

Attention Holdup^{*}

Francesco Fabbri[†]

August 31, 2025

Abstract

A novel form of the holdup problem emerges in market interactions where a consumer invests in learning about product quality set by a producer. Flexible information acquisition, modeled using rational inattention, mitigates this holdup problem by enabling positive quality provision: in equilibrium, the producer randomizes between zero and a higher-than-price quality. A trade-off arises. When attention costs decrease, high quality rises as the consumer learns about quality more precisely. However, by acquiring better information, the consumer loses commitment power and needs to trade less often to discourage producer deviations. When attention costs vanish, the holdup becomes inevitable: to deter deviations, the consumer never trades, causing the market to fail. Following this trade-off, high information or production costs and low prices increase profits by providing commitment benefits that enhance trade efficiency. A new refinement, preventing information from being acquired for free, uniquely selects this binary-quality equilibrium.

KEYWORDS: holdup problem, rational inattention, quality provision.

JEL CODES: D11, D21, D42, D83.

^{*}I am indebted to Pietro Ortoleva, Faruk Gul, and Wolfgang Pesendorfer for their guidance and encouragement throughout this project. I am grateful to Pierpaolo Battigalli, Roland Benabou, Simone Cerreia-Vioglio, Sylvain Chassang, Daniel Chen, Mark Dean, Laura Doval, Agata Farina, Mira Frick, Yannai Gonczarowski, Ryota Iijima, Zi Yang Kang, Alessandro Lizzeri, Leonardo Maltoni, Massimo Marinacci, Moritz Meyer-ter-Vehn, Sofia Moroni, Paola Moscarrello, Hervé Moulin, Xiaosheng Mu, Jacopo Peregó, Giulio Principi, Doron Ravid, Michael Richter, Fedor Sandomirskiy, Carol Shou, Ali Shourideh, Erik Snowberg, João Thereze, Can Urgan, Maren Vairo, Quitzé Valenzuela-Stookey, Alexander Wolitsky, Darwin Yang, Leeat Yariv, and Qiyuan Zheng for their insightful comments. I also thank seminar participants at the 2025 BFI Conference on Rational Inattention, Bocconi University, Boston University, Columbia University, the CSEF Capri, in Theory, Georgetown University, Kellogg School of Management, and Princeton University.

[†]University of California, Berkeley, Department of Economics, ffabbri@berkeley.edu.

1 Introduction

Relation-specific investments enable opportunistic behavior, giving rise to a key issue: the holdup problem. Once investment costs are sunk, investors become vulnerable to opportunistic actions by other parties who have no incentives to compensate them for their costs. No compensation leads to underinvestments and inefficiencies. As highlighted by a broad body of literature (Grout, 1984; Grossman and Hart, 1986; Aghion et al., 1994), contracts play a crucial role in disciplining opportunistic behavior. However, contracting does not always suffice, for instance, when investments are difficult to observe due to their complexity (Segal, 1999; Battigalli and Maggi, 2002; Tirole, 2009). In such cases, parties may need to process information to understand the investment decisions made by their counterparts. When information processing is costly, a second holdup problem emerges concerning the investment in information. This is the object of our analysis.

To exemplify, consider the following situation. A researcher invests in the quality of her research before seeking an academic position. Research quality is difficult to observe. While it can be communicated by the researcher or her advisers, research institutions may still want to process further information before hiring, for instance, by examining the candidate's work. Information acquisition is itself a costly investment that research institutions undertake only if, on average, it compensates its costs by improving hiring decisions. With little incentive to process this information, research quality, being costly to improve and hard to assess, would decline. In particular, without processing information, quality may fall to a level that makes hiring nonviable, causing this market to fail.

In the context of a quality provision game, we show that when information is acquired flexibly, information investments partially compensate for the failure of contractual commitment. This occurs since flexible information simultaneously provides the right incentives to invest in quality and information. But, these incentives depend significantly on processing information costs. When processing costs decrease, not only is quality learned more precisely, but learning occurs in a way that forces only higher qualities to be produced. We refer to this as the *learning effect*. However, as the consumer acquires better information, she is subject to the opportunistic behavior of the producer, who can deviate by producing lower quality levels. To discourage the producer's deviations, the consumer participates in the market less often. This is the *participation effect*. As a consequence of lower participation, high qualities are produced less frequently.

The interplay between the learning and participation effect becomes particularly

stark when information costs vanish. In this case, the learning effect forces the seller to produce the maximum feasible quality. However, at this quality, the producer's incentives to deviate to lower quality levels are very high; by the participation effect, the consumer prevents these deviations by never engaging in trade. As a result, the market fails. Therefore, while information investments can be held up when costs are too high, causing market failure, approximating the case in which these investments are free is also problematic as information can be exploited to force unrealistically high-quality levels. The takeaway is that the market as a whole benefits from moderate information frictions.

