action, the consumer pays for information not being used (c.f. Matějka and McKay (2015)).

[13] For beliefs in the boundary of $\triangle(\Omega)$, we let $0 \cdot \ln(0) \equiv \lim_{x\to 0^+} x \cdot \ln x = 0$.

[14] If a_i is chosen with positive probability, and $\pi(a_i|\omega)$ denotes the conditional probability $\mathbb{P}(a_i|\omega)$ for a

An increase in λ makes information more expensive and impacts the consumer's probability of taking a certain action. The firm will therefore choose the value of λ that maximizes the odds of its favorite action being played. Note that when the firm chooses λ, it does not know the true state of the world yet. Thus, this setting is more suited to applications in which there is uncertainty on both sides: think of a new firm interacting with a possible consumer for the first time.[15]

2.1 The consumer's Rational Inattention Problem

Since the firm has to anticipate what the consumer will do given a value of λ, it is important to first understand how the consumer will optimally choose to learn. After receiving an action recommendation a_i, Bayesian updating leads to a posterior μ^{a_i} and the consumer chooses a_i. Two opposing forces are in play: the consumer wants to learn the state of nature to choose the best possible action; the consumer wants to avoid learning costs. In what follows, let $U(\mu) \equiv \mathbb{E}_{w\sim\mu}[u(a^*(\mu), \omega)]$ be the indirect utility for a belief $\mu \in \triangle(\Omega)$.

Definition 1. *For given μ_0 and $\lambda \in \mathbb{R}_+$, the consumer's Rational Inattention Problem is*

$$\max_{\tau\in\triangle(\triangle(\Omega))} \mathbb{E}_\tau \left[U(\mu^{a_i})\right] - \lambda(H(\mu_0) - \mathbb{E}_\tau[H(\mu^{a_i})]), \tag{4}$$
$$s.t.\ \mathbb{E}_\tau[\mu^{a_i}] = \mu_0.$$

The restriction is the usual requirement of Bayes Plausibility. We denote by τ_λ the solution to (4) for a given λ, which identifies the posteriors and the probability that each action is taken.

given signal π,

$$\mu^{a_i}(\omega) = \frac{\pi(a_i|\omega)\mu_0(\omega)}{\sum_{j=1}^{2}\pi(a_i|\omega_j)\mu_0(\omega_j)}. \tag{3}$$

[15] In a similar fashion, Li and Shi (2017) provides a model in which a seller can allow the buyer to obtain more information about the valuation of a good. Neither the seller nor the buyer have previous perfect information about the valuation.

2.1.1 Solving the consumer's Problem

In this section, we go straight to describing a functional form of the solution to the consumer's problem based on the characterization from Caplin et al. (2019). We hold on to the details until we go over this solution more explicitly in Section 4, where we also look at the comparative statics for λ and discuss when obfuscation is possible.

The characterization of Caplin et al. (2019)[16] allows us to find the posteriors that are reached following a given signal realization for any value of λ (see Lemma 3). With the posteriors in hand, we then find τ_λ through Bayes' plausibility. In the 2×2 case, we can simplify the characterization of the solution from Caplin et al. (2019) by using the following auxiliary functions:[17]

$$\underline{\mu}(\lambda) \equiv \frac{1 - exp(\frac{u(a_H,\omega_L)-u(a_L,\omega_L)}{\lambda})}{exp(\frac{u(a_H,\omega_H)-u(a_L,\omega_H)}{\lambda}) - exp(\frac{u(a_H,\omega_L)-u(a_L,\omega_L)}{\lambda})}, \tag{5}$$

$$\overline{\mu}(\lambda) \equiv \frac{1 - exp(\frac{u(a_L,\omega_L)-u(a_H,\omega_L)}{\lambda})}{exp(\frac{u(a_L,\omega_H)-u(a_H,\omega_H)}{\lambda}) - exp(\frac{u(a_L,\omega_L)-u(a_H,\omega_L)}{\lambda})}. \tag{6}$$

For a given λ, $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ represent the threshold beliefs for which actions a_L and a_H, respectively, are taken with probability 1,[18] which we illustrate for simplicity in Figure 1. To understand how these thresholds characterize the solution to the consumer's problem for a value of λ, we can go case by case.

If (i) $\mu_0 \in [0, \underline{\mu}(\lambda))$, the cost of obtaining more information outweighs the benefits of learning - she knows enough, - and the consumer prefers to purchase a non-informative signal that costs 0. Therefore, the posterior is equal to the prior. Since $\mu_0 \in [0, \underline{\mu}(\lambda))$, acting according to the prior means choosing a_L with probability 1 (remember that μ_0 is the probability that the state is H, so the closer μ_0 is to 0, the more likely it is that action a_L is optimal). Analogously, if (ii) $\mu_0 \in (\overline{\mu}(\lambda), 1]$, the consumer acquires a completely

[16]To give an intuition, the characterization comes from applying the KKT conditions to Equation (4). Matějka and McKay (2015) and Ravid (2020) also discuss how to solve this problem.

[17]The exponential function arises from the use of the Entropy as a cost function.

[18]For more details see Caplin et al. (2019); Matysková and Montes (2023) and Section 4 in this paper.

non-informative signal, but now she always chooses a_H. If, however, (iii) $\mu_0 \in [\underline{\mu}(\lambda), \overline{\mu}(\lambda)]$, then it is profitable to acquire new information. The consumer will choose a distribution of posteriors with support $\mu^{a_L} = \underline{\mu}(\lambda)$ and $\mu^{a_H} = \overline{\mu}(\lambda)$. Thus, if the recommended action is a_L the induced posterior will be $\underline{\mu}(\lambda)$ and she chooses a_L; if the realization is a_H the posterior will be $\overline{\mu}(\lambda)$ and she chooses a_H, so both actions are chosen with a positive probability.

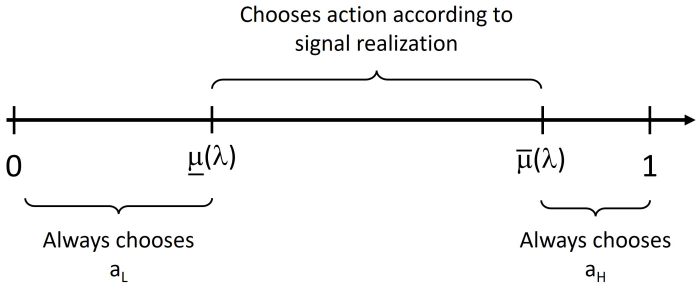

Figure 1: How the functions $\underline{\mu}$ and $\overline{\mu}$ define the choice by the consumer for a given λ depending on where the prior is in the $[0, 1]$ interval.

To see how the firm will choose λ is then sufficient to understand how the thresholds $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ change with λ. To ensure continuity, we define

$$\underline{\mu}(0) \equiv \lim_{\lambda \to 0^+} \underline{\mu}(\lambda) = 0, \text{ and } \overline{\mu}(0) \equiv \lim_{\lambda \to 0^+} \overline{\mu}(\lambda) = 1.$$

We get these limits since $u(a_i, \omega_{-i}) - u(a_i, \omega_i) < 0$ and $u(a_i, \omega_i) - u(a_{-i}, \omega_i) > 0$. By calculating the limit when $\lambda \to \infty$, we get

$$\lim_{\lambda \to \infty} \underline{\mu}(\lambda) = \lim_{\lambda \to \infty} \overline{\mu}(\lambda) = \frac{u(a_L, \omega_L) - u(a_H, \omega_L)}{u(a_H, \omega_H) - u(a_H, \omega_L) + u(a_L, \omega_L) - u(a_L, \omega_H)} = \mu_{ind}. \quad (7)$$

When $\lambda = 0$, the thresholds are in the extremities 0 and 1; as λ increases, they move toward μ_{ind} in the interior. As a consequence of these properties, we also have that $\mu_{ind} \in \left[\underline{\mu}(\lambda), \overline{\mu}(\lambda)\right]$ for all $\lambda \in [0, \infty)$. This makes sense: as λ grows and new information becomes more expensive, the regions for which buying new information is beneficial diminish. At the limit, only when the consumer is very unsure (μ_0 close to μ_{ind}) will she buy any information - a fact we discuss in more detail in Section 4. Finally, we highlight that $\underline{\mu}$ / $\overline{\mu}$ are strictly increasing/decreasing

respectively[19] and continuous.[20] Figure 2 illustrates how the threshold beliefs $\underline{\mu}$ and $\overline{\mu}$ change with λ in an example where μ_{ind}.

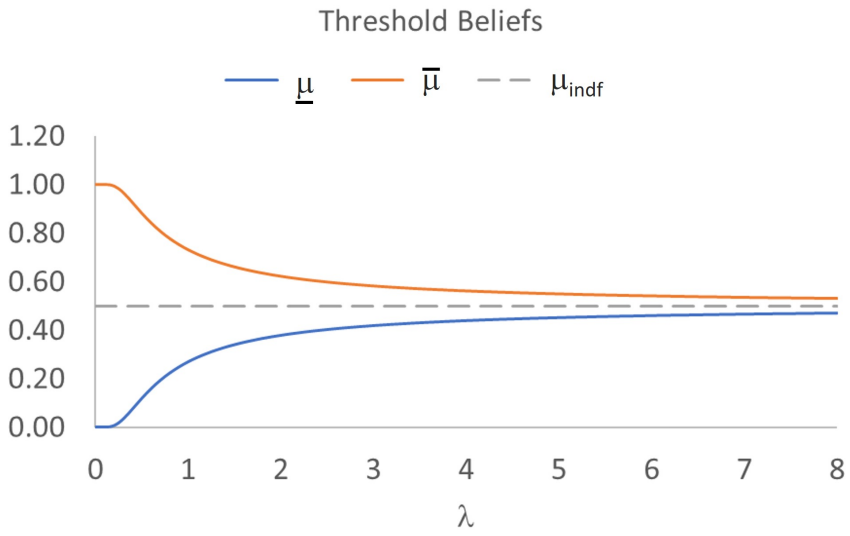

Figure 2: The value of $\overline{\mu}$ and $\underline{\mu}$ as a function of lambda when $\mu_{ind} = 0.5$.

2.2 The Firm's Problem

Now that we know how the consumer will react to the cost of information, we are better equipped to understand how the firm's utility changes with λ, which in turn lets us predict how much the firm will choose to obfuscate. We define the firm's indirect expected utility given a belief μ as $V(\mu) \equiv \mathbb{E}_{\mu}[v(a^*(\mu))]$. Anticipating the optimal choice τ_λ by the consumer (which is unique in the 2×2 case), we can define the firm's indirect expected utility from obfuscation as $\tilde{V}_{\mu_0}(\lambda) \equiv E_{\tau_\lambda}[V(\mu^{a_i})]$. We use the prior as an index to highlight that, because of Bayes' Plausibility, it affects the optimal choice of λ.

Intuitively, the firm will choose a λ that maximizes the probability that its favorite actions are played. For instance, if the firm agrees with the consumer in every state, it will choose $\lambda = 0$; if the firm wants to maximize the chance of mistakes, it will choose full obfuscation.

[19]For more details, see Matysková and Montes (2023). We can also check that the derivative is always positive/negative.

[20]Continuity is ensured because both $\underline{\mu}(\cdot)$ and $\overline{\mu}(\cdot)$ are composites of exponential functions, which are continuous.

Lemma 1. *In the* 2×2 *model, there exists* $\overline{\lambda} \in \mathbb{R}_{++}$ *such that*

$$\begin{cases} \underline{\mu}(\overline{\lambda}) = \mu_0, \text{ if } \mu_0 < \mu_{ind}, \\ \overline{\mu}(\overline{\lambda}) = \mu_0, \text{ if } \mu_0 > \mu_{ind}. \end{cases} \tag{8}$$

Proof. We first highlight three facts mentioned before: (i) $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ are continuous and strictly increasing/decreasing respectively; (ii) $\underline{\mu}(0) = 0$, $\overline{\mu}(0) = 1$, $\lim_{\lambda\to\infty} \underline{\mu}(\lambda) = \lim_{\lambda\to\infty} \overline{\mu}(\lambda) = \mu_{ind}$; (iii) $\mu_0 \neq \mu_{ind}$ by assumption. Now, consider the case where $\mu_0 < \mu_{ind}$. Since $\underline{\mu}$ is continuous, $\underline{\mu}(0) = 0$ and $\lim_{\lambda\to\infty} \underline{\mu}(\lambda) = \mu_{ind}$, by the Intermediate Value Theorem, we deduce the existence of $\overline{\lambda} \in (0, \infty)$ such that $\underline{\mu}(\overline{\lambda}) = \mu_{ind}$. The case where $\mu_0 > \mu_{ind}$ is analogous. □

We formalize this result for multiple states and actions in Section 4. If $\lambda = 0$, the consumer acquires a fully informative signal, which we call *full transparency*; if $\lambda \in (0, \overline{\lambda})$, then $\mu_0 \in [\underline{\mu}(\lambda), \overline{\mu}(\lambda)]$ and the consumer acquires some information, which we call *partial obfuscation*; if $\lambda \geqslant \overline{\lambda}$ information is too expensive and the consumer plays according to the prior, which we call *full obfuscation*. It is then clear that $\tilde{V}_{\mu_0}(\lambda) = V(\mu_0)$ for any $\lambda \geqslant \overline{\lambda}$, so $\tilde{V}_{\mu_0}(\lambda)$ it becomes constant. Therefore, we consider that the firm chooses $\lambda \in [0, \overline{\lambda}]$. We are now ready to state the firm's problem.

Definition 2. *The firm's Obfuscation Problem is*

$$\max_{\lambda\in[0,\overline{\lambda}]} \tilde{V}_{\mu_0}(\lambda). \tag{9}$$

Proposition 6, stated in Section 5, ensures that a solution to the firm's problem exists as long as $\mu_0 \neq \mu_{ind}$ (the proof itself is in Appendix A). For an intuition, $\mu_0 \neq \mu_{ind}$ ensures that exists $\overline{\lambda}$ in $\mathbb{R}_{++}$ that induces full obfuscation, so the firm maximizes an upper continuous function over a compact set. We denote the solution to the firm's problem by λ^*.

2.3 A Geometric Interpretation for The Firm's Problem

Here we develop a geometric intuition for the firm's problem and provide a functional form for $\tilde{V}_{\mu_0}$. This functional form has a similar interpretation as the concavification result in Bayesian Persuasion (BP). Given that the consumer is subject to Bayes' Plausibility, we can find the value of $\tilde{V}_{\mu_0}(\lambda)$ by evaluating at the prior μ_0 the line in $\mathbb{R}^2$ that connects the points $(\underline{\mu}(\lambda), V(\underline{\mu}))$ and $(\overline{\mu}(\lambda), V(\overline{\mu}))$.

This may be more easily seen on Figure 3, where we plot the firm's payoff when $v(a_H) = 1$, $v(a_L) = 0$, $\mu_0 = 0.75$, and $\mu_{ind} = \frac{1}{2}$. Note that μ_{ind} is the only information we need from the consumer to define the payoffs of the firm as a function of λ. The left-hand side of Figure 3 shows the split of posteriors under full-information ($\lambda = 0$). Then, $\tilde{V}_{\mu_0}(\lambda)$ is found as the y-coordinate of the point at the intersection of the vertical line passing through the prior and the steep line connecting $(\underline{\mu}(\lambda), V(\underline{\mu}))$ and $(\overline{\mu}(\lambda), V(\overline{\mu}))$. On the right side of Figure 3 we see what happens when a non-informative signal is acquired ($\lambda = \overline{\lambda}$). The posterior will equal the prior, and we can find $\tilde{V}_{\mu_0}(\overline{\lambda})$ by looking at the y-coordinate of the intersection between the vertical line crossing the prior and the step function $V(\mu)$.

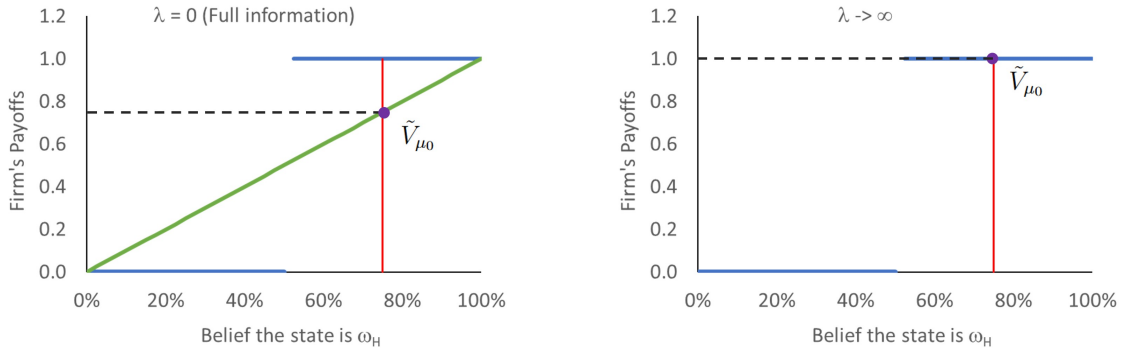

Figure 3: Left: the expected payoff of full information is at the intersection of the steep line with the vertical line of the prior. Right: the posterior is equal to the prior with no new information. In both cases, we set $v(a_H) = 1$, $v(a_L) = 0$, $\mu_0 = 0.75$, and $\mu_{ind} = \frac{1}{2}$.

To illustrate the dynamics of how $\tilde{V}_{\mu_0}$ changes with λ, consider Figure 4. An increase in λ moves the beliefs $\underline{\mu}$ and $\overline{\mu}$ closer to μ_{ind}. The value of $\tilde{V}_{\mu_0}$ then changes as the intersection of the vertical prior line with the line segment connecting V evaluated at the posteriors

moves up or down. At the top row of Figure 4, the value of $\tilde{V}_{\mu_0}$ is monotonically increasing as the point of intersection goes up. In the middle row, we have an example where $\tilde{V}_{\mu_0}$ is monotonically decreasing as the point of intersection goes down. At the bottom row, $\tilde{V}_{\mu_0}$ changes non-monotonically: it increases a bit at first, but after a point it decreases with λ.

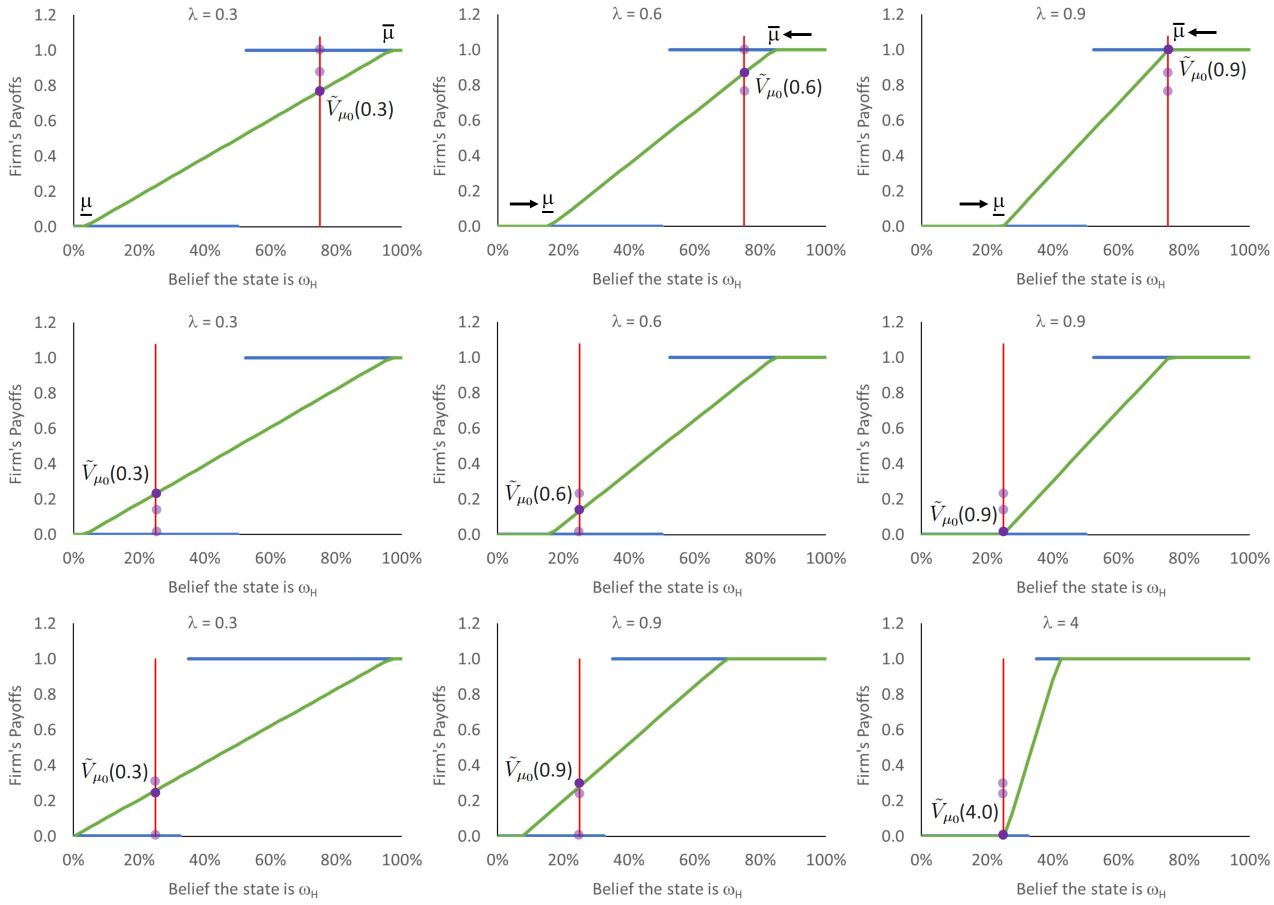

Figure 4: Dynamics of an increase in λ on the firm's payoff when $\tilde{V}_{\mu_0}$ is increasing (top row); decreasing (middle row); non-monotonic (bottom row). We set $v(a_H) = 1$ and $v(a_L) = 0$ in all cases. In the top row: $\mu_0 = 0.75, \mu_{ind} = 0.5$. In the middle row: $\mu_0 = 0.25, \mu_{ind} = 0.5$. In the bottom row: $\mu_0 = 0.25, \mu_{ind} = 1/3$.

Define $\mathcal{O} \equiv [0, \overline{\lambda}]$ and, to ease notation, let $m(\cdot)$ be given by[21]

$$m(\lambda) := \left(\frac{V(\underline{\mu}(\lambda)) - V(\overline{\mu}(\lambda))}{\underline{\mu}(\lambda) - \overline{\mu}(\lambda)} \right).$$

[21]The function $m(\cdot)$ gives us the tangent of the steep line connecting points $(\underline{\lambda}, \tilde{V}_{\mu_0}(\underline{\mu}(\lambda))$ and $(\overline{\lambda}, \tilde{V}_{\mu_0}(\overline{\mu}(\lambda))$.

The functional form for $\tilde{V}_{\mu_0}$ is thus

$$\tilde{V}_{\mu_0}(\lambda) = \begin{cases} m(\lambda)(\mu_0 - \underline{\mu}(\lambda)) + V(\underline{\mu}(\lambda)) \text{ if } \lambda \in \mathcal{O}, \\ V(\mu_0), \text{ otherwise.} \end{cases} \tag{10}$$

Note that $m(\overline{\lambda})(\mu_0 - \underline{\mu}(\overline{\lambda}) + V(\underline{\mu}(\overline{\lambda})) = V(\mu_0)$ and so $\tilde{V}_{\mu_0}$ is continuous. Indeed, finding $\overline{\lambda}$, and therefore $\mathcal{O}$, is just a matter of discovering the unique λ that satisfies

$$m(\lambda)\left(\mu_0 - \underline{\mu}(\lambda)\right) + V(\underline{\mu}(\lambda)) = V(\mu_0).$$

3 Transparency, Obfuscation, and Partial Obfuscation

In this section, we aim to predict the level of obfuscation resulting from the interaction between the consumer and the firm. This matters if we care about the consumer's welfare, which is decreasing in the marginal cost of information as mistakes become more prevalent. There is also a sense of efficiency when decisions are made with perfect information. Therefore, understanding the incentives to obfuscate may lead researchers and policy-makers to a better understanding of which policies are more conducive to full transparency.

There are three possibilities: (i) full transparency - in which the consumer fully learns the state of the world; (ii) full obfuscation - in which the consumer chooses a completely uninformative experiment; (iii) partial obfuscation - in which the consumer acquires an informative signal, but still makes mistakes with positive probability. The following provides a formal definition, which we carry on to the model with multiple actions and states.

Definition 3. *The firm's obfuscation problem admits a* partial obfuscation solution *if neither* $\lambda = 0$ *nor* $\lambda = \bar{\lambda}$ *solves the firm's optimization problem. Otherwise, the equilibrium lies on the boundary, and we refer to it as a* boundary obfuscation solution*, which can take two forms:*

(i) full transparency, when the optimal choice is $\lambda = 0$*;*

(ii) full obfuscation, when the optimal choice is $\lambda = \bar{\lambda}$*.*

We begin by providing a sufficient condition for a problem to have a boundary obfuscation solution, which can be determined by examining the shape of the firm's indirect utility function V.

Proposition 1. *Given a Firm's Problem, the following holds:*

(i) if V *is concave, then* $\tilde{V}_{\mu_0}$ *is increasing in* λ*;*

(ii) if V *is convex, then* $\tilde{V}_{\mu_0}$ *is decreasing in* λ*.*

The result applies even when the firm has state-dependent preferences. The intuition follows directly from Jensen's inequality. When the firm's indirect utility $V(\mu)$ is concave, the expected value of V under any distribution of posteriors is below $V(\mu_0)$. In other words, allowing the consumer to learn, thereby spreading her beliefs around the prior, reduces the firm's expected payoff. The firm therefore prefers higher information costs and less learning. Conversely, if V is convex, the inequality reverses: the firm's expected payoff increases when the consumer becomes better informed, so it favors transparency.

Does partial obfuscation ever occur? At first sight, one might expect only boundary outcomes in the binary environment: with two states and two actions, the firm should either choose full transparency ($\lambda^* = 0$) or full obfuscation ($\lambda^* = \bar{\lambda}$). Yet the model shows that partial obfuscation can emerge even in this simple setting and under fairly general conditions. This case is particularly relevant in real-world situations where frictions or additional costs influence the firm's choice of λ.[22] We illustrate this with an example.

[22]The version of the model with an explicit cost of obfuscation is detailed in Appendix C.

Example 1 (Online T-shirt Market). Consider an online market for T-shirts. We will interpret setting $\lambda = 0$ as the seller sending the T-shirt for free to the consumer, who can try it for a week. If the consumer likes it, she pays the seller; if not, she gets a free return. Consider now a researcher studying this market who wrongfully believes that only boundary obfuscation is possible. If the researcher observes an online retailer that emails promotional material with pictures and sizes of the T-shirts it sells, it would signal to the researcher that the retailer is going out of its way to make it easier for the consumer to get better information. In a "boundary obfuscation only" world, this would therefore signal that the retailer would prefer to set $\lambda = 0$ (send the T-shirts for free for the consumer to try), but chooses not to due to unobservable costs, for instance, logistics. The researcher could then conclude that imposing a policy of free returns on online sales may be inefficient, as the firm would do it if it could afford it.

If, however, the firm was optimally choosing partial obfuscation, it is possible that the seller was purposely avoiding the free returns to prevent the consumer from acquiring too much information, and the researcher's conclusion about the free returns policy being inefficient would be invalid. □

Partial obfuscation appears when $\tilde{V}_{\mu_0}$ is non-monotonic. The prevalence of examples of partial obfuscation abounds: costs of add-ons, which are not completely hidden but may be in fine print; a firm that reveals the attributes of a mobile phone and allows consumers to compare different models but includes dozens of features to be checked, and so on. In addition, when a firm wants to induce full obfuscation, but that is prohibitively costly, it may choose to stop at partial obfuscation. Our base model shows that partial obfuscation is more general: even absent obfuscation costs, partial obfuscation may be optimal.

Our main result for this section, Theorem 1, provides a characterization of cases of partial obfuscation. Our main example, presented in Appendix B.1, details the case of a firm trying to induce the consumer to buy a product that may be good or bad. Before going into the theorem itself, we state the following result, which will prove useful.

Lemma 2. *If the consumer has Symmetric Preferences (i.e., $\mu_{ind} = \frac{1}{2}$), then we have a boundary obfuscation solution.*

In such scenarios, the firm chooses a boundary solution for any prior. Appendix A gives a detailed proof, but the main idea is that, in this environment, any increase in λ will always result in symmetric movements by $\underline{\mu}$ and $\overline{\mu}$. Since V is a step function when the firm has state-independent preferences, the resulting $\tilde{V}_{\mu_0}$ will be monotonic. This result provides some insight into how the consumer optimally learns based on the belief of indifference, given by the ratio $\frac{u(a_L,\omega_L)-u(a_H,\omega_L)}{u(a_H,\omega_H)-u(a_H,\omega_L)+u(a_L,\omega_L)-u(a_L,\omega_H)}$, which in some sense can be thought of as the ratio of the stakes of each state. If the stakes are the same in both states, then the consumer will learn about both states equally. But if one state has higher stakes, say ω_H, she will want to be more certain that she is in state ω_H when it happens to avoid the costlier mistakes, even if it means choosing the wrong action more often when ω_L happens instead.[23]

From this result, we get the motivation for looking into the consumer from the perspective of the belief that makes her indifferent between the two actions.

Definition 4. *We say that we have a* skeptical consumer *if:*

(i) *The firm's payoff is higher under action a_H, the indifference belief is lower than $\frac{1}{2}$ and the initial prior is even lower, i.e., $v(a_H) > v(a_L)$, with $\mu_0 < \mu_{ind} < 1/2$; or*

(ii) *The firm's payoff is higher under action a_L, the indifference belief is higher than $\frac{1}{2}$ and the initial prior is even higher, i.e., $v(a_H) < v(a_L)$, with $\mu_0 > \mu_{ind} > 1/2$.*

We have a *skeptical consumer* whenever there is a large range of beliefs for which the optimal action of the firm coincides with the consumer's optimal action and yet the consumer has such an extreme prior belief that her optimal action under the prior is different from the firm's optimal action. Another way to interpret such agent is one in which even a low (high) belief would be sufficient to convince her to take the action that is optimal for the firm, but

[23]Brown and Jeon (2024) provides empirical evidence of such behavior. They document that, when faced with choosing between different types of health insurance, patients with a higher likelihood of needing insurance tend to acquire more information about the insurance plans before making a decision.

her belief is even lower (higher) than this already low (high) threshold. The above definition helps us get to the main result for this section.

Theorem 1. *Given a Firm's Obfuscation Problem,*

(i) *we have a Partial Obfuscation solution if, and only if, we have a Skeptical consumer;*

(ii) *we have a Full Obfuscation solution if, and only if, we have Agreement at the Prior, i.e.* $a^*(\mu_0) = \arg\max_{a_i \in A} v(a_i)$*;*

(iii) *otherwise, we have a Full Transparency solution.*

The proof for the theorem can be found in Appendix A. For part (i), we first prove that skeptical consumer indeed yields Partial Obfuscation, then we proceed to show that all other possible cases result in Boundary Obfuscation. Part (ii) shows that the full obfuscation happens if, and only if, absent new information, the consumer would choose the firm's preferred action. Full transparency encompasses all the remaining cases.

Note that state-independent preferences are not necessary for partial obfuscation to occur. We illustrate this in another application in Appendix B.2 where the firm is a monopolist who wants to deter the entry of a possible competitor in the market.

Leveraging the results from Theorem 1, we can shed some light on policy implications that seek to improve the consumer's welfare by inducing full transparency. The next result can help us see how. Let $u_0^* \in \mathbb{R}$ be the consumer's expected payoff if the firm were to set $\lambda = 0$. This is the highest payoff the consumer can achieve as it is the full information payoff.[24]

Theorem 2. *Assume we do not have Full Transparency. Then, there exists a number* $\underline{\varepsilon} > 0$ *and an* $i \in \{H, L\}$ *such that reducing* $u(a_i, \omega_{-i})$ *by at least* $\underline{\varepsilon}$*, induces a full information equilibrium (the firm sets* $\lambda^* = 0$*) that achieves* u_0^**.*

It may at first be counterintuitive that strictly reducing the consumer's payoff from an action in a state can improve the consumer's equilibrium welfare. We are effectively increasing

[24]We have that $u_0^* = \mu_0 u(a_H, \omega_H) + (1 - \mu_0) u(a_L, \omega_L)$.

the costs of mistakes. However, the intuition is straightforward: increasing the costs of mistakes ensures that the consumer will only buy when she has enough information, leaving the firm no choice but to reduce λ. Moreover, the high cost of a mistake has no direct impact on welfare because under the induced full information, the consumer never makes a mistake.

Applying Theorem 2 only requires knowing things that are common knowledge: payoffs and the prior. It does not require experimentation on the part of the policy-maker and still leaves learning in the hands of the consumer. It is a useful example to see that understanding the payoff structure and obfuscation incentives is vital when examining policies aimed at improving the consumer's welfare.

Before moving forward, one might wonder if we could use a First Order Condition to study the optimal value of λ for the firm. Indeed, that can be done.

Proposition 2. *Given a firm's Obfuscation Problem, there exists at most one unique* $\lambda^* \in (0, \overline{\lambda})$ *such that* $\frac{d\tilde{V}_{\mu_0}(\lambda^*)}{d\lambda} = 0$.

Proposition 2 implies two things. First, any $\lambda \in (0, \overline{\lambda})$ that satisfies $\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} = 0$ is not only a local minimum/maximum, but a global minimum/maximum. Therefore, we have Partial Obfuscation only if there exists a $\lambda^* \in (0, \overline{\lambda})$ such that $\frac{d\tilde{V}_{\mu_0}(\lambda^*)}{d\lambda} = 0$. To see if the candidate λ^* is a global minimum instead of a global maximum, it suffices to check if $\tilde{V}_{\mu_0}(\lambda^*) > \tilde{V}_{\mu_0}(0)$. Second, it implies that a Partial Obfuscation solution to the firm's problem is unique.

Since the functional form of $\tilde{V}_{\mu_0}$ is rather convoluted - and its derivative even more so - it may be quite impractical to check the FOC.[25] On the other hand, checking for a Skeptical consumer takes very little effort.

3.1 Comparison to Communication Games

It is natural to ask how our model compares to other communication models. In particular, how it compares to the canonical models of Bayesian Persuasion (Kamenica and Gentzkow

[25] One could also check the SOC to separate between the minimum and maximum, but higher order derivatives of $\tilde{V}_{\mu_0}$ are even more complicated.

(2011)) and Cheap Talk (Crawford and Sobel (1982)). In a standard Bayesian Persuasion problem, the sender can choose any feasible information structure to influence the receiver's beliefs and actions. This gives the sender the highest possible expected payoff, which acts as an upper bound for any other model where the sender influences information.

Our strategic obfuscation model imposes a cost on the receiver for acquiring information. This means the sender's ability to influence the receiver is constrained by the receiver's willingness to pay for information. Consequently, the sender's expected payoff in a strategic obfuscation scenario will be weakly lower than what they could achieve in a pure Bayesian Persuasion setting.

The comparison with Cheap Talk is more nuanced. Cheap Talk involves costless, non-binding communication. A key difference lies in the sender's ability to commit. In a Cheap Talk game, the sender has no commitment power. In our strategic obfuscation model, the sender imposes a restriction on information through the cost of acquisition. While the sender doesn't directly commit to a specific message, the cost structure itself shapes what information the receiver will realistically choose to get. This gives the sender a form of indirect control or commitment that is absent in Cheap Talk.

However, Cheap Talk also offers a type of flexibility that strategic obfuscation lacks. In a Cheap Talk game, a sender could potentially use a wider range of communication experiments or signals that might not be possible to replicate using the rational inattention framework with information acquisition cost. Indeed, we show in Example 2 below that Cheap Talk can outperform strategic obfuscation when we allow for more than 2 actions. Example 3 then shows that it is not always the case.

Example 2. [Cheap Talk outperforming Strategic Obfuscation] There are two players, player 1 (sender or firm), and player 2 (receiver or consumer). Player 2 chooses an action from the set $A = \{a, b, c, d\}$, which determines both players' payoffs. The state $\omega \in \{High, Low\}$ is drawn with prior $\mu_0 = \Pr(\omega = High)$. Table 1 shows the payoffs for each player (the rows represent the state of the world).

Table 1: Payoffs for the players.

		Player 2			
		a	b	c	d
Player 1	High	1, -2	0, -1	1, 0	0, 1
	Low	1, 1	0, 0.9	1, 0	0, -1

We can see that player 1 has state-independent preferences, but player 2 wants to know the state of the world to choose the optimal action. Given the optimal action for player 2 at each belief, $a^*(\mu) = \arg\max_{x \in A} E_\mu[u(x, \omega)]$, Figure 5 plots the indirect utility for Player 1, $V(\mu) = \mathbb{E}_\mu[v(a^*(\mu)]$.

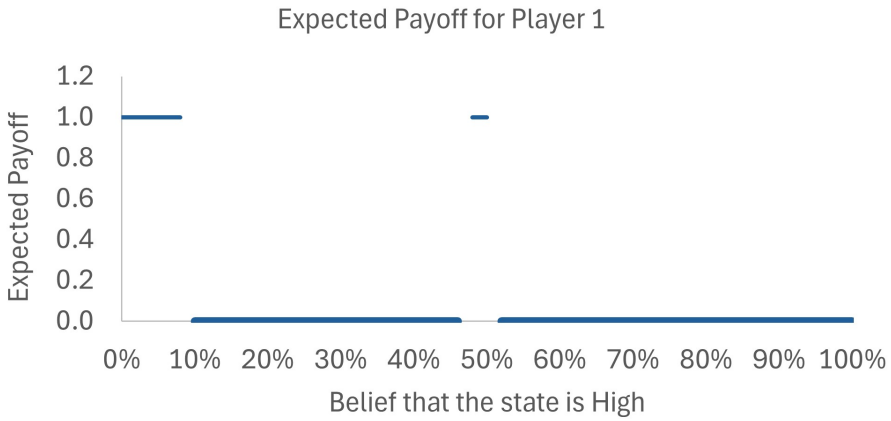

Figure 5: Expected payoffs for player 1.

Fix the prior $\mu_0 = 0.4$. We can now compare the outcomes under Cheap Talk and under Strategic Obfuscation. Under Cheap Talk, we can leverage the results from Lipnowski and Ravid (2020), and we know that there exists an equilibrium in which the sender (player 1) induces the distribution of posteriors τ given by

$$\tau(0.05) = \frac{2}{9};\ \tau(0.5) = \frac{7}{9},$$

which is Bayes' Plausible and yields a payoff of 1. The strategy is to induce posteriors that fall in the small areas where the payoff for player 1 is 1. Under Strategic Obfuscation, however, a payoff of 1 for player 1 is not possible as there is no $\lambda \in \mathbb{R}_+$ such that $\underline{\mu}(\lambda) < 9.09\%$ and

$\overline{\mu}(\lambda) \in (0.474, 0.500)$ simultaneously. Indeed, recall that $\underline{\mu}(\cdot)$ is strictly increasing while $\overline{\mu}(\cdot)$ is strictly decreasing and note that $\underline{\mu}(1.5) \approx 20\%$, while $\overline{\mu}(1.5) \approx 80\%$.

This example provides an intuition for when Cheap Talk yields a better payoff than obfuscation. It is when the actions that yield the best payoffs for the sender/firm happen in regions in the belief space that cannot all be reached by obfuscation simultaneously. □

Example 3. [Cheap Talk underperforming Strategic Obfuscation]

Consider now that Player 1 chooses an action from the set $A = \{a, b, c\}$, and that the payoffs are given by Table 2.

Table 2: Payoffs for the players.

		Player 2		
		a	b	c
Player 1	High	0.5, -1	0, 0.5	1, 1
	Low	0.5, 1	0, 0.5	1, -1

The corresponding expected utility for Player 2 is plotted in Figure 6. For a prior of 40%, a payoff of 0.5 is the maximum equilibrium payoff that a sender can get under cheap talk (Lipnowski and Ravid, 2020). However, under SO, the firm can choose $\lambda = 0$, which leads to full information and an expected payoff of 0.7.

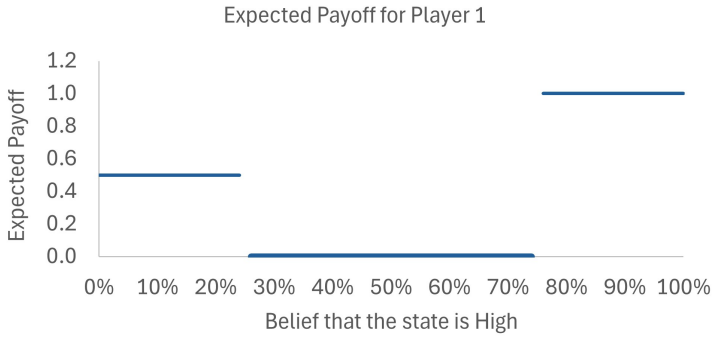

Figure 6: Expected payoffs for player 1.

□

For the remainder of this section, we keep the same environment of binary actions and

binary states of the world. Formally, let there be two players: player 1 (the firm, or the sender) wants to induce player 2 (the consumer, or the receiver) to take a certain action. Let $\Omega = \{\omega_H, \omega_L\}$ denote the possible states of the world, and $A = \{a_H, a_L\}$ denote the set of actions of player 2. Players 1 and 2 have utility functions $v, u : \Omega \times A \to \mathbb{R}$ respectively, and a common prior $\mu_0 \in \triangle(\Omega)$. Fixing these common objects, we can define the corresponding Bayesian Persuasion, Cheap Talk, and Obfuscation settings.[26]

Definition 5. Let s^P be the optimal payoff for the Sender under Bayesian Persuasion; let s^C be an equilibrium payoff for the Sender under Cheap Talk; and let s^O be the optimum payoff for the firm under Strategic Obfuscation.

With these in hand, we claim the next result.

Proposition 3. *The following holds:*

(i) $s^O \leqslant s^P$. *Under State-Independent Preferences for the sender/firm,* $s^O = s^P$ *if, and only if, we have Agreement at the Prior, i.e.,* $a^*(\mu_0) \subseteq \arg\max_{a_i \in A} v(a_i)$.

(ii) Under State-Independent Preferences for the sender/firm, $s^C \leqslant s^O$.

The highest outcome that the Sender can achieve under Bayesian Persuasion is an upper bound to what can be achieved under Strategic Obfuscation. The optimal experiment under both models will coincide only when the consumer chooses the firm-preferred action absent new information (Agreement at the Prior). For an intuition, note that the firm can only control the intensity of the cost of information, so we are essentially restricting the ability of the firm to induce signal structures. The only signals the firm can induce are the ones whose posteriors end up as $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ for some λ. In Bayesian Persuasion, however, the Sender is free to choose any signal structure as long as it satisfies Bayes plausibility.

In a simplified setting of 2 x 2, the comparison with Cheap Talk is rather straightforward (Cheap Talk would lead to Babbling Equilibria), but we kept the comparison in Proposition 3

[26]We can assume that the set of messages M for Cheap Talk is rich enough, i.e., $|M| \geqslant |A|$.

for completeness. In a more complex setting, as the previous examples show, the comparison between Cheap Talk and SO requires a more careful eye.

4 The General Model

Here, we expand the 2×2 model to allow for multiple states and actions, and we let the firm have state-dependent preferences. The rest of the setting (players, information, timing, etc...) remains the same. Formally, the consumer (she) now chooses an action from a finite set $A = \{a_1, \ldots, a_N\}$. We denote by $\Omega = \{\omega_1, \ldots, \omega_M\}$ the finite set of payoff-relevant states of nature, where the realization of a state is drawn from the prior distribution $\mu_0 \in int(\triangle(\Omega))$, known by both individuals.

The consumer and the firm have material utility functions $u : A \times \Omega \to \mathbb{R}$ and $v : A \times \Omega \to \mathbb{R}$ respectively. We let $a^*(\mu) \equiv \arg\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)]$ represent the set of actions that the consumer may take given a belief $\mu \in \triangle(\Omega)$, again breaking ties in favor of the firm. This will help ensure that a solution to the firm's problem exists. We restrict A to contain only actions that are optimal for the consumer for at least one prior belief.[27] We further assume that no two actions give the same payoffs in every state. Finally, to ensure that full obfuscation is possible (as we will see later), we assume that there is only one optimal action at the prior, so that $a^*(\mu_0)$ is a singleton[28] (as in the 2×2 case). This property holds almost-everywhere in the space of beliefs.[29]

As before, the consumer buys information $\pi(\cdot|\omega)$ (signal structure or Blackwell experiment) with a cost given by the Shannon Entropy, where the marginal cost of information is

[27] The consumer would never choose any other actions, and they can be removed without loss of generality.

[28] We could also get rid of this assumption if we assume that there is a cost to obfuscation (no matter how small), as in the extension in Appendix C.

[29] To give an intuition why, think of an example with 2 states and 2 actions: the resulting belief of indifference can be thought of as a single point in the $[0, 1]$ interval.

parametrized by λ.[30] We can again think in terms of the induced distribution of posteriors $\tau \in \triangle(\triangle(\Omega))$. The firm chooses the value of $\lambda \in \mathbb{R}_+$ and, for a given cost λ, the consumer chooses τ. When a signal a_i is realized, she takes action a_i, and utilities $u(a_i, \omega)$ and $v(a_i, \omega)$ are realized (given that the state is ω).

4.1 The consumer's Rational Inattention Problem

Here, we go into more detail over the solution to the consumer's problem. The statement of the consumer's problem is the same as given in Equation (4).

As mentioned before, we will use the characterization of the solution for the consumer's problem in Caplin et al. (2019), which can be called the posterior approach. However, when we get to the firm's problem, it will also be worth thinking of the solution in terms of the induced action probabilities. Indeed, Matějka and McKay (2015) show that we can alternatively write the consumer's problem in terms of the unconditional probabilities of taking a certain action. More formally, it is as if the consumer selects a stochastic choice function $P_\lambda : \Omega \to \triangle(A)$ from a set $\mathcal{P} = \{P : \Omega \to \triangle(A)\}$ (Corollary 1 from Matějka and McKay (2015)). For a signal π_λ that induces a posterior distribution τ_λ, the conditional probability that an action a is taken at state ω is $P_\lambda(a|\omega) = \pi_\lambda(a|\omega)$, and the unconditional probability of a is $P_\lambda(a) = \sum_{m=1}^{M} \mu_0(\omega_m)\pi_\lambda(a|\omega_m)$, so $P_\lambda(a) = \tau_\lambda(\mu^a)$ for any action played with positive probability. We let $A(\tau_\lambda) = \{a \in A : P_\lambda(a) > 0\}$ denote the set of actions played with positive probability under a solution τ_λ.[31]

[30] In the general case, the form of the Entropy function is given by

$$H(\mu) = -\sum_{m=1}^{M} \mu(\omega_m) \ln(\mu(\omega_m)).$$

[31] Caplin et al. (2019) calls this the consideration set.

4.1.1 Solving the consumer's Problem

The solution to the consumer's problem, for any $\lambda > 0$, has been characterized in Caplin et al. (2019), which we rewrite here for completeness.

Lemma 3. *(Proposition 2, Caplin et al. (2019).) A distribution of posteriors τ_λ and posteriors $\{\mu^a : a \in A(\tau_\lambda)\}$ is a solution to the consumer's Rational Inattention Problem (with $\lambda > 0$) if and only if, $\sum_{a \in A(\tau_\lambda)} \tau_\lambda(\mu^a)\mu^a(\omega) = \mu_0(\omega)$, and:*

1. *For any $\omega \in \Omega$ and any $a_i, a_j \in A(\tau_\lambda)$,*

$$\frac{\mu^{a_i}(\omega)}{\exp\left(\frac{u(a_i,\omega)}{\lambda}\right)} = \frac{\mu^{a_j}(\omega)}{\exp\left(\frac{u(a_j,\omega)}{\lambda}\right)}. \tag{11}$$

2. *Given $a_i, a_j \in A$ such that $a_i \in A(\tau_\lambda)$, and $a_j \notin A(\tau_\lambda)$,*

$$\sum_{m=1}^{M} \mu^{a_i}(\omega_m) \exp\left(\frac{u(a_j,\omega_m) - u(a_i,\omega_m)}{\lambda}\right) \leqslant 1. \tag{12}$$

We denote by $T_{\mu_0}(\lambda)$ the set of solutions to the consumer's problem for a given λ. Again, let $\tau_\lambda \in T_{\mu_0}(\lambda)$ denote an arbitrary solution to the consumer's problem.

4.1.2 The Impact of the Cost of Information on the consumer's Solution

In this section, we analyze how the solution τ_λ for the consumer's Problem changes as we vary λ. To find τ_λ, we must solve the system of equations defined by (11) and (12), which does not provide an easy functional form. Therefore, to think about the dynamics of the solution as λ increases, it is helpful to define, for each $a_i \in A$, the set

$$I(\lambda; a_i) \equiv \left\{ \mu \in \triangle(\Omega) : \sum_{m=1}^{M} \mu(\omega_m) \exp\left(\frac{u(a_j,\omega_m) - u(a_i,\omega_m)}{\lambda}\right) \leqslant 1, \forall j \neq i \right\}. \tag{13}$$

This set characterizes, for a given λ, all prior beliefs in which buying no information and

playing a_i with probability 1 is optimal.[32]

We define $I(0, a_i) \equiv \lim_{\lambda \to 0^+} I(\lambda; a_i)$, which has a simple characterization. Note that, as λ approaches $0, I(\lambda; a_i)$ only includes beliefs that put a very low probability on states in which a_i is not the optimal action. Indeed, let $\Omega(a_i) = \{\omega \in \Omega : u(a_i, \omega) \geqslant u(a_j, \omega)\ \forall j \neq i\}$ denote the set of all states in which a_i is an optimal action.

Lemma 4. *We have that* $\lim_{\lambda \to 0^+} I(\lambda; a_i) = \{\mu \in \triangle(\Omega) : \text{supp}(\mu) \subset \Omega(a_i)\}$.

Note that it is possible for $I(0, a_i) = \emptyset$ for some i, but not all, and that if $I(0, a_i) \neq \emptyset$, it is contained in the boundary of $\triangle(\Omega)$ as we assume that no action is optimal at all states.

If $\lambda_1 < \lambda_2$, then $I(a, \lambda_1) \subset I(a, \lambda_2)$, which intuitively implies that the regions where not learning is optimal grow (in the set inclusion sense) as information becomes more expensive.[33] Let $\delta_\mu \in \triangle(\triangle(\Omega))$ be the degenerate distribution of posteriors that puts probability 1 on belief μ, and the following Proposition will help us understand the consumer's decision when λ is large enough that $\mu_0 \in I(\lambda; a_i)$ for some a_i.

Proposition 4. *We have that* $\delta_{\mu_0} \in T_{\mu_0}(\lambda)$ *if and only if* $\mu_0 \in I(\lambda; a_i)$ *for some* $a_i \in A$. *Moreover, if* $\mu_0 \in int(I(\lambda; a_i))$, *then* $T_{\mu_0}(\lambda) = \{\delta_{\mu_0}\}$.

Indeed, if $\mu_0 \in I(\lambda; a_i)$, setting $\mu^a = \mu_0$, and $\tau_\lambda(\mu^a) = 1$ automatically satisfies the sufficient and necessary conditions of the solution to the consumer's problem. New information is expensive, so the consumer prefers to act according to the already available information: her prior. In the full proof, we also formalize the intuition that, at any belief $\mu \in I(\lambda; a_i)$, action a_i is optimal, and if μ is also interior, the consumer cannot be indifferent between any two actions at μ (i.e., $a^*(\mu) = \{a_i\}$).

The main contribution of Proposition 4 lies in showing that acquiring no information is the unique solution whenever μ_0 is in the interior of $I(\lambda; a_i)$. This allows us to know the

[32]This set is used in Caplin et al. (2019) and Matysková and Montes (2023) to define the (in their case, exogenous) set of prior beliefs in which no-learning is a solution for the consumer's problem. Our Proposition 4 shows when it is the unique solution to the consumer's problem. We use uniqueness to guarantee that the objective function of the firm's problem is well behaved. We cannot restrict our attention to the conditions for uniqueness specified in Caplin et al. (2019) Section 4 as those hold only for a fixed λ.

[33]For a formal proof, check Proposition 3 in Matysková and Montes (2023).

values of the cost of information that induces no learning as a function of the utility of each choice. When $\mu_0 \in int(I(\lambda; a_i))$ for some $\lambda > 0$, the firm knows the consumer will take action a_i with probability 1, and we do not need to consider tie-breaking rules.[34]

We are now equipped to understand how the consumer acquires information for any $\lambda \in \mathbb{R}_+$. At $\lambda = 0$, the consumer buys a fully informative signal, learns the true state, and takes the optimal action (this means $\mu^{a_i} \in I(a_i, 0)$ for all a_i s.t. $P_\lambda(a_i) > 0$). If instead $\mu_0 \notin \cup_i I(\lambda; a_i)$, the consumer will purchase an informative signal and follow the action recommended by the signal realization. Lemma 3 yields the posteriors that follow a signal realization and the unconditional probability that an action is recommended. No matter what happens, for any $\lambda \geqslant 0$, the solution τ_λ satisfies $\text{supp}(\tau_\lambda) \subset \cup_i I(\lambda; a_i)$. Finally, if λ is high enough such that and $\mu_0 \in int(I(\lambda; a_i))$ for some action a_i, the consumer acquires a completely uninformative signal and chooses action a_i with probability 1.

The last step is to understand what happens as λ grows arbitrarily large. Surprisingly, there are prior beliefs for which, no matter how high the marginal cost of information λ gets, the consumer will always optimally choose to acquire some information. Let $C_\mu(A) = \arg\max_{a_i \in A} \mathbb{E}_\mu[u(a_i, \omega)]$ denote the set of optimal actions for the Consumer given a belief μ.

Proposition 5. *For any $\mu \in int(\triangle(\Omega))$, the following statements are equivalent:*

(i) $a^*(\mu)$ is a singleton;

(ii) there exists $\overline{\lambda} \in \mathbb{R}_{++}$ s.t. $\mu \in int(I(\overline{\lambda}; a_i))$ for some $a_i \in C_\mu(A)$.

Proposition 5 characterizes the prior beliefs for which the consumer always learns, no matter the cost. They are exactly the beliefs in which she has more than one optimal action available. We can see why: when she is indifferent between actions, any information, no matter how coarse, will help break the tie. Therefore, even if λ is arbitrarily large, the

[34] Alternatively, there could be other "firm preferred" solutions where the consumer is indifferent between not learning or learning even with a very high λ. Given our tie-breaking rule that favors the firm, solving the firm's problem would be much more complicated.

consumer can afford a cheap enough signal (by making it less informative) that is still useful. Consequently, inducing total obfuscation is impossible in these cases, and only in these cases.

Corollary 1. *Full obfuscation is possible, i.e.* $\delta_{\mu_0} \in T_{\mu_0}(\lambda)$ *for some* $\lambda \in \mathbb{R}_+$*, if and only if,* $a^*(\mu_0)$ *is a singleton.*

An important property of the consumer's solution is that of Locally Invariant Posteriors (Caplin et al., 2022): the support of the distribution of induced posteriors is the same for any $\mu_0 \notin \cup_i I(\lambda; a_i)$. The prior only affects the solution by altering the distribution of posteriors through Bayes' plausibility. While this is surely still true in our model when looking from the consumer's side - i.e., for a fixed λ, - when we consider the firm, the prior has a much bigger impact on equilibrium outcomes. As we will see in Section 4.2, two different priors, even if both are in $\cup_i I(\lambda; a_i)$, may lead to wildly different obfuscation levels and, consequently, different supports for the distribution of induced posteriors.

4.2 The firm's Problem

It may be helpful to see that we can rewrite the firm's indirect utility function $\tilde{V}_{\mu_0}$ as

$$\tilde{V}_{\mu_0}(\lambda) = \sum_{i=1}^{N} P_\lambda(a_i) \sum_{m=1}^{M} \mu^{a_i}(\omega_m) v(a_i, \omega_m). \tag{14}$$

The statement of the firm's problem then stays the same as in Equation (9). When the firm has state-independent preferences, we can write $\tilde{V}_{\mu_0}(\lambda) = \sum_{i=1}^{N} P_\lambda(a_i) v(a_i)$.
Intuitively, the firm chooses the λ that maximizes the probability its favorite actions are played, which becomes clearer in Equation (14).

At this point, the assumption that $a^*(\mu_0)$ is a singleton comes into play. If the firm wants to induce full obfuscation and $a^*(\mu_0)$ is not a singleton, by Proposition 5 that would not be possible. However, since we assume $a^*(\mu_0)$ is a singleton, there exists $\overline{\lambda}$ such that $\mu_0 \in int(I(\lambda; a_i))$ for some $a_i \in A$ and we can then restrict the firm's choice set to $\mathcal{O} \equiv [0, \overline{\lambda}]$. We are ready to state our next result.

Proposition 6. *The Firm's Obfuscation Problem has a solution.*

The formal proof is in Appendix A, but it follows the usual arguments. The assumption that $a^*(\mu_0)$ is a singleton ensures that $\mathcal{O}$ is compact and the assumption that ties are broken in the firm's favor ensures that the maximizing function is upper semicontinuous.

4.3 Transparency and Obfuscation in the General Case

We begin by discussing three classes of games where we can easily determine the equilibrium level of obfuscation. In what follows, let $a_F^*(\mu) = argmax_{a\in A}\mathbb{E}_\mu[v(a,\omega)]$ denote the optimal action for the firm given a belief μ.

Aligned firm: $\arg\max_{a_i\in A} u(a_i,\omega_m) \subseteq \arg\max_{a_i\in A} v(a_i,\omega_m)$ for all $\omega_m \in \Omega$. The firm's utility is decreasing in λ for any prior as the consumer's mistakes are sub-optimal for the firm. The solution is $\lambda = 0$ and we have *full transparency*. By contraposition, obfuscation can arise only if there exists at least one belief μ in which $a_F^*(\mu) \neq a^*(\mu)$.

Totally misaligned firm: $u(a_i,\omega_m) + v(a_i,\omega_m) = c,\ \forall a_i,\omega_m$, where c is a constant. We essentially have a zero-sum game: the consumer's utility is weakly decreasing in λ, while the firm's utility is weakly increasing. Therefore, we have full obfuscation.

Agreement at the Prior: the firm has state-independent preferences ($v(a_i,\omega_m) = v(a_i)$ for all $m = 1,\dots,M$) and $a^*(\mu_0) \subseteq \arg\max_{a_i\in A} v(a_i)$. Then $\lambda = \overline{\lambda}$ is necessarily optimal, and we have a full obfuscation problem. This case has been mentioned before throughout the paper.

As a final remark, consider a paternalistic principal (for instance, when $v = u$), so that the principal and the agent agree on the optimal action in every state. Lipnowski et al. (2020) show that full disclosure is optimal if and only if there are two possible states of nature. When the agent exhibits rational inattention, however, the principal cannot fully disclose the state, as the agent would be unable to process all the available information. In contrast, our model allows the principal (or firm) to effectively reduce the cost of information acquisition to zero, making full disclosure optimal even when there are more than two states.

4.4 Strategic Obfuscation and Consideration Sets

This section applies our general framework to a well-known example from Caplin et al. (2019, Example 3.1) that illustrates the formation of *consideration sets*. The environment features a consumer who must choose one good from a menu of N options, where exactly one option is good and the remaining $N-1$ are bad. Identifying the good option requires costly information acquisition, and we interpret this cost as being endogenously set by the firm through its choice of λ. The firm has a single seller-preferred option,[35] which may or may not coincide with the consumer's uniquely good option.

In Caplin et al. (2019), the information cost parameter λ is exogenous. Here, we make it a firm choice variable, allowing us to study how strategic obfuscation influences which alternatives the consumer actually considers. When λ is small, learning is cheap and the consumer examines many options; as λ rises, learning becomes harder and her consideration set shrinks. Eventually, when λ is sufficiently high, she focuses only on the option with the highest prior probability of being good and may stop learning altogether.

Formally, for a fixed λ, the consideration set consists of the actions chosen with positive probability, $\{a \in A : P_\lambda(a) > 0\}$. In this example, the consumer must choose the single good option from a menu $A = \{a_1, \ldots, a_N\}$ that also includes $N-1$ bad options. The state space is $\Omega = \{\omega_1, \ldots, \omega_N\}$, where the realization of ω_m indicates that a_m is the good option. The consumer's utility is given by $u(a, \omega)$, where a is the chosen action and ω the realized state of the world. Choosing the good (bad) option yields utility u_G (u_B), with $u_G > u_B$. Formally,

$$u(a_i, \omega_m) = \begin{cases} u_G, & \text{if } i = m, \\ u_B, & \text{otherwise.} \end{cases}$$

Given the prior μ_0, we order the states without loss of generality so that $\mu_0(\omega_i) \geqslant \mu_0(\omega_{i+1})$ for all i; that is, option 1 is the most likely to be good, while option N is the least likely.

[35] We can think of an ice cream truck in which the price for all flavors is the same, but the vanilla flavor (the blandest) has a lower marginal cost of production.

Theorem 1 in Caplin et al. (2019) characterizes the optimal choice for the rationally inattentive consumer given λ, which in turn defines the corresponding consideration set. The authors highlight the impact of λ on the consideration set: if the cost is too high, some options are never considered and the consumer learns nothing about them; even when the optimal good is in the consideration set, the consumer makes mistakes and may choose sub-optimally. To illustrate, Caplin et al. provide the following numerical example: let $u_G = 1, u_B = 0$, and $N = 10$. Moreover, let the prior distribution over the states of the world be $\mu_0(\omega_m) = \alpha 0.8^{m-1}, m = 1, \ldots, 10$ and $\alpha = \frac{1-0.8}{1-0.8^{10}}$. Figure 7 shows how the size of the consideration set varies with the information cost λ.

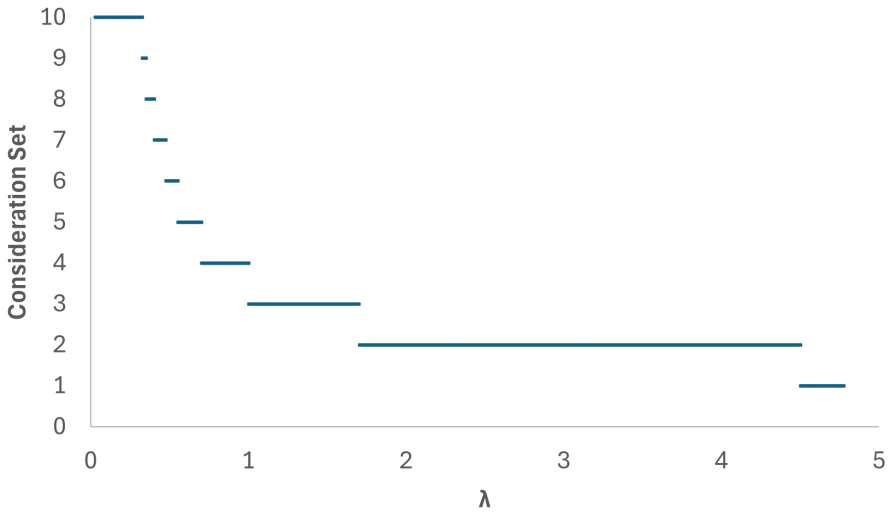

Figure 7: The number of options in the consideration set (vertical axis) as a function of λ. For low λ, all 10 options are considered; as λ increases, less and less options are considered, until around $\lambda \approx 4.5$ when only option 1 is considered and it is bought with probability 1.

When the seller-preferred option is 1, the optimal λ induces no-learning: with no new information, the consumer goes for the most likely best option according to her prior and buys option 1 with probability 1. When the seller-preferred option is 10, the optimal λ induces full learning. When the firm wants to sell a product that a consumer perceives to be bad, it wants to induce learning to increase the chances of changing her mind. Figure 8 illustrates these results.

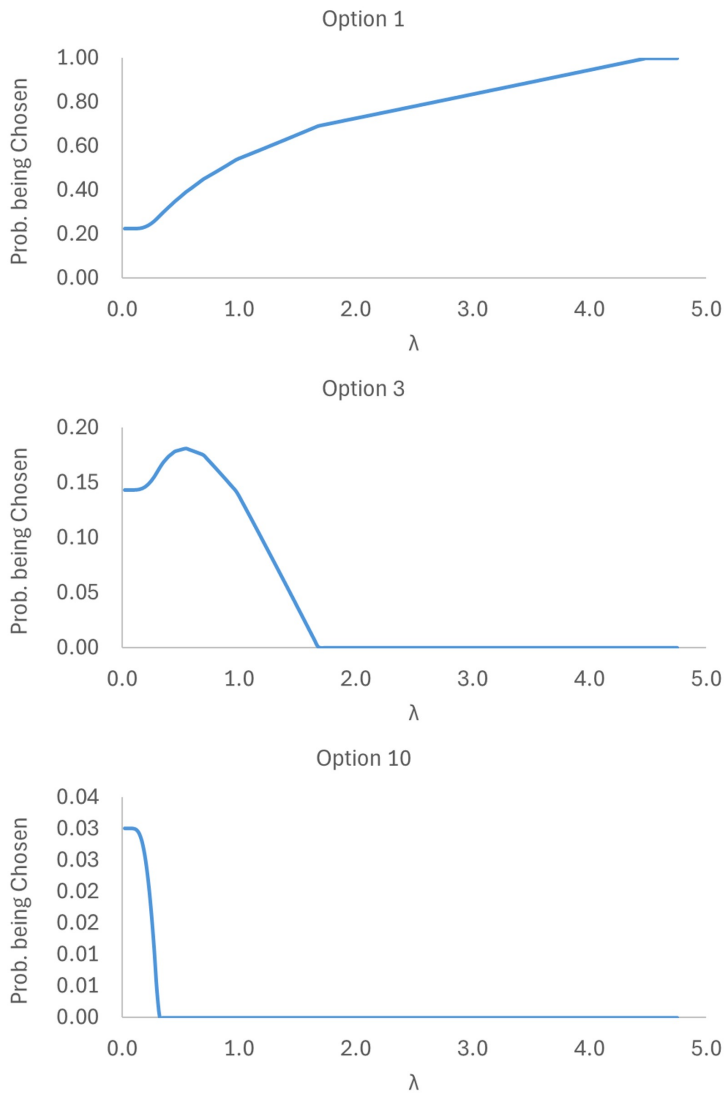

Figure 8: The probability of the consumer choosing: (a) option 1, which is increasing in λ; (b) option 3, which is non-monotonic in λ; (c) option 10, which is decreasing in λ.

The cases of options 1 and 10 illustrate the main forces at play. Learning promotes experimentation and is therefore beneficial for ex-ante worse goods. At the same time, it enlarges the consideration set and increases the likelihood that other goods are chosen. This tradeoff generates the non-monotonic relationship between the probability that option 3 is selected and the information cost λ. At first glance, one might expect the optimal λ for maximizing sales of option 3 to be just low enough to include it in the consideration set while keeping all options $i > 3$ excluded. Interestingly, this is not the case. The value

$\lambda \approx 0.55$ that maximizes the probability of selling option 3 induces a consideration set that also includes options 4 and 5, as shown in Figure 7. For option 3, allowing for slightly more experimentation—even if this means that additional goods enter the consideration set—remains beneficial.

Two last remarks. First, if the prior was uniform ($\mu_0(a_i) = 1/N$ for all a_i), the consideration set would remain constant and include all options for all values of λ. Proposition 5 tells us that, in this case, the firm could not induce full obfuscation. Second, from a consumer welfare perspective, the best case is when option 10 is seller-preferred as it induces full information and the consumer always gets the good option.

This setup reveals how the firm's preferred product determines its optimal degree of obfuscation. If the firm's favored product already enjoys a high prior probability, raising λ—and thus discouraging search—benefits the firm. In contrast, if the preferred product has a low prior probability, the firm gains from lowering λ to encourage information acquisition. In some intermediate cases, the relationship between λ and the probability of sale is non-monotonic: a moderate amount of obfuscation can increase sales, but excessive obfuscation drives the consumer to rely on her prior and ignore other options. This analysis connects strategic obfuscation to the emergence of limited consideration and choice overload.

5 Conclusion

This paper develops a model in which a firm can induce a rationally inattentive consumer to take particular actions by choosing the cost of processing information. The framework endogenizes information costs and shows how persuasion can operate through the design of learning frictions rather than through direct message transmission.

Our first contribution is to provide a theoretical foundation for why agents may rationally shroud information. In line with empirical evidence on strategic obfuscation, we show that the optimal degree of obfuscation is often neither zero nor extreme, but moderate: the

firm engages in what we call *partial obfuscation*. We illustrate several mechanisms through which this outcome arises—for example, asymmetries in consumer preferences—and discuss economic settings where such behavior is likely to occur, including competition, insurance, and politics.

Second, the model highlights the welfare consequences of obfuscation. By raising information costs, firms can induce consumers to make mistakes, and policy interventions that overlook these incentives may be counterproductive. Our Good or Bad Product example illustrates how policies that appear to harm consumers in one dimension may improve welfare by reducing firms' incentives to obfuscate.

Finally, the paper contributes to the literature on information design by providing a framework in which a firm influences decisions solely by controlling the consumer's access to information. Persuasion, in this view, concerns the design of information costs rather than commitment to disclosure. The framework extends naturally to environments with multiple actions and states, and offers a basis for future research on the interaction between rational inattention and strategic communication.

References

Alonso, R. and O. Camara (2024). Organizing data analytics. *Management Science 70*(5), 3123–3143.

Asriyan, V., D. Foarta, and V. Vanasco (2023). The good, the bad, and the complex: Product design with imperfect information. *American Economic Journal: Microeconomics 15*(2), 187–226.

Barros, L. (2021). Information acquisition design. *Master Thesis, Sao Paulo School of Economics - FGV*.

Ben-Shahar, O. and C. E. Schneider (2014, apr). *More Than You Wanted to Know*. Princeton University Press.

Bergemann, D. and S. Morris (2019, March). Information design: A unified perspective. *Journal of Economic Literature 57*(1), 44–95.

Bizzotto, J., J. Rüdiger, and A. Vigier (2020). Testing, disclosure and approval. *Journal of Economic Theory 187*, 105002.

Bloedel, A. W. and I. Segal (2021). Persuading a rationally inattentive agent. *Available at SSRN*.

Brown, Z. Y. and J. Jeon (2024). Endogenous information and simplifying insurance choice. *Econometrica 92*(3), 881–911.

Caplin, A., M. Dean, and J. Leahy (2019). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies 86*(3), 1061–1094.

Caplin, A., M. Dean, and J. Leahy (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy 130*(6), 1676–1715.

Crawford, V. P. and J. Sobel (1982, nov). Strategic Information Transmission. *Econometrica 50*(6), 1431.

Cusumano, C. M., F. Fabbri, and F. Pieroth (2024). Competing to commit: Markets with rational inattention. *American Economic Review 114*(1), 285–306.

De Clippel, G. and K. Rozen (2025). Communication, Perception and Strategic Obfuscation. working paper.

De Oliveira, H., T. Denti, M. Mihm, and K. Ozbek (2017). Rationally inattentive preferences and hidden information costs. *Theoretical Economics 12*(2), 621–654.

Dean, M. and N. Neligh (2023). Experimental tests of rational inattention. *Journal of Political Economy 131*(12), 3415–3461.

Edmond, C. and Y. K. Lu (2021). Creating confusion. *Journal of Economic Theory 191*, 105145.

Ellis, A. (2018). Foundations for optimal inattention. *Journal of Economic Theory 173*, 56–94.

Ellison, G. (2005, may). A Model of Add-On Pricing. *The Quarterly Journal of Eco-*

nomics 120(2), 585–637.

Ellison, G. and A. Wolitzky (2012). A search cost model of obfuscation. *RAND Journal of Economics 43*(3), 417–441.

Fabbri, F. (2024). Attention holdup. *Working Paper*.

Fréchette, G. R., A. Lizzeri, and J. Perego (2022). Rules and commitment in communication: An experimental analysis. *Econometrica 90*(5), 2283–2318.

Gabaix, X. and D. Laibson (2006, may). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *The Quarterly Journal of Economics 121*(2), 505–540.

Gentzkow, M. and E. Kamenica (2014). Costly Persuasion. *American Economic Review 104*(5), 457–462.

Jin, G. Z., M. Luca, and D. Martin (2022). Complex disclosure. *Management Science 68*(5), 3236–3261.

Kalayci, K. and J. Potters (2011). Buyer confusion and market prices. *International Journal of Industrial Organization 29*(1), 14–22.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics 11*(1).

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Le Treust, M. and T. Tomala (2019). Persuasion with limited communication capacity. *Journal of Economic Theory 184*, 104940.

Li, H. and X. Shi (2017, November). Discriminatory information disclosure. *American Economic Review 107*(11), 3363–85.

Lin, X. and C. Liu (2024). Credible persuasion. *Journal of Political Economy 132*(7), 2228–2273.

Lipnowski, E., L. Mathevet, and D. Wei (2020, mar). Attention Management. *American Economic Review: Insights 2*(1), 17–32.

Lipnowski, E. and D. Ravid (2020). Cheap talk with transparent motives. *Econometrica 88*(4), 1631–1660.

Lipnowski, E., D. Ravid, and D. Shishkin (2022). Persuasion via weak institutions. *Journal of Political Economy 130*(10), 2705–2730.

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

Matějka, F. and A. McKay (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review 105*(1), 272–298.

Matysková, L. and A. Montes (2023). Bayesian persuasion with costly information acquisition. *Journal of Economic Theory 211*, 105678.

Mensch, J. and D. Ravid (2022). Monopoly, product quality, and flexible learning. *arXiv preprint arXiv:2202.09985*.

Min, D. (2021). Bayesian persuasion under partial commitment. *Economic Theory 72*, 743–764.

Nguyen, A. and T. Y. Tan (2021). Bayesian persuasion with costly messages. *Journal of Economic Theory 193*, 105212.

Nocke, V. and P. Rey (2024). Consumer search, steering, and choice overload. *Journal of Political Economy 132*(5), 1684–1739.

Perez-Richet, E. and V. Skreta (2022). Test design under falsification. *Econometrica 90*(3), 1109–1142.

Ravid, D. (2020). Ultimatum bargaining with rational inattention. *American Economic Review 110*(9), 2948–63.

Richards, T. J., G. J. Klein, C. Bonnet, and Z. Bouamra-Mechemache (2020). Strategic obfuscation and retail pricing. *Review of Industrial Organization 57*, 859–889.

Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal 27*(4), 623–656.

Sims, C. A. (2003, apr). Implications of rational inattention. *Journal of Monetary Eco-*

nomics 50(3), 665–690.

Wei, D. (2020). Persuasion under costly learning. *Journal of Mathematical Economics*, 1–26.

Appendix

A Proofs

Proof of Proposition 1 *Given a firm's problem, the following holds:*

(i) *if V is concave, then $\tilde{V}_{\mu_0}$ is increasing in λ;*

(ii) *if V is convex, then $\tilde{V}_{\mu_0}$ is decreasing in λ.*

Proof of Proposition 1. We shall detail the proof only for the case in which V is concave, as the proof for V convex is analogous. First, we know that there exists $\overline{\lambda}$ large enough such that, for every $\lambda \geqslant \overline{\lambda}$, $\tilde{V}_{\mu_0}(\lambda) = V(\mu_0)$. As a result, we can restrict our analysis to $\lambda \in [0, \overline{\lambda}]$. In this interval, $\tilde{V}_{\mu_0}(\lambda) = m(\lambda) * \mu_0 + V(\underline{\mu}(\lambda)) - m(\lambda) * \underline{\mu}(\lambda)$. We also know that:

(i) $\tilde{V}_{\mu_0}$ valued at a specific $\lambda \in [0, \overline{\lambda}]$, is represented by the *y-coordinate* of the point given by the intersection of the line segment connecting the points $(\underline{\mu}(\lambda), V(\underline{\mu}(\lambda)))$ and $(\overline{\mu}(\lambda), V(\overline{\mu}(\lambda)))$ and the vertical line at μ_0;

(ii) $\underline{\mu}$ and $\overline{\mu}$ are strictly increasing and decreasing respectively. This means that, for $0 \leqslant \lambda_1 < \lambda_2 \leqslant \overline{\lambda}$, we have $0 \leqslant \underline{\mu}(\lambda_1) < \underline{\mu}(\lambda_2) \leqslant \mu_0 \leqslant \overline{\mu}(\lambda_2) < \overline{\mu}(\lambda_1) \leqslant 1$, with at most one equality for $\underline{\mu}(\lambda_2) \leqslant \mu_0 \leqslant \overline{\mu}(\lambda_2)$.

Thus, if we prove the general result that, for $a < c \leqslant \mu_0 \leqslant d < b$ (again, with at least one strict inequality), and a concave function $f : \mathbb{R}_+ \to \mathbb{R}$, where $[a, b] \subset \mathbb{R}_+$, the intersection of the line segment connecting $(a, f(a))$ and $(b, f(b))$, when evaluated at μ_0, has an equal or lower value than the line segment connecting $(c, f(c))$ and $(d, f(d))$ when evaluated at μ_0, we are done. To put it simply, we are considering that $[a, b]$ represents $[\underline{\mu}(\lambda_1), \overline{\mu}(\lambda_1)]$ and $[c, d]$ represents $[\underline{\mu}(\lambda_2), \overline{\mu}(\lambda_2)]$ for arbitrary $\lambda_1 < \lambda_2$.

Figure 9 gives a visual representation. Points A and B represent the intersection of different line segments. What we wish to prove is that Point B is lies at least as high as Point

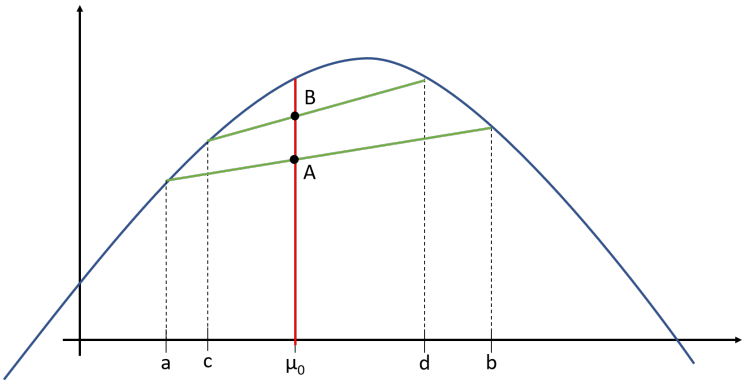

Figure 9: A concave function f and the intersection of the two line segments with the vertical line at μ_0.

A, as Point B is the intersection of the line $y = \mu_0$ with the innermost line segment (which would correspond to $\tilde{V}_{\mu_0}$ at a given cost λ).

With basic geometry, we can write the formulas for each line segment $C1$ and $C2$, for $x_1 \in [a, b],\ x_2 \in [c, d]$ as the following functions

$$C1(x_1) : \left(\frac{f(b) - f(a)}{b - a}\right) x_1 + f(a) - \left(\frac{f(b) - f(a)}{b - a}\right) a,$$
$$C2(x_2) : \left(\frac{f(d) - f(c)}{d - c}\right) x_2 + f(c) - \left(\frac{f(d) - f(c)}{d - c}\right) c.$$

Now, evaluating $C1$ and $C2$ at c, we get

$$C1(c) : (f(b) - f(a)) \frac{(c - a)}{(b - a)} + f(a),$$
$$C2(c) : f(c).$$

But since $c \in (a, b)$, there exists $t_c \in (0, 1)$ such that $c = t_c b + (1 - t_c)a$. We then get that

$$\frac{(c - a)}{(b - a)} = \frac{t_c b + a - t_c a - a}{b - a} = t_c \frac{b - a}{b - a} = t_c,$$

and so:

$$(f(b) - f(a)) \frac{(c - a)}{(b - a)} + f(a) = (f(b) - f(a))\, t_c + f(a). \tag{15}$$

Finally, since f is concave,

$$C2(c) = f(c) = f(t_c b + (1 - t_c)a) \geqslant t_c f(b) + (1 - t_c) f(a) = (f(b) - f(a))\, t_c + f(a) = C1(c).$$

The same argument can be made for the point d. Since $d \in (a, b)$, there exists $t_d \in (0, 1)$ such that $d = t_d b + (1 - t_d)a$, and we get that:

$$C2(d) = f(d) = f(t_d b + (1 - t_d)a) \geqslant t_d f(b) + (1 - t_d) f(a) = (f(b) - f(a))\, t_d + f(a) = C1(d)$$

And so, $C1(c) \leqslant C2(c)$ and $C1(d) \leqslant C2(d)$. Since the line segments are by definition linear, they are on the same line or they intersect at most once: $C2$ is never below $C1$ in the interval $[c, d]$. Finally, as $\mu_0 \in [c, d]$, then $C2(\mu_0) \geqslant C1(\mu_0)$. As mentioned before, $C2(\mu_0)$ represent the value of $\tilde{V}_{\mu_0}$ at λ_2 and $C1(\mu_0)$ the value of $\tilde{V}_{\mu_0}$ at λ_1, and we get that $\tilde{V}_{\mu_0}$ is increasing in λ. □

Lemma 2 *If the consumer has Symmetric Preferences (i.e., $\mu_{ind} = \frac{1}{2}$), then we have a boundary obfuscation solution.*

Proof of Lemma 2. First, note that since the consumer's preferences are symmetric, $\underline{\mu}(\lambda) = 1 - \overline{\mu}(\lambda)$ for all λ. This can be checked by noting that $\mu_{ind} = 0.5$ implies that $u(a_H, \omega_H) - u(a_L, \omega_H) = u(a_L, \omega_L) - u(a_H, \omega_L)$ and using substitution in Equations 5 and 6. Given that the firm's preferences are state-independent, for any belief $\mu < \mu_{ind}$, $V(\mu) = V(0)$ and for any belief $\mu > \mu_{ind}$, $V(\mu) = V(1)$. Assume that $V(0) > V(1)$ (the other case follows a similar argument). Now, fix $\lambda_1, \lambda_2 \in \mathcal{O} = [0, \overline{\lambda}]$ (as outside this interval the function is constant) with $\lambda_1 < \lambda_2$, and let $d \coloneqq \mu_0 - \underline{\mu}(\lambda_1)$, $D \coloneqq \overline{\mu}(\lambda_1) - \mu_0$. These last variables represent the distance between the induced posteriors and the prior.

By triangle similarity, we have

$$\frac{d}{V(\underline{\mu}(\lambda_1)) - \tilde{V}_{\mu_0}(\lambda_1)} = \frac{D}{\tilde{V}_{\mu_0}(\lambda_1) - V(\overline{\mu}(\lambda_1))}.$$

Since $\underline{\mu}(\lambda_1) < \mu_{ind}$ and $\overline{\mu}(\lambda_1) > \mu_{ind}$, we have $V(\underline{\mu}(\lambda_1)) = V(0)$ and $V(\overline{\mu}(\lambda_1)) = V(1)$. We can then write

$$\frac{d}{V(0) - \tilde{V}_{\mu_0}(\lambda_1)} = \frac{D}{\tilde{V}_{\mu_0}(\lambda_1) - V(1)} \implies \tilde{V}_{\mu_0}(\lambda_1) = \frac{D}{d+D}V(0) + \frac{d}{d+D}V(1). \tag{16}$$

Now, let $\delta \coloneqq \underline{\mu}(\lambda_2) - \underline{\mu}(\lambda_1)$. Since consumer's preferences are Symmetric, $\overline{\mu}(\lambda_1) - \overline{\mu}(\lambda_2) = \delta$, and we can write

$$\tilde{V}_{\mu_0}(\lambda_2) = \frac{D-\delta}{d+D-2\delta}V(0) + \frac{d-\delta}{d+D-2\delta}V(1). \tag{17}$$

But note that, by construction, $\delta < \frac{d+D}{2}$, and so the following holds:

$$\frac{D-\delta}{d+D-2\delta} > \frac{D}{d+D} \iff d < D.$$

We now have two cases:

(i) $D > d$: This happens when $\mu_0 < \mu_{ind}$. This also implies that from Equation 16 to Equation 17 we are increasing the proportion of $V(0)$ and since we assume that $V(0) > V(1)$, this means that $\tilde{V}_{\mu_0}(\lambda_1) < \tilde{V}_{\mu_0}(\lambda_2)$ and the problem has a boundary solution as $\tilde{V}_{\mu_0}$ is increasing.

(ii) $d > D$: This happens when $\mu_0 > \mu_{ind}$. This also implies that from Equation 16 to Equation 17 we are decreasing the proportion of $V(0)$ and since we assume that $V(0) > V(1)$, this means that $\tilde{V}_{\mu_0}(\lambda_1) > \tilde{V}_{\mu_0}(\lambda_2)$ and the problem has a boundary solution as $\tilde{V}_{\mu_0}$ is decreasing.

□

Theorem 1 *Given a Firm's Obfuscation Problem,*

(i) *we have a Partial Obfuscation solution if, and only if, we have a Skeptical consumer;*

(ii) *we have a Full Obfuscation solution if, and only if, we have Agreement at the Prior, i.e.* $a^*(\mu_0) = \arg\max_{a_i \in A} v(a_i)$*;*

(iii) *otherwise, we have a Full Transparency solution.*

Proof of Theorem 1. For the proof of part (i), we start by showing sufficiency. Our goal is to find a λ that yields a higher payoff for the firm than either $\lambda = 0$ or $\lambda = \overline{\lambda}$. We focus here on the case in which $v(a_L) > v(a_H)$, or equivalently, $V(0) > V(1)$, as the other case follows a symmetric argument. Then, if we have a skeptical agent, $\mu_{ind} > 1/2$ and $\mu_0 > \mu_{ind}$, which implies that the consumer would choose a_H if they act according to their prior. In other words, $\tilde{V}_{\mu_0}(\overline{\lambda}) = V(1)$.

As usual, $\mathcal{O} = [0, \overline{\lambda}]$, and by triangle similarity, $\tilde{V}_{\mu_0}(0) = V(0)(1-\mu_0) + V(1)\mu_0$. Given our assumption that $V(0) > V(1)$, for any $\lambda \in int(\mathcal{O})$, $\tilde{V}_{\mu_0}(\lambda) > \tilde{V}_{\mu_0}(\overline{\lambda})$. So we must check if there exists a λ in $int(\mathcal{O})$ such that $\tilde{V}_{\mu_0}(\lambda) > \tilde{V}_{\mu_0}(0)$. Indeed, by triangle similarity, we have that for any $\lambda \in int(\mathcal{O})$,

$$\tilde{V}_{\mu_0}(\lambda) = V(0)\frac{\overline{\mu}(\lambda) - \mu_0}{\overline{\mu}(\lambda) - \underline{\mu}(\lambda)} + V(1)\frac{\mu_0 - \underline{\mu}(\lambda)}{\overline{\mu}(\lambda) - \underline{\mu}(\lambda)}.$$

Thus, if we find a λ such that

$$1 - \mu_0 < \frac{\overline{\mu}(\lambda) - \mu_0}{\overline{\mu}(\lambda) - \underline{\mu}(\lambda)}, \tag{18}$$

we are done. Now, note that both $\underline{\mu}$ and $(1-\overline{\mu})$ are positive, continuous and strictly increasing functions with $\lim_{\lambda \to 0^+} \underline{\mu}(\lambda) = \lim_{\lambda \to 0^+} 1 - \overline{\mu}(\lambda) = 0$ and define $q \coloneqq \frac{\mu_0}{1-\mu_0}$, so that Equation

(18) gives

$$1-\mu_0 < \frac{\overline{\mu}(\lambda)-\mu_0}{\overline{\mu}(\lambda)-\underline{\mu}(\lambda)} \iff \underline{\mu}(\lambda) > q(1-\overline{\mu}(\lambda)) \iff 1 > q\frac{(1-\overline{\mu}(\lambda))}{\underline{\mu}(\lambda)}. \tag{19}$$

Since $\mu_{ind} > 1/2$, we have that $D \coloneqq u(a_L,\omega_L) - u(a_H,\omega_L) > d \coloneqq u(a_H,\omega_H) - u(a_L,\omega_H)$, and this together with Equations (5) and (6) yields

$$\lim_{\lambda\to 0^+} \frac{1-\overline{\mu}(\lambda)}{\underline{\mu}(\lambda)} = 0 \implies \lim_{\lambda\to 0^+} \lambda\frac{1-\overline{\mu}(\lambda)}{\underline{\mu}(\lambda)} = 0. \tag{20}$$

To help us understand the mechanics of this limit, note that using the functional form for $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$, the limit becomes

$$\lim_{\lambda\to 0^+} \frac{1-\overline{\mu}(\lambda)}{\underline{\mu}(\lambda)} = \lim_{\lambda\to 0^+} \frac{1-\frac{1-e^{\frac{D}{\lambda}}}{e^{\frac{-d}{\lambda}}-e^{\frac{D}{\lambda}}}}{\frac{1-e^{\frac{-D}{\lambda}}}{e^{\frac{d}{\lambda}}-e^{\frac{-D}{\lambda}}}} = \lim_{\lambda\to 0^+} \frac{e^{\frac{-d}{\lambda}}-1}{e^{\frac{-d}{\lambda}}-e^{\frac{D}{\lambda}}}\frac{e^{\frac{d}{\lambda}}-e^{\frac{-D}{\lambda}}}{1-e^{\frac{-D}{\lambda}}} = 0.$$

The final equality is a direct consequence of $D > d$. Therefore, we deduce the existence of $\lambda \in (0,\overline{\lambda})$ small enough such that Equation 19 is satisfied, and we have a partial obfuscation problem.

We now prove necessity by checking all possible cases we can have with state-independent preferences. Let $\mu_{ind} > \frac{1}{2}$. We then have the following possible cases:

1. $v(a_L) > v(a_H)$:

   (i) $\mu_0 > \mu_{ind}$: this is a skeptical agent case and it is, therefore, a partial obfuscation problem.

   (ii) $\mu_0 < \mu_{ind}$: the optimum for the firm is to choose $\lambda = \overline{\lambda}$ given that $a^*(\mu_0) = a_L$. The problem is thus of boundary obfuscation.

2. $v(a_L) < v(a_H)$:

   (i) $\mu_0 > \mu_{ind}$: the optimum for the firm is to choose $\lambda = \overline{\lambda}$ given that $a^*(\mu_0) = a_H$.

The problem is thus of boundary obfuscation. Note, however, that for analogous reasons to when we have a skeptical agent, $\tilde{V}_{\mu_0}$ is non-monotonic in this case. The difference here is that the λ that induces the lowest payoff is now interior. Indeed, confusing "just a little" might be the worst decision for the firm.

(ii) $\mu_0 < \mu_{ind}$: the optimum for the firm is to choose $\lambda = 0$ as in this case $\tilde{V}_{\mu_0}$ is decreasing in λ, and so the problem is of boundary obfuscation. To see this, note that $\underline{\mu}(\lambda) > 1 - \overline{\mu}(\lambda)$. Building on the construction for the symmetric case proof and using Equation 16, we see that as λ rises we are increasing the proportion of $V(0) = v(a_L, \cdot)$, which in this case we assume is the worst option.

The cases where $\mu_{ind} < \frac{1}{2}$ are analogous. Since we've already shown that when $\mu_{ind} = \frac{1}{2}$ the problem is of boundary obfuscation, the proof of part (i) is complete.

We now prove part (ii). Assume we have agreement at the prior. With a slight abuse of notation, let $\overline{a} = a^*(\mu_0)$ represent the favorite action for the firm, and we have

$$\tilde{V}_{\mu_0}(\lambda) = P_\lambda(a_H)v(a_H) + P_\lambda(a_L)v(a_L) \leqslant v(\overline{a}) \text{ for every } \lambda.$$

Since we have agreement at the prior, $\tilde{V}_{\mu_0}(\overline{\lambda}) = v(\overline{a})$, and $\overline{\lambda}$ solves the firm's problem. Now, assume we have full obfuscation, i.e., $\overline{\lambda}$ is a solution to the firm's problem. As before, assume $v(a_L) > v(a_H)$, or equivalently, that $V(0) > V(1)$ as the other case is analogous. Then, $\{a_L\} = \arg\max_{a \in A} v(a)$. We need to show that $\{a_L\} = a^*(\mu_0) = \arg\max_{a \in A} \mathbb{E}_{\mu_0}[u(a, \omega)]$.

Since we have full obfuscation,

$$\begin{aligned} & \tilde{V}_{\mu_0}(\overline{\lambda}) \geqslant \tilde{V}_{\mu_0}(\lambda) \ \forall \lambda \in \mathbb{R}_+ \\ \implies & V(\mu_0) \geqslant \tilde{V}_{\mu_0}(\lambda) \ \forall \lambda \in \mathbb{R}_+ \\ \implies & V(\mu_0) \geqslant \tilde{V}_{\mu_0}(0) = \mu_0 V(1) + (1 - \mu_0)V(0) \end{aligned} \tag{21}$$

Since the firm has state-independent preferences, $V(\mu) \in \{V(0), V(1)\}$ for all $\mu \in \triangle(\Omega)$,

so $V(\mu_0) \in \{V(0), V(1)\}$. As μ_0 is interior and $V(0) > V(1)$, Equation (21) then implies that $V(\mu_0) = V(0) = v(a_L)$, which in turn implies that $a^*(\mu_0) = \{a_L\}$.

Note that the only other possible case is of full transparency, which gives us part (iii). □

Theorem 2. *Assume we do not have Full Transparency. Then, there exists a number $\underline{\varepsilon} > 0$ and an $i \in \{H, L\}$ such that reducing $u(a_i, \omega_{-i})$ by at least $\underline{\varepsilon}$, induces a full information equilibrium (the firm sets $\lambda = 0$) that achieves u_0^*.*

Proof of Theorem 2. First, assume we have a type (i) Skeptical consumer from Definition 4. The other case is analogous. Now, define $\underline{\varepsilon} = (u(a_H, \omega_H) + u(a_H, \omega_L)) - (u(a_L, \omega_L) + u(a_L, \omega_H))$ and note that since $\mu_{ind} < \frac{1}{2}$, $\underline{\varepsilon} > 0$. If we decrease $u(a_H, \omega_L)$ by any $\varepsilon \geqslant \underline{\varepsilon}$, Theorem 1 tells us that we no longer have a partial obfuscation problem as now $\mu_{ind} \geqslant \frac{1}{2}$. Thus, we only need to check the boundary cases. If the firm sets $\lambda = \overline{\lambda}$, the consumer will choose the firm's least favorite action, a_L, with probability 1 as $\mu_0 < \mu_{ind}$. Setting $\lambda = 0$ yields a strictly better expected payoff for the firm: $\mu_0(\omega_H)v(a_H) + \mu_0(\omega_L)v(a_L)$, so it chooses full transparency. Under full transparency, the consumer's ex-ante expected payoff in this new game is $\mu_0(\omega_H)u(a_H, \omega_H) + \mu_0(\omega_L)u(a_L, \omega_L)$, which is the same payoff as the first-best payoff in the original SO game.

Now, assume we have Agreement at the Prior with $v(a_H) > v(a_L)$. Then, $\mu_{ind} < \mu_0$. Now, take

$$\underline{\varepsilon} = \frac{\mu_0}{1 - \mu_0}[u(a_h) - u(a_L, \omega_H)] - [u(a_L, \omega_L) - u(a_H, \omega_L)] \geqslant 0,$$

and since $\mu_0 > \mu_{ind}$, $\underline{\varepsilon} > 0$. Since the new belief of indifference is higher than μ_0, we do not have either agreement at the prior, not a skeptical consumer, so we get full transparency and the consumer's new equilibrium payoff is u_0^*. The case where $v(a_L) > v(a_H)$ is analogous. □

Proposition 2 *Given a firm's Obfuscation Problem, there exists at most one $\lambda^* \in (0, \overline{\lambda})$ such that $\frac{d\tilde{V}_{\mu_0}(\lambda^*)}{d\lambda} = 0$.*

Proof of Proposition 2. For $\lambda \in (0, \overline{\lambda})$, can rewrite Equation (10) as

$$\tilde{V}_{\mu_0}(\lambda) = \frac{\mu_0 - \overline{\mu}(\lambda)}{\underline{\mu}(\lambda) - \overline{\mu}(\lambda)}[v(a_L) - v(a_H)] + v(a_H).$$

We assume, without loss, that $v(a_H) > v(a_L)$. Considering Equations (5) and (6), we get that

$$\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} = \frac{(v(a_H) - v(a_L))}{4\lambda^2}\Bigg[(-1+\mu_0)(u(a_H,\omega_H)\text{Csch}^2\left(\frac{u(a_H,\omega_H) - u(a_L,\omega_H)}{2\lambda}\right) + \mu_0\left(-u(a_H,\omega_L) + u(a_L,\omega_L)\right)\text{Csch}^2\left(\frac{u(a_H,\omega_L) - u(a_L,\omega_L)}{2\lambda}\right)\Bigg].$$

where Csch denotes the hyperbolic cosecant. Note that the term outside the brackets is always positive and that $\text{Csch}(-x) = -\text{Csch}(x)$. Rearranging the equation, we get that

$$\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} > 0 \iff \text{Csch}^2\left(\frac{u(a_L,\omega_L) - u(a_H,\omega_L)}{2\lambda}\right) > \frac{1-\mu_0}{\mu_0}\frac{u(a_H,\omega_H) - u(a_L,\omega_H)}{u(a_L,\omega_L) - u(a_H,\omega_L)}\text{Csch}^2\left(\frac{u(a_H,\omega_H) - u(a_L,\omega_H)}{2\lambda}\right).$$

In what follows, let

$$G(\lambda) \equiv \frac{\text{Csch}^2\left(\frac{u(a_L,\omega_L) - u(a_H,\omega_L)}{2\lambda}\right)}{\text{Csch}^2\left(\frac{u(a_H,\omega_H) - u(a_L,\omega_H)}{2\lambda}\right)}.$$

We now go case by case.

1. $u(a_H,\omega_H) - u(a_L,\omega_H) = u(a_L,\omega_L) - u(a_H,\omega_L)$ (so $\mu_{ind} = \frac{1}{2}$). Then,

$$\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} > 0 \iff 1 > \frac{1-\mu_0}{\mu_0}.$$

Therefore, if $\mu_0 > \frac{1}{2}$ then $\tilde{V}_{\mu_0}$ is strictly increasing in $[0, \overline{\lambda})$; if $\mu_0 < \frac{1}{2}$ it is strictly

decreasing.[36]

2. $u(a_H,\omega_H)-u(a_L,\omega_H) > u(a_L,\omega_L)-u(a_H,\omega_L)$. In this case, we get that $G(\lambda)$ is strictly decreasing for $\lambda \in \mathbb{R}_+$ and

$$\lim_{\lambda\to 0} G(\lambda) = +\infty, \ \lim_{\lambda\to+\infty} G(\lambda) = 0.$$

Therefore, by the Intermediate Value Theorem, there exists $\lambda^* \in \mathbb{R}_+$ such that $\frac{d\tilde{V}_{\mu_0}(\lambda^*)}{d\lambda} = 0$. Moreover, since $G(\lambda)$ is strictly decreasing, λ^* is unique. Since there exists at most one such $\lambda^* \in \mathbb{R}_+$, then there exists at most one $\lambda^* \in (0,\overline{\lambda}) \subset \mathbb{R}_+$.

3. $u(a_H,\omega_H)-u(a_L,\omega_H) < u(a_L,\omega_L)-u(a_H,\omega_L)$. In this case, we get that $G(\lambda)$ is strictly increasing for $\lambda \in \mathbb{R}_+$ and

$$\lim_{\lambda\to 0} G(\lambda) = 0, \ \lim_{\lambda\to+\infty} G(\lambda) = \frac{(u(a_H,\omega_H)-u(a_L,\omega_H))^2}{(u(a_L,\omega_L)-u(a_H,\omega_L))^2}.$$

Therefore, for λ close enough to 0, $\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} < 0$. We then have two cases: if $\frac{(u(a_H,\omega_H)-u(a_L,\omega_H))}{(u(a_L,\omega_L)-u(a_H,\omega_L))} > \frac{1-\mu_0}{\mu_0}$, then by the Intermediate Value Theorem, there exists λ^* such that $\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} = 0$, and since $G(\lambda)$ is strictly increasing, λ^* is unique. Otherwise, $\frac{d\tilde{V}_{\mu_0}(\lambda)}{d\lambda} < 0$ for all $\lambda \in \mathbb{R}_+$.

We then get that there exists at most one $\lambda^* \in (0,\overline{\lambda})$ such that $\frac{d\tilde{V}_{\mu_0}(\lambda^*)}{d\lambda} = 0$. □

Proposition 3 *The following holds:*

(i) $s^O \leqslant s^P$. *Under State-Independent Preferences for the sender/firm,* $s^O = s^P$ *if, and only if, we have Agreement at the Prior, i.e.,* $a^*(\mu_0) \subseteq \arg\max_{a_i\in A} v(a_i)$.

(ii) *Under State-Independent Preferences for the firm,* $s^C \leqslant s^O$.

[36] Since we assume $\mu_0 \neq \mu_{ind}$, we do not include the case where $\mu_0 = \frac{1}{2}$.

Proof of Proposition 3. We start by proving (i). Note that in Bayesian Persuasion, the firm maximizes its utility by choosing a distribution of posteriors in the set

$$BP(\mu_0) \coloneqq \{\tau \in \triangle(\triangle(\Omega)) : \sum_{supp(\tau)} \mu\tau(\mu) = \mu_0\}.$$

So, given commitment, the Sender is only restricted by Bayes plausibility. Under Strategic Obfuscation, the firm maximizes its utility in the set

$$SO(\mu_0) \coloneqq \{\tau \in \triangle(\triangle(\Omega)) : \exists \lambda \in \mathcal{O} \text{ with } supp(\tau) \subseteq \{\overline{\mu}(\lambda), \underline{\mu}(\lambda)\} \wedge \sum_{supp(\tau)} \mu\tau(\mu) = \mu_0\}.$$

We can clearly see that $SO(\mu_0) \subseteq BP(\mu_0)$, and so the firm in Strategic Obfuscation cannot do better than the Sender in Bayesian Persuasion.

For the second statement of part (i), suppose we have Agreement at the Prior, with $v(a_H) > v(a_L)$. Then, $s^o = \tilde{V}_{\mu_0}(\overline{\lambda}) = V(\mu_0) = v(a_H)$. Under BP, the optimal signal is fully uninformative, so $s^P = V(\mu_0) = v(a_H)$. Now, assume we do not have Agreement at the Prior. Then, $V(\mu_0) = v(a_L)$. Under BP, the optimal split of posteriors would put positive weight on the belief of indifference, which, as we have seen, is not reached under SO for any $\lambda \in \mathbb{R}_+$. Therefore, the concavization result is not reached.

For the proof of (ii), we go back to the case of state-independent preferences for the firm. We repeat the claims in Proposition 2, and for any belief $\mu < \mu_{ind}$, $V(\mu) = V(0)$ and for any belief $\mu > \mu_{ind}$, $V(\mu) = V(1)$. Assume that $V(0) > V(1)$ (the other case follows a similar argument). From Lipnowski and Ravid (2020), we know that an outcome (τ, s) is an equilibrium outcome if, and only if, $\tau \in \triangle(\triangle(\Omega))$ satisfies Bayes plausibility and $s \in \cap_{\mu \in supp(\tau)} V(\mu)$. But, given the state-independent preferences, we have that $\forall F \in 2^{\triangle(\Omega)}$,

the following holds:

$$\begin{cases} \cap_{\mu \in F} V(\mu) & = \{V(0)\} \text{ if } F \subseteq [0, \mu_{ind}], \\ \cap_{\mu \in F} V(\mu) & = \{V(1)\} \text{ if } F \subseteq (\mu_{ind}, 1], \\ \cap_{\mu \in F} V(\mu) & = \emptyset \text{ otherwise.} \end{cases}$$

And so the only possible payoffs for the Sender in a Cheap Talk equilibrium are $V(0)$ or $V(1)$. When adding Bayes Plausibility, we have one of two cases. If $\mu_0 < \mu_{ind}$, the only possible payoff is $V(0)$, and this can be achieved under Strategic Obfuscation by setting $\lambda = \overline{\lambda}$ (since $\mu_0 < \mu_{ind}$, the decision maker will take the optimum action for the firm if no new information is acquired); if $\mu_0 > \mu_{ind}$, $V(1)$ is the only achievable payoff in equilibrium for the Sender, i.e., the consumer chooses the worst action for the Sender with probability 1. But in Strategic Obfuscation the firm can always get a better payoff by setting $\lambda = 0$. Since μ_0 is interior, by fully revealing the state the optimal action for the firm is expected to be chosen with strictly positive probability. □

The next results pertain to the more general case of the model, as described in Section 4. The following Lemmas help understand how the comparative statics of the solution to the consumer's problem change as λ varies.

Lemma 4. *We have that* $\lim_{\lambda \to 0^+} I(\lambda; a_i) = \{\mu \in \triangle(\Omega) : \text{supp}(\mu) \subset \Omega(a_i)\}$.

Proof of Lemma 4. Fix $\tilde{\mu} \in \triangle(\Omega)$ such that $\text{supp}(\tilde{\mu}) \subset \Omega(a_i)$. Then,

$$\sum_{m=1}^{M} \tilde{\mu}(\omega_m) \exp\left(\frac{u(a_j, \omega_m) - u(a_i, \omega_m)}{\lambda}\right) = \sum_{\omega \in \Omega(a_i)} \tilde{\mu}(\omega) \exp\left(\frac{u(a_j, \omega) - u(a_i, \omega)}{\lambda}\right).$$

But, for any $\omega \in \Omega(a_i)$,

$$\exp\left(\frac{u(a_j, \omega_m) - u(a_i, \omega_m)}{\lambda}\right) \leqslant 1 \text{ for any } \lambda > 0.$$

Therefore, $\tilde{\mu} \in I(\lambda, a_i)$ for all $\lambda > 0$, which implies $\tilde{\mu} \in \lim_{\lambda \to 0^+} I(\lambda, a_i)$.

Now, fix $\tilde{\mu} \in \lim_{\lambda \to 0^+} I(\lambda, a_i)$ and assume, by way of contradiction, that $\tilde{\mu} \notin \{\mu \in \triangle(\Omega) : supp(\tilde{\mu}) \subset \Omega(a_i)\}$. Then, there exists $\omega_m \in \Omega$ and $a_j \in A$ such that $\tilde{\mu}(\omega_m) > 0$, and $u(a_j, \omega_m) > u(a_i, \omega_m)$. But then,

$$\lim_{\lambda \to 0^+} \exp\left(\frac{u(a_j, \omega_m) - u(a_i, \omega_m)}{\lambda}\right) > 1. \text{ Contradiction.}$$

□

Lemma 5. *Fix $a_i \in A$ and $\lambda \geqslant 0$. Then, $\mu \in I(\lambda, a_i) \implies a_i \in a^*(\mu)$.*

Proof of Lemma 5. The result when $\lambda = 0$ is immediate, so let $\lambda > 0$. We prove by taking the contrapositive. Let $\mu \in \triangle(\Omega)$ be such that exists action a_j with $\sum_{m=1}^{M} \mu(\omega_m)(u(a_j, \omega_j) - u(a_i, \omega_j)) > 0$, so that $a_i \notin a^*(\mu)$. To ease notation, let

$$g_\lambda(a_i, a_j, \omega_j) \equiv \frac{u(a_j, \omega_j) - u(a_i, \omega_j)}{\lambda}, \text{ and}$$
$$\varphi(x) \equiv e^x, \text{ for } x \in \mathbb{R}.$$

We then have that $\mathbb{E}_\mu[g_\lambda(a_i, a_j, \omega_j)] > 0 \implies \varphi(\mathbb{E}_\mu[g_\lambda(a_i, a_j, \omega_j)]) > 1$. By Jensen's Inequality,

$$\sum_{j=1}^{M} \mu(\omega_j) \exp(g_\lambda(a_i, a_j, \omega_j)) = \mathbb{E}_\mu[\varphi(g_\lambda(a_i, a_j, \omega_j))] \geqslant \varphi(\mathbb{E}_\mu[g_\lambda(a_i, a_j, \omega_j)]) > 1.$$

Therefore, $\mu \notin I(\lambda; a_i)$. □

Lemma 6. *For any prior μ_0 and $\lambda \geqslant 0$, the set of solutions for the consumer's problem is convex.*

Proof of Lemma 6. Fix μ_0, $\lambda \geqslant 0$, and the corresponding induced set of solutions to the consumer's problem. If the set is a singleton we are done. So let τ^1 and τ^2 be different

solutions, and fix $\alpha \in (0,1)$. We define $\tau^3 \in \triangle(\triangle(\Omega))$ as $\tau^3(\mu) = \alpha\tau^1(\mu) + (1-\alpha)\tau^2(\mu)$. Clearly,

$$\mu \in \text{supp}(\tau^3) \iff \tau^3(\mu) > 0 \iff \alpha\tau^1(\mu) + (1-\alpha)\tau^2(\mu) > 0 \iff [\tau^1(\mu) > 0] \vee [\tau^2(\mu) > 0],$$

and therefore $\text{supp}(\tau^3) = \text{supp}(\tau^1) \cup \text{supp}(\tau^2)$. We then get

$$\begin{aligned}
\sum_{a \in A(\tau^3)} \tau^3(\mu^a)\mu^a(\omega) &= \sum_{a \in A(\tau^3)} \left(\alpha\tau^1(\mu^a) + (1-\alpha)\tau^2(\mu^q)\right)\mu^a(\omega) \\
&= \alpha \sum_{a \in A(\tau^1)} \tau^1(\mu^a)\mu^a(\omega) + (1-\alpha) \sum_{a \in A(\tau^2)} \tau^2(\mu^a)\mu^a(\omega) \\
&= \alpha\mu_0(\omega) + (1-\alpha)\mu_0(\omega) \text{ (since } \tau^1 \text{ and } \tau^2 \text{ are solutions)} \\
&= \mu_0(\omega) \text{ for all } \omega.
\end{aligned}$$

Therefore, τ^3 satisfies Bayes' plausibility. Furthermore,

$$\begin{aligned}
&\mathbb{E}_{\tau^3}[U(\mu^a)] - \lambda(H(\mu_0) - \mathbb{E}_{\tau^3}[H(\mu^a)]) \\
&\quad= \sum_{a \in A(\tau^3)} \tau^3(\mu^a)U(\mu^a) - \lambda\left(H(\mu_0) - \sum_{a \in A(\tau^3)} \tau^3(\mu^a)H(\mu^a)\right) \\
&\quad= \alpha\left(\sum_{a \in A(\tau^1)} \tau^1(\mu^a)U(\mu^a) - \lambda\left(H(\mu_0) - \sum_{a \in A(\tau^1)} \tau^1(\mu^a)H(\mu^a)\right)\right) \\
&\qquad + (1-\alpha)\left(\sum_{a \in A(\tau^2)} \tau^2(\mu^a)U(\mu^a) - \lambda\left(H(\mu_0) - \sum_{a \in A(\tau^2)} \tau^2(\mu^a)H(\mu^a)\right)\right) \\
&\quad= \alpha\left(\mathbb{E}_{\tau^1}[U(\mu^a)] - \lambda(H(\mu_0) - \mathbb{E}_{\tau^1}[H(\mu^a)])\right) \\
&\qquad + (1-\alpha)\left(\mathbb{E}_{\tau^2}[U(\mu^a)] - \lambda(H(\mu_0) - \mathbb{E}_{\tau^2}[H(\mu^a)])\right) \\
&\quad= \mathbb{E}_{\tau^1}[U(\mu^a)] - \lambda(H(\mu_0) - \mathbb{E}_{\tau^1}[H(\mu^a)]).
\end{aligned}$$

The last equality follows from the fact that both τ^1 and τ^2 are solutions. Therefore, τ^3 is

also a solution to the consumer's problem. □

Corollary 2. *Let τ^1 and τ^2 solve the consumer's Problem. Then, Equation* (11) *holds for all $a_i, a_j \in A(\tau^1) \cup A(\tau^2)$.*

Proof of Corollary 2. Fix $\alpha \in (0,1)$. By Lemma 6, τ^3 defined by $\tau^3(\mu) = \alpha\tau^1(\mu)+(1-\alpha)\tau^2(\mu)$ for all μ is also a solution. Necessity implies that Equation (11) holds for any $a_i, a_j \in A(\tau^3) = A(\tau^1) \cup A(\tau^2)$. □

Proposition 4. *We have that $\delta_{\mu_0} \in T_{\mu_0}(\lambda)$ if and only if $\mu_0 \in I(\lambda; a_i)$ for some $a_i \in A$. Moreover, if $\mu_0 \in int(I(\lambda; a_i))$, then $T_{\mu_0}(\lambda) = \{\delta_{\mu_0}\}$.*

Proof of Lemma 4. That $\delta_{\mu_0} \in T_{\mu_0}(\lambda)$, or equivalently, that $\tau_\lambda(\mu_0) = 1$ is a solution if and only if $\mu_0 \in I(\lambda; a_i)$ is a direct application of Lemma 3: by construction $\sum_{\mu^a \text{supp } \tau_\lambda} \tau_\lambda(\mu_0)\mu^a(\omega) = \mu_0(\omega)$; then, for the necessary part, if $\mu_0 \notin \cup_i I(\lambda; a_i)$, it cannot satisfy (12); for the sufficiency, $\mu_0 \in I(\lambda; a_i)$ implies that Equation (12) is satisfied, and $\tau_\lambda(\mu_0) = 1$ implies that Equation (11) is vacuously satisfied as $A(\tau_\lambda) = \{a_i\}$. Lemma 5 shows that a_i is the recommended action for the uninformative signal, and a_i is chosen with probability 1.

The rest of the proof shows that the solution is unique whenever $\mu_0 \in int(I(a, \lambda)$, so suppose this statement holds. Since μ_0 is assumed to be interior, $\lambda > 0$. Now, suppose towards a contradiction, that there exists another solution $\tau \in \triangle(\Omega)$ such that $\tau(\mu_0) < 1$. Therefore, there must be another action $a_j \neq a_i$ that is recommended with positive probability ($a_j \in A(\tau)$) and induces posterior μ^{a_j}. By Corollary 2, we must have

$$\frac{\mu_0(\omega)}{\exp\left(\frac{u(a_i,\omega)}{\lambda}\right)} = \frac{\mu^{a_j}(\omega)}{\exp\left(\frac{u(a_j,\omega)}{\lambda}\right)}, \text{ for all } \omega. \tag{22}$$

But $\mu_0 \in int(I(\lambda; a_i))$, and thus

$$\sum_{m=1}^{M} \mu_0(\omega_m) \frac{\exp(u(a_j, \omega_m)/\lambda)}{\exp(u(a_i, \omega_m)/\lambda)} < 1.$$

Substituting according to Equation (22), gives us

$$\sum_{m=1}^{M} \mu^{a_j}(\omega_m) \frac{\exp(u(a_i,\omega_m)/\lambda)}{\exp(u(a_j,\omega_m)/\lambda)} \frac{\exp(u(a_j,\omega_m)/\lambda)}{\exp(u(a_i,\omega_m)/\lambda)} < 1,$$

$$\implies \sum_{m=1}^{M} \mu^{a_j}(\omega_m) < 1. \text{ Contradiction.}$$

□

Proposition 5. *For any $\mu \in int(\triangle(\Omega))$, the following statements are equivalent:*

(i) $a^*(\mu)$ is a singleton;

(ii) there exists $\overline{\lambda} \in \mathbb{R}_{++}$ s.t. $\mu \in I(a_i, \overline{\lambda})$ for some $a_i \in C_\mu(A)$.

Proof of Proposition 5. Fix $\mu \in int(\triangle(\Omega))$. We begin by showing that $(ii) \implies (i)$. Assume $a^*(\mu)$ is a singleton, and let $i \in \{1, \ldots, N\}$ be such that $a^*(\mu) = \{a_i\}$. Therefore,

$$\sum_{m=1}^{M} \mu(\omega_m)[u(a_j,\omega_m) - u(a_i,\omega_m)] < 0 \text{ for all } j \neq i. \tag{23}$$

For the rest of the proof, we fix $j \neq i$ and let $z(m) = u(a_j,\omega_m) - u(a_i,\omega_m)$. Now, let $\overline{m} = \arg\min_m z(m)$. We can rewrite Equation (23) as

$$\sum_{m \neq \overline{m}} \mu(\omega_m) z(m) + \mu(\omega_{\overline{m}}) z(\overline{m}) < 0 \iff \sum_{m \neq \overline{m}} \mu(\omega_m) z(m) + \left(1 - \sum_{m \neq \overline{m}} \mu(\omega_m)\right) z(\overline{m}) < 0$$

$$\iff \sum_{m \neq \overline{m}} \mu(\omega_m)[z(m) - z(\overline{m})] < -z(\overline{m}).$$

Take $\tilde{m} \neq \overline{m}$ such that $z(\tilde{m}) > z(\overline{m})$.[37] We then get

$$\sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m)[z(m) - z(\overline{m})] + \mu(\omega_{\tilde{m}})[z(\tilde{m}) - z(\overline{m})] < -z(\overline{m})$$

$$\iff \mu(\omega_{\tilde{m}}) < \underbrace{\frac{-z(\overline{m})}{[z(\tilde{m}) - z(\overline{m})]}}_{T_0} - \sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m) \underbrace{\frac{[z(m) - z(\overline{m})]}{[z(\tilde{m}) - z(\overline{m})]}}_{T_m}. \tag{24}$$

Now take a belief $\gamma \in \triangle(\Omega)$. We then have that, for any $\lambda \in \mathbb{R}_{++}$,

$$\gamma \in I(\lambda; a_i) \iff \sum_{m=1}^{M} \gamma(\omega_m) \exp(z(m)/\lambda) \leqslant 1$$

$$\iff \sum_{m \neq \overline{m}} \gamma(\omega_m) \exp(z(m)/\lambda) + \left(1 - \sum_{m \neq \overline{m}} \gamma(\omega_m)\right) \exp(z(\overline{m})/\lambda) \leqslant 1$$

$$\iff \sum_{m \neq \overline{m}} \gamma(\omega_m)[\exp(z(m)/\lambda) - \exp(z(\overline{m})/\lambda)] \leqslant 1 - \exp(z(\overline{m})/\lambda)$$

$$\iff \sum_{m \neq \tilde{m}, \overline{m}} [\exp(z(m)/\lambda) - \exp(z(\overline{m})/\lambda)] + \gamma(\omega_{\tilde{m}})[\exp(z(\tilde{m})/\lambda) - \exp(z(\overline{m})/\lambda)] \leqslant 1 - \exp(z(\overline{m})/\lambda)$$

$$\iff \gamma(\omega_{\tilde{m}}) \leqslant \underbrace{\frac{1 - \exp(z(\overline{m})/\lambda)}{[\exp(z(\tilde{m})/\lambda) - \exp(z(\overline{m})/\lambda)]}}_{\Gamma_0} - \sum_{m \neq \tilde{m}, \overline{m}} \gamma(\omega_m) \underbrace{\frac{[\exp(z(m)/\lambda) - \exp(z(\overline{m})/\lambda)]}{[\exp(z(\tilde{m})/\lambda) - \exp(z(\overline{m})/\lambda)]}}_{\Gamma_m}.$$

But, note that

$$\frac{1 - \exp(z(\overline{m})/\lambda)}{[\exp(z(\tilde{m})/\lambda) - \exp(z(\overline{m})/\lambda)]} \xrightarrow{\lambda \to \infty} \frac{-z(\overline{m})}{[z(\tilde{m}) - z(\overline{m})]},$$

and, for any $m \neq \tilde{m}, \overline{m}$ we get,

$$\frac{[\exp(z(m)/\lambda) - \exp(z(\overline{m})/\lambda)]}{[\exp(z(\tilde{m})/\lambda) - \exp(z(\overline{m})/\lambda)]} \xrightarrow{\lambda \to \infty} \frac{[z(m) - z(\overline{m})]}{[z(\tilde{m}) - z(\overline{m})]}.$$

[37] We know such $\tilde{m}$ exists since $\overline{m}$ minimizes $z(\cdot)$ and we cannot have $z(m) = z(\overline{m})$ for all m, otherwise a_j would be strictly dominated by a_i in all states, which is ruled out by assumption.

We then deduce the existence of $\overline{\lambda}$ large enough such that

$$|T_0 - \Gamma_0| < \frac{\delta}{M-1},$$

and, for $m \neq \tilde{m}, \overline{m}$

$$|T_m - \Gamma_m| < \frac{\delta}{M-1}.$$

Going back to Equation (24), there exists $\delta > 0$ such that $\mu(\omega_{\tilde{m}}) + \delta < T_0 + \sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m) T_m$. We then get

$$\begin{aligned}
\mu(\omega_{\tilde{m}}) &< T_0 - \frac{\delta}{M-1} - \sum_{m \neq \tilde{m}, \overline{m}} \left[\mu(\omega_m) T_m + \frac{\delta}{M-1} \right] \\
&= T_0 - \frac{\delta}{M-1} - \sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m) \left[T_m + \frac{\delta}{M-1} \right] - \sum_{m \neq \tilde{m}, \overline{m}} (1 - \mu(\omega_m)) \frac{\delta}{M-1} \\
&\leqslant \left(T_0 - \frac{\delta}{M-1} \right) - \sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m) \left[T_m + \frac{\delta}{M-1} \right] \\
&\leqslant \Gamma_0 - \sum_{m \neq \tilde{m}, \overline{m}} \mu(\omega_m) \Gamma_m.
\end{aligned}$$

Therefore, $\mu \in I(a_i, \overline{\lambda})$.

We show that $(i) \implies (ii)$ by taking the contrapositive. Assume $a^*(\mu)$ is not a singleton, so that for two different actions a_i, a_j we get

$$\mathbb{E}_\mu[u(a_j, \omega) - u(a_i, \omega)] = 0 \implies \mathbb{E}_\mu \left[\frac{u(a_j, \omega) - u(a_i, \omega)}{\lambda} \right] = 0 \ \ \forall \lambda > 0.$$

We again denote by $\varphi : \mathbb{R} \to \mathbb{R}_+$ the exponential function: $\varphi(x) = e^x$ for any $x \in \mathbb{R}$. Then, by strict Jensen's inequality

$$\mathbb{E} \left[\varphi \left(\frac{u(a_j, \omega) - u(a_i, \omega)}{\lambda} \right) \right] > \varphi \left(\mathbb{E}_\mu \left[\frac{u(a_j, \omega) - u(a_i, \omega)}{\lambda} \right] \right) = \varphi(0) = 1 \ \ \forall \lambda > 0.$$

Therefore, $\mu \notin I(a, \lambda)$ for any $a \in a^*(\mu)$ and $\lambda > 0$. By Lemma 5, $\mu \notin I(a_k, \lambda)$ for any $a_k \in A \setminus a^*(\mu)$. □

Proposition 6 *The firm's Obfuscation Problem has a solution.*

Proof of Proposition 6. Let $T(\lambda)$ denote the set of solutions for the consumer's problem at λ. For any $\lambda \geqslant \overline{\lambda} \equiv \inf\{\lambda \in \mathbb{R}_{++} : \mu_0 \in int(I(\lambda; a_i)) \text{ for some } a_i \in A\}$ the consumer's unique solution is buying a non-informative signal, so $\tilde{V}_{\mu_0}(\lambda) = V(\mu_0)$ and is constant. If $T(\overline{\lambda})$ is not a singleton, breaking ties in favor of the Principle ensures upper-semicontinuity at this point. The rest of the proof assumes $\lambda < \overline{\lambda}$.

By Corollary (2), we can consider a solution τ_λ such that $A(\tau_\lambda) = \bigcup_{\tau \in T(\lambda)} A(\tau)$. Fix $a_i \in A(\tau_\lambda)$, let $z^i_{jm}(\lambda) \coloneqq exp(\frac{u(a_j,\omega_m)-u(a_i,\omega_m)}{\lambda})$, and we can write the inequalities that a posterior μ^{a_i} must satisfy as

$$\sum_{m=1}^{M} \mu^{a_i}(\omega_m) z^i_{jm}(\lambda) \leqslant 1 \text{ for all } j \neq i, \tag{25}$$

with equality for any j such that $a_j \in A(\tau_\lambda)$. Now, let us define an $N \times M$ matrix $\Psi^i(\lambda)$ by

$$\Psi^i_{jm}(\lambda) = z^i_{jm}(\lambda).$$

We can rewrite the inequalities described by Equation (25) in matrix form as $\Psi^i(\lambda)\mu^{a_i} \leqslant e$, where the $N \times 1$ column vector e has all entries 1. Since τ_λ is a solution to the consumer's problem, it must be that $K \geqslant 2$ equations from $\Psi^i(\lambda)\mu^{a_i} \leqslant e$ bind.[38] We then define a $K \times M$ matrix $\Phi^i(\lambda)$ formed by the K rows of $\Psi^i(\lambda)$ in which the equations hold with equality and let $\tilde{e}$ be a $K \times 1$ vector of 1s.

Now, fix an interval $[\lambda_1, \lambda_2] \subseteq [0, \overline{\lambda}]$ such that $A(\tau_\lambda)$ does not change for any $\lambda \in [\lambda_1, \lambda_2]$. If $K = M$, we can find the corresponding solution by setting $\mu^{a_i} = \Phi^i(\lambda)^{-1}\tilde{e}$, which is a continuous function of λ. Bayes' plausibility then ensures the distribution of posteriors also

[38] Since $\lambda < \overline{\lambda}$, at least 2 actions must be played with positive probability.

varies continuously, and therefore $\tilde{V}_{\mu_0}$ is continuous.

If $K < M$ then there may be infinite solutions for the system of equations $\Phi^i(\lambda)\mu^{a_i} = \tilde{e}$. We have two cases: (i) if $M \geqslant K(M-K)$, the requirement of Bayes Plausibility[39] pins down the possible posteriors and we have a unique solution for the consumer's problem. The system of equations that bind do not change in this interval, and a similar argument as for the $K = M$ case follows; (ii) if $M < K(M-K)$ there may be multiple solutions to the consumer's problem. In that case, we call upon the assumption that the consumer chooses the optimal choice for the principal. This pins down τ_λ, and as the set of actions taken with positive probability does not change in the interval, the binding equations do not change, and the pinned down solution must change continuously with λ. If $K > M$, we again have multiple solutions for the consumer's problem. We turn once more to the assumption that ties are broken in favor of the firm. This pins down a solution, which changes continuously with λ as long as $A(\tau_\lambda)$ does not change.

Therefore, discontinuities of $\tilde{V}_{\mu_0}$ must occur at values of λ^* in which $A({\tau_\lambda}^*)$ changes, or in other words, $\lim_{\lambda \to \lambda^{*-}} A(\tau_\lambda) \neq \lim_{\lambda \to \lambda^{*+}} A(\tau_\lambda)$. The consumer's problem necessarily has multiple solutions in those cases, and breaking ties in favor of the principle ensures that $\tilde{V}_{\mu_0}$ is upper semicontinuous at λ^*.

Since we assume that $a^*(\mu_0)$ is a singleton, by Proposition 5 we know that $\overline{\lambda}$ is finite so $\mathcal{O}$ is a compact set. Therefore $\tilde{V}_{\mu_0}$ attains a maximum in $\mathcal{O}$. □

B Applications

This section provides two applications of our model, but note that many other economically relevant contexts - insurance, politics, etc. - can be inserted in the framework. The first example is the previously mentioned good/bad product case. The second example illustrates that partial obfuscation can appear even under state-dependent preferences.

[39]The prior must be in the convex hull of the support of the solution τ_λ.

B.1 A Good or Bad Product

Producer, Inc. wants to sell its product to the Customer, who may choose between two actions: *buy* or *not buy*. The quality of the product can be either *Good* or *Bad* when compared to another similar product on the market (made by a different firm) that the Customer is already familiar with. The common prior that the product is *Bad* is given by $\mu_0 \in (0,1)$. In accordance with our assumption, the prior is interior for both players because the Customer is not acquainted with all the features of the product, and the Producer does not have all the information on the Customer's preferences yet. The prior then represents the fraction of consumers in society that prefer the competitor's products. One can also consider the case of a reseller who was not present in the production process and so is unsure of the product's quality for the final buyer.

The Customer may conduct a private investigation to obtain more information on the quality of the product: cost-benefit analysis, price comparison, search for online reviews, etc. The Producer may facilitate or hinder the quest for information by selling the product in different sizes or batches than competitors, putting the specifications in fine print, or not including the price for add-ons.[40]

We normalize the payoff to zero for both players if the Customer chooses *not buy*. If the Customer chooses *buy*, the Customer's payoff will be based on how much better the Product is compared to the competitors'. If it is *Good*, the benefit is given by $b > 0$, and if it is *Bad*, the payoff is given by $-c$, where $c > 0$. The Producer gets a benefit of $d > 0$ independently of the state of the world. The payoffs are presented in Table 3, where rows denote the state of the world, not actions.

For a more concrete illustration, we set $\mu_0 = 0.75$, $d = b = 2$ and $c = 1$, which gives $\mu_{ind} \approx 0.667$. We then get the payoffs in Table 4. Note that we have a certain asymmetry considering the payoffs for when the Customer chooses *buy*. This can be interpreted as either

[40]Kalayci and Potters (2011) describes the ordeal of buying a mobile phone. With more than 30 listed technical attributes (such as weight, memory size, and battery capacity), comparing options becomes an arduous task.

Table 3: Payoffs table for the Good or Bad Product game.

		Customer	
		not buy	buy
Producer	Good	$0, 0$	d, b
	Bad	$0, 0$	$d, -c$

the Customer not getting very angry when buying a worse product, or as there being in place a recycling policy that returns some benefit for the Customer if they choose to recycle the bad product.

Table 4: Payoffs table with numerical values.

		Customer	
		not buy	buy
Producer	Good	$0, 0$	2, 2
	Bad	$0, 0$	$2, -1$

Now, to solve the problem, we first plot U, the indirect utility for the Customer as a function of the beliefs, as seen in Figure 10. The horizontal axis denotes the probability that the state is *Bad*. On the right of Figure 10, it is shown the split of posteriors when $\lambda = 0$. At this value of λ the Customer will acquire a fully informative signal, and thus, the posterior will be either 1 or 0. If, on the other hand, the cost is too high, the Customer will purchase a non-informative signal and will act according to the prior.

Since the beliefs are restricted by Bayes' Plausibility, we can obtain the expected utilities by looking at where the vertical line of the prior intersects with the posteriors induced by the signal. When $\lambda = 0$, have full information and the expected payoff is 0.5. When λ is "too high", i.e., $\lambda > \overline{\lambda}$, we have full obfuscation and the expected payoff is 0.

Now that we know how the Customer will act at each belief, Figure 11 shows the expected utility V for the Producer. Again, on the right, we can find the expected payoffs for a given signal at the point where the vertical line of the prior intersects the induced split of posteriors. Under full information, the expected payoff for the Producer is 0.5; under no new information, the expected payoff is 0. Figure 12 plots the expected utility for the Customer on the left and the cost of information at each belief on the right (remember that $\lambda H(\mu_0)$ is a constant

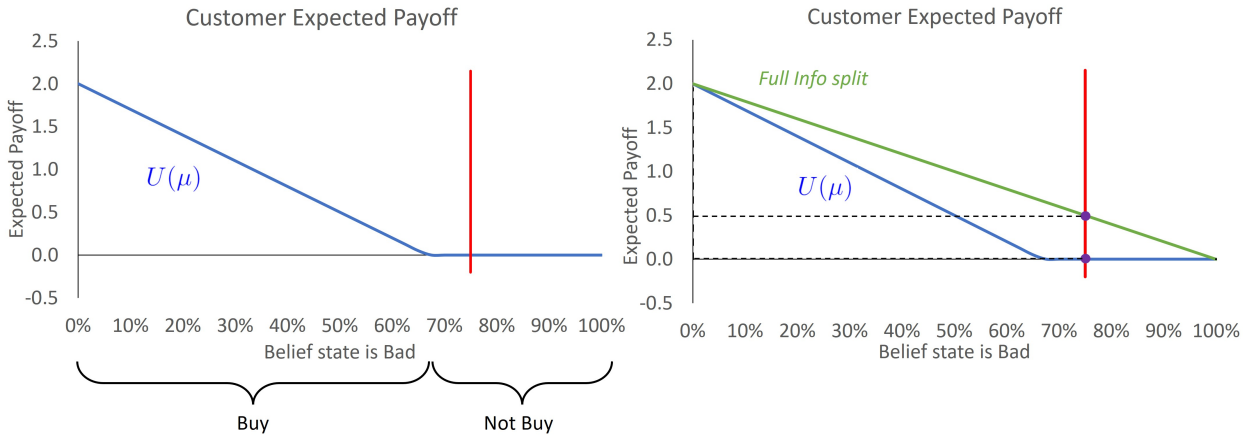

Figure 10: Payoffs as a function of beliefs for the Customer. The vertical line represents the prior. Left: Curly brackets indicate the optimal action at that belief. Right: the diagonal green line indicates the split of posteriors and resulting expected payoffs when $\lambda = 0$.

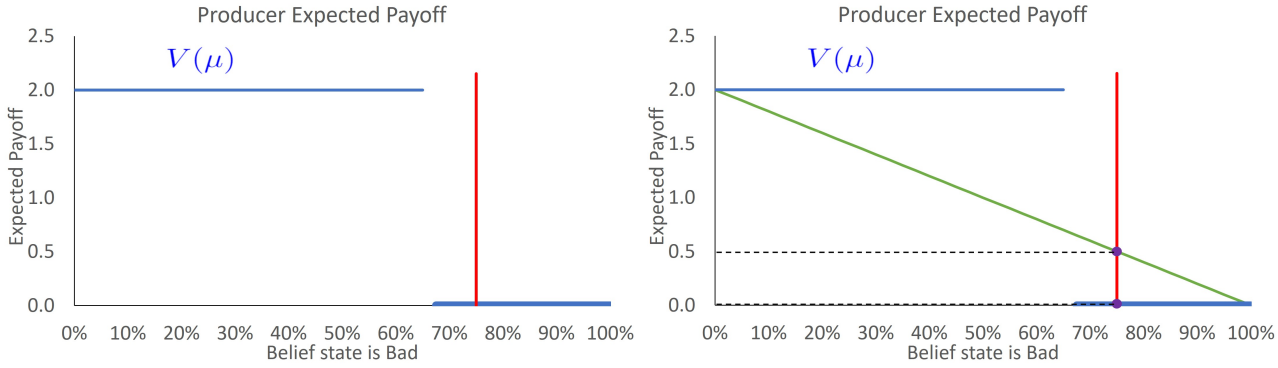

Figure 11: Payoffs as a function of λ for the Producer. Right: the diagonal green line indicates the split of posteriors for $\lambda = 0$.

from the point of view of the Consumer). The sum of these two functions results in the final objective function that the Customer will maximize.

We then can plot the function the Customer will maximize when $\lambda = 1$ as shown in Figure 13. The Costumer chooses a distribution of posteriors with support at $\underline{\mu}(1) = 0.335$ and $\overline{\mu}(1) = 0.91$. Since the prior is between these two values, the Customer will learn by purchasing a signal that, when the realization is *buy*, updates her beliefs to $\underline{\mu}(1)$, and when the realization is *not buy*, updates her beliefs to $\overline{\mu}(1)$, so her expected payoff is 0.11. Had the prior fallen outside this interval, the Customer would have chosen a fully uninformative signal and acted according to the prior. Figure 13 also shows that $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ indicate

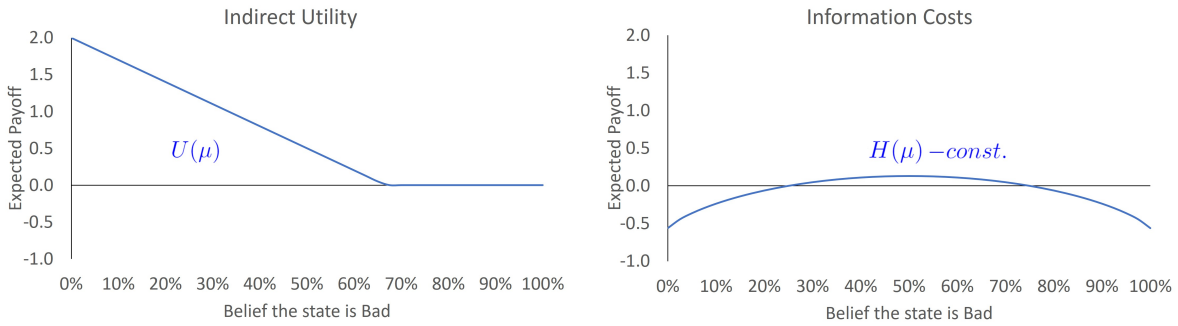

Figure 12: The Consumer's objective function is given by the sum of the indirect utility U (plotted on the left) and the entropy learning costs (plotted on the right). We set $\lambda = 1$ and $\mu_0 = 0.75$.

where the concavication detaches from the Customer's objective function.

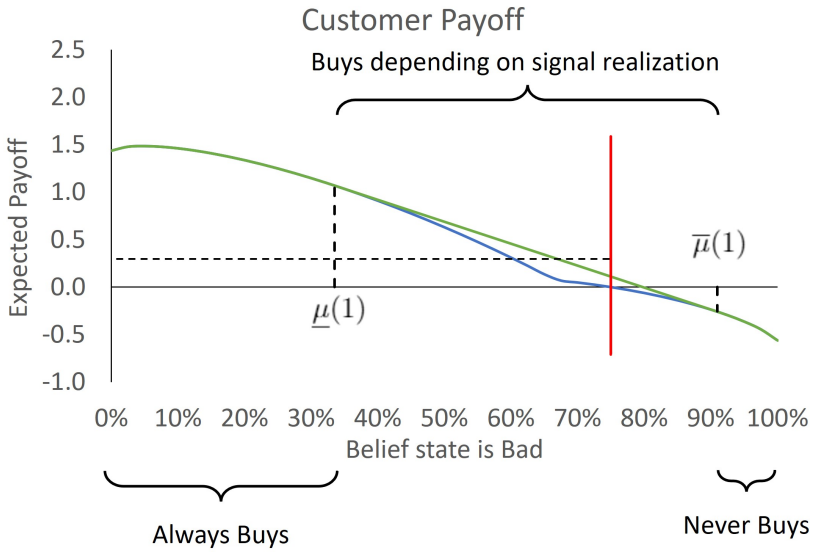

Figure 13: The maximizing function and its concavication for the Customer for $\lambda = 1$ and $\mu_0 = 0.75$.

As an illustration of how λ affects the maximizing function and the solution, Figure 14 shows on the left the maximizing function for $\lambda = 0.5$, in which the solution yields a payoff of 0.243 and on the right the maximizing function for $\lambda = 2$, which the solution yields a payoff of 0.02. It is visible that on the right, the concavity of the Entropy has a higher impact on the final function. We can also see that as λ increases, $\underline{\mu}(\lambda)$ and $\overline{\mu}(\lambda)$ approach μ_{ind}.

Now we return to the Producer and try to understand how intermediate values of λ affect

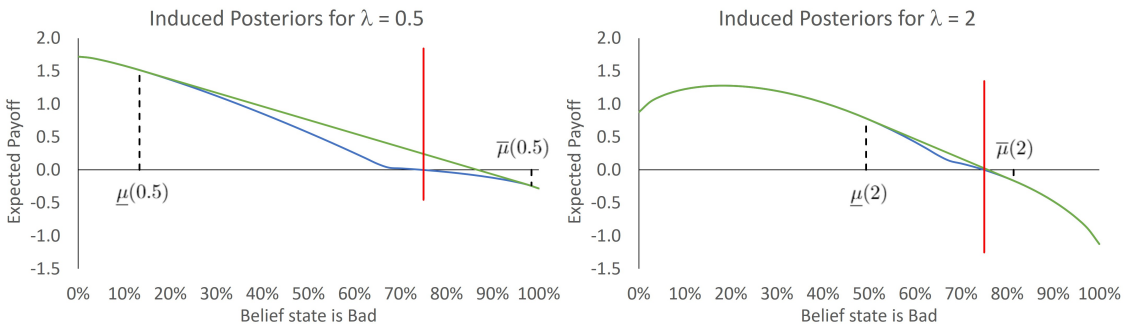

Figure 14: Customer's maximizing function and solution. Left: Case for $\lambda = 0.5$. Right: Case for $\lambda = 2$.

its payoffs. Figure 15 plots the expected payoff for the case in which $\lambda = 1$. The payoff now is 0.556, higher than both the full information and no new information cases.

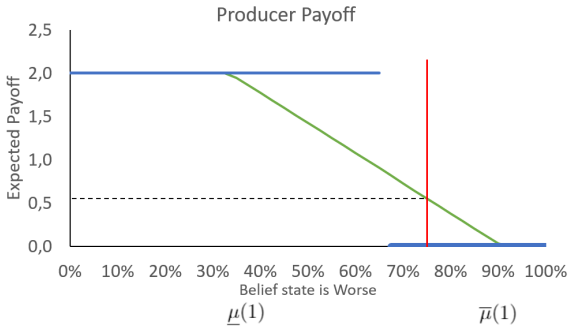

Figure 15: Producer's indirect utility and the chosen split for $\lambda = 1$.

In fact, since we know the functional form for $\tilde{V}_{\mu_0}$ from Equation (10), we can plot it for the different values of λ. We do this in Figure 16. Our numerical solution indicates that, at the optimal level of obfuscation ($\lambda^* \approx 0.75$), the Producer should try to engage in partial obfuscation. Therefore, the Producer may divulge all product specifications, but amongst irrelevant information that the Customer has to parcel through. In equilibrium, the Customer gets an expected utility of 0.162 while the Producer gets an expected utility of 0.556.

Now, let's see what happens if we tweak the parameters slightly, setting $c = 2$. The new payoff table is shown in Table 5, and we can see that the Customer has symmetric preferences, i.e., $\mu_{ind} = 0.5$.

The steps to the solution in this new case are the same. We already know the behavior of

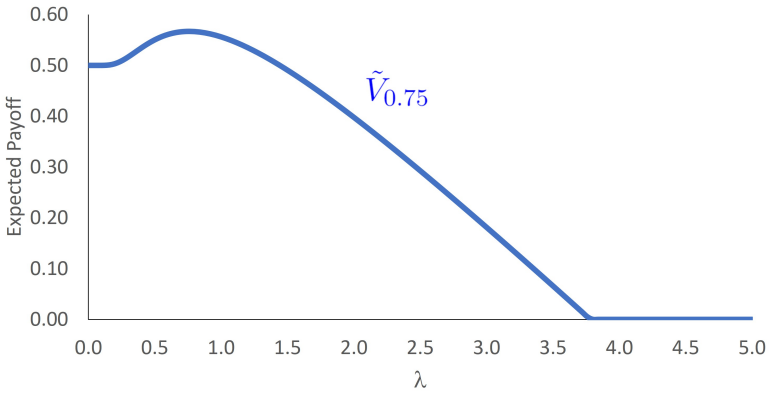

Figure 16: Producer's indirect utility from λ.

Table 5: Payoffs table with new value of c.

		Customer	
		not buy	buy
Producer	Good	$0, 0$	$2, 2$
	Bad	$0, 0$	$2, -2$

the Customer as λ increases, and we have the more general functional form of $\tilde{V}_{\mu_0}$ as given by Equation 10. Hence, we can plot the indirect utility for the Producer in this new context. Figure 17 shows the result. Now, the indirect utility is actually decreasing, and the optimal choice for the Producer is to engage in no obfuscation ($\lambda = 0$), which clearly makes the Customer better-off.

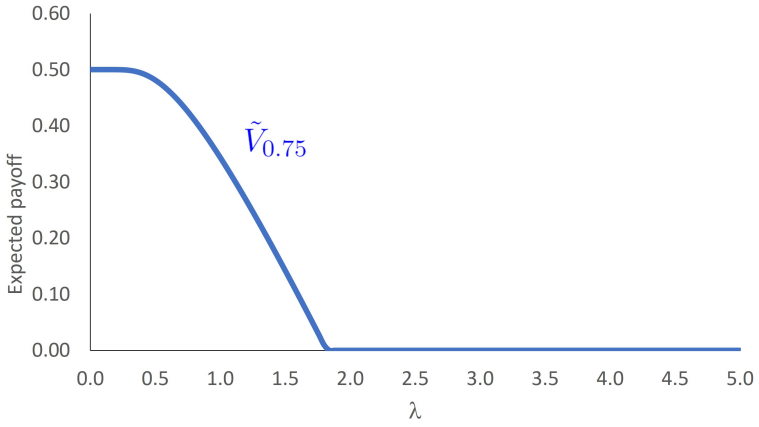

Figure 17: Producer's indirect utility from λ for the new case in which $c = 2$.

Indeed, let λ^* denote the equilibrium value of λ, and we have the following expected

payoffs in equilibrium:

(i) with $u(buy, Bad) = -1$, $\lambda^* \approx 0.75$, the Producer has a payoff of 0.556 and the Customer has a payoff of 0.162;

(ii) with $u(buy, Bad) = -2$, $\lambda^* = 0$, and the Producer has a payoff of 0.50, while the Customer has a payoff of 0.50.

When disregarding the strategic obfuscating behavior, it would seem that revoking the recycling policy would only make the Customer worse. We effectively reduce the agent's utility in a given state with no apparent trade-off. However, this induces full transparency and, consequently, the Customer is actually better off in equilibrium. Another way of interpreting this is that if the consumer could commit to "hurting" themselves if they are in the *Bad* state, they could also induce the firm to choose full disclosure.

In this example, the mechanism behind this phenomenon is the asymmetry in the Customer's payoffs and how it affects the dynamics of the problem. It implies that the Customer's response to higher information costs will not be symmetrical when gathering information. When we set $c = 1$, we make the Customer less averse to buying. As a result, the Customer is more inclined to purchase a signal structure that recommends *buy* more often, even if it leads to more mistakes. While λ is still low, the Customer has enough leeway to choose this sort of signal. The Producer takes advantage of this fact and induces the Customer to *buy* more. As λ increases, however, this kind of signal becomes impracticable because, to recommend *buy* more and more frequently, the signal must also recommend *not buy* less and less, which requires accuracy, and that's expensive. At a certain point, the costs outweigh the benefits, and the Customer will be inclined to follow the prior, in which *not buy* is the better option. By setting $c = 2$, we are effectively making the Customer more "afraid of mistakes", and she will be less inclined to choose a signal biased towards recommending *buy*: there is no space for exploiting the Customer.[41] The best course of action for the Producer is

[41]This is an analogous result to the one in Martin (2017), albeit not in a pricing context, in which an increase in search costs simply induces the buyer to take the outside option, which is bad for the seller.

then full transparency to at least induce *buy* when the state is actually *Good*.

Finally, we also point out that depending on the parameters of the model, we can have a boundary obfuscation problem. For instance, if the prior was 0.25 (instead of 0.75), Theorem 1 says that the best result is total obfuscation, or $\lambda = \overline{\lambda}$. What changed? The idea is that at the new prior of 0.25, the default action (i.e. if the consumer plays according to the prior) is *buy*. Thus, the firm (the Producer) does not want the consumer to learn anything new and chooses $\lambda = \overline{\lambda}$.

B.2 Incumbent vs. Entrant

We now provide a 2 x 2 case obfuscation case with state-dependent preferences for both parties that still results in a partial obfuscation problem. In the context of this example, the economic agent that we have so far called a consumer is now an entrant firm.[42]

Company Monopolistic Enterprises (MON - the incumbent) is a monopolist. Realizing the potential for profits, Competition Group (COM - the entrant) is deciding on whether they should make a play to enter this market and compete with MON. As a result, COM has to choose between two possible actions: *compete* and *not compete*. The success of COM's endeavor hinges on how strong MON's hold on the market is. We can interpret this as product quality, customer loyalty, capacity for further investment, etc. We denote by *Stronger* the state of the world in which MON has a good hold on the market and COM's efforts will be unsuccessful; and by *Weaker* the state of the world in which MON does not have a good hold, and COM will be able to completely kick it off the market. The common prior[43] that the state is *Weaker* is 0.7.

The payoff of being a monopolist in this market is given by $b > 0$, while the payoff of not entering the market is normalized to 0. Furthermore, we assume that COM has to pay a cost

[42]We agree that the term consumer is not the most general when thinking of different applications. We used it, however, because it was the more natural term for the simplest application that comes to mind: that of a firm selling a product that may be good or bad to a consumer.

[43]MON may not know the state because it depends on the quality of the products that COM will bring to the market, which is unknown. Alternatively, we can interpret this as MON committing to an obfuscation strategy.

of $d_C > 0$ to enter the market(for instance, investing in marketing, stores, hiring employees, etc.). To deal with the extra competition, MON incurs a cost $0 < d_M < d_C$ if COM decides to enter.

COM would like to gather information on whether MON is indeed able to fight back or not. Being rationally inattentive, COM must bear a cost to do so. MON can affect the cost of obtaining information: how much they disclose in quarterly results, consumer reports, or if they are soon to receive new investments, etc.

For a more concrete example, in Table 6 we present the payoff table for the case in which we set $b = 1$, $d_C = 0.5$, and $d_M = 0.1$. Under these values, $\mu_{ind} = 0.5$.

Table 6: Payoffs table for Incumbent vs. Entrant.

		Competition Group	
		Not Compete	Compete
Monopolistic Enterprises	Stronger	1, 0	0.9, -0.5
	Weaker	1, 0	-0.1, 0.5

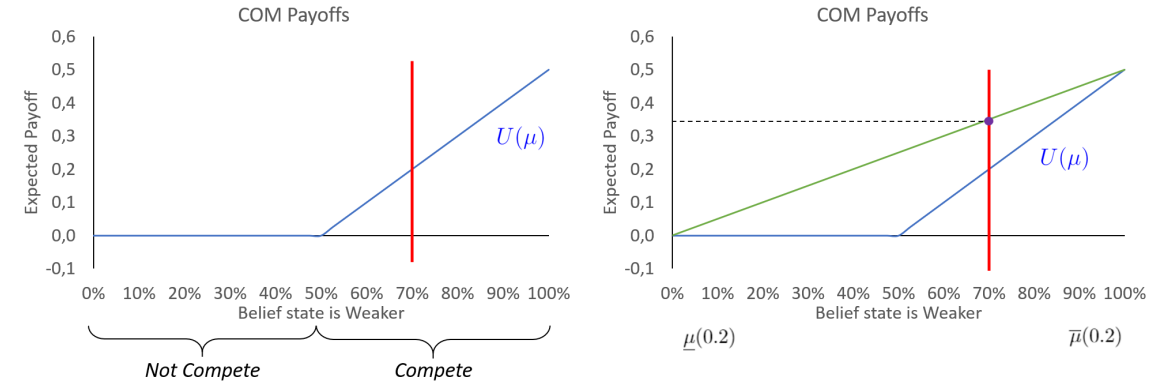

Figure 18: Payoffs as a function of beliefs for COM. The vertical red lines indicate the prior. Left: Curly brackets indicate the optimal action at that belief. Right: the split of posteriors for $\lambda = 0$.

For the payoff structure in Table 6, Figure 18 shows the expected payoffs for COM as a function of the beliefs. The horizontal axis represents the belief that the state is *Weaker*. Since the agent chooses the action that maximizes the expected payoff at each belief, we

were able to plot the indirect utility U for COM. On the right side of Figure 18 we also plot the split of posteriors under full information ($\lambda = 0$) and it is straightforward that in this case the posterior will be either at 0 or at 1 (COM will try to know the state of the world with certainty). Given Bayes' Plausibility, we can find the payoff for full information at the intersection of the line connecting the value of U at the extremes and the vertical line at the prior. In the case of no new information, the posterior equals the prior, and at the intersection of U with the vertical line at the prior we get the expected payoff with a non-informative signal. The payoff under full information is 0.35 and under no new information is 0.2. Analogously, Figure 19 shows V for MON, whose expected utility under full information is 0.23, and with no new information, the expected payoff is 0.2.

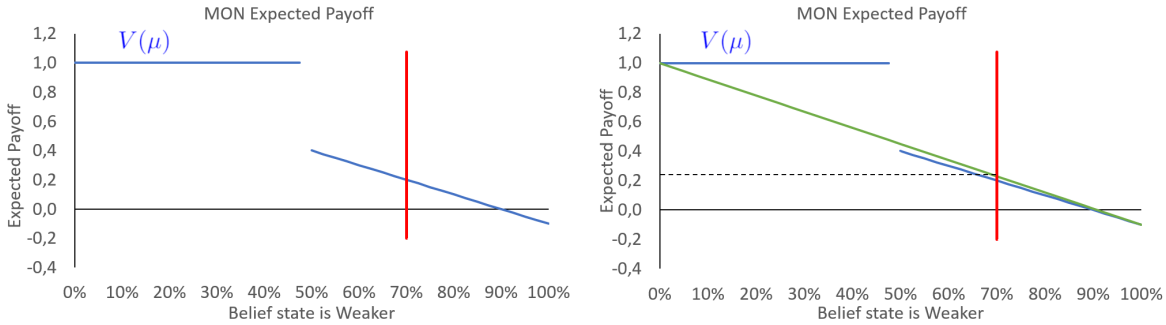

Figure 19: Payoffs as a function of beliefs for MON. The vertical red lines indicate the prior. On the right, the diagonal green line indicates the split of posteriors for $\lambda = 0$.

We then plot $\tilde{V}_{0.7}$. As Figure 20 evidences, the numerical solution for the optimal λ for MON is approximately 0.3. By setting $\lambda > 0$, MON induces the competitor to enter the market when the state is *Stronger*, which is better than when they compete only when the state is *Weaker*. This also implies that the competitor will choose *not compete* when the state is *Weaker*, but that is not a problem for the monopolist, as in this action the payoff does not depend on the state. But if MON keeps increasing the value of λ, CON will *compete* more often, as the *compete* signal will appear with higher probability. After a certain point, costs will outweigh the benefits as the monopolist expects to be at the *Weaker* state most of the time (the prior is higher than 0.5). And so, in equilibrium, MON will seek to confuse

CON, but "just a little bit".

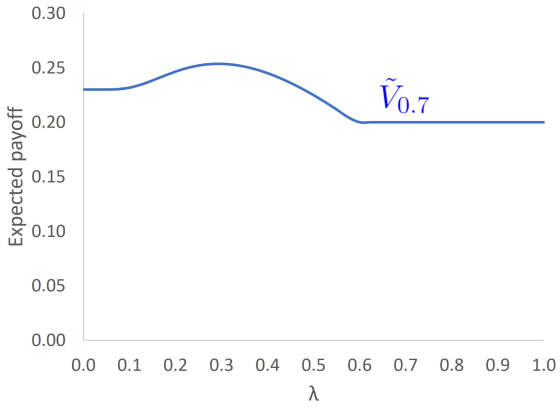

Figure 20: Plot of $\tilde{V}_{0.7}$.

C Costly Obfuscation

In this section, we extend the model to allow for costly obfuscation and briefly discuss the conclusions. It is important to note that it may be easier to hide or to reveal information depending on the available technology and the status quo. Think of a couple planning where to go out for dinner at the cheapest restaurant in the neighborhood. In the past, finding the cheapest restaurant required making a list of local places, then calling each in turn; now, all it requires is a quick search on Google Maps.

To allow for the fact that it may be easier or harder to withhold information, we add a parameter $\lambda_0 \in \mathbb{R}_+$, which we call the *natural* marginal cost of information. Obfuscation (or revealing) becomes costlier the further away the chosen λ is from λ_0. In the restaurant example, before Google Maps was invented, we had a high λ_0. After Google Maps, we have a low λ_0.

Definition 6. *The firm's Costly Obfuscation Problem is*

$$\max_{\lambda \in [0,\overline{\lambda}]} \tilde{V}_{\mu_0}(\lambda) - \kappa(\lambda - \lambda_0)^2. \tag{26}$$

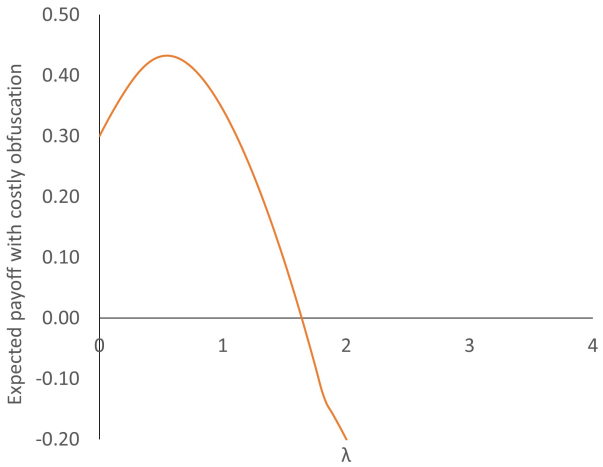

Figure 21: The firm's payoffs as a function of λ under costly obfuscation. We set $\lambda_0 = 1$ and $\kappa = 0.2$. There is a kink around $\lambda = 1.75$ because that is the point where $\tilde{V}_{\mu_0}$ becomes constant.

Here, $\kappa \geqslant 0$ represents the marginal cost of pushing obfuscation away from the natural level and $\overline{\lambda}$ is defined as before. It is easy to see that the existence of a solution for the costly version is guaranteed by the existence of a solution for the original problem. We can further remove the assumption that $a^*(\mu_0)$ is a singleton whenever $\kappa > 0$, and the costly version of the problem will still always have a solution. Indeed, as long as $\lambda_0 \in \mathbb{R}_+$, setting $\lambda \to \infty$ will never be optimal for the firm.

We can look at the impact of λ_0 on the equilibrium level of obfuscation. Let λ^* be the solution to the firm's problem when $\kappa = 0$. The closer λ_0 is to λ^*, the less of an effect κ has on the equilibrium level of obfuscation. Indeed, if $\lambda^* = \lambda_0$, the equilibrium level of obfuscation does not change with adding a cost of obfuscationκ.

Let us look at a simple example of costly obfuscation building upon the application in Section B.1. We will start with the payoff structure in Table 5 and the same prior of $\mu_0 = 0.75$. We know from our Theorem 1 that, with $\kappa = 0$, we have a full transparency problem. Now, let us set $\kappa = 0.2$ and $\lambda_0 = 1$, so both full transparency and full obfuscation are costly for the firm. Figure 21 shows the payoffs for the firm as a function of λ. We can see that we do not get full transparency anymore (compare Figure 21 with Figure 17).

In this example, where the firm would choose full transparency if it could, policies that reduce κ and λ_0 would result in more transparency, but that is not always the case. It is easy to see that, with costly obfuscation, inducing full transparency needs to take into account the values of both κ and λ_0 simultaneously.