The Good, the Bad and the Complex: Product Design with Imperfect Information

Vladimir Asriyan, Dana Foarta and Victoria Vanasco*

March 20, 2019

Abstract

This paper explores the incentives of product designers to complexify products, and the resulting implications for overall product quality. In our model, a consumer can accept or reject a product proposed by a designer, who can affect the quality and the complexity of the product. While the product's quality determines the direct benefits of the product to the consumer, the product's complexity affects the information a Bayesian consumer can extract about the product's quality. Examples include policymakers who propose policies for approval by voters, or banks that design financial products that they later offer to retail investors. We find that complexity is not necessarily a feature of low quality products. For example, while an increase in alignment between the consumer and the designer leads to more complex but better quality products, higher product demand or lower competition among designers leads to more complex and lower quality products. Our findings can help rationalize the observed trends in quality and complexity of regulatory policies and financial products.

Keywords: product design; information frictions; information transmission; signaling; complexity; regulation; financial products.

JEL Codes: D82, D83, G18, P16, D78.

*Asriyan: CREi, BGSE and CEPR, vasriyan@crei.cat; Foarta: Stanford Graduate School of Business, ofoarta@stanford.edu; Vanasco: CREi and BGSE, vvanasco@crei.cat. We thank Anat Admati, Larbi Alaoui, Fernando Broner, Steve Callander, Peter DeMarzo, Ben Hèbert, Jason Roderick Donaldson, Deeksha Gupta, Arvind Krishnamurthy, Andrey Malenko, Alberto Martin, Giorgia Piacentino, Giacomo Ponzetto, Andy Srkzypacz, Jaume Ventura, Jeffrey Zwiebel and seminar/conference participants at Bank of Italy, CREi-UPF, EIEF Rome, Essex, Stanford University, Stockholm School of Economics, Yale SOM, 11th Annual Paul Woolley Centre Conference at LSE, Sixth Oxford Financial Intermediation Theory Conference, IX Workshop on Institutions, Individual Behavior and Economic Outcomes, Warwick/Princeton/Utah Political Economy Conference in Venice, and Individual Behavior and Economic Outcomes in Sardinia for helpful comments and feedback. Vanasco acknowledges financial support from the Barcelona GSE Seed Grant.

1 Introduction

Rapidly increasing complexity has been an issue of concern for policymakers and financial market participants alike. In the regulatory sphere, for instance, it has become more common to encounter policy proposals lacking specificity, with broadly worded or ambiguous provisions (McMillan, 2014; Davis, 2017). Similarly, in the financial industry, it has been argued that the products sold to retail investors have become increasingly complex – with product descriptions that contain jargon and complicated or vague explanations (Célérier and Vallée, 2017; Ghent et al., 2019). Such an increase in complexity could be a concern if it prevents agents from accurately evaluating policies or judging the quality of the products that they buy, which could in turn foster the proliferation of undesirable policies and products.

In many situations, agents evaluate a product before deciding whether to accept it or not. Some examples include the median voter who evaluates a policy proposal, retail investors who evaluate financial products, the editor of a journal who evaluates a paper. To make a decision, the agent gathers information from several sources, such as the description of the product attributes, reviews, media reports, etc. The product designer, who chooses the product attributes, can in turn influence the quality of the information that the agent receives by complexifying or simplifying his product.[1] For example, a product can be complexified by adding unnecessary attributes and contingencies, or by opting for complicated jargon and lengthy and ambiguous descriptions.

A growing literature suggests that product complexity can result from purposeful obfuscation of bad product attributes, such as high prices, by firms that are faced with unsophisticated consumers (see Spiegler (2016) for a survey of this literature). In that view, unsophistication is essential since otherwise consumers would eventually stop demanding products whose attributes they do not understand well (Milgrom, 1981). By contrast, in this paper we develop

[1]We view policymakers as designers of legislation and policy proposals, academics as designers of research papers, and financial intermediaries as designers of financial products.

a framework that rationalizes the proliferation of complex products in a rational setting. We show that when the information that the consumer receives about product attributes is imperfect –i.e. there is risk of miscommunication,- the designers of good products may choose to complexify them in order to reduce the risk of miscommunication. Alternatively, the designers of bad products may choose to simplify them in order to exploit this risk. We thus break the link between bad product attributes and complexity; as a result, our model suggests novel drivers of product complexity and shows that complexification can be accompanied by either a decrease or an increase in overall product quality.

We consider a setting where a consumer demands a product that is supplied by a designer, who takes private actions that affect the product's output and complexity. Whereas a product's output (good or bad) determines the direct payoff to the consumer, a product's complexity affects her information set. In particular, when the designer complexifies (simplifies) the product, the consumer is more likely to extract less (more) precise information about the product's output. The objective of the designer is to get the consumer to accept his product, whereas the consumer only wants to accept a good product. For example, a policymaker wants to obtain the median voter's approval for his tax reform, while a bank wants to convince a retail investor to accept a savings account.[2] Finally, we suppose that the product designer is misaligned with the consumer; that is, he receives a higher payoff from having a bad product accepted. Such misalignment aims to capture conflicts of interest stemming from good products being more costly to design, career concerns, ideological preferences, or privately negotiated sales commission incentives. The timing of the game is as follows. First, the designer privately takes actions to affect the product's output and its complexity, and then he proposes the product to the consumer. Second, the consumer obtains noisy information about the product's output and decides whether to accept the product. If the consumer accepts, product payoffs are realized; otherwise, everyone gets their outside options.

The product designer in our setting takes actions to separately affect the product's output

[2] Our objective as paper designers is to have the Editor accept our proposed paper for publication.

and its complexity.[3] For example, the output of a financial product is determined by the net present value (NPV) that it generates to an investor. There are, however, many financial contracts that generate the same NPV. Thus, for a given NPV, a financial product can be complexified by adding contingencies that generate zero NPV to the investor, by using ambiguous words in the product's description, by linking payments to financial indices that the consumer is unlikely to know, etc. A similar argument can be made for policymakers in charge of writing policy proposals. By studying both attributes separately, we gain a better understanding of the incentives to produce good/bad quality vs. complex/simple products.

An essential ingredient of our model is that the information received by the consumer is imperfect: in particular, the product designer cannot verify the product's output to the consumer.[4] By simplifying his product, the designer can only increase the likelihood that the consumer's information is more precise, but he cannot make it perfectly precise. This is natural, since in many settings a designer's actions cannot fully determine the information set of the consumer. For instance, even though a policymaker may take actions to simplify a proposed regulation, the median voter may still find it complex due to features of the environment — e.g., the proposed regulation may be complex because it addresses a complicated topic–, or due to other exogenous factors — e.g., amendments added by another agency.

In our setting, the consumer disregards sufficiently imprecise information when making her acceptance decision. It follows that when the consumer would accept the product in the absence of new information –i.e., her belief about the product's output being good is sufficiently high or her outside option is sufficiently low,- the designer has an incentive to complexify, as more precise information could make the consumer change her mind and reject the product. In this case, good products may be complexified in order to reduce the risk of miscommunication. For example, a policy maker who has enough support to pass a tax reform

[3] In a different context, Bar-Isaac et al. (2010) study incentives to produce different product attributes by exploring a firm's integrated strategy for marketing, pricing, and investment in quality, where marketing affects the information consumers receive about the product's other attributes.

[4] As we show in Section 2.1, when the designer can verify the product's output, there is a unique equilibrium where only good products are designed, and their quality is fully revealed to the consumer.

may have little incentives to provide detailed information about the reform to his supporters. Analogously, when the consumer would reject the product in the absence of new information, the designer has an incentive to simplify, as more precise information could make the consumer change her mind and accept the product. In this case, bad products may be simplified in order to exploit the risk of miscommunication. For example, a policy maker with no support for his tax reform can only obtain such support by providing additional information about it. A related result is present in Perez-Richet and Prady (2011), who show that a privately informed sender may "complicate to persuade" a receiver in order to obtain a certification.[5]

In equilibrium, the consumer's belief about the product's output being good must be consistent with the designer's decision to produce a good product. When making his decision, the designer faces a tradeoff between a higher probability of acceptance (obtained with a good product) versus a higher payoff conditional on acceptance (obtained with a bad product). The equilibrium is in mixed strategies, in which the designer produces good products with some probability. We show that a product is more likely to be of bad quality when the alignment between the consumer and the designer is low, or when the product demand, interpreted as a relatively low outside option for the consumer, is high.

Our model suggests that in order to understand the relationship between quality and complexity of products, it is essential to understand the underlying drivers of product heterogeneity. For instance, we show that as the alignment between the consumer and the designer increases, products become of better quality and more complex, implying a positive relationship between quality and complexity. On the other hand, as the consumer's demand for the product increases, products become of worse quality and more complex, implying a negative relationship between quality and complexity.

The result that complexity is optimal when alignment between the designer and the consumer is sufficiently high is closely related to findings in Dewatripont and Tirole (2005), who

[5] In a different context, Zwiebel (1995) studies managerial project selection in an environment with reputational concerns and asymmetric information. He shows that the manager may choose conservative projects that are less profitable than innovative ones, because they are easier to evaluate by outside parties, i.e. there is a lower risk of miscommunication.

study optimal communication strategies in a sender-receiver game with moral hazard. They find that when the decision of the receiver is *supervisory*, i.e. whether to take a given action or not (as in our paper), the sender may choose to put zero effort to communicate, and the receiver to interpret, when the sender-receiver alignment (*congruence* in their paper) is sufficiently high. Similar incentives have been highlighted, albeit in a different context, by Bond and Zeng (2018), who find that silence may be optimal when the sender of information is risk-averse and faces uncertainty about the receiver's preferences.

Our results are complementary to the findings in Gabaix and Laibson (2006), who study the incentives of firms to shroud certain product attributes from consumers.[6] A crucial ingredient of their model is that there is a fraction of unsophisticated consumers who make no inferences from the fact that they do not observe a certain attribute. In a related paper, Pagano and Volpin (2012) show that securitizers may have incentives to increase the opacity of their products in order to draw unsophisticated investors into the market. In these papers, complexity allows sellers to extract rents from unsophisticated buyers, a mechanism that is not present in our setting where there is one type of rational consumer. Instead, we show that complexity could be a way for designers of good quality products to induce product acceptance by reducing the risk of miscommunication, a mechanism that is not present in these papers. Analogously, we show that simplicity could be a way for designers of bad quality products to induce product acceptance by exploiting the risk of miscommunication.

Our model belongs to the class of signaling games, where a sender (product designer) sends a signal to a receiver (consumer) who makes a decision based on the information she receives, which affects both agents' payoffs. This naturally relates our paper to the literature on strategic information transmission in games (Milgrom, 1981; Grossman, 1981; Crawford and Sobel, 1982; Kartik, 2009) and especially to the recent applied work on the value of ignorance or opacity in incentive provision schemes (Brocas and Carrillo, 2007; Ederer et al., 2018). We differ from this literature in two respects. First, in our model, the designer is restricted in

[6]Relatedly, Auster and Pavoni (2018) consider an optimal investment delegation problem to study the incentives of the financial intermediaries to shroud investment opportunities from investors.

his information choice as he must choose among signal distributions that are neither fully informative nor uninformative; hence, there is risk of miscommunication. Second, in our model, the designer takes an action which determines the underlying "state" about which information is transmitted. Our work also relates to the growing literature on Bayesian persuasion (Rayo and Segal, 2010; Kamenica and Gentzkow, 2011) and on information design more broadly (Bergemann and Morris, 2016; Taneva, 2019). This literature focuses on the designer's ex-ante problem, whereas we study a setting in which the information structure is chosen ex-post, *after* the designer has privately chosen the underlying "state."

We show that our results are robust to several extensions and alternative specifications. First, we extend the model to study the effect of competition among designers (Section 6.1). To do so, we consider a sequential search setting whereby, if the consumer rejects a given product, then she meets another designer with some probability, capturing the presence of search frictions. We show that as search frictions decrease, which we interpret as competition among designers intensifies, product quality increases and complexity decreases. This result is in contrast to the literature on obfuscation or price complexity, which typically finds that competition leads to more obfuscation (Spiegler, 2006; Ellison and Ellison, 2009; Carlin, 2009; Ellison and Wolitzky, 2012).[7] The reason for our opposing results is that in these papers, by obfuscating, a firm reduces the consumer's ability to uncover the attributes of competing products, e.g. by increasing search costs. Instead, in our paper, by complexifying, a designer reduces the precision of the information that a rational consumer extracts about his own product. As a result, by effectively increasing the consumer's outside option, more intense competition incentivizes each designer to supply products that are better for the consumer, i.e., those that are good and simple.

Second, we show that our model can be recast within a rational inattention framework (Sims, 2003; Wiederholt et al., 2010), in which the consumer can reduce her uncertainty about the

[7]Some papers in this literature also highlight that obfuscation gives room to price dispersion, which in turn allows firms to price-discriminate among different consumer-types (e.g. fast vs. slow searchers, sophisticated vs. unsophisticated), a force that is present even in the absence of competition (Salop, 1977). This mechanism is not present in our model, where there is one type of rational consumer.

product's output by incurring an entropy-reduction cost (Section 6.2).[8] By complexifying a product, the designer makes it more costly for the consumer to extract information about the product's output. Though this approach is less tractable, we show that allowing for this more flexible information acquisition technology does not change the qualitative insights derived from our baseline model. Through this lens, our model relates to Roesler and Szentes (2017), who study buyer-optimal learning in a similar context, but where the seller cannot affect the buyers' learning process, and to Oehmke and Zawadowski (2019), who focus on sellers' incentives to complexity products in a setting where complex products give more value to buyers but are costlier to understand.

Finally, we consider other extensions that may be relevant for particular applications, such as (i) introducing designers who are aligned with the consumer (Section 6.3), (ii) allowing the designer to set a price for his product (Section 6.4), and (iii) introducing direct costs to the consumer from unnecessary complexification or simplification by the designer (Section 6.5).

Even though our model is stylized and abstracts from many institutional details of real-world settings, we discuss its main predictions within the context of two concrete applications: financial products and regulatory policy. In financial markets, intermediaries design financial products to offer to retail investors, such as savings accounts and asset-backed securities. The design of a product consists of determining a set of cash flows for different states of the world (price, future payments, fees, contingencies, etc.) and of writing these contract terms down. The investor evaluates a product and decides whether to invest or not. In this context, our model suggests that the increase in the complexity of financial products documented by Célérier and Vallée (2017) could be driven by (i) an increase in investor's trust in financial advisors, or (ii) an increased demand for financial products. Both of these features were characteristics of financial markets prior to the 2008/09 crisis, and while the trust in the financial system may have fallen in response to the crisis, the high demand for relatively safe financial products still persists. This suggests that the observed proliferation of worse and

[8] Another related setting is that of computational complexity, as in Aragones et al. (2005).

more complex products could be an endogenous response of product designers to an increasing demand for relatively safe financial products.

In the political sphere, politicians are the designers in charge of proposing policies, such as plans for taxation or regulation. We can interpret the consumer in such a setup as the median voter, from whom the politician must obtain approval for policy proposals. A more complex policy proposal by a politician may, for instance, take the form of less specific promises and hazy details on the exact implementation of the policy goal. This idea is present in the literature on strategic ambiguity, deliberate vagueness or noise by politicians (Alesina and Cukierman, 1990; Aragones and Postlewaite, 2002; Espinosa and Ray, 2018). Our model suggests that policy proposals are more likely to be complex when (i) public opinion about the politician's alignment with the median voter is high, or (ii) there is urgency to pass a given policy; that is, the status-quo is costly. Both of these features were present when complex policies, such as the Affordable Care Act and the Dodd-Frank Act were passed. Moreover, in policy areas where public opinion of politicians is low, policy proposals tend to be simpler with subsequent details and lawmaking authority being delegated to federal agencies, for example to the FDA in the case of the pharmaceutical industry.

The rest of the paper is organized as follows. In Section 2, we setup our baseline model. In Section 3, we present our main results. We conduct comparative statics in Section 4, which we discuss in the context of applications in Section 5. In Section 6, we explore several extensions of the baseline model. We conclude in Section 7. All proofs are relegated to the Appendix.

2 The Model

We consider the following interaction between a consumer and a product designer. The consumer needs a product, that only the designer can produce. The designer privately takes two actions $\{y, \kappa\}$, where $y \in \{\mathbf{G}ood, \mathbf{B}ad\}$ affects the product's output, and $\kappa \in \{\underline{\kappa}, \bar{\kappa}\}$ affects the product's complexity. The designer then proposes the product to the consumer, who evaluates it and decides whether to accept it ($a = 1$) or take an outside option ($a = 0$).

A Product's Output. By taking action y, the designer affects the output of the proposed product. The payoff to the consumer from accepting a product with output y (which we refer to as a y-product) is $w(y)$, and her outside option if no product is accepted is w_0. The designer receives payoff $v(y)$ from having a y-product accepted, and zero otherwise. We make the following assumptions on the payoffs:

Assumption 1 *The payoffs satisfy the following properties:*

1. $w(G) > w_0 > w(B) \geq 0$.

2. $v(B) > v(G) > 0$.

3. $w(G) + v(G) > w_0 > w(B) + v(B)$.

The first assumption states that the consumer wants to accept a G-product but reject a B-product. The second assumption states that the designer is *misaligned* with the consumer, as he prefers to have a B-product accepted.[9] Since in practice not all designers may be misaligned with the consumer, we extend our analysis and introduce aligned designers in Section 6.3. The third assumption states that only G-products are efficient to produce. In Section 6.4, we show that in some applications our payoffs structure can also be rationalized by introducing prices and costs of production, whereby G-products are costlier to produce.

A Product's Complexity. By taking action κ, the designer affects how difficult it is for the consumer to understand the product's output, y. Formally, we suppose that after the designer proposes product (y, κ), the consumer is able to extract a binary signal $S \in \{b, g\}$ about the product's output with some noise $z \equiv \mathbb{P}(y = G|S = b) = \mathbb{P}(y = B|S = g)$, where $z \sim F(\cdot|\kappa)$ and $F(\cdot|\kappa)$ has a corresponding pdf $f(\cdot|\kappa)$ with full support on $\left[0, \frac{1}{2}\right]$, satisfying the property

[9] In the policy sphere, misalignment of policymakers vis-à-vis the public may arise due to ideological differences, lobbying, or career concerns. In the financial products industry, misalignment may arise due to financial advisors receiving higher fees for selling products that are not necessarily the best fit for their clients (i.e., fixed vs. adjustable-rate mortgages).

that $\frac{f(z|\bar{\kappa})}{f(z|\underline{\kappa})}$ increasing in z (MLRP). Thus, the consumer is more likely to receive a noisier signal about the product's output if the designer chooses $\kappa = \bar{\kappa}$. In what follows, we will refer to z as a product's complexity, and we will say that the designer *complexifies* a product when he increases the product's expected complexity, i.e., he chooses $\kappa = \bar{\kappa}$. Otherwise, we will say that the designer *simplifies* a product. This setup captures the intuitive notion that the consumer is able to extract less information from more complex products.[10]

Remarks on our Modeling Approach. The imperfect link between the action κ and observed complexity z essentially works as a refinement and rules out multiplicity of equilibria, typical of signaling games, that arise from the freedom in specifying off-equilibrium beliefs. This approach allows us to obtain a unique equilibrium and facilitates rich comparative statics (Matthews and Mirman, 1983). Moreover, it has the natural interpretation that the consumer is unable to perfectly observe the underlying actions that the designer takes towards complexification. This may be due to features of the environment unknown to the consumer, e.g. a proposed regulation may be complex because it addresses a complicated topic or because the legislator intended to make it difficult to decipher; or due to other exogenous factors in the production process, e.g. amendments added to a proposed regulation by another agency.

In our baseline setup, we assume that product complexity does not directly affect the consumer's payoffs. This is convenient, as it allows us to isolate the strategic role of complexity in deterring information acquisition or learning by consumers. Nevertheless, in some applications of interest, unnecessary complexity can be costly to the consumer, e.g. costs of complying with multiple provisions. We address such considerations in Section 6.5 where we incorporate direct costs to the consumer from the designer's action to complexify or simplify his product.

The Consumer's Problem. The consumer has to decide whether to accept the designer's product or not. Before making her decision, the consumer observes the product's complexity

[10] We formalize this interpretation in Section 6.2 where we consider the alternative approach of rational inattention, in which the consumer can choose how much to learn about the product's output subject to an entropy-reduction cost. In that setting, higher entropy-reduction cost corresponds to higher product complexity, and the designer can take an action to increase this cost. Even though this alternative specification is less tractable, we show that the equilibrium of this alternative model resembles that of our baseline model.

z and signal s, and forms her posterior belief about the output y being good, denoted by $\mu(s,z) \equiv \mathbb{P}(y = G|s,z)$. The consumer's expected payoff is given by

$$W(a|s,z) \equiv a \cdot [\mu(s,z) \cdot w(G) + (1 - \mu(s,z)) \cdot w(B)] + (1-a) \cdot w_0. \tag{1}$$

The consumer chooses $a \in \{0,1\}$ to maximize (1).

The Designer's Problem. The designer's expected payoff is given by

$$V(y,\kappa) \equiv \mathbb{P}(a = 1|y,\kappa) \cdot v(y) \tag{2}$$

where $\mathbb{P}(a = 1|y,\kappa)$ denotes the probability that product $\{y,\kappa\}$ is accepted by the consumer. The designer chooses $y \in \{G,B\}$ and $\kappa \in \{\underline{\kappa},\bar{\kappa}\}$ to maximize (2). We denote the designer's strategy by $\{m,\sigma_G,\sigma_B\}$, where $m = \mathbb{P}(y = G)$ is the probability that the designer chooses a G-product and $\sigma_y = \mathbb{P}(\kappa = \bar{\kappa}|y)$ is the probability that he chooses to complexify a y-product.

Equilibrium Concept. We use Perfect Bayesian Equilibrium (PBE) as our equilibrium concept. This has the following implications. First, given her beliefs, the consumer's acceptance strategy must maximize her expected payoff (*Consumer Optimality*). Second, the designer's strategy must maximize his expected payoff, given the consumer's strategy (*Designer Optimality*). Finally, the consumer's beliefs must be consistent with the designer's strategy and updated using Bayes' rule when possible (*Belief Consistency*).

2.1 Benchmark with Perfect Communication

Before we proceed to the equilibrium analysis, it is useful to establish a benchmark against which our results can be compared. To highlight the role of miscommunication, we consider a benchmark in which the designer can perfectly communicate (or verify) his product's output to the consumer. More concretely, suppose that the designer can directly choose the complexity of the product, $z \in [0,\frac{1}{2}]$. Thus, after choosing output, y, the designer can potentially verify it to the consumer by setting $z = 0$.[11] The rest of the setup is unchanged.

The following proposition states that if perfect communication is available to the designer, then the efficient outcome is attained.

Proposition 1 *If the designer can perfectly communicate his product output to the consumer, then in equilibrium only the G-product is produced and it is accepted with probability one.*

Intuitively, when perfect communication is available, the designer of a G-product never wants to expose himself to the risk of miscommunication; thus, he chooses to perfectly reveal his product output to the consumer. As a result, the designer of a B-product cannot exploit miscommunication to get his product accepted, since the consumer rationally infers that the designer has chosen a B-product if its output is not perfectly communicated, and thus rejects it with probability one (w.p.1). Therefore, in equilibrium, only G-products are produced and they are accepted w.p.1, i.e., the efficient outcome is attained. This result is reminiscent of the famous unraveling result in the persuasion game of Milgrom (1981), where a privately informed salesman always chooses to truthfully report the quality of his product to the uninformed buyer. Finally, it is easy to see that the efficient outcome would also be attained if the product output were perfectly observable to the consumer, since then the consumer would simply reject any product that she observes has a B-output. Therefore, in what follows, it will be essential that *both* (i) the consumer does not observe the product output, y, chosen by the designer, *and* (ii) the designer cannot perfectly communicate it to the consumer.

3 Equilibrium

In this section, we characterize the equilibria of our game. First, we consider the consumer's optimal acceptance strategy, given her beliefs about the product proposed by the designer (Section 3.1). Second, we analyze the designer's strategy: his optimal choice of output (Section

[11] To capture perfect communication, we set $\mu(g, 0) = 1$ and $\mu(b, 0) = 0$, i.e., the designer has the option to verify his product output.

3.3) and of complexification (Section 3.2), given the consumer's acceptance strategy. Finally, we impose belief consistency (Section 3.4).

By inspection of our game, it is immediate that there is always a trivial equilibrium with zero trade in which (i) the consumer correctly believes that the designer has chosen a B-product and thus she rejects it w.p.1, and (ii) the designer indeed chooses a B-product w.p.1 as he is indifferent to producing a G- vs. a B-product (both yield a zero expected payoff). In what follows, we focus on the more interesting equilibria with positive trade, where the designer chooses a G-product with positive probability.

3.1 The Consumer's Acceptance Strategy

From the consumer's problem, as given in (1), we see that she follows a threshold strategy: she accepts the product, $a(s, z) = 1$, if and only if her posterior belief about the product having good output is sufficiently high, $\mu(s, z) \geq \omega$, where $\omega \equiv \frac{w_0 - w(B)}{w(G) - w(B)}$ captures the relative value of the consumer's outside option.[12]

Thus, in order to understand the consumer's acceptance strategy, we need to analyze the determinants of her posterior belief. Let $\mu \equiv \mathbb{P}(y = G)$ denote the consumer's prior belief, which must be positive as we are looking at equilibria with positive trade. After the designer proposes his product, the consumer observes the product complexity z and signal s about the product's output. Since complexity is informative about the designer's action κ, it may contain information about output y. Let $\mu(z)$ denote the consumer's interim belief upon observing z, which is given by

$$\mu(z) \equiv \mathbb{P}(y = G|z) = \frac{\mu}{\mu + (1 - \mu)\, \ell(z)}, \tag{3}$$

where the prior belief μ and the likelihood ratio $\ell(z) \equiv \frac{\mathbb{P}(z|y=B)}{\mathbb{P}(z|y=G)}$ will need to be consistent with the designer's equilibrium strategy, which the consumer takes as given. As a result, the

[12]If the consumer is indifferent, we assume without loss of generality that she accepts the product. Since such an indifference will arise with probability zero, what happens in that event will not affect the equilibrium.

consumer's posterior belief, $\mu(s,z)$, after observing complexity z and signal s is

$$\mu\left(s,z\right)=\frac{\mathbb{P}\left(S=s|y=G\right)\cdot\mu\left(z\right)}{\mathbb{P}\left(S=s|y=G\right)\cdot\mu\left(z\right)+\mathbb{P}\left(S=s|y=B\right)\cdot\left(1-\mu\left(z\right)\right)}. \tag{4}$$

The consumer's acceptance strategy is contingent on the observed signal whenever she accepts the product if she observes a good signal, $S=g$, but rejects it if she observes a bad signal, $S=b$. For this to be optimal, the signal has to be informative enough so that

$$\mu\left(b,z\right)<\omega\leq\mu\left(g,z\right). \tag{5}$$

Figure 1 illustrates the behavior of the posterior $\mu(s,z)$ as it depends on the observed signal s and complexity z, for a given prior belief $\mu\in(0,1)$ and likelihood ratio $\ell(z)$. The figure highlights several key features of the consumer's problem. First, with zero complexity, the signal is perfectly informative, so $\mu(g,0)=1$ and $\mu(b,0)=0$. Second, with maximal complexity, the consumer gains no new information upon observing signal s, so $\mu(g,\frac{1}{2})=\mu(b,\frac{1}{2})=\mu(\frac{1}{2})\in(0,1)$. Finally, the posterior belief $\mu(g,z)$ is decreasing in complexity, while the posterior belief $\mu(b,z)$ is increasing in complexity.[13]

The following definition will be useful in characterizing the consumer's acceptance strategy.

Definition 1 *We say that the consumer is* **optimistic** *if her interim belief satisfies* $\mu\left(\frac{1}{2}\right)\geq\omega$, *whereas we say that she is* **pessimistic** *if* $\mu\left(\frac{1}{2}\right)<\omega$.

The consumer is optimistic when, after observing an uninformative signal, her posterior belief is higher than her relative outside option. This case is depicted in Figure 1(a): here, the consumer always accepts the product after a good signal (i.e., $\mu(g,z)>\omega$ for all z), while after a bad signal she accepts the product whenever its complexity z is above threshold $\bar{z}$ (i.e., $\mu(b,z)\geq\omega$ iff $z\geq\bar{z}$). Instead, Figure 1(b) depicts the case of a pessimistic consumer, who

[13] In the Appendix, we impose a regularity condition on the conditional distributions of complexity z, which ensures that for each s the posterior $\mu\left(s,z\right)$ is monotonic in z (see Condition A.1).

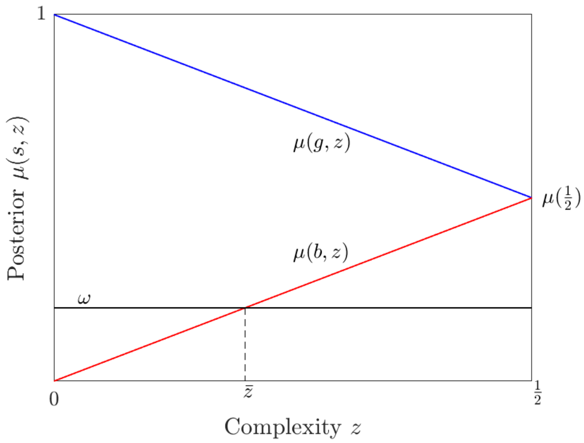

(a) Optimistic Consumer: $\mu(\frac{1}{2}) \geq \omega$.

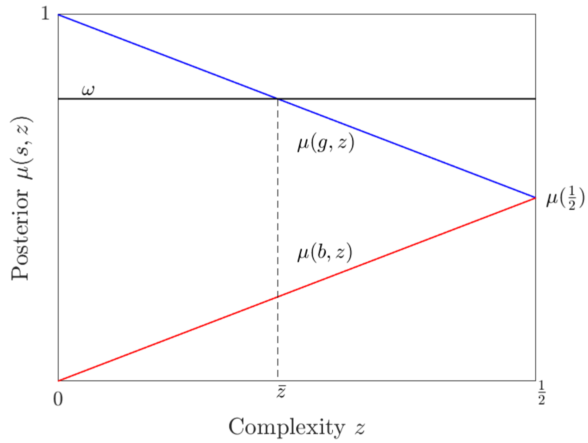

(b) Pessimistic Consumer: $\mu(\frac{1}{2}) < \omega$.

Figure 1: This figure illustrates the posterior belief $\mu(s, z)$ as it depends on the signal, s, and complexity, z.

always rejects the product after a bad signal (i.e., $\mu(b, z) < \omega$ for all z), while after a good signal she accepts the product whenever its complexity z is below threshold $\bar{z}$ (i.e., $\mu(g, z) \geq \omega$ iff $z \leq \bar{z}$). To summarize, when complexity is relatively high ($z > \bar{z}$), the consumer disregards her signal, and she accepts the product if she is optimistic, but rejects it if she is pessimistic. Conversely, when complexity is relatively low ($z \leq \bar{z}$), the consumer makes her acceptance decision contingent on the signal, and she accepts the product after a good signal, but rejects it after a bad signal. We formalize this discussion in the following lemma.

Lemma 1 *When the consumer is optimistic, her acceptance strategy is:*

$$a(s, z) = \begin{cases} \mathcal{I}_{\{S=g\}} & \text{if} \quad z \leq \bar{z} \\ 1 & \text{if} \quad z > \bar{z} \end{cases}, \tag{6}$$

where $\mu(b, \bar{z}) = \omega$ and $\mathcal{I}_{\{S=g\}}$ is the indicator equal to one when the signal is good. When the consumer is pessimistic, her acceptance strategy is:

$$a(s, z) = \begin{cases} \mathcal{I}_{\{S=g\}} & \text{if} \quad z \leq \bar{z} \\ 0 & \text{if} \quad z > \bar{z} \end{cases} \tag{7}$$

where $\mu(g, \bar{z}) = \omega$.[14]

Importantly, both the value of threshold $\bar{z}$ and whether the consumer is optimistic or pessimistic are endogenous to equilibrium, since the prior belief μ and the likelihood ratio $\ell(z)$, which determine the beliefs $\mu(\frac{1}{2})$ and $\mu(s, z)$, will need to be consistent with the designer's strategy and Bayes' rule.

3.2 The Designer's Complexification Strategy

We next consider the designer's choice of κ, given his choice of output, y, and the consumer's acceptance strategy described in the previous section. From the designer's objective in (2), it follows that a designer who chooses a y-product also (weakly) prefers to simplify, $\underline{\kappa}$, whenever

$$\mathbb{P}(a = 1|y, \underline{\kappa}) \geq \mathbb{P}(a = 1|y, \bar{\kappa}). \tag{8}$$

Otherwise, the designer prefers to complexify, $\bar{\kappa}$. Using Lemma 1, we can compute the probability of acceptance of a y-product conditional on the product's complexity z. When the consumer is optimistic,

$$\mathbb{P}(a = 1|G, z) = \begin{cases} 1 - z & \text{if } z < \bar{z} \\ 1 & \text{if } z \geq \bar{z} \end{cases} \quad \text{and} \quad \mathbb{P}(a = 1|B, z) = \begin{cases} z & \text{if } z < \bar{z} \\ 1 & \text{if } z \geq \bar{z}, \end{cases} \tag{9}$$

whereas when the consumer is pessimistic,

$$\mathbb{P}(a = 1|G, z) = \begin{cases} 1 - z & \text{if } z \leq \bar{z} \\ 0 & \text{if } z > \bar{z} \end{cases} \quad \text{and} \quad \mathbb{P}(a = 1|B, z) = \begin{cases} z & \text{if } z \leq \bar{z} \\ 0 & \text{if } z > \bar{z}. \end{cases} \tag{10}$$

[14]When $\mu = 1$ or $\mu = 0$, we set without loss of generality $\bar{z} = 0$. As we will see, however, in any positive trade equilibrium $\mu \in (0, 1)$.

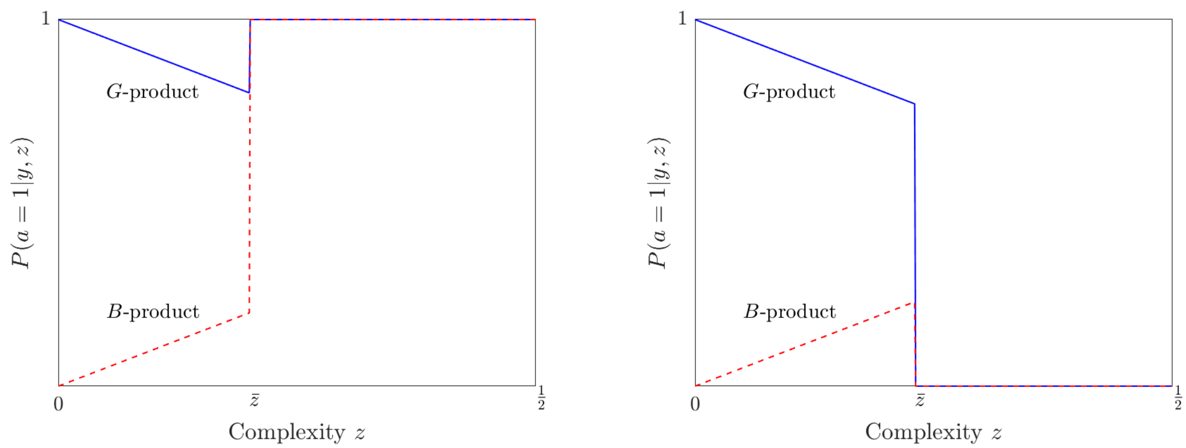

(a) Optimistic Consumer: $\mu(\frac{1}{2}) > \omega$.

(b) Pessimistic Consumer: $\mu(\frac{1}{2}) < \omega$.

Figure 2: This figure illustrates the probability of acceptance of a y-product as a function of complexity z.

Therefore, the designer's expected probability of a (y, κ) product being accepted is

$$\mathbb{P}\left(a=1|y,\kappa\right) = \int_0^{\frac{1}{2}} \mathbb{P}\left(a=1|y,z\right) \cdot f\left(z|\kappa\right) \cdot dz, \tag{11}$$

where $\mathbb{P}\left(a=1|y,z\right)$ is given by (9) if the consumer is optimistic and by (10) if she is pessimistic.

The following proposition characterizes the optimal complexification strategy of a designer who has produced a y-product.

Proposition 2 *Let $\widehat{z}$ denote the unique solution to $\int_0^{\widehat{z}} z \cdot f(z|\underline{\kappa})dz = \int_0^{\widehat{z}} z \cdot f(z|\bar{\kappa})dz$. Then, when the consumer is optimistic,*

$$\sigma_B = 1 \quad \textit{and} \quad \sigma_G \begin{cases} = 1 & \textit{if } \bar{z} < \widehat{z} \\ \in [0,1] & \textit{if } \bar{z} = \widehat{z} \\ = 0 & \textit{if } \bar{z} > \widehat{z}, \end{cases} \tag{12}$$

whereas, when the consumer is pessimistic,

$$\sigma_B \begin{cases} = 0 & \textit{if } \bar{z} < \widehat{z} \\ \in [0,1] & \textit{if } \bar{z} = \widehat{z} \\ = 1 & \textit{if } \bar{z} > \widehat{z} \end{cases} \quad \textit{and} \quad \sigma_G = 0. \tag{13}$$

Proposition 2 says that when the consumer is optimistic, the designer has a tendency to complexify, and vice versa. The intuition for this result can be obtained from Figure 2, which illustrates the acceptance probability $\mathbb{P}(a = 1|y, z)$ as it depends on the product output y and complexity z. First, a B-product's probability of acceptance increases in complexity when the consumer is optimistic (see Figure 2(a)). As a result, the designer of a B-product strictly prefers to complexify. Second, a G-product's probability of acceptance decreases in complexity when the consumer is pessimistic (see Figure 2(b)). As a result, the designer of a G-product strictly prefers to simplify.

The effect of complexity on the probability of acceptance, however, is non-monotonic for a B-product when the consumer is optimistic and for a G-product when the consumer is pessimistic. As a result, the designer's optimal complexification strategy depends critically on the consumer's acceptance strategy. If the consumer is optimistic and $\bar{z}$ is sufficiently low, the G-product designer complexifies his product in order to reduce the risk of miscommunication. On the other hand, if the consumer is pessimistic and $\bar{z}$ is sufficiently low, the B-product designer simplifies his product in order to exploit such risk.

3.3 The Designer's Output Strategy

When choosing the product's output, the designer faces a trade-off between increasing the product's acceptance probability (by choosing $y = G$) or increasing his payoff conditional on acceptance (by choosing $y = B$). Given the consumer's acceptance strategy, the net expected payoff to the designer from choosing the G-product over the B-product is

$$\gamma \equiv \max_{\kappa} \mathbb{P}(a = 1|G, \kappa) \cdot v(G) - \max_{\kappa} \mathbb{P}(a = 1|B, \kappa) \cdot v(B). \tag{14}$$

The first term is the expected payoff from choosing the G-product given the corresponding optimal choice of complexification/simplification action κ, as characterized in Proposition 2. The second term is the expected payoff from choosing the B-product given the corresponding optimal choice of κ. The probabilities in each scenario are computed as in equation (11). The next result then follows immediately.

Proposition 3 *Given the consumer's acceptance strategy, the designer chooses the G-product with probability*

$$m = \begin{cases} = 1 & \text{if } \gamma > 0 \\ \in [0,1] & \text{if } \gamma = 0 \\ = 0 & \text{if } \gamma < 0, \end{cases}$$

where γ is given by (14).

3.4 Characterization of Equilibria

In Section 3.1, we characterized the consumer's acceptance strategy given her beliefs. In Sections 3.2 and 3.3, we characterized the designer's output and complexification strategy given the consumer's acceptance strategy. To characterize the equilibria of our model, we now require that the consumer's beliefs be consistent with the designer's strategy and Bayes' rule.

First, we take the equilibrium prior belief μ as given, and we find the designer's equilibrium complexification strategy $\{\sigma_y\}$ by requiring that the consumer's interim belief, $\mu(z)$, be consistent with this strategy and Bayes' rule.

Proposition 4 *Suppose that in equilibrium the consumer's prior belief is $\mu \in (0,1)$, then there exist thresholds $\mu_1 - \mu_4$ such that:*[15]

1. *If $\mu \in (0, \mu_1]$, all products are simplified, $\sigma_B = \sigma_G = 0$.*

[15] The expressions for these thresholds are given by equations (33), (35), (38) and (41) in the Appendix.

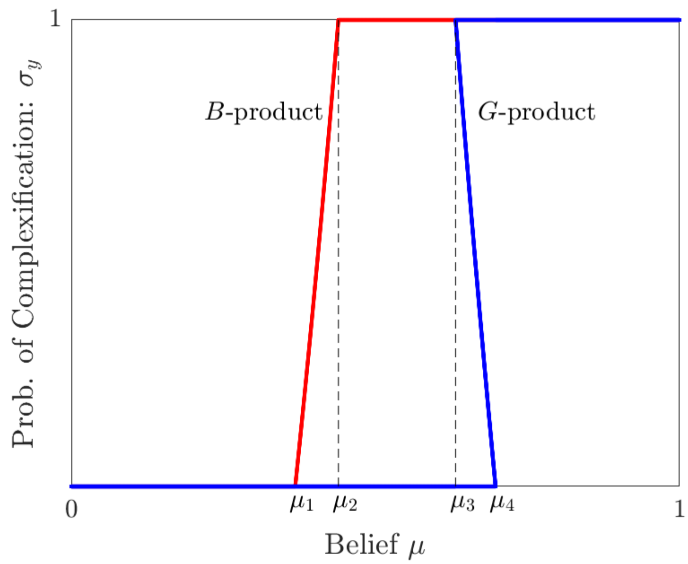

Figure 3: This figure illustrates how the complexification strategy of the designer who chooses a y-product varies with the consumer's belief μ.

2. *If $\mu \in (\mu_1, \mu_2]$, G-products are simplified, $\sigma_G = 0$, whereas B-products are complexified with probability*

$$\sigma_B = \left(\frac{f(\widehat{z}|\bar{\kappa})}{f(\widehat{z}|\underline{\kappa})} - 1\right)^{-1} \left(\frac{1-\widehat{z}}{\widehat{z}}\frac{\mu}{1-\mu}\frac{1-\omega}{\omega} - 1\right).$$

3. *If $\mu \in (\mu_2, \mu_3]$, G-products are simplified, $\sigma_G = 0$, whereas B-products are complexified, $\sigma_G = 1$.*

4. *If $\mu \in (\mu_3, \mu_4)$, G-products are complexified with probability*

$$\sigma_G \in \left\{0,\ 1 - \left(1 - \frac{f(\widehat{z}|\underline{\kappa})}{f(\widehat{z}|\bar{\kappa})}\right)^{-1} \left(1 - \frac{1-\widehat{z}}{\widehat{z}}\frac{1-\mu}{\mu}\frac{\omega}{1-\omega}\right),\ 1\right\},$$

whereas B-products are complexified, $\sigma_B = 1$.

5. *If $\mu \in [\mu_4, 1)$, all products are complexified, $\sigma_B = \sigma_G = 1$.*

The results of this proposition are illustrated in Figure 3. When μ is sufficiently low, all products are simplified, as the consumer is pessimistic and rejects sufficiently complex products. Instead, when μ is sufficiently high, all products are complexified, as the consumer is optimistic and accepts sufficiently complex products. Since complexification is always more beneficial for a B-product designer, the B-product is more likely to be complex than the G-product for intermediate values of μ.

Second, using Proposition 5, we find the designer's output strategy m by requiring that the consumer's prior belief, μ, be consistent with this strategy, i.e. $\mu = m$. It is straightforward to rule out equilibrium with $m = \mu = 1$, since in that case the consumer would accept the product w.p.1, making it optimal for the designer to only produce a B-product. Therefore, in any equilibrium with positive trade, it must be that $m = \mu \in (0, 1)$. The following proposition shows that there is a unique μ consistent with equilibrium.

Proposition 5 *There is generically a unique equilibrium with positive trade. In it, there exists a $\psi \in (0, 1)$ such that the designer produces a G-product with probability $m = \psi$, and he complexifies (simplifies) as described in Proposition 4 when $\mu = \psi$.*

An equilibrium with $m = \mu \in (0, 1)$ exists if there is an interior belief at which the designer is indifferent between producing G- and B-products. Figure 4 illustrates the designer's payoff from producing a G-product net of the payoff from producing a B-product as it depend on μ, which we denote by $\Gamma(\mu)$.[16] Observe that $\Gamma(\mu)$ is positive when μ is small and becomes negative when μ is large. Intuitively, when μ is small, the consumer is pessimistic and, thus, accepts products with low probability. The designer then expects a higher payoff from producing a G-product, given its higher probability of acceptance. As μ increases, the probability of acceptance increases for both products. At first, the difference between the probabilities of acceptance for a G- versus a B-product increases to further favor producing the G-product, as the consumer become more likely to make her acceptance decision contingent on the signal. As μ increases further, however, the gap between the probabilities of acceptance shrinks,

[16]Formally, Γ is a correspondence which maps belief μ to the set of net payoffs γ defined in (14) that are consistent with the designer's complexification strategy $\{\sigma_y\}$ (characterized in Proposition 4) and Bayes' rule. In Appendix A, we show that the properties of the correspondence Γ are as depicted in Figure 4, with the exception that ψ may lie in either of the five regions.

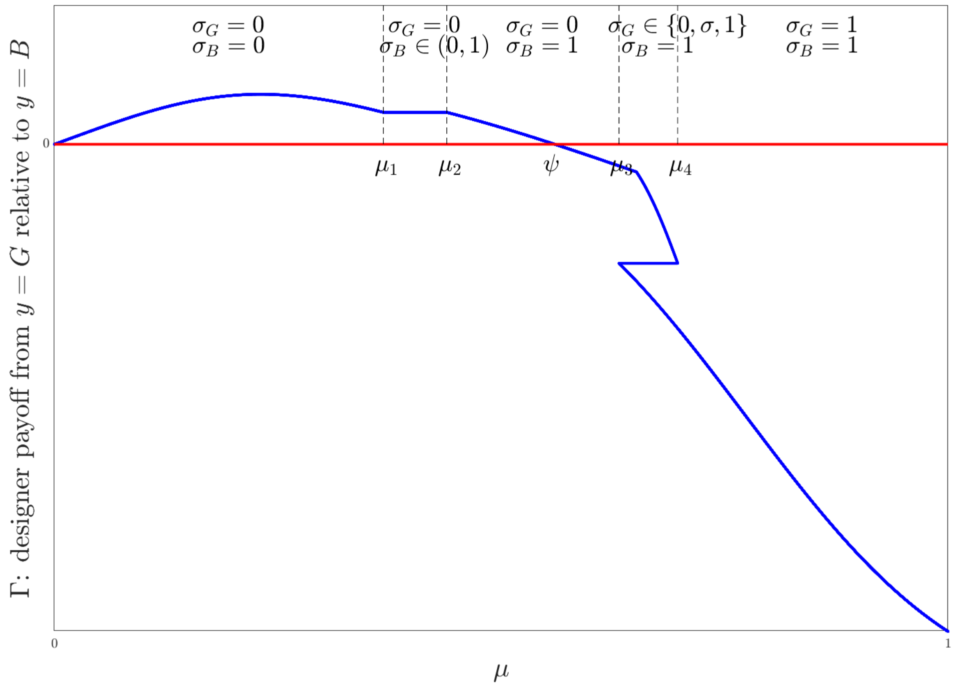

Figure 4: This figure illustrates how the designer's net payoff from choosing the G-product varies with the consumer's belief μ.

since consumer becomes more optimistic and is thus more likely to disregard information and accept all products, regardless of their output. Eventually, as μ becomes sufficiently large, the designer obtains a higher expected payoff from a B-product compared to a G-product. In the next section, we explore the properties of equilibrium using comparative statics.

4 Comparative Statics

In this section, we study the comparative statics of our model, which we later interpret in the context of two applications (Section 5). We say that a product's expected quality increases when the probability of a product having a good output, $\mu = \mathbb{P}(y = G)$, increases. We say that a product's expected complexity increases when the probability that a product is complexified, $\mathbb{P}(\kappa = \bar{\kappa}) = \mu\sigma_G + (1 - \mu)\sigma_B$, increases, which is natural since an increase in $\mathbb{P}(\kappa = \bar{\kappa})$ is equivalent to an increase in $\mathbb{E}\{z\}$.

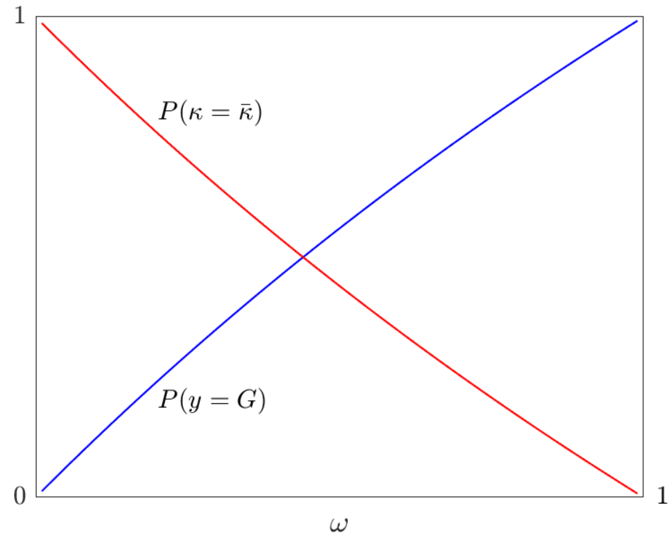

Figure 5: This figure illustrates how expected product quality and complexity vary with the consumer's relative outside option.

We begin by considering the effect of a decrease in the consumer's relative outside option. A decrease in ω could result from a decrease in the consumer's outside option, w_0, or from an increase in the consumer's payoff when the product is accepted, i.e., either $w(G)$ or $w(B)$. Thus, a decrease in ω increases the net payoff of the product to the consumer.

Proposition 6 *As ω decreases, μ decreases while σ_G and σ_B do not change. Thus, a decrease in the consumer's relative outside option decreases expected quality and increases expected complexity.*

Figure 5 illustrates the effect of a change in the consumer's relative outside option on expected product quality and complexity. Intuitively, as ω decreases, the consumer becomes less selective in accepting a product. As this makes the B-product more likely to be accepted, the designer has a lower incentive to produce a G-product, resulting in lower expected product quality. The designer's equilibrium complexification strategy, however, is independent of ω. While the decrease in the relative outside option pushes the consumer to be less selective, this is fully offset by the resulting decrease in expected product quality that makes the consumer more selective. As a result, the consumer's acceptance strategy remains unchanged, together

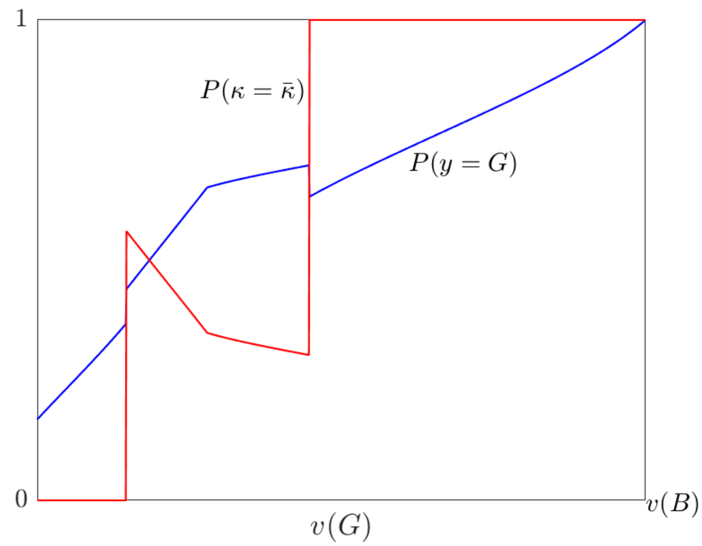

Figure 6: This figure illustrates how expected product quality and complexity vary with the designer's payoff from producing a G-product.

with the incentives of the designer to complexify his product. Finally, as complexification is always weakly higher for a B- than a G-product, a decrease in expected product quality results in an increase in expected complexity.

We next analyze the effect of an increase in the designer's alignment with the consumer, which we capture by an increase in the designer's payoff from producing a G-product, $v(G)$.[17]

Proposition 7 *As $v(G)$ increases, μ increases for large changes in $v(G)$, though it may be non-monotonic locally, while σ_G and σ_B increase. Thus, a large increase in the designer's alignment with the consumer increases both expected quality and expected complexity.*

Figure 6 illustrates the effect of a change in the designer's alignment with the consumer on expected product quality and complexity. Intuitively, an increase in $v(G)$ increases the net payoff to the designer from producing a G-product. Therefore, unsurprisingly, expected product quality increases. Since the consumer's prior belief μ must increase as well, she becomes less selective in accepting products, which increases the designer's incentive to complexify. In the figure, we see that (i) for low $v(G)$ all products are simplified, (ii) for intermediate $v(G)$,

[17]Note that the effect of an increase in $v(G)$ is the same as the effect of a decrease in $v(B)$.

the B-products are complexified while the G-products are simplified (implying that expected complexity falls as μ increases), and (iii) for $v(G)$ large all products are complexified.

The results in Propositions 6 and 7 have implications for the relationship between expected product quality and complexity, and highlight the importance of understanding the underlying drivers of product heterogeneity, i.e., what primitives drive the variation of product attributes across time, industries, regulatory environments, etc. We next interpret these results in the context of two concrete applications.

5 Applications

In this section, we consider two applications that are at the center of policy and academic discussion on issues related to complexity: the design of financial products and of regulatory policies. We discuss how to map our framework to each of these applications, and the resulting implications. Even though our model is stylized and does not capture the richness of the institutional details of these environments, we use results from our extensions in Section 6 to address particular issues of interest that arise within each application.

5.1 Financial Products

Banks and other financial intermediaries design financial products that they offer to retail investors, such as savings and retirement accounts, mortgages, credit lines, etc. When a retail investor (i.e., the consumer in our model) approaches a financial advisor (i.e., the product designer), the latter chooses which product to offer to the consumer. In practice, financial advisors may receive different payments from selling a given financial product to investors, and the products that give the financial advisor a higher commission need not be the best ones for the retail investor. When incentives are not aligned, the financial advisor faces a trade-off as the one faced by the designer in our model: to increase the probability of acceptance by offering the product that best suits the needs of the investor, or to increase his payoff conditional on acceptance by choosing the product with a higher commission. In turn, the choice of product design within financial institutions will crucially depend on the financial advisor's ability to sell different types of financial products, such as more or less complex. In what follows, we examine the drivers of financial product quality (e.g., the net present value of a financial product to the investor) and complexity (e.g., the multidimensionality of the contract) through the lens of our model.[18]

Demand for financial products. An increase in the consumer's demand for a given financial product is captured in our model by a decrease in the consumer's relative outside option, ω. As shown in Proposition 6, our model suggests that as retail investors' demand for a given financial product increases, product quality falls while product complexity increases. These predictions are consistent with the trend leading up to the financial crisis of an increasing demand for "safe" financial products, such as mortgage-backed securities, which also exhibited worsening quality and increased complexity.[19]

Competition in financial markets. To capture competition in financial markets, we extend our model to a dynamic search setting in Section 6.1. We find that higher competition, as captured by smaller search frictions, results in an increase in product quality as well as a decrease in complexity. These results suggest that efforts to increase competition among financial services providers may be effective in improving the quality of financial products that the consumers choose while also reducing their complexity. Note that this prediction is in contrast to those in the literature on obfuscation (Spiegler, 2006; Ellison and Ellison, 2009; Carlin, 2009; Ellison and Wolitzky, 2012), where a common finding is that higher competition leads to more obfuscation, as obfuscation effectively increases the consumer's search cost.

Compensation structures. Our model can also be used to examine the role played by the financial advisors' compensation structures. If the compensation of financial advisors is linked

[18] While the net present value of a financial product may also be related to its dimensionality, it is possible for two products to yield the same net present value to an investor but vary in their dimensionality.

[19] For evidence on the increasing demand for safe products, see Bernanke (2005); on the worsening quality of securitized products, see Jaffee et al. (2009); and on the increasing complexity of financial products, see Célérier and Vallée (2017).

to the volume or the characteristics of financial products that they sell, then the advisors may be misaligned with retail investors. In our model, we interpret this friction as the difference between the payoff associated with selling a bad vs. good product, $v(B) - v(G)$. As shown in Proposition 7, as financial advisors become more aligned with the investor, there is an increase in the provision of good financial products. Interestingly, higher alignment could result in more complexity as investors become more optimistic.

Trust in financial advisors. The determinants of misalignment between consumers and financial product designers have been discussed in the context of policy proposals to expand the categories of product designers bound by fiduciary duty to their customers. For instance, the Securities and Exchange Comission (SEC) has increased its efforts to forbid the use of the term "financial advisor" for those managing brokerage accounts (particularly retirement funds) unless the broker has formally accepted a fiduciary duty to act in the investor's best interest.[20] To analyze policies such as this one, in Section 6.3 we extend our framework to introduce designers who are aligned with the consumer. Our results suggest that policies that aim to improve the alignment between financial advisors and retail investors would increase the quality of financial products, together with their complexity.

In our model, retail investors are rational and they do not misperceive the distribution of financial advisors. However, if we allowed for such deviations, an unjustified increase in the trust in financial advisors would actually generate a decrease in product quality accompanied by an increase in complexity. This observation is consistent with the behavior of several financial intermediaries who during the 2000s were allegedly designing and selling increasingly lower quality and more complex financial products to overly optimistic retail investors, a behavior that contributed to the financial crisis and resulted in multiple lawsuits.[21]

[20] *"Fiduciary Rule" Poised for Second Life Under Trump Administration*, article by Dave Michaels on the Wall Street Journal, January 10th, 2018: https://www.wsj.com/articles/fiduciary-rule-poised-for-second-life-under-trump-administration-1515580200.

[21] In 2011, the Federal Housing Finance Agency filed lawsuits against some of the largest US financial institutions, involving allegations of securities law violations and fraud in the packaging and sale of mortage-backed-securities. For a detailed description, see https://www.fhfa.gov/SupervisionRegulation/LegalDocuments/Pages/Litigation.aspx.

5.2 Regulatory Policies

Politicians design and propose policies to achieve a policy agenda. In doing so, they balance their own preferences (e.g., determined by ideology, lobbying) with the need to obtain voter approval for that policy. In this environment, politicians (i.e., product designers) propose policies to the median voter (i.e., the consumer) who may accept or reject them. For instance, if the policy in question is a tax reform, the policy's quality is given by whether it implies "higher or lower taxes" or "more or less redistribution." By contrast, a policy is more likely to be complex if it contains unnecessarily complicated wording or many unlikely contingencies. An illustrative example of such complexity comes from the regulatory framework proposed by the Basel Committee on Banking Supervision. An analysis of its text has shown that an average sentence in the Basel documents consists of 25.7 words, significantly longer than the average 21 words in a sentence of the British National Corpus.[22] Moreover, the second sentence of the very first document published by the Basel Committee on Banking supervision spans over 72 words.[23]

Public opinion. Public opinion data provides politicians with real time information about voters' support. In our model, public opinion is captured by the median voter's belief that the politician in place will propose policies that are in the voter's interest. Through the lens of our model, high public opinion should be associated with more complex policy proposals, and vice-versa. Indeed, legal scholars have argued that in policy domains where public opinion is low (e.g., financial services or pharmaceuticals), policy proposals from the US Congress tend to be simpler, leaving it to federal agencies to draft additional rules (Stiglitz, 2017).

Urgency. Urgency to pass a given policy varies depending on the reform under discussion. For example, there was a strong sense of urgency to pass financial regulation reform after the 2008-09 crisis. One possible reason is that the public would not trust the financial system otherwise. In our model, urgency is captured as the median voter's outside option, a measure

[22] The British National Corpus is a collection of texts covering a broad range of modern British English.

[23] Analysis performed by Neue Zurcher Zeitung, as cited by Marie-Jose Kolly and Jurg Muller, https://www.endofbanking.org/2018/05/22/how-banking-regulation-has-grown-out-of-all-proportions

of the status-quo. Thus, our model suggests that when urgency to pass a reform is higher, policies will tend to be of lower quality and more complex, and vice-versa. This is consistent with the observation that in the U.S., reforms passed in times of urgency, such as the Dodd-Frank Act or the Affordable Care Act, have been described as overly complex.

Direct costs of complex or simple rules. In the context of regulation, it is natural to think that the actions politicians take to complexify or simplify policies may have a direct impact on the voter. On the one hand, actions that are expected to make policies more complex may include redundant requirements or contingencies that come with high costs of compliance. On the other hand, adding provisions or contingencies may be desired when regulating complex systems. In view of this, in Section 6.5 we extend our model to introduce a direct cost to the consumer due to complexification or simplification. While the introduction of such direct costs does not change our qualitative results, we do find unsurprisingly that regulatory complexity should be lower (higher) when it generates a direct cost (benefit) to the median voter.

6 Robustness and Extensions

In this section, we explore several extensions and alternative specifications of our model.

6.1 Sequential Search and Competition

We now study the effects of competition among designers. In particular, we suppose that if the consumer rejects a product, she searches for a new designer whom she finds with probability $\beta \in (0, 1)$. The new designer proposes a product to the consumer, and the game repeats until the consumer accepts an offered product. In this setting, higher β corresponds to lower search frictions and, hence, more fierce competition among designers.

In a stationary equilibrium, in which U denotes the consumer's equilibrium value, we have

$$U = \mathbb{E}\left[\max_{a\in\{0,1\}}\{a \cdot [\mu(s, z) \cdot w(G) + (1 - \mu(s, z)) \cdot w(B)] + (1 - a) \cdot \beta U\}\right]. \tag{15}$$

Note that for $w_0 = \beta U$, the equilibrium is fully characterized in Section 3.4. Thus, the main difference with our baseline model is that the consumer's outside option is now endogenous.

Proposition 8 *An equilibrium exists, and in it $\beta U \in (w(B), w(G))$ provided that β is not too low. Furthermore, βU is increasing in β. Thus, an increase in competition among designers increases expected quality and decreases expected complexity.*

Comparative statics with respect to β are qualitatively similar to those with respect to the relative outside option ω in Proposition 6. This is because, in a search environment, βU is the consumer's effective outside option, which we have shown increases in β. Therefore, competition in our framework has desirable effects for the consumer and for overall efficiency.

6.2 Interpretation through Rational Inattention

We now that our baseline setup can be interpreted through the lens of a rational inattention framework, where the consumer is able to extract less information from a more complex product. We suppose that the consumer can choose how much uncertainty about the product output to reduce, subject to an entropy-reduction cost, where entropy measures the consumer's uncertainty (Sims, 2003). Within this setup, we say that a product is more complex if it has a higher entropy-reduction cost.[24] Although this approach allows for a more flexible information acquisition technology, it has the drawback that we can no longer obtain as sharp of an equilibrium characterization. Nevertheless, we show next that the model's main mechanisms remain robust to this alternative specification. Here, we discuss our results at an informal level, but we refer the reader for details and derivations to Appendix B.2.

In the new setup, the consumer chooses the distribution of her signal about product output subject an entropy-reduction cost, χ, which is higher for a more complex product. Since the consumer's action is binary, i.e., accept or reject, it is without loss of generality to focus on binary signals $S \in \{b, g\}$ (Woodford, 2009; Yang, 2015), where the consumer accepts the

[24] As in our baseline model, we suppose that the mapping between the designer's action, κ, and the product's complexity is imperfect.

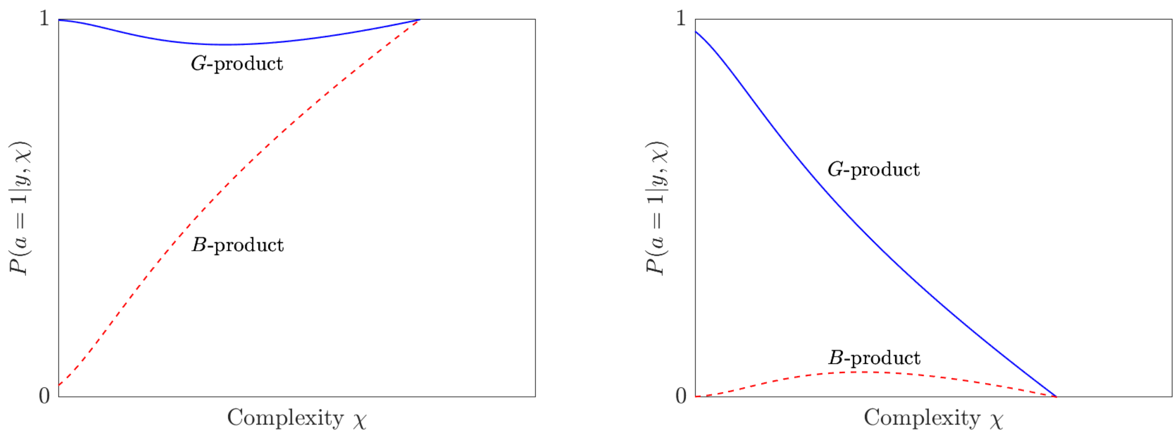

(a) Optimistic Consumer

(b) Pessimistic Consumer

Figure 7: This figure illustrates the probability of acceptance of a y-product as a function of the product's complexity, χ.

product if and only if she receives a g signal. The resulting information structure differs from the one in our baseline setup in two ways. First, the signals are no longer symmetric, as the consumer allocates "precision" optimally between the g and the b signals, trading off the costs of rejecting a G-product (type I error) with the costs of accepting a B-product (type II error). Second, the amount of information that the consumer extracts overall depends not only on the product's complexity but also on the consumer's prior belief μ, which is endogenous.

Equipped with the optimal information structure, we can compute the probability of acceptance of a y-product, as it depends on the entropy-reduction cost χ. These probabilities are depicted in Figure 7, which we can see closely resemble those in our baseline model (see Figure 2). When the product's complexity is low, the consumer extracts an informative signal and makes her decision contingent on its realization. Otherwise, the consumer accepts the production w.p.1 if she is optimistic, and she rejects it w.p.1 if she is pessimistic.

Although a full analytical characterization of the equilibrium set is difficult to obtain, we show (numerically) that the equilibrium set of the model with optimal information extraction resembles closely that of our baseline model. Figure 8 depicts (a) the complexification strategy

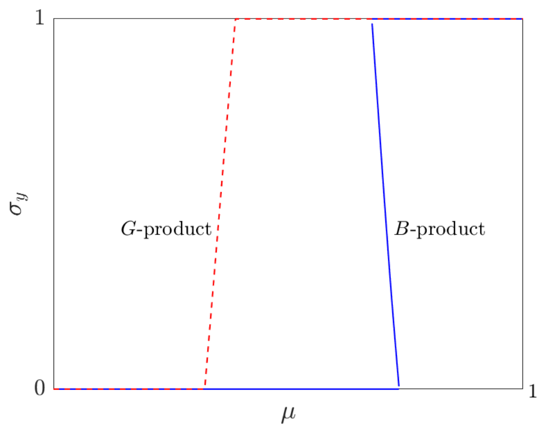

(a) Equilibrium complexification

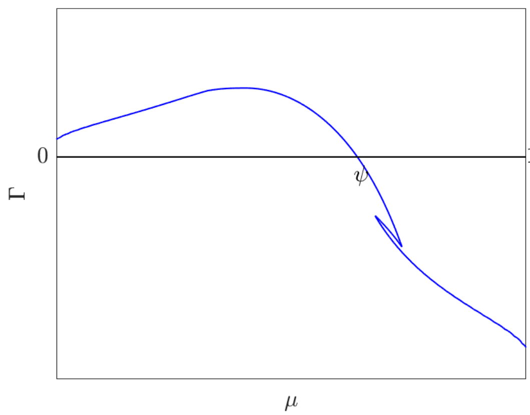

(b) Equilibrium output determination

Figure 8: The left panel illustrates how the complexification strategy of the designer who produces a y-product varies with the prior belief μ. The right panel illustrates the designer's net payoff from choosing the G-product, given belief μ.

$\{\sigma_y\}$ of the designer that is consistent with an equilibrium prior belief μ, and (b) the designer's net payoff from producing a G- vs. a B- product as it depends on μ. As in the baseline model, the right panel determines the expected product quality, $\mu = \psi$, whereas the left panel determines the equilibrium complexification of a y-product, given that $\mu = \psi$.

6.3 Introducing an Aligned Designer

In this section, we study the implications of introducing a designer whose payoffs are *aligned* with the consumer's. We suppose that with probability $q \in [0, 1]$ the consumer encounters the aligned designer, who obtains a higher payoff from having a G-product being accepted, $\bar{v}(G) > \bar{v}(B) > 0$. With probability $1-q$, however, the consumer meets the misaligned designer as in the baseline model. Observe that the case of $q = 0$ corresponds to our baseline model. Finally, whether the designer is aligned or misaligned is not observable to the consumer.

An aligned designer takes private actions $\{y, \kappa\}$ to maximize his expected payoff, $\mathbb{P}(a = 1|y, \kappa) \cdot \bar{v}(y)$. As both his probability of acceptance and payoff conditional on acceptance are higher for a G-product, the aligned designer always produces a G-product. The complexifica-

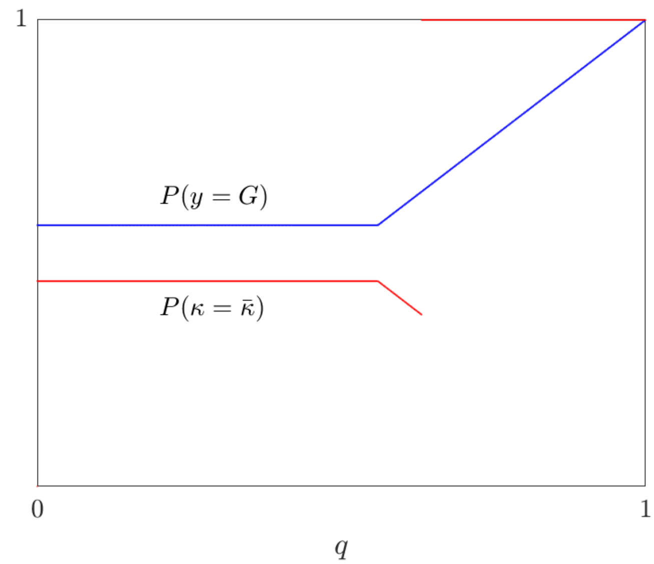

Figure 9: This figure illustrates how expected product quality and complexity vary with the probability that the consumer meets the aligned designer.

tion strategy for the G-product is as in our baseline model given by Proposition 2.

The presence of an aligned designer only affects the equilibrium analyzed in Section 3 through the belief that a G-product is offered, which is now given by $\mu = q + (1 - q) \cdot m$. As before, m is the probability with which the misaligned designer produces a G-product. The following proposition characterizes the main effects of introducing an aligned designer.

Proposition 9 *The presence of an aligned designer need not change the equilibrium outcome unless q is sufficiently large, in which case μ, σ_G and σ_B increase.*

When the probability of meeting an aligned designer is sufficiently small, there is no effect on equilibrium outcomes, as the presence of an aligned designer is fully offset by an increase in the misaligned designer's incentives to produce a B-product. As a result, expected product quality and complexity are as in the baseline model, with $\mu = \psi$. When q is sufficiently large, the misaligned designer only produces a B-product, $m = 0$, and further increases in q lead to higher product quality and, thus, to more complexification (see Proposition 4). These effects are depicted in Figure 9.

6.4 Introducing Prices

For some applications, e.g. financial products, it is natural to assume that a designer not only proposes a product to the consumer but that he also sets a price that is observable to the consumer. To analyze the role of such transfers in our environment, we modify the agents' payoff as follows. If a product is accepted, the designer's payoff is given by the price he charges the consumer minus the cost of production, $p - c(y)$.[25] In turn, the consumer's payoff from accepting a y-product is given by her valuation minus the price she pays, $\widetilde{w}(y) - p$. The following assumption replaces Assumption 1.

Assumption 1' *The payoffs satisfy the following properties:*

1. $\widetilde{w}(G) > w_0 > \widetilde{w}(B) \geq 0$.

2. $c(G) > c(B) \geq 0$.

3. $\widetilde{w}(G) - c(G) > w_0 > \widetilde{w}(B) - c(B)$.

As prices are set by the designer, the consumer makes inferences not only from the product's complexity, z, but also from its price, p; so the consumer's posterior belief is now denoted by $\mu(s, z, p)$. It is easy to see that the consumer will accept the product if and only if her posterior belief is greater than a price-adjusted relative outside option

$$\mu(s, z, p) \geq \frac{\widetilde{w}_0 - \widetilde{w}(B) + p}{\widetilde{w}(G) - \widetilde{w}(B)}. \tag{16}$$

Given the consumer's acceptance strategy, the designer chooses $\{y, \kappa, p\}$ to maximize his expected payoff

$$\mathbb{P}(a = 1|y, \kappa, p) \cdot (p - c(y)). \tag{17}$$

For simplicity, we focus on equilibria in which the designer has a pure strategy over the price. The following proposition summarizes the main results of this section.

[25] We assume that the cost of production is incurred upon product acceptance in order to stay close to the payoff structure of our baseline model.

Proposition 10 *In any positive trade equilibrium, the price set by the designer is independent of the product's output. Moreover, any price* $p^* \in (c(G), \widetilde{w}(G) - w_0)$ *can be supported in equilibrium. The expected product quality and complexity are determined as in the baseline model with payoffs given by* $w(y) \equiv \widetilde{w}(y) - p^*$ *and* $v(y) \equiv p^* - c(y)$*.*

The result that separation through prices is not possible is intuitive. As a B-product is cheaper to produce, the designer of such a product is willing to set any price the G-product designer is willing to set. As a result, the B-product designer always mimics the pricing strategy of a G-product designer in order to avoid being identified. Due to the freedom in specifying off equilibrium beliefs, multiple prices can be supported as an equilibrium. The bounds on possible prices are due to the fact that in any positive trade equilibrium a G-product designer will not post a price below his cost of production, or a price high enough for the product to be rejected w.p.1.

6.5 Direct Costs or Benefits of Complexity

We have assumed that the designer's action to complexify only affects the information that the consumer receives, but that it does not directly enter the consumer's payoff. In some applications, however, these actions can be an important determinant of the product's value to the consumer. For instance, in the context of regulation, complexification may involve introducing provisions with high cost of compliance. Such provisions might nevertheless be beneficial when regulating complex systems. To address these considerations, we now extend the baseline model to incorporate direct costs or benefits of complexification.

We suppose that the consumer's payoff from accepting a (y, κ) product is now given by $w(y) - c(\kappa)$. We make the normalization that $c(\underline{\kappa}) = 0$ but we let $c(\bar{\kappa}) = c$ to be either positive or negative, in order to capture situations in which complexification is costly ($c > 0$) or beneficial ($c < 0$). Observe that the case of $c = 0$ corresponds to our baseline model. We make the following additional assumption on the consumer's payoff.

Assumption 2 $w(G) - c > w_0 > w(B) - c$.

This assumption guarantees that, as in our baseline setting, the consumer wants to accept a G-product and to reject a B-product. Given information (s, z), the consumer's payoff is now

$$a \cdot \mathbb{E}\left[w(y) - c(\kappa)|s, z\right] + (1 - a) \cdot w_0 \tag{18}$$

The following proposition summarizes the effect of direct costs or benefits of complexification in our setting.

Proposition 11 *As c increases, μ increases while σ_G and σ_B do not change. Thus, an increase in the direct cost from complexification increases expected product quality and decreases expected complexity.*

The intuition behind this proposition is as the one described in Section 4, where we analyzed the effect of changes in the relative outside option, ω. This is because an increase in c effectively increases the consumer's relative outside option, since more complex products are now more costly and therefore less attractive to the consumer.

7 Concluding Remarks

In this paper, we explore the incentives of product designers to produce complex products, and the resulting implications for overall product quality. We find that product complexity is not necessarily a feature of low quality products. In particular, complexification or simplification may be used strategically by designers of both high and low quality products. We derive implications for the relationship between quality and complexity of products, and highlight the importance of understanding the underlying drivers of product heterogeneity when conducting empirical work. Our model provides a new rationale for the observed proliferation of complex regulatory policies and financial products. In particular, we argue that high demand for safe assets may have been an important driver of the increasing complexity and worsening quality of structured products. We view these results as complementary to those proposed by the literature on obfuscation and shrouded attributes. In addition, our model sheds light on the increased complexity of regulatory policies, such as Dodd-Frank and Basel, by linking it to the urgency of passing banking regulatory reforms in the years following the financial crisis.

References

Alesina, A. and A. Cukierman (1990). The politics of ambiguity. *The Quarterly Journal of Economics 105*(4), 829–850.

Aragones, E., I. Gilboa, A. Postlewaite, and D. Schmeidler (2005). Fact-free learning. *American Economic Review 95*(5), 1355–1368.

Aragones, E. and A. Postlewaite (2002). Ambiguity in election games. *Review of Economic Design 7*(3), 233–255.

Auster, S. and N. Pavoni (2018). Optimal delegation and limited awareness, with an application to financial intermediation.

Bar-Isaac, H., G. Caruana, and V. Cuñat (2010). Information gathering and marketing. *Journal of Economics & Management Strategy 19*(2), 375–401.

Bergemann, D. and S. Morris (2016). Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review 106*, 586–591.

Bernanke, B. (2005). The global saving glut and the u.s. current account deficit. Speech 77, Board of Governors of the Federal Reserve System (U.S.).

Bond, P. and Y. Zeng (2018). Silence is safest: Non-disclosure when the audience's preferences are uncertain.

Brocas, I. and J. D. Carrillo (2007). Influence through ignorance. *The RAND Journal of Economics 38*(4), 931–947.

Carlin, B. I. (2009). Strategic price complexity in retail financial markets. *Journal of financial Economics 91*(3), 278–287.

Célérier, C. and B. Vallée (2017). Catering to investors through security design: Headline rate and complexity. *The Quarterly Journal of Economics 132*(3), 1469–1508.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica*, 1431–1451.

Davis, S. (2017). Regulatory complexity and policy uncertainty: headwinds of our own making. *Becker Friedman Institute for Research in Economics Working Paper No. 2723980*.

Dewatripont, M. and J. Tirole (2005). Modes of communication. *Journal of Political Economy 113*(6), 1217–1238.

Ederer, F., R. Holden, and M. Meyer (2018). Gaming and strategic opacity in incentive provision. *The RAND Journal of Economics 49*(4), 819–854.

Ellison, G. and S. F. Ellison (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica 77*(2), 427–452.

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

Espinosa, F. and D. Ray (2018, May). Noisy agents. Working Paper 24627, National Bureau of Economic Research.

Gabaix, X. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets*. *The Quarterly Journal of Economics 121*(2), 505–540.

Ghent, A. C., W. N. Torous, and R. I. Valkanov (2019). Complexity in structured finance. *The Review of Economic Studies*.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics 24*(3), 461–483.

Jaffee, D., A. Lynch, M. Richardson, and S. Van Nieuwerburgh (2009). Mortgage origination and securitization in the financial crisis. *Financial Markets, Institutions & Instruments 18*(2), 141–143.

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

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies 76*(4), 1359–1395.

Matthews, S. A. and L. J. Mirman (1983). Equilibrium limit pricing: The effects of private information and stochastic demand. *Econometrica: Journal of the Econometric Society*, 981–996.

McMillan, J. (2014). *The end of banking: money, credit, and the digital revolution.* Zero/One Economics.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 380–391.

Oehmke, M. and A. Zawadowski (2019). The tragedy of complexity.

Pagano, M. and P. Volpin (2012). Securitization, transparency, and liquidity. *The Review of Financial Studies 25*(8), 2417–2453.

Perez-Richet, E. and D. Prady (2011). Complicating to persuade. *SSRN Working Paper*.

Rayo, L. and I. Segal (2010). Optimal information disclosure. *Journal of political Economy 118*(5), 949–987.

Roesler, A.-K. and B. Szentes (2017). Buyer-optimal learning and monopoly pricing. *American Economic Review 107*(7), 2072–80.

Salop, S. (1977). The noisy monopolist: Imperfect information, price dispersion and price discrimination. *The Review of Economic Studies 44*(3), 393–406.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics 50*(3), 665–690.

Spiegler, R. (2006). Competition over agents with boundedly rational expectations. *Theoretical Economics 1*(2), 207–231.

Spiegler, R. (2016). Choice complexity and market competition. *Annual Review of Economics 8*, 1–25.

Stiglitz, E. H. (2017). Delegating for trust. *U. Pa. L. Rev. 166*, 633.

Taneva, I. A. (Forthcoming, 2019). Information design. *AEJ Microeconomics*.

Wiederholt, M. et al. (2010). Rational inattention. *The New Palgrave Dictionary of Economics (Online Edition ed.)*.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics 56*, S100–S124.

Yang, M. (2015). Coordination with flexible information acquisition. *Journal of Economic Theory 158*, 721–738.

Zwiebel, J. (1995). Corporate conservatism and relative compensation. *Journal of Political economy 103*(1), 1–25.

A Proofs for Sections 2-4

We impose the following regularity condition on the distribution of complexity $z \in [0, \frac{1}{2}]$.

Condition A.1 *The function* $\frac{f(z|\underline{\kappa})}{f(z|\bar{\kappa})} \cdot \frac{z}{1-z}$ *is monotonic in* z.

Since $f(z|\kappa) > 0\ \forall z, \kappa$, given a prior belief $\mu \in (0,1)$, the posterior beliefs satisfy the following properties: $\mu(g,0) = 1$, $\mu(b,0) = 0$, and $\mu(g,\frac{1}{2}) = \mu(b,\frac{1}{2}) = \mu(\frac{1}{2}) \in (0,1)$. That is, the signal perfectly reveals output when complexity is zero, $z = 0$, and it is uninformative when complexity is maximal, $z = \frac{1}{2}$. Condition A.1 simply ensures that, for each s, the posterior $\mu(s,z)$ is monotonic in complexity z, as shown next.

Lemma A.1 *Suppose that Condition A.1 holds. Then, in equilibrium, (i) the posterior belief* $\mu(g,z)$ *is increasing in* z*, (ii) the posterior belief* $\mu(b,z)$ *is increasing in* z*, and (iii) the threshold* $\bar{z}$ *defined in Definition 1 is unique.*

Proof. From (4), the posterior belief $\mu(b,z)$ is increasing in z if and only if the likelihood ratio $\frac{\sigma_G f(z|\bar{\kappa})+(1-\sigma_G)f(z|\underline{\kappa})}{\sigma_B f(z|\bar{\kappa})+(1-\sigma_B)f(z|\underline{\kappa})} \cdot \frac{z}{1-z}$ is increasing in z. But the latter follows from MLRP and Condition A.1. Analogously, the posterior belief $\mu(g,z)$ is decreasing in z if and only if the likelihood ratio $\frac{\sigma_B f(z|\bar{\kappa})+(1-\sigma_B)f(z|\underline{\kappa})}{\sigma_G f(z|\bar{\kappa})+(1-\sigma_G)f(z|\underline{\kappa})} \cdot \frac{z}{1-z}$ is increasing in z. But the latter also follows from MLRP and Condition A.1. Finally, the uniqueness of the threshold $\bar{z}$ follows from the monotonicity of the posteriors combined with the facts that $\mu(g,0) = 1$, $\mu(b,0) = 0$, and $\mu(g,\frac{1}{2}) = \mu(b,\frac{1}{2}) = \mu(\frac{1}{2}) \in (0,1)$. ■

We will use the result in Lemma A.1 in the proofs that follow.

Proof of Proposition 1. Note that, in any equilibrium, a G-product must be accepted w.p.1, since perfect communication is available to the designer and the consumer accepts a G-product w.p.1. Next, let z_y denote the complexity chosen by the designer of a y-product, and suppose for contradiction that in equilibrium the designer of a G-product does not perfectly communicate his product output to the consumer, i.e., he chooses $z_G > 0$ with positive probability. Since the product has to be accepted w.p.1, it must be that the consumer accepts it independently of the signal, i.e., $a(b, z_G) = a(g, z_G) = 1$. But then, the designer of a B-product can set $z_B = z_G$ and get his product accepted w.p.1. If in equilibrium, however, both G- and B-products were accepted w.p.1, then the designer would only produce a B-product, since he gets a higher payoff with that product, $v(B) > v(G)$. But then, the consumer would reject all products w.p.1, a contradiction. Thus, in any equilibrium, the designer of a G-product must perfectly communicate his product output to the consumer, i.e., $z_G = 0$. Hence, there does not exist an equilibrium in which the designer produces a B-product with positive probability, since the consumer would know this and would reject such a product w.p.1. ■

Proof of Lemma 1. See text. ■

Proof of Proposition 2. We begin by studying the designer's optimal choice of κ in the case when the consumer is optimistic (see Definition 1).

Case 1 (consumer is optimistic). In this case, the designer's product is accepted w.p.1 when complexity is high enough, $z > \bar{z}$. So, his optimal choice of κ solves:

$$\max_{\kappa \in \{\underline{\kappa}, \bar{\kappa}\}} \int_0^{\bar{z}} P(s = g|y) \cdot f(z|\kappa) dz + \int_{\bar{z}}^{1/2} f(z|\kappa) dz. \tag{19}$$

Thus, it is optimal for the designer of B-product to choose $\bar{\kappa}$ if

$$\int_0^{\bar{z}} z \cdot f(z|\bar{\kappa}) dz + \int_{\bar{z}}^{1/2} f(z|\bar{\kappa}) dz \geq \int_0^{\bar{z}} z \cdot f(z|\underline{\kappa}) dz + \int_{\bar{z}}^{1/2} f(z|\underline{\kappa}) dz, \tag{20}$$

and it is uniquely optimal if the inequality is strict. This is equivalent to:

$$\int_0^{\bar{z}} (1 - z) \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa})) dz \geq 0. \tag{21}$$

But, note that we have:

$$\int_0^{\bar{z}} (1 - z) \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa})) dz > (1 - \bar{z})(F(\bar{z}|\underline{\kappa}) - F(\bar{z}|\bar{\kappa})) > 0. \tag{22}$$

for $\bar{z} > 0$, as will be the case in equilibrium. Thus, condition (20) is satisfied with strict inequality, and it is uniquely optimal for the designer of the B-product to choose $\bar{\kappa}$.

On the other hand, it is optimal for the designer of G-product to choose $\bar{\kappa}$ if

$$\int_0^{\bar{z}} (1 - z) \cdot f(z|\bar{\kappa}) dz + \int_{\bar{z}}^{1/2} f(z|\bar{\kappa}) dz \geq \int_0^{\bar{z}} (1 - z) \cdot f(z|\underline{\kappa}) dz + \int_{\bar{z}}^{1/2} f(z|\underline{\kappa}) dz, \tag{23}$$

and it is uniquely optimal if the inequality is strict. This is equivalent to:

$$\int_0^{\bar{z}} z \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa})) dz \geq 0. \tag{24}$$

Condition (23) is satisfied if $\bar{z} \leq \widehat{z}$, and holds with strictly inequality if $\bar{z} < \widehat{z}$. Thus, if $\bar{z} < \widehat{z}$, it is uniquely optimal for the designer of G-product to choose $\bar{\kappa}$. Otherwise, if $\bar{z} = \widehat{z}$, the designer is indifferent to the choice of κ, and if $\bar{z} > \widehat{z}$, it is uniquely optimal to choose $\underline{\kappa}$.

Next, we study the designer's choice of κ in the case when the consumer is pessimistic.
Case 2 (consumer is pessimistic). In this case, the designer's product is rejected if complexity is too high, $z > \bar{z}$. So, his optimal choice of κ solves:

$$\max_{\kappa \in \{\underline{\kappa}, \bar{\kappa}\}} \int_0^{\bar{z}} P(s = g|y) \cdot f(z|\kappa) dz. \tag{25}$$

Thus, it is optimal for the designer of B-product to choose $\underline{\kappa}$ if

$$\int_0^{\bar{z}} z \cdot f(z|\bar{\kappa}) dz \leq \int_0^{\bar{z}} z \cdot f(z|\underline{\kappa}) dz, \tag{26}$$

and it is uniquely optimal if the inequality is strict. This is equivalent to:

$$\int_0^{\bar{z}} z \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa}))dz \geq 0. \tag{27}$$

Condition (26) is satisfied if $\bar{z} \leq \widehat{z}$, and holds with strict inequality if $\bar{z} < \widehat{z}$. Thus, if $\bar{z} < \widehat{z}$, it is uniquely optimal for the designer of B-product to choose $\underline{\kappa}$. Otherwise, if $\bar{z} = \widehat{z}$, the designer is indifferent to the choice of κ, and if $\bar{z} > \widehat{z}$, it is uniquely optimal to choose $\bar{\kappa}$.

On the other hand, it is optimal for the designer of the G-product to choose $\underline{\kappa}$ if

$$\int_0^{\bar{z}} (1-z) \cdot f(z|\bar{\kappa})dz \leq \int_0^{\bar{z}} (1-z) \cdot f(z|\underline{\kappa})dz, \tag{28}$$

and it is uniquely optimal if the inequality is strict. This is equivalent to:

$$\int_0^{\bar{z}} (1-z) \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa}))dz \geq 0. \tag{29}$$

Re-writing the above condition, we have

$$\int_0^{\bar{z}} (f(z|\underline{\kappa}) - f(z|\bar{\kappa}))dz > \int_0^{\bar{z}} z \cdot (f(z|\underline{\kappa}) - f(z|\bar{\kappa}))dz, \tag{30}$$

which immediately implies that condition (28) is satisfied for all $\bar{z} > 0$, as will be the case in equilibrium, and it is uniquely optimal for the designer of the G-product to choose $\underline{\kappa}$. ■

Proof of Proposition 3. The designer's net benefit from choosing the G-product is:

$$\gamma(\mu) = \max_{\kappa} \mathbb{P}_{\mu}(a=1|G,\kappa) \cdot v(G) - \max_{\kappa} \mathbb{P}_{\mu}(a=1|B,\kappa) \cdot v(B). \tag{31}$$

Since $\mathbb{P}_{\mu}(a=1|G,\kappa) \geq \mathbb{P}_{\mu}(a=1|B,\kappa) > 0$ and $v(G) > v(B)$, we can have $\gamma(\mu) \lesseqgtr 0$. Thus, the designer chooses the G-product whenever $\gamma(\mu) > 0$, the B-product whenever $\gamma(\mu) < 0$, and he is indifferent whenever $\gamma(\mu) = 0$. ■

Proof of Proposition 4. Suppose that, in equilibrium, the consumer's belief that the designer has produced a G-product is $\mu \in (0,1)$.

Pooling on κ*.* Consider first the candidate equilibrium in which $\sigma_B = \sigma_G = 0$. By Proposition 2, this requires that $\mu \leq \omega$ and $\bar{z} \leq \widehat{z}$. On equilibrium path, the consumer does not update upon observing z and, thus, threshold $\bar{z}$ is given by $\mu(g,\bar{z}) = \omega$, which is equivalent to:

$$\bar{z} = \frac{(1-\omega)\cdot\mu}{(1-\omega)\cdot\mu + \omega\cdot(1-\mu)}. \tag{32}$$

This is an equilibrium if and only if $\bar{z} \leq \widehat{z}$, which is equivalent to:

$$\mu \leq \frac{\omega \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega} \equiv \mu_1. \tag{33}$$

Consider next the candidate equilibrium in which $\sigma_B = \sigma_G = 1$. By Proposition 2, this requires that $\mu \geq \omega$ and $\bar{z} \leq \widehat{z}$. On equilibrium path, the consumer does not update upon observing z and, thus, threshold $\bar{z}$ is given by $\mu(b, \bar{z}) = \omega$, which is equivalent to:

$$\bar{z} = \frac{\omega \cdot (1-\mu)}{(1-\omega) \cdot \mu + \omega \cdot (1-\mu)}. \tag{34}$$

This is an equilibrium if and only if $\bar{z} \leq \widehat{z}$, which is equivalent to:

$$\mu \geq \frac{\omega \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega} \equiv \mu_3. \tag{35}$$

Therefore, $\sigma_B = \sigma_G = 0$ is an equilibrium if and only if $\mu \in (0, \mu_1]$, whereas $\sigma_B = \sigma_G = 1$ is an equilibrium if and only if $\mu \in [\mu_3, 1)$.

Separation on κ. Consider the candidate equilibrium in which $\sigma_B = 1$ and $\sigma_G = 0$. There are two cases to consider, depending on whether the consumer is optimistic or pessimistic.

First, suppose that

$$\mu\left(g, \frac{1}{2}\right) = \mu\left(b, \frac{1}{2}\right) = \frac{\mu}{\mu + (1-\mu) \cdot \ell\left(\frac{1}{2}\right)} \leq \omega, \tag{36}$$

where $\ell(\cdot) = \frac{f(\cdot|\bar{\kappa})}{f(\cdot|\underline{\kappa})}$. Then, the consumer is pessimistic (see Definition 1). On equilibrium path, the consumer updates upon observing z, and thus threshold $\bar{z}$ is given by

$$\mu\left(g, \bar{z}\right) = \frac{\mu}{\mu + (1-\mu) \cdot \ell(\bar{z}) \cdot \frac{\bar{z}}{1-\bar{z}}} = \omega. \tag{37}$$

This is an equilibrium if and only if $\bar{z} \geq \widehat{z}$, i.e.

$$\mu_2 \equiv \frac{\omega \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega} \leq \mu \leq \frac{\omega \cdot \ell\left(\frac{1}{2}\right)}{\omega \cdot \ell\left(\frac{1}{2}\right) + 1 - \omega} \equiv \tilde{\mu}. \tag{38}$$

Second, suppose that

$$\mu\left(g, \frac{1}{2}\right) = \mu\left(b, \frac{1}{2}\right) = \frac{\mu}{\mu + (1-\mu) \cdot \ell\left(\frac{1}{2}\right)} > \omega, \tag{39}$$

Then, the consumer is optimistic. The threshold $\bar{z}$ is now given by

$$\mu(b, \bar{z}) = \frac{\mu}{\mu + (1-\mu) \cdot \ell(\bar{z}) \cdot \frac{1-\bar{z}}{\bar{z}}} = \omega. \tag{40}$$

This is an equilibrium if and only if $\bar{z} \geq \widehat{z}$, i.e.

$$\tilde{\mu} = \frac{\omega \cdot \ell\left(\frac{1}{2}\right)}{\omega \cdot \ell\left(\frac{1}{2}\right) + 1 - \omega} < \mu \leq \frac{\omega \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega} \equiv \mu_4. \tag{41}$$

Therefore, $\sigma_B = 0$ and $\sigma_G = 1$ is an equilibrium if and only if $\mu \in [\mu_2, \mu_4]$.

Semi-separation on κ. Consider the candidate equilibrium, in which $\sigma_B \in (0,1)$. By Proposition 2, such an equilibrium requires that the consumer be pessimistic and so $\sigma_G = 0$. On equilibrium path, there is updating from observation of z, and threshold $\bar{z}$ must exactly equal $\widehat{z}$ so that the designer of B-product is indifferent to the choice of κ (Proposition 2) and is willing to mix:

$$\mu\left(g, \widehat{z}\right) = \frac{\mu}{\mu + (1-\mu) \cdot (\sigma_B \cdot \ell(\widehat{z}) + 1 - \sigma_B) \cdot \frac{\widehat{z}}{1-\widehat{z}}} = \omega, \tag{42}$$

which in turn implies that:

$$\sigma_B = \frac{\frac{1-\widehat{z}}{\widehat{z}} \cdot \frac{\mu}{1-\mu} \cdot \frac{1-\omega}{\omega} - 1}{\frac{f(\widehat{z}|\bar{\kappa})}{f(\widehat{z}|\underline{\kappa})} - 1}. \tag{43}$$

Since the posterior belief $\mu\left(g, \widehat{z}\right)$ is continuous and decreasing in σ_B (MLRP implies that $\ell(\widehat{z}) > 1$), this equilibrium exists if and only if:

$$\mu(g, \widehat{z})|_{\sigma_B = 1} < \omega < \mu(g, \widehat{z})|_{\sigma_B = 0}, \tag{44}$$

which is equivalent to:

$$\mu_1 = \frac{\omega \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega} < \mu < \frac{\omega \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega} = \mu_2. \tag{45}$$

Therefore, $\sigma_G = 0$ and $\sigma_B \in (0,1)$ is an equilibrium if and only if $\mu \in (\mu_1, \mu_2)$.

Consider the candidate equilibrium in which $\sigma_G \in (0,1)$. By Proposition 2, such an equilibrium requires that the consumer be optimistic and so $\sigma_B = 1$. On equilibrium path, there is updating from observation of z, and threshold $\bar{z}$ must exactly equal $\widehat{z}$ so that the designer of G-product is indifferent to the choice of κ and is willing to mix::

$$\mu\left(b, \widehat{z}\right) = \frac{\mu}{\mu + (1-\mu) \cdot \frac{1-\widehat{z}}{\widehat{z}} \cdot \frac{1}{\sigma_G + (1-\sigma_G) \cdot \ell(\widehat{z})^{-1}}} = \omega. \tag{46}$$

which in turn implies that

$$\sigma_G = 1 - \frac{1 - \frac{1-\widehat{z}}{\widehat{z}} \cdot \frac{1-\mu}{\mu} \cdot \frac{\omega}{1-\omega}}{1 - \frac{f(\widehat{z}|\underline{\kappa})}{f(\widehat{z}|\bar{\kappa})}}. \tag{47}$$

Since the posterior belief $\mu\left(b, \widehat{z}\right)$ is continuous and increasing in σ_G, this equilibrium exists if and only if:

$$\mu(b, \widehat{z})|_{\sigma_G = 0} < \omega < \mu(b, \widehat{z})|_{\sigma_G = 1}, \tag{48}$$

which is equivalent to:

$$\mu_3 = \frac{\omega \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega} < \mu < \frac{\omega \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega} = \mu_4. \tag{49}$$

Therefore, $\sigma_G \in (0,1)$ and $\sigma_B = 1$ is an equilibrium if and only if $\mu \in (\mu_3, \mu_4)$.

We have thus characterized all the possible equilibrium $\{\sigma_y\}$, as a function of belief μ:

1. If $\mu \in (0, \mu_1]$, then $\sigma_G = \sigma_B = 0$.
2. If $\mu \in (\mu_1, \mu_2)$, then $\sigma_G = 0$ and $\sigma_B = \frac{\frac{1-\widehat{z}}{\widehat{z}} \cdot \frac{\mu}{1-\mu} \cdot \frac{1-\omega}{\omega} - 1}{\frac{f(\widehat{z}|\bar{\kappa})}{f(\widehat{z}|\underline{\kappa})} - 1}$.
3. If $\mu \in [\mu_2, \mu_3]$, then $\sigma_G = 0$ and $\sigma_B = 1$.
4. If $\mu \in (\mu_3, \mu_4)$, then $\sigma_G \in \left\{0,\ 1 - \frac{1 - \frac{1-\widehat{z}}{\widehat{z}} \cdot \frac{1-\mu}{\mu} \cdot \frac{\omega}{1-\omega}}{1 - \frac{f(\widehat{z}|\underline{\kappa})}{f(\widehat{z}|\bar{\kappa})}}\ , 1\right\}$ and $\sigma_B = 1$.
5. If $\mu \in [\mu_4, 1]$, then $\sigma_G = \sigma_B = 1$.

This establishes the stated result. ∎

Proof of Proposition 5. The designer's net expected payoff from choosing G-product relative to B-product is defined in (14). Define the correspondence $\Gamma : [0, 1] \to 2^{\mathbb{R}}$, where $\Gamma(\mu)$ is the set of designer net payoffs $\gamma(\mu)$ implied by all the possible $\{\sigma_y\}$, given in Proposition 4, which are consistent with the consumer's prior belief μ, where we now make explicit the dependence of the net payoff on the consumer's belief μ.

First, note that $0 \in \Gamma(0)$, and that $\Gamma(0)$ is a singleton, since the product is rejected w.p.1 when the consumer's belief is $\mu = 0$. Second, consider $\mu \in (0, 1]$. Note that $\Gamma(\mu)$ is a singleton for $\mu \notin (\mu_3, \mu_4)$, since $\{\sigma_y\}$ corresponding to such μ are unique. On the other hand, $\Gamma(\mu)$ consists of three elements if $\mu \in (\mu_3, \mu_4)$, since then either (i) $\sigma_G = 0$ and $\sigma_B = 1$, (ii) $\sigma_G \in (0, 1)$ and $\sigma_B = 1$, or (iii) $\sigma_G = 1$ and $\sigma_B = 1$. We consider each case next.

Case $\mu \in (0, \mu_1]$. By Proposition 4, equilibrium must have $\sigma_G = \sigma_B = 0$, and it must be that the consumer is pessimistic, since $\mu_1 < \tilde{\mu}$. Furthermore, $\Gamma(\mu)$ is a singleton with:

$$\gamma(\mu) = v(G) \cdot \int_0^{\bar{z}(\mu)} (1 - z) f(z|\underline{\kappa})\, dz - v(B) \cdot \int_0^{\bar{z}(\mu)} z f(z|\underline{\kappa})\, dz, \tag{50}$$

since the product is rejected whenever $z > \bar{z}(\mu)$. Therefore:

$$\gamma'(\mu) = [v(G) - (v(G) + v(B)) \cdot \bar{z}(\mu)] \cdot f(z|\underline{\kappa}) \cdot \frac{d\bar{z}(\mu)}{d\mu}. \tag{51}$$

where $\bar{z}(\mu)$ is given by (32) and, thus, satisfies $\frac{d\bar{z}(\mu)}{d\mu} > 0$, $\bar{z}(0) = 0$, and $\bar{z}(\mu_1) = \widehat{z}$. As a result, for μ sufficiently small, $\gamma'(\mu) > 0$ and $\gamma(\mu) > 0$. Next, consider the value μ_v such that:

$$\bar{z}(\mu_v) = \frac{v(G)}{v(G) + v(B)} \implies \mu_v \equiv \frac{v(G) \cdot \omega}{v(G) \cdot \omega + v(B) \cdot (1 - \omega)}. \tag{52}$$

If $\mu_v > \mu_1$, then $\gamma'(\mu) > 0\ \forall \mu \in (0, \mu_1)$. Otherwise, $\gamma'(\mu) > 0$ for $\mu \in (0, \mu_v)$ and $\gamma'(\mu) < 0$ for $\mu \in (\mu_v, \mu_1)$.

Case $\mu \in (\mu_1, \mu_2]$. Equilibrium must have $\sigma_G = 0$ and $\sigma_B \in (0,1)$, and it must be that the consumer is pessimistic, since $\mu_2 < \tilde{\mu}$. In this case, $\bar{z}(\mu) = \widehat{z}$ and $\Gamma(\mu)$ is singleton with:

$$\gamma(\mu) = v(G) \cdot \int_0^{\widehat{z}} (1-z) f(z|\underline{\kappa})\, dz - v(B) \cdot \int_0^{\widehat{z}} z f(z|\underline{\kappa})\, dz, \tag{53}$$

since the product is rejected whenever $z > \bar{z}(\mu) = \widehat{z}$. It therefore follows that $\gamma(\mu)$ is constant on interval (μ_1, μ_2) and equal to $\gamma(\mu_1)$.

Case $\mu \in (\mu_2, \mu_3]$. Equilibrium must have $\sigma_G = 0$ and $\sigma_B = 1$. The consumer is pessimistic if $\mu < \tilde{\mu}$, and she is optimistic otherwise. Here, again $\Gamma(\mu)$ is a singleton.

Suppose that $\mu < \tilde{\mu}$. Then, the consumer is still pessimistic, and we have:

$$\gamma(\mu) = v(G) \cdot \int_0^{\bar{z}(\mu)} (1-z) f(z|\underline{\kappa})\, dz - v(B) \cdot \int_0^{\bar{z}(\mu)} z f(z|\bar{\kappa})\, dz, \tag{54}$$

since the product is rejected whenever $z > \bar{z}(\mu)$. Therefore:

$$\gamma'(\mu) = [v(G) \cdot (1-\bar{z}(\mu)) \cdot f(\bar{z}(\mu)|\underline{\kappa}) - v(B) \cdot \bar{z}(\mu) \cdot f(\bar{z}(\mu)|\bar{\kappa})] \cdot \frac{d\bar{z}}{d\mu}. \tag{55}$$

where $\bar{z}(\mu)$ is now given by (37). Since $\bar{z}(\mu_2) = \widehat{z}$, $\gamma(\mu)$ is contiuous at μ_2. Furthermore, $\gamma'(\mu) \geq 0$ iff

$$\bar{z}(\mu) \leq \frac{v(G)}{v(G) + v(B) \cdot \ell(\bar{z}(\mu))} \iff \frac{(1-\mu)\cdot\omega}{\mu\cdot(1-\omega)} \geq \frac{v(B)}{v(G)} \iff \mu \leq \mu_v, \tag{56}$$

with strict inequalities iff $\mu < \mu_v$. Furthermore, since $\frac{(1-\mu)\cdot\omega}{\mu\cdot(1-\omega)} \geq \frac{v(B)}{v(G)}$ is decreasing in μ and equal to $\ell(\frac{1}{2}) < 1$ when $\mu = \tilde{\mu}$, it follows that $\mu_v < \tilde{\mu}$ and thus $\gamma'(\tilde{\mu}) < 0$.

Suppose that $\mu > \tilde{\mu}$. Now, the consumer is optimistic and $\bar{z}(\mu) \geq \widehat{z}$ is given by (40), with $\frac{d\bar{z}(\mu)}{d\mu} < 0$. Therefore, we have:

$$\gamma(\mu) = v(G) \cdot \left[\int_0^{\bar{z}(\mu)} (1-z) f(z|\underline{\kappa})\, dz + 1 - F(\bar{z}|\underline{\kappa})\right] - v(B) \cdot \left[\int_0^{\bar{z}(\mu)} z f(z|\bar{\kappa})\, dz + 1 - F(\bar{z}|\bar{\kappa})\right],$$

since the product is now accepted whenever $z > \bar{z}(\mu)$. Thus:

$$\gamma'(\mu) = [v(B) \cdot (1-\bar{z}(\mu)) \cdot f(\bar{z}(\mu)|\bar{\kappa}) - v(G) \cdot \bar{z}(\mu) \cdot f(\bar{z}(\mu)|\underline{\kappa})] \cdot \frac{d\bar{z}}{d\mu} < 0, \tag{57}$$

where the inequality follows from the observation that $\ell(\bar{z}(\mu)) \geq \ell(\widehat{z}) > 1$. Recall that $\tilde{\mu}$ is the threshold between the region where the consumer is pessimistic and the region where she is optimistic. Since $\bar{z}(\tilde{\mu}) = \frac{1}{2}$ and thus $F(\bar{z}|\kappa) = 1$ for $\kappa \in \{\underline{\kappa}, \bar{\kappa}\}$, it is easy to check that $\gamma(\mu)$ is continuous at $\tilde{\mu}$.

Case $\mu \in [\mu_4, 1]$. Equilibrium must have $\sigma_G = \sigma_B = 1$, and it must be that the consumer is optimistic, since $\mu_4 > \tilde{\mu}$. Here, again $\Gamma(\mu)$ is a singleton. Moreover, $\bar{z}(\mu)$ given by (34),

$\frac{d\bar{z}(\mu)}{d\mu} < 0$, and:

$$\gamma\left(\mu\right) = v(G) \cdot \left[\int_0^{\bar{z}(\mu)} \left(1-z\right) f\left(z|\bar{\kappa}\right) dz + 1 - F\left(\bar{z}|\bar{\kappa}\right)\right] - v(B) \cdot \left[\int_0^{\bar{z}(\mu)} z f\left(z|\bar{\kappa}\right) dz + 1 - F\left(\bar{z}|\bar{\kappa}\right)\right],$$

since the product is accepted whenever $z > \bar{z}(\mu)$. It follows that $\gamma(\mu)$ is decreasing in μ since:

$$\gamma'\left(\mu\right) = \left[v(B)\cdot(1-\bar{z}) - v(G)\cdot\bar{z}\right]\cdot f\left(z|\bar{\kappa}\right)\cdot\frac{d\bar{z}}{d\mu} < 0.$$

Finally, note that $\gamma(1) = v(G) - v(B) < 0$.

Case $\mu \in (\mu_3, \mu_4)$. By Proposition 4, equilibrium must have either: (i) $\sigma_G = 0$ and $\sigma_B = 1$; (ii) $\sigma_G \in (0,1)$ and $\sigma_B = 1$; and (iii) $\sigma_G = \sigma_B = 1$; and it must be that the consumer is optimistic, since $\mu_3 > \tilde{\mu}$. Thus, $\Gamma(\mu)$ consists of three elements, and we let $\gamma^j(\mu) \in \Gamma(\mu)$ for $j \in \{1,2,3\}$ denote the net expected payoff to the L-type of choosing G-product, when the equilibrium $\{\sigma_y\}$ is in region (i), (ii), and (iii) respectively.

We have already shown that the functions $\gamma^1(\mu)$ and $\gamma^3(\mu)$ are decreasing in μ (see Case $\mu \in [\mu_2, \mu_3)$ when $\mu > \tilde{\mu}$, and Case $\mu \in [\mu_4, 1]$). Let us consider $\gamma^2(\mu)$, which is given by:

$$\gamma^2\left(\mu\right) = v(G) - v(G)\cdot\int_0^{\widehat{z}} z f\left(z|\bar{\kappa}\right) dz - v(B)\cdot\int_0^{\widehat{z}} z f\left(z|\bar{\kappa}\right) dz - v(B)\cdot\left(1 - F\left(\widehat{z}|\bar{\kappa}\right)\right), \quad (58)$$

and is thus constant on the interval (μ_3, μ_4). Furthermore, it is easy to check that $\lim_{\mu\downarrow\mu_3}\gamma^1(\mu) = \gamma(\mu_3)$ where $\gamma(\mu_3)$ is defined in Case $\mu \in (\mu_2, \mu_3]$, $\lim_{\mu\uparrow\mu_4}\gamma^1(\mu) = \lim_{\mu\uparrow\mu_4}\gamma^2(\mu) = \lim_{\mu\downarrow\mu_3}\gamma^2(\mu) = \lim_{\mu\downarrow\mu_3}\gamma^3(\mu)$, and $\lim_{\mu\uparrow\mu_4}\gamma^3(\mu) = \gamma(\mu_4)$ where $\gamma(\mu_4)$ is defined in Case $\mu \in [\mu_4, 1]$.

Therefore, we have shown that (i) $\Gamma(\mu)$ is a singleton for $\mu \in [0, \mu_3]$, with $\lim_{\mu\to 0}\gamma(\mu) = \gamma(0) = 0$ where $\gamma(\mu)$ is continuous, and increasing on $[0, \mu_v]$, but decreasing on $[\mu_v, \mu_3]$; (ii) $\Gamma(\mu)$ is a singleton, where $\gamma(\mu)$ continuous and decreasing on $[\mu_4, 1]$, with $\gamma(1) < 1$; (iii) finally, $\Gamma(\mu)$ has three elements on (μ_3, μ_4), where $\{\gamma^j(\mu)\}$ are continuous and (weakly) decreasing, with $\lim_{\mu\downarrow\mu_3}\gamma^1(\mu) = \gamma(\mu_3)$, $\lim_{\mu\uparrow\mu_4}\gamma^1(\mu) = \lim_{\mu\uparrow\mu_4}\gamma^2(\mu)$, $\lim_{\mu\downarrow\mu_3}\gamma^2(\mu) = \lim_{\mu\downarrow\mu_3}\gamma^3(\mu)$, $\lim_{\mu\uparrow\mu_4}\gamma^3(\mu) = \gamma(\mu_4)$, and where $\gamma^2(\mu)$ is constant (see Figure 4 for illustration). Hence, it must be that (generically) there is a unique μ on $(0,1)$, denoted by ψ, such that $0 \in \Gamma(\psi)$. We conclude that there is (generically) a unique positive trade equilibrium. In it, the designer produces the G-product w.p.$\psi \in (0,1)$ and his complexification strategy $\{\sigma_y\}$ is given by Proposition 4, where the consumer's prior belief is $\mu = \psi$. ■

Proof of Proposition 6. By inspection of the designer's net payoff γ from producing a G-product relative to a B-product, we see that μ and ω only affect it through their effect on threshold complexity $\bar{z}$ (see proof of Proposition 5), which determines whether the consumer's acceptance decision is contingent on the signal or not (see Lemma 1). As a result, any change in ω must be fully offset by a corresponding change in $\mu = \psi$ so as to keep the designer indifferent between producing a G- vs. a B-product. It is easy to check that an increase in ω increases the thresholds $\mu_1 - \mu_4$. Therefore, in equilibrium, it must be that ψ increases in ω, but the designer's complexification strategy $\{\sigma_y\}$ does not change. As a result, an increase in ω increases expected quality, and it decreases expected complexity since $\sigma_G \leq \sigma_B$. ■

Proof of Proposition 7. An increase in $v(G)$ affects the equilibrium complexification strategy $\{\sigma_y\}$ only to the extent that it affects the consumer's equilibrium belief μ (see proof of Proposition 4 and note that the thresholds $\mu_1 - \mu_4$ are independent of $v(G)$). Now, consider the designer's net payoff γ from producing a G-product relative to producing a B-product, as given by:

$$\gamma = \max_{\kappa} \mathbb{P}(a = 1|G, \kappa; \{\sigma_y\}, \mu) \cdot v(G) - \max_{\kappa} \mathbb{P}(a = 1|B, \kappa; \{\sigma_y\}, \mu) \cdot v(B), \tag{59}$$

where we now make explicit that the equilibrium probability of acceptance of a y-product, $\mathbb{P}(a = 1|y, \kappa; \{\sigma_y\}, \mu)$, will depend on the equilibrium belief μ and complexification strategy $\{\sigma_y\}$. By our previous argument, for a given belief μ, $\max_\kappa \mathbb{P}(a = 1|y, \kappa; \{\sigma_y\}, \mu)$ is independent of $v(G)$ and therefore γ must be increasing in $v(G)$. Thus, if equilibrium features $\mu = \psi \notin [\mu_3, \mu_4]$, it must be that ψ increases with $v(G)$ (see proof of Proposition 5 and Figure 4). The same holds if $\psi \in [\mu_3, \mu_4]$ and the change in $v(G)$ is large enough so that the new equilibrium ψ is greater than μ_4. In particular, it is easy to check that $\mu = \psi$ goes to 1 as $v(G)$ goes to $v(B)$. However, when $\psi \in (\mu_3, \mu_4)$, it is possible that an increase in $v(G)$ implies that ψ falls as the equilibrium jumps from separation on κ to pooling at $\kappa = \bar{\kappa}$ (see Figure 6). From Proposition 4, since the thresholds $\mu_1 - \mu_4$ are unchanged, as $\mu = \psi$ increases, $\{\sigma_y\}$ increase. As a result, a large enough change in $v(G)$ increases both expected quality and expected complexity, though locally the effect may be non-monotonic. ■

B Derivations and Proofs for Section 6

In this Appendix, we provide derivations and proofs for the extensions to our baseline model.

B.1 Sequential Search and Competition

Proof of Proposition 8. For each $U \in [w(B), w(G)]$, consider the map $T_\beta : U \mapsto \mathbb{R}$ defined by:

$$T_\beta(U) = \mathbb{E}\left\{\max_{a\in\{0,1\}} \{a \cdot (\mu(s,z) \cdot w(G) + (1-\mu(s,z)) \cdot w(B)) + (1-a) \cdot \beta U\}\right\}, \tag{60}$$

where recall $\mu(s, z)$ is the consumer's equilibrium belief that the proposed product has output G, given signal s and complexity z. For an exogenously given value of U, which pins down the consumer's outside option $w_0 = \beta U$, this map gives us the consumer's ex-ante value $T_\beta(U)$. An equilibrium is a fixed point of this map, and we denote it by U^*.

Since, in equilibrium, the consumer receives informative signals with positive probability, it must be that $T_\beta(w(B)) > w(B)$. Since, in equilibrium, B-products are produced with positive probability and the consumer's signals are imperfectly informative, it must be that $T_\beta(w(G)) < w(G)$. Thus, in equilibrium, the outside option $w_0 = \beta U^*$ must satisfy Assumption 1, i.e. $\beta U^* \in (w(B), w(G))$, provided that β is not too low.[26] Therefore, to show that an equilibrium

[26]Otherwise, when β is small enough, then the consumer will accept all products and the equilibrium strategies and payoffs will be independent of β in that region.

exists, it suffices to show that $T_\beta(\cdot)$ is increasing. But note that an increase in the outside option increases the consumer's ex-ante welfare directly and indirectly through its effects on equilibrium μ and $\{\sigma_y\}$. The latter follows from Proposition 6, where we have shown that $\mu = \psi$ increases in the outside option, whereas $\{\sigma_y\}$ are independent of it.

For comparative statics, note that, for a given U, an increase from β to some β' is equivalent to an increase in the consumer's outside option. Thus, it must be that $T_\beta(U) < T_{\beta'}(U)$. The fixed point must therefore be higher at β' than at β, since $T_\beta(\cdot)$ is increasing. If there are multiple fixed points, then the statement holds locally and for the maximal one. ■

B.2 Interpretation through Rational Inattention

In this Appendix, we adjust our setting to allow the consumer to optimally extract information, i.e. reduce her uncertainty about the product's output, subject to an entropy-reduction cost, as in the literature on rational inattention (Sims, 2003). Within this framework, a more complex product is one that is associated with a high entropy-reduction cost for the consumer.

The uncertainty faced by the consumer with belief $\tilde{\mu} = P(y = G)$ is measured by the entropy function

$$H(\tilde{\mu}) = -\left(\tilde{\mu} \cdot log(\tilde{\mu}) + (1 - \tilde{\mu}) \cdot log(1 - \tilde{\mu})\right), \tag{61}$$

which reaches a minimum of zero at $\tilde{\mu} \in \{0, 1\}$ and a maximum of $-log\left(\frac{1}{2}\right)$ at $\tilde{\mu} = \frac{1}{2}$. As before, we let S denote the signal observed by the consumer and s denote its realization. The signal has a distribution conditional on the product's output, $\pi(s|y) \equiv \mathbb{P}(S = s|y)$, which determines the consumer's posterior belief:

$$\tilde{\mu}(s) \equiv P(y = G|s) = \frac{\pi(s|G) \cdot \tilde{\mu}}{\pi(s|G) \cdot \tilde{\mu} + \pi(s|B) \cdot (1 - \tilde{\mu})}. \tag{62}$$

The entropy associated with the posterior belief is $H(\tilde{\mu}(s))$.

We measure the amount of information that the consumer obtains from a particular information structure π as the expected reduction in entropy:

$$I(\pi) = H(\tilde{\mu}) - \int_s H(\tilde{\mu}(s)) \cdot \pi(s) \cdot ds, \tag{63}$$

and we assume that the consumer faces a cost $\chi \cdot I(\pi)$ of entropy-reduction for some $\chi > 0$. Since a higher entropy-reduction cost makes it harder for the consumer to extract information about the product's output, we will refer to χ as the product's complexity.

As in our baseline model, the mapping from the designer's action κ to the product's complexity is imperfect. We suppose that the complexity χ is random and satisfies $\chi \sim F(\chi|\kappa)$ with pdf $f(\chi|\kappa)$ that has full support on $[0, \infty)$ and where $\frac{f(\chi|\bar{\kappa})}{f(\chi|\underline{\kappa})}$ is increasing in χ (MLRP). When complexity is minimal, $\chi \to 0$, it is essentially costless for the consumer to find out the product's output. When complexity is maximal, $\chi \to \infty$, extracting any information about the product's output becomes prohibitively costly.

Since the consumer's action is binary, i.e. she chooses to accept or reject the product, it is without loss of generality to restrict attention to information structures that consist of binary signals $S \in \{b, g\}$ such that the consumer accepts the product if and only if $S = g$ (Woodford, 2009; Yang, 2015). Let π_y denote the probability that the consumer accepts the product, conditional on the designer producing a y-product. Let $\mu(\chi)$ be the consumer's interim belief after observing the product's complexity χ. For a given χ, the consumer's problem is then reduced to choosing π_G and π_B in order to maximize her expected payoff:

$$\mu(\chi) \cdot \pi_G \cdot (w(G) - w_0) + (1 - \mu(\chi)) \cdot \pi_B \cdot (w(B) - w_0) - \chi \cdot I(\pi) \tag{64}$$

where

$$I(\pi) = H(\mu(\chi) \cdot \pi_G + (1 - \mu(\chi)) \cdot \pi_B) - \mu(\chi) \cdot H(\pi_G) - (1 - \mu(\chi)) \cdot H(\pi_B). \tag{65}$$

Figure 7 depicts the solution to the consumer's problem for a given prior belief $\mu \in (0, 1)$, for the case where the consumer's interim belief satisfies $\mu(\chi) = \mu$, i.e. when the equilibrium features pooling on complexification. As we can see, the resulting probabilities of acceptance closely resemble those in our baseline model, as depicted in Figure 2. When complexity is low, the consumer extracts an informative signal and makes her decision contingent on its realization. Otherwise, when complexity is high, the consumer makes her decision solely based on her interim belief (which in this case equals the prior). Finally, observe that when complexity is high enough, then the consumer either accepts the product w.p.1 (left panel) or she rejects it w.p.1 (right panel). As in our baseline model, which of the two scenarios arises depends on whether the consumer is optimistic or pessimistic; that is, what she would do in the absence of an informative signal.[27]

Naturally, an equilibrium requires that the consumer's prior belief μ and her interim belief $\mu(\chi)$ be consistent with the designer's strategy $\{m, \sigma_G, \sigma_B\}$ and Bayes' rule. Although a full analytical characterization of the equilibrium set is now difficult to obtain, we check numerically that the equilibrium set of the model with optimal information extraction resembles closely that of our baseline model. As we discussed in the text, Figure 8 is the analogue of the Figures 3 and 4. As in our baseline model, an equilibrium is found by requiring that the belief μ equals ψ, so that the designer is indifferent to producing a G- or a B-product, and then reading off the equilibrium complexification strategy of the designer from the left panel, given that $\mu = \psi$.

B.3 Introducing an Aligned Designer

Proof of Proposition 9. Note that the aligned designer's net benefit from choosing the G-product is

$$\gamma = \max_{\kappa} \mathbb{P}(a = 1|G, \kappa) \cdot \bar{v}(G) - \max_{\kappa} \mathbb{P}(a = 1|B, \kappa) \cdot \bar{v}(B). \tag{66}$$

Since the probability of acceptance is always higher for a G-product, i.e. $\mathbb{P}(a = 1|G, \kappa) \geq \mathbb{P}(a = 1|B, \kappa) > 0$ for all κ and $\bar{v}(G) > \bar{v}(B)$, the aligned designer produces a G-product w.p.1. In turn, this designer's complexification strategy is given by σ_G, as characterized in Proposition 2. Thus, the presence of an aligned designer will affect the equilibrium outcomes

[27] When the equilibrium features some separation on κ, a regularity condition similar to Condition A.1 needs to be imposed to ensure that there is a unique threshold complexity level above which the consumer either accept w.p.1 or rejects w.p.1.

only by affecting the probability of a G-product being produced, captured by the fact that belief consistency now requires that $\mu = q + (1-q) \cdot m$.

Consider the case of $q < \psi$. First, note that in equilibrium $\mu \geq q$, since $m \geq 0$. Second, note that it must be that $m = 0$, since the misaligned designer's net payoff from producing a G-product relative to a B-product is always negative for $\mu > \psi$ (see proof of Proposition 5). Thus, we have that in equilibrium $\mu = q \geq \psi$ and, thus, expected quality is higher in the presence of an aligned designer. It follows from Proposition 4 that $\{\sigma_y\}$ are higher as well, since the presence of an aligned designer simply increases μ.

Next, consider the case of $q \leq \psi$. We now show that the equilibrium μ and $\{\sigma_y\}$ need not change in the presence of an aligned designer; their presence is simply offset by the misaligned designer producing a G-product with smaller probability. If $q < \mu_3$, then the misaligned designer's payoff from producing a G-product is strictly positive if equilibrium had $\mu = q$ (see proof of Proposition 5 and Figure 4), which is inconsistent with an equilibrium; thus, it must be that the misaligned designer produces a G-product with positive probability $m = \frac{\psi - q}{1-q}$ so that the equilibrium belief is $\mu = \psi$ and he is indifferent to producing a G- vs. B-product. Now, suppose that $q \in [\mu_3, \psi]$. If $\min\{\gamma : \gamma \in \Gamma(\mu_3)\} > 0$, then the equilibrium is as the one described above since the misaligned designer's net payoff from producing a G-product is still strictly positive if equilibrium had $\mu = q$, which cannot be consistent with equilibrium. If, instead, $\min\{\gamma : \gamma \in \Gamma(\mu_3)\} \leq 0$, then multiple equilibria exist. In particular, the equilibrium where the misaligned designer produces a G-product with probability $m = \frac{\psi - q}{1-q}$ (so that $\mu = \psi$) still exists, since $0 \in \Gamma(\psi)$. However, there is also an equilibrium where $\mu = q \in [\mu_3, \psi]$, since there exists a $\gamma < 0$ such that $\gamma \in \Gamma(q)$. ■

B.4 Introducing Prices

Proof of Proposition 10. Consider an equilibrium in which the y-product has price p_y, with $p_G \neq p_B$. It is straightforward that the consumer would then reject all products with price p_B, and thus the B-product designer would expect to make profits by deviating to price p_G. Thus, in equilibrium, different output products cannot be offered at different prices. Consider a candidate equilibrium in which p^* is the price set by the designer, with the off-equilibrium belief that the designer has produced a B-product if he sets any other price. Then, deviations by the designer to any other price have a payoff of zero, as a product with price different than p^* is rejected w.p.1. Also, note that $p^* < c(G)$ cannot be an equilibrium, as the G-product designer would make losses. Thus, for $p^* < c(G)$, it must be that $\mu = 0$ as only B-products are produced and rejected w.p.1, i.e. there is no trade. Similarly, any price $p^* > \widetilde{w}(G) - w_0$ would induce a rejection w.p.1 by the consumer, as the product generates losses to the consumer, i.e. again there is no trade. For any price $p^* \in (c(G), \widetilde{w}(G) - w_0)$, define payoffs $w(y) \equiv \widetilde{w}_y - p^*$ and $v(y) \equiv p^* - c(y)$, and note that they satisfy Assumption 1. As a result, a positive trade equilibrium exists by Proposition 5. ■

B.5 Direct Costs or Benefits of Complexity

Proof of Proposition 11. Fix the designer's equilibrium strategy, then the consumer's expected payoff conditional on signal s and complexity z is now given by:

$$a \cdot \{\mu(s,z) \cdot [w(G) - \sigma_G(z) \cdot c] + (1 - \mu(s,z)) \cdot [w(B) - \sigma_B(z) \cdot c]\} + (1 - a) \cdot w_0, \quad (67)$$

where

$$\sigma_y(z) \equiv \mathbb{P}(\kappa = \bar{\kappa}|y, z) = \frac{\sigma_y f(z|\bar{\kappa})}{\sigma_y f(z|\bar{\kappa}) + (1 - \sigma_y) f(z|\underline{\kappa})}, \quad (68)$$

and $\sigma_y = \mathbb{P}(\kappa_y = \bar{\kappa}|y)$. Thus, the consumer's relative outside option now depends on the equilibrium complexification strategy of the designer. In particular, the consumer accepts the product if and only if her posterior belief exceeds an adjusted relative outside option:

$$\mu(s,z) \geq \frac{w_0 - w(B) + \sigma_B(z) \cdot c}{w(G) - w(B) + (\sigma_B(z) - \sigma_G(z)) \cdot c} \equiv \omega(z; \sigma_G, \sigma_B). \quad (69)$$

The left-hand side of condition (69) is the consumer's posterior belief, which is the same as in (4). The right-hand side is the relative outside option, which is now adjusted to also incorporate the fact that complexification by the designer is now directly costly (if $c > 0$) or beneficial (if $c < 0$) to the consumer.

The construction of equilibria is analogous to our baseline model, with the exception that the threshold complexity $\bar{z}$, defining whether the consumer's acceptance decision is contingent on the realization of the signal, and thus the thresholds for prior beliefs $\mu_1 - \mu_4$, defining the equilibrium complexity, must now be adjusted. Letting $\widehat{z}$ and $\omega(z; \sigma_G, \sigma_B)$ be defined as in Proposition 2 and equation (69) respectively, and following the steps in the proof of Proposition 4, we can show that:

1. If $\mu \leq \mu_1$, then $\sigma_G = \sigma_B = 0$, where

$$\mu_1 = \frac{\omega(\cdot; 0, 0) \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega(\cdot; 0, 0) \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega(\cdot; 0, 0)}. \quad (70)$$

2. If $\mu \in (\mu_1, \mu_2)$, then $\sigma_G = 0$ and $\sigma_B \in (0, 1)$, where

$$\mu_2 = \frac{\omega(\cdot; 0, 1) \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}}}{\omega(\cdot; 0, 1) \cdot \ell(\widehat{z}) \cdot \frac{\widehat{z}}{1-\widehat{z}} + 1 - \omega(\cdot; 0, 1)}, \quad (71)$$

   and where σ_B solves

$$\frac{\mu}{\mu + (1 - \mu) \cdot (\sigma_B \cdot \ell(\widehat{z}) + 1 - \sigma_B) \cdot \frac{\widehat{z}}{1-\widehat{z}}} = \omega(\widehat{z}, 0, \sigma_B). \quad (72)$$

3. If $\mu \in [\mu_2, \mu_4]$, then $\sigma_G = 0$ and $\sigma_B = 1$, where

$$\mu_4 = \frac{\omega(\cdot; 0, 1) \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega(\cdot; 0, 1) \cdot \ell(\widehat{z}) \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega(\cdot; 0, 1)}. \quad (73)$$

4. If $\mu \in (\mu_3, \mu_4)$, then $\sigma_G \in \{0, \tilde{\sigma}, 1\}$ and $\sigma_B = 1$, where

$$\mu_3 = \frac{\omega(\cdot; 1, 1) \cdot \frac{1-\widehat{z}}{\widehat{z}}}{\omega(\cdot; 1, 1) \cdot \frac{1-\widehat{z}}{\widehat{z}} + 1 - \omega(\cdot; 1, 1)} \tag{74}$$

and where $\tilde{\sigma}$ solves

$$\frac{\mu}{\mu + (1-\mu) \cdot \frac{\tilde{\sigma} \cdot \ell(\widehat{z})}{\tilde{\sigma} \cdot \ell(\widehat{z}) + 1 - \tilde{\sigma}} \cdot \frac{1-\widehat{z}}{\widehat{z}}} = \omega(\widehat{z}; \tilde{\sigma}, 1). \tag{75}$$

5. If $\mu \geq \mu_4$, then $\sigma_G = \sigma_B = 1$.

Finally, observe that any change in c is equivalent to a change in the relative outside option, which we already argued only affects the equilibrium μ but leaves the equilibrium complexification strategy $\{\sigma_y\}$ unchanged (see Proposition 6). And, it is clear that a higher c implies a higher relative outside option and therefore higher equilibrium μ. Thus, as complexification becomes more costly (beneficial), the expected product quality increases (decreases), while expected product complexity decreases (increases). ■