We study the holdup problem in information investments by incorporating rational inattention into a quality provision game. A producer invests in the quality of its product before offering it to a unit-demand consumer at a fixed price. Production costs are linear in quality and sunk at the time of trade. The consumer learns about quality by selecting an information structure. Following Sims (2003) and the related literature, we assume that the information processing cost, which we interpret as an attention cost, is proportional to the average entropy reduction of consumer beliefs. Once information is realized, the consumer decides whether to purchase, and the game ends.

A known issue arises when entropic costs are used in strategic settings. Because entropy reduction is unaffected by zero-probability events, Nash equilibrium does not restrict consumer's off-equilibrium behavior. To fix this problem, we adopt Ravid's (2020) *credibility* notion, a trembling-hand-like refinement selecting best responses in which the consumer treats every quality as if it occurred in equilibrium. We identify an additional equilibrium with deterministic quality that, although credible, is unappealing as it relies on entropy becoming degenerate when the consumer beliefs' are deterministic. To exclude it, we introduce *Shannon best responses*, which complement credibility by requiring attention strategies to remain optimal to arbitrarily small information costs. By Proposition 1, Shannon best responses ensure positive entropic costs whenever the consumer is paying attention. We call the related equilibrium concept *Shannon-Nash*.

We focus on Shannon-Nash equilibria with *positive quality provision*, where the producer provides positive quality with positive probability. Notice that market failure is always possible in our setting: there is an equilibrium with *pessimistic expectations*, where the consumer believes quality is too low to pay attention and does not buy and, as a consequence, the producer sets a low quality to save on production costs.

We highlight two relevant features of Shannon-Nash equilibria with positive quality: *inefficient trade* and *binary quality*. In particular, the producer randomizes between zero

and a higher-than-price quality. Trade is inefficient since the consumer cannot accept the producer's offer for sure without accepting every offer for sure. But if the consumer always accepts, the producer deviates to the lowest quality possible. Furthermore, quality provision must be random, with some qualities below and above the price level: when best responses are Shannon, the consumer never wastes attention if it does not affect her purchasing decision. But without paying attention, the producer would again deviate to lower qualities. Quality provision is binary since, to prevent deviations, the linear isoprofit curve must lie above the consumer's best response, which is s-shaped.

Our main result shows that flexible information acquisition partially compensates for the failure of contractual commitment if the consumer is sufficiently attentive. Specifically, a unique Shannon-Nash equilibrium with positive quality provision exists if and only if the parameter that governs the attention cost, λ , is lower than a threshold, $\bar{\lambda}$. To build intuition, as attention costs increase, the consumer becomes less attentive and makes mistakes more often. By the learning effect, this implies that the high-quality level decreases. In particular, as $\lambda \rightarrow \bar{\lambda}$, high quality converges to its minimum level, the price. To sustain trade in this region, the producer provides high quality with a frequency close to one, approximating the perfect information outcome. By the participation effect, the consumer trades with a high frequency. However, at $\lambda = \bar{\lambda}$, the market fails: the producer is indifferent between setting quality equal to zero or the price, providing no incentives for the consumer to pay attention. When $\lambda \geq \bar{\lambda}$, the consumer makes mistakes too frequently to discourage the producer from deviating to quality levels that are lower than the price. Therefore, the attention investment is no longer held up.

As previewed, when the attention cost vanishes, the market fails since the consumer no longer trades. If attention is free, $\lambda = 0$, there are two Pareto-ranked Shannon-Nash equilibria: the pessimistic expectations equilibrium and the *perfect information* equilibrium—the consumer pays full attention for free, and the producer sets quality equal to the price. In the limit, the pessimistic expectations equilibrium is selected, the worst of the two. This outcome occurs not because the quality of information declines, for instance, as in Ravid et al. (2022); on the contrary, as attention increases, by the learning effect, the consumer distinguishes among high and low-quality levels with greater precision. It occurs because, since uncertainty is endogenous in our model, the consumer discourages deviations by recurring to the no-trade outside option more frequently, lowering market participation. As $\lambda \rightarrow 0$, participation vanishes and hence the pessimistic expectation equilibrium is selected.

The discussion above already suggests how surplus varies in the attention cost. When costs are low, the consumer makes fewer mistakes—which helps the consumer but hurts the producer—and trade inefficiency is high—which hurts both. For this reason, profits are increasing in the attention cost and converge to zero as $\lambda \rightarrow 0$. The consumer utility, instead, is non-monotonic: it converges to zero at both extremes ($\lambda \rightarrow 0$ or $\lambda \rightarrow \bar{\lambda}$), while it is strictly positive in equilibrium as the consumer extracts information rents. Therefore, if the consumer could manipulate her attention, she would prefer to be more inattentive for values of λ close to zero and less inattentive for values of λ close to $\bar{\lambda}$. The producer, instead, would always prefer to face a more inattentive consumer. As a consequence, a bit of inattention is Pareto-efficient. In particular, increasing inattention benefits the market as a whole when the consumer is particularly attentive, while it benefits only the producer when inattention is already sufficiently high.

Our analysis so far has assumed a fixed price. How does the provision of quality change as the price changes? First, we show that a Shannon-Nash equilibrium with positive quality provision exists if and only if the price, p , is higher than a certain threshold, \bar{p} . To build intuition, when $p = \bar{p}$, we have that $\lambda = \bar{\lambda}$, which implies that the producer is indifferent between providing a quality equal to zero or the price. If $p < \bar{p}$, the attention and production costs become too high to sustain positive quality provision: the producer finds it optimal to set quality to zero, causing the market to break down.

We further show that profits decrease in the price. Intuitively, low prices offer *commitment benefits* for the producer by increasing the average trade probability and strengthening the incentives to provide the high-quality product more frequently. This follows from the incentives in the quality provision game: as $p \rightarrow \bar{p}$, the threshold $\bar{\lambda} \rightarrow \lambda$, implying that the high-quality level is close to the price. Therefore, on the one hand, lower prices induce equilibria with lower high-quality levels and higher average trade probability, increasing profits. On the other, a decrease in price lowers the producer's gains from trade. We show that the combined positive effects of lower high quality and higher trade efficiency always outweigh the negative effect of decreased gains from trade.

Next, we study what happens when production costs vary. We show that profits are higher when costs are higher—the opposite of what happens in typical markets. The reason is that high production costs serve as a *commitment device* for the producer by strengthening the incentives to produce lower high-quality levels more frequently. In particular, we show that the positive effects of a higher production cost—a lower high-quality level and a higher probability of trade—always overcome the negative effect of making production more expensive.

We conclude by discussing alternative assumptions. First, we consider a variant of our quality provision game in which costs are *not sunk*. We show that the market fails in this case even with flexible information. The intuition is that, as above, sunk costs provide commitment benefits to the producer. Then, we show that the binary-quality structure of our equilibrium is common to other information acquisition environments that satisfy *convex-concave* information—a generalization of normal additive noise. We also show that the investment in information still generates a holdup problem when the consumer has *private information* about her preference intensity for quality. Finally, we drop the assumption of credibility and consider only best responses robust to non-degenerate entropic costs. We show that, due to the failure of a tangency condition, equilibria may be supported by *three* quality levels.

1.1 Related Literature

This paper contributes directly to the literature studying the role of information in markets, in particular in producer-consumer interactions. The key observation that imperfect information creates commitment power is not new and is due to Kessler (1998). Roesler and Szentes (2017) explore the role of imperfect information in a setting in which the seller observes the buyers signal structure before setting the price. In our framework, the producer quality provision needs to incentivize consumer attention; consequently, free learning does not produce commitment benefits. Furthermore, when attention costs vanish, our model selects the no-learning equilibrium, which implies zero surplus for both consumer and producer. This result is reminiscent of Ravid et al. (2022), which show that to sustain an equilibrium when information costs are positive, the buyer has to ignore a large amount of information, selecting the worst free-learning equilibrium when costs vanish. Our result differs in two ways. First, when information costs vanish, our market fails. Second, our uncertainty is endogenous: the consumer exploits imperfect information to prevent producer deviations but gathers precise information about the equilibrium qualities. As costs vanish, the consumer acquires more precise information about equilibrium qualities, but the market fails since the consumer participates too rarely to sustain positive quality provision.

Our paper fits in a recent literature documenting trade inefficiencies when the consumer observes the producer's offer imperfectly (Ravid, 2020; Denti et al., 2022; Wolitzky, 2023; Cusumano, Fabbri, and Pieroth, 2024). Inefficiencies arise since the producer, lack-

ing commitment power, cannot refrain from making bad offers when the consumer is inattentive. In our equilibrium, while inefficiencies are present, inattention helps trade: a high attention cost disciplines the producer by strengthening the incentives to provide the high quality more frequently. To refine our equilibrium, we adopt the credibility notion by Ravid (2020), which requires the consumer's best response to be robust to small mistakes made by the producer. We complement it by imposing robustness to non-degenerate entropic costs, which selects credible strategies where attention entails positive costs. This refinement is related to the experimental version of entropy-based costs proposed by Denti et al. (2022) as both prevent degenerate beliefs in equilibrium.

We connect to the classical literature on the holdup problem by interpreting information acquisition as a relation-specific and non-contractible investment (Grout, 1984; Grossman and Hart, 1986; Aghion et al., 1994). In Gul (2001) and Lau (2008), the buyer can generate asymmetric information by randomizing over different investment decisions. In both models, the consumer has no information rents and obtains zero utility in equilibrium. Since our framework considers an investment in information, which is by nature stochastic, the consumer does not need to randomize to generate asymmetric information. This approach is shared with Condorelli and Szentes (2020) in which the consumer's investment results in a distribution of valuations for the good. For this reason, as in Condorelli and Szentes, the consumer has information rents and positive utility.

Our findings relate to the large literature studying how costly information acquisition about exogenous traits shapes market outcomes (Matějka, 2016; Martin, 2017; Boyacı and Akçay, 2018; Yang, 2020; Mensch, 2022).¹ Thereze (2024) and Mensch and Ravid (2024) study a model of monopolistic screening and find lower-than-efficient qualities for all types. In our setting, quality is stochastic and may be higher than the efficient level to incentivize information investments. This aspect is absent in their model as the monopolist commits to a menu of goods before learning occurs.

Our producer faces a moral hazard problem in product quality, linking us to the moral hazard literature (Mirrlees, 1976; Holmström, 1979; Georgiadis et al., 2024). Similarly to us, Georgiadis and Szentes (2020) studies a model with flexible information acquisition in which a principal monitors the unobserved effort of an agent by observing the outcome of a diffusion process. The substantive difference is that their principal can commit to a path-contingent stopping rule and a path-contingent wage scheme, while our consumer lacks commitment power and acquires it only through imperfect information.

¹See Maćkowiak et al. (2023) for a detailed literature review.

1.2 Information holdup: an illustration

We conclude the introduction by illustrating in a simple framework the essential ingredients that may generate a holdup problem in markets with information investments. The holdup arises due to the nature of costly information acquisition, which is relation-specific and non-contractible. A formal analysis of this section is in Appendix A.

Consider a simplified version of our quality provision game. A producer sets the quality of its product before offering it to a consumer who cannot observe quality directly but can learn about it by acquiring information. Price is fixed, production costs are increasing in the quality level and sunk at the time of trade, and the information technology is perfect monitoring—the consumer incurs a fixed cost to learn about quality perfectly. Notice that, as in many market interactions, the consumer does not have *commitment power*, and decides whether to buy following her beliefs about quality.

An equilibrium with no trade, which we refer to as the *pessimistic expectations* equilibrium, always exists. The consumer believes quality is low and leaves the market without monitoring or buying. As a result, the producer sets the lowest possible quality level to save on production costs. Are there other equilibria?

It is clear that no other deterministic equilibrium exists: since monitoring is costly, when quality is deterministic, the consumer has no incentive to monitor. If the consumer does not buy, we fall again into the pessimistic expectations equilibrium; if the consumer buys with positive probability without learning the quality, the producer profitably deviates by providing the lowest possible quality. Therefore, an equilibrium with positive quality exists only if the producer randomizes across different quality levels. In turn, the monitoring strategy must ensure the producer is willing to randomize across different quality levels without deviating to other qualities.

Our first observation is that an equilibrium with positive quality cannot exist under perfect monitoring. Intuitively, perfect monitoring implies perfect information about quality, which, in the absence of commitment power, does not provide any instrument to prevent the producer from deviating to lower qualities. To understand why, notice that a consumer who observes quality will buy the product whenever quality exceeds the price. In particular, the consumer cannot refrain from buying a quality lower than expected as long as it is above the price. But then, the producer has no incentives to produce a higher-than-price quality. As a result, the consumer does not monitor, and the market unravels. This is the holdup problem in information investments.

Our next observation is that an equilibrium with positive quality can exist when

information is imperfect. This is because imperfect information provides commitment power to the consumer. For instance, suppose the information technology detects only whether the quality is above or below a certain fixed threshold, higher than the price. In this case, we have an equilibrium with positive quality since the producer has no incentive to deviate to a quality lower than the detection threshold. In particular, two equilibria with positive quality exist: in both, the consumer randomly monitors to make the producer indifferent between producing quality levels equal to zero and the threshold. In one equilibrium, the threshold-level quality is produced sufficiently rarely that the consumer never buys without monitoring; in the other, this quality is produced more often, and the consumer always buys even without monitoring. This difference is particularly stark when the monitoring cost approaches zero: in the former equilibrium, the probability that the threshold-level quality is produced converges to zero, approximating market failure, while in the latter, it converges to one, approximating the efficient market outcome.


This solution is unsatisfactory as it does not enable any prediction about quality provision. The frequency in which the threshold-level quality is produced, and, as a consequence, trade efficiency, vastly differs in the two equilibria described above. Furthermore, these equilibria tightly depend on the fixed threshold of quality detection. If we take a flexible approach and let the consumer set this threshold freely, everything goes. Depending on the consumer's beliefs about the producer's behavior, any quality level can be sustained in equilibrium by suitably selecting the right threshold.

For the remainder of our analysis, we study a model of flexible information acquisition, which abstracts away from particular information structures by assuming that the consumer has access to all of them. By resorting to the acquisition of imperfect information, this model mitigates the holdup problem in information investment while allowing for sharp predictions about quality provision in this market.

2 Inattention to quality provision

A *producer* (he) sets the quality of his product before offering it for a fixed price to a unit-demand *consumer* (she), who can either accept or reject the offer. The consumer does not observe quality directly but can process information about it at a cost before making her purchasing decision. We assume the consumer is *rationally inattentive* (Sims, 2003) and optimally trades off the benefits of obtaining more precise information about the product against the costs of acquiring information, which we assume are entropy-based.

Market interaction. The following timeline summarizes how the agents in this market interact. Each element is described below.



- (i) First, the *price* $p > 0$ at which the market operates realizes. For now, price is exogenously fixed; we investigate the consequences of endogenous pricing in Section 3.1.
- (ii) After observing the price, the producer sets the *quality* level $q \in Q := [0, \infty)$ of its product. We denote by $\sigma \in \Delta(Q)$ the producer's strategy.²
- (iii) The consumer can gather information about the realized quality. This is equivalent to selecting an *attention strategy* $\beta : Q \rightarrow [0, 1]$, which specifies the probability at which the consumer accepts the producer's offer for every quality level. More formally, one can construct attention strategies from binary signal spaces, one signal recommending to accept the offer and the other to reject it, and interpret $\beta(q)$ as the probability of receiving the recommendation to accept the offer when quality is q . Once a recommendation is received, the consumer follows it, and the game ends.³

Payoffs. When the game ends, payoffs are obtained. Given a consumer's attention strategy $\beta : Q \rightarrow [0, 1]$, the producer sets product quality $\sigma \in \Delta(Q)$ to maximize *profits*

$$\Pi(\sigma, \beta) = \int (p \cdot \beta(q) - c(q)) d\sigma(q),$$

where $c : Q \rightarrow [0, \infty)$ denotes the cost of quality provision. Throughout the analysis, we assume that c is linear in the quality level and normalize the cost of producing zero quality to zero, that is, $c(q) = \alpha \cdot q$, where $\alpha \in (0, 1)$. This assumption imposes constant marginal production costs, simplifying the study of the producer incentives while preserving their key trade-off—higher quality results in a higher probability of the offer being accepted

²We endow $[0, \infty)$ with the Borel σ -algebra, and $\Delta([0, \infty))$ with the topology of strong convergence. All functions are assumed to be measurable with respect to the relevant Borel σ -algebra.

³Restricting the consumer's strategy space to attention strategies is without loss under entropy-based information costs. See Matějka and McKay (2015) and Ravid (2019).

but also higher costs.⁴ With a slight abuse of notation, we denote by $\Pi(q, \beta)$, or simply $\Pi(q)$, the producer profits when the realized quality is $q \in Q$.

Production costs naturally impose an upper bound on the quality level. Specifically, for every $\alpha \in (0, 1)$, it is without loss to consider qualities in the compact set $Q = [0, p/\alpha]$, as producing any quality greater than p/α would result in negative profits. Therefore, p/α represents the *maximum feasible quality* the producer can provide.

In our model, production costs are *sunk*: producing a positive quality entails positive costs regardless of whether a trade occurs. Therefore, the producer faces a moral hazard problem while setting product quality. Later, in Section 4, we consider the case of recoverable costs and show that not even flexible information acquisition can prevent the holdup problem. Intuitively, when costs are sunk, the producer's potential losses facilitate information investments by providing more incentives for randomization. When costs are recoverable instead, the producer does not face any loss, is never indifferent across different quality levels, and the consumer has no incentives to acquire information.

The consumer processes information at a cost, which we model using *mutual information*. The mutual information associated with attention strategy $\beta : Q \rightarrow [0, 1]$, when the consumer's beliefs about quality are $\mu \in \Delta(Q)$, is given by

$$I(\mu, \beta) = H(\mathbb{E}_\mu[\beta]) - \mathbb{E}_\mu[H(\beta)], \quad (1)$$

where $H : [0, 1] \rightarrow [0, 1]$ is the Shannon entropy, $H(x) = -x \log(x) - (1-x) \log(1-x)$.⁵ Mutual information quantifies the difference in entropy between the unconditional probability of accepting an offer and the probability of accepting it conditional on each quality level. In simpler terms, it captures how much the consumer's behavior, as driven by attention, varies across different quality levels.

The consumer *utility* equals her gains from trade net of the costs of processing information. For every producer's strategy $\sigma \in \Delta(Q)$, the consumer selects an attention strategy $\beta : Q \rightarrow [0, 1]$ to maximize

$$U(\sigma, \beta) = \int \beta(q) \cdot (q - p) d\sigma(q) - \lambda \cdot I(\sigma, \beta),$$

where the parameter $\lambda > 0$ represents the consumer's unit cost of information processing. As λ increases, the consumer finds it more expensive, in mutual information terms, to

⁴We discuss this assumption in more detail when we sketch the proofs of our results in Section 3.3.

⁵We follow the usual convention, that is, $H(0) = H(1) = 0$.

differentiate her behavior across different quality levels. Again, slightly abusing notation, we write $U(q, \beta)$, or simply $U(q)$, to denote utility when the realized quality is $q \in Q$.

Equilibrium. We adopt *Nash equilibrium* as the underlying solution concept to analyze our quality provision game. Recall that a *strategy profile* (σ, β) is a Nash equilibrium (NE) if β is a best response to σ , and σ is a best response to β .

The *pessimistic expectations* equilibrium, discussed in Section 1.2, remains a valid NE under rational inattention. In this equilibrium, the consumer expects quality to be zero and therefore never purchases, $\beta = 0$; the producer, in turn, optimally responds by setting the quality to zero, $\sigma = \mathbb{1}_{\{q=0\}}$. In the remainder, we focus on equilibria involving *positive quality provision*, where the producer provides positive quality with positive probability.

Before, we need to discuss a known concern related to rational inattention in strategic settings. The issue arises because mutual information is determined by prior beliefs and, therefore, unaffected by off-equilibrium contingencies. As a result, the consumer can pay attention to off-equilibrium qualities (i) arbitrarily, and (ii) at no additional costs. This is problematic since the producer’s behavior depends on the consumer’s purchasing decision at all qualities—on and off-equilibrium. In particular, it generates equilibria with deterministic and positive quality provision, where the consumer selects a non-constant attention strategy for free, allowing her to react differently to different quality levels without incurring any cost to process information about the realized quality.⁶

To solve this issue, we propose a new refinement of Nash equilibrium which complements Ravid’s (2020) notion of *credible best responses*. Credibility tackles (i) by selecting the consumer’s best responses that treat every quality as if it occurred in equilibrium. We additionally impose *robustness to non-degenerate entropic costs*, which requires attention strategies to remain optimal even when processing costs are not zero, but arbitrarily small. A best response that satisfies both credibility and our robustness notion, which we refer to as a *Shannon best response*, is optimal for reasons that extend beyond (i) and (ii).

2.1 Credible best responses

Ravid (2020) identifies a class of equilibria sustained by non-credible threats—such as the consumer threatening not to buy unless a specific quality level is met— which illustrates part of the issue previously discussed.

⁶This issue does not occur when the information cost is defined in terms of experiments rather than divergence of beliefs. See Denti et al. (2022) for a detailed discussion.

Non-credible threats. Consider the following strategy profile: $\sigma = \beta = \mathbb{1}_{\{q=\gamma \cdot p\}}$, where $\gamma \in [1, 1/\alpha]$. That is, the consumer's attention strategy recommends purchasing with probability 1 when quality equals $q = \gamma \cdot p$, and rejecting otherwise. As a consequence, the producer has no choice but to meet this quality level; σ is best response to β . Therefore, every quality $q \neq \gamma \cdot p$ occurs with zero probability, and, by equation (1), the consumer utility is unaffected by zero-probability events. Since β is best response to qualities that occur with positive probability, it is a best response to σ . Notice that γ is set such trade yields a non-negative payoff to both players, and, by varying γ , every quality level in the interval $[p, p/\alpha]$ can be sustained as an equilibrium outcome.

The credibility refinement. As Ravid (2020) shows, equilibria involving non-credible threats are fragile to the producer making arbitrarily small mistakes. Intuitively, an attention strategy recommending purchase only if the producer provides a certain quality level becomes suboptimal when the producer plays every quality with positive probability. To exclude these equilibria, we adopt Ravid's credibility notion, which requires the consumer's attention strategy to be robust to small mistakes in the producer's behavior.

DEFINITION 1 (Ravid, 2020). *We say that β is a credible best response to σ if, for every $\tilde{q} \in Q$, there exists a sequence $(\tilde{\sigma}_n)_{n \in \mathbb{N}}$ such that, for every $n \in \mathbb{N}$, $q \in Q$,*

- $\tilde{\sigma}_n(\tilde{q}) > 0$, and $\tilde{\sigma}_n(q) \rightarrow \sigma(q)$ strongly,⁷
- β is a best response to $\tilde{\sigma}_n$.

Furthermore, (σ, β) is a credible equilibrium if is NE and β is a credible best response to σ .

Credibility is similar to, but weaker than, Selten's (1975) trembling-hand perfection as it allows perturbations in the producer's strategy to vary with off-equilibrium quality levels, while trembling-hand perfection does not.⁸

Lemma 1 characterizes credible best responses by showing they follow the adjusted multinomial logit formulation typical of rational inattention problems *everywhere* instead of almost surely (Matějka and McKay, 2015).

⁷A sequence $(\nu_n)_{n \in \mathbb{N}}$ of Borel probability measures over $Y \subseteq \mathbb{R}$ converges strongly to a Borel probability measure ν if $\nu_n(A) \rightarrow \nu(A)$ for every Borel $A \subseteq Y$.

⁸Formally, trembling-hand perfection would require β to be a best response to every element in a sequence of vanishing producer's strategies that are full-support.

LEMMA 1 (Ravid, 2020). *The attention strategy β is a credible best response to $\mu \in \Delta(Q)$ if and only if it is a best response, and*

$$\beta(q) = \frac{\pi \cdot e^{\frac{q-p}{\lambda}}}{\pi \cdot e^{\frac{q-p}{\lambda}} + 1 - \pi}, \quad (2)$$

for every $q \in Q$, where $\pi := \mathbb{E}_\mu[\beta]$.

To clarify the role of credibility, recall that a rational inattentive consumer follows the logit formula (2) μ -almost surely. Specifically, when choosing the probability of accepting an offer of quality $q \in \text{supp } \mu$, the consumer weights the gains from trade, $q - p$, against the attention costs, λ , and adjusts by the *average trade probability*, $\pi := \mathbb{E}_\mu[\beta]$.⁹

Credibility ensures that the logit formula (2) holds for *every* $q \in Q$, requiring the consumer to treat all qualities as if they occur with positive probability. Without credibility, the consumer would fail to internalize the incentives associated with zero-probability qualities and set an arbitrary purchasing probability for them. Notice that this is what enables non-credible threats—the consumer threatens to never buy any off-equilibrium quality—which are thus prevented by credibility.

2.2 Shannon-Nash equilibrium

In the context of our quality provision game, we identify an additional equilibrium that, although credible, is unappealing. This equilibrium relies on mutual information becoming degenerate when the producer sets quality deterministically, allowing the consumer to react differently to different quality levels without incurring any cost.

An artificial equilibrium. Credibility does not exclude degenerate entropic costs, allowing for the existence of a credible equilibrium with deterministic and positive quality. In this credible equilibrium, the producer sets the quality equal to the price making the consumer indifferent between accepting and rejecting the offer. As a result, the consumer accepts the offer with an intermediate probability $\pi \in (0, 1)$. Since the consumer's best response is credible, π induces an attention strategy that follows the logit formula (2) for every quality level $q \in Q$. In particular, the consumer sets π such that it is optimal

⁹In addition to Ravid (2020), see Csiszár (1974), Matějka and McKay (2015), Yang (2015), and Denti et al. (2020) for further properties of best responses in rational inattention problems.

for the producer to play $\sigma = \mathbb{1}_{\{q=p\}}$: higher or lower qualities would result in higher than profitable costs and lower than profitable average trade probabilities, respectively.¹⁰

Why is this credible equilibrium unappealing? Since no quality level different than the price is produced in equilibrium, the consumer's unconditional and conditional behavior coincides, implying that mutual information equals zero. In other words, the consumer is not acquiring any information almost surely. Still, the consumer processes information *everywhere* by forming a credible attention strategy that specifies a different probability of buying for *every* quality level. This attention strategy determines the equilibrium as it disciplines the producer's quality provision, which depends on the consumer's purchasing decision for all qualities, including those not part of the equilibrium. On the one hand, if the consumer were truly not processing any information, for instance, by setting the same purchasing probability for all quality levels, then the producer would deviate and produce quality equal to zero. On the other, if this attention strategy were minimally costly, it would not be worth it for the consumer, as it would result in a negative utility.

Shannon best responses. To exclude the artificial equilibrium resulting from degenerate entropic costs, we impose a mild adjustment to the notion of credibility.

DEFINITION 2. *Let β be a best response to $\mu \in \Delta(Q)$. We say that β is robust to non-degenerate entropic costs if there exists an $\varepsilon > 0$ such that*

$$\beta \in \operatorname{argmax} \int \beta(q)(q-p) d\sigma(q) - \lambda \cdot I_\varepsilon(\mu, \beta),$$

where $I_\varepsilon(\mu, \beta) = \max\{I(\mu, \beta), \varepsilon\}$ if β is non-constant and zero otherwise.

This property ensures that information is costly to process or, put differently, that the consumer behaves *as if* paying attention to quality always entails some positive, albeit infinitely small, cost. When β is constant, the consumer is not paying attention and incurs no cost. When β is non-constant, regardless of whether qualities belong to the support of μ , the consumer pays attention to distinguish across different quality levels and must incur some costs. The notion of robustness to non-degenerate entropic costs is particularly weak as it does not constrain how costly attention should be when mutual information is zero as long as β can be justified by some positive costs.¹¹

¹⁰See Proposition 8 and Corollary 3 in Appendix B for a formal treatment of the artificial equilibrium.

¹¹An alternative way to formulate a similar robustness notion is to consider consumer best responses that are optimal to entropic costs plus some arbitrarily small fixed costs, i.e., if β is non-constant, then it

As Proposition 1 below shows, the combination of credibility and robustness to non-degenerate entropic costs rules out precisely the cases where a non-constant, credible attention strategy is optimal due to degenerate beliefs.

DEFINITION 3. *We say that β is a Shannon best response to $\mu \in \Delta(Q)$ if β is a credible best response to μ that is robust to non-degenerate entropic costs.*

PROPOSITION 1. *The following holds:*

- (i) *Let $\mu \in \Delta(Q)$ with $|\text{supp}\mu| > 1$. Then, β is a Shannon best response to μ if and only if it is a credible best response to μ .*
- (ii) *Let $\mu \in \Delta(Q)$ with $|\text{supp}\mu| = 1$. Then, β is a Shannon best response to μ if and only if it is a credible best response to μ with $\mathbb{E}_\mu[\beta] := \pi \in \{0, 1\}$.*

As a consequence of Proposition 1, Shannon best responses select only credible best responses in which the consumer incurs positive information costs when paying attention. Specifically, when $|\text{supp}\mu| = 1$, this selection refines away credible best responses with $\pi \in (0, 1)$, which generate the artificial equilibrium previously discussed.

The intuition behind the proposition is as follows. If $\pi \in \{0, 1\}$, then, by credibility, β is constant everywhere, and points (i) and (ii) are straightforward. Let $\pi \notin \{0, 1\}$. For point (i), if $|\text{supp}\mu| > 1$, then, by optimality, β is non-constant μ -almost surely. This implies that the mutual information is positive, $I(\mu, \beta) > 0$, and therefore that β is robust to non-degenerate entropic costs. For point (ii), assume that β is a credible best response with $\pi \in (0, 1)$. In this case, $\mu = \mathbb{1}_{\{q=p\}}$, as for quality levels that are higher or lower than the price the consumer would always or never buy, respectively. However, when $\mu = \mathbb{1}_{\{q=p\}}$, the consumer obtains zero utility and cannot incur any additional information cost.

Denti et al. (2022) propose an alternative approach to address the issue of degenerate entropic costs. To reconcile entropy-based costs with information costs defined in terms of Blackwell experiments, Denti et al. compute mutual information using as reference a full-support probability distribution. This solution prevents degenerate beliefs, ensuring there is no free information. However, it also modifies the consumer's best response, which no longer follows the logit formula typical of rational inattention problems (2), but

must remain optimal when costs are $I(\mu, \beta) + \epsilon$, for some $\epsilon > 0$, where $\mu \in \Delta(Q)$. The two notions yield the same predictions in our quality provision game. As we show in Proposition 3, the consumer has positive utility in equilibrium, which implies that the equilibrium strategy is robust to some additional fixed costs. Similarly, the artificial equilibrium is excluded as additional fixed costs would imply negative utility.

depends on the reference measure.¹² Shannon best responses instead preserve the logit formula by selecting credible best responses (see Lemma 1).

Shannon meets Nash. The previous analysis motivates the following solution concept.

DEFINITION 4. *The profile (σ, β) is a Shannon-Nash equilibrium if it is a Nash equilibrium and β is a Shannon best response.*

Shannon-Nash equilibria are closely related to the credible equilibria with trade in Ravid (2020) and Cusumano, Fabbri, and Pieroth (2024). They study, respectively, monopolistic and competitive pricing when the consumer is rational inattentive and simultaneously learns about the price and quality of the good, which is exogenous in their settings. When quality is deterministic, each of these games have a unique credible equilibrium with trade which follows the same logic of the artificial equilibrium with deterministic and positive quality and, therefore, is not Shannon-Nash. However, they consider random qualities: firms price deterministically given the quality level, which is stochastic. As a result, the consumer faces positive information costs to learn about quality, which implies, by Proposition 1, that these equilibria are Shannon-Nash.

3 A quality-efficiency trade-off

This section presents the main result of our analysis, which consists of the trade-off between the provision of high quality and trade efficiency. We start by discussing the market consequences of Shannon best responses: inefficient trade and binary quality. Then, we present our main result, discussing the existence of an equilibrium with positive quality, the learning and participation effect, and the limiting cases of attention costs. Finally, we study how the surplus vary as the attention cost varies.

Inefficiency and binary quality. The following implications rely on the notion of Shannon-Nash equilibrium.

PROPOSITION 2. *If (σ, β) is a Shannon-Nash eqm with positive quality provision, then:*

- *Trade is inefficient: the average trade probability $\pi := \mathbb{E}_\sigma[\beta] \in (0, 1)$;*
- *The producer randomizes between two quality levels: $\text{supp}\sigma = \{0, q^*\}$, where $q^* > p$.*

¹²See Proposition 7 in Denti et al. (2022) for further details.

Shannon-Nash equilibria predict trade inefficiency and random quality provision, which are natural features of markets where quality is observable only by acquiring information. To incentivize information acquisition, which is necessary to prevent market failure as discussed in Section 1.2, the producer must randomize between qualities below and above the price level, creating inefficiencies. Specifically, in this case, the producer randomizes between two quality levels only: zero and a higher-than-price quality.

The intuition is as follows. Trade inefficiency is implied by credibility alone. By Lemma 1, under efficient trade, $\pi = 1$, the consumer buys each quality level with certainty. But this cannot occur in equilibrium: if the consumer buys regardless of the quality level, the producer will set quality equal to zero to save on production costs. As illustrated by the artificial equilibrium of Section 2.2, random quality provision is a consequence of Shannon best responses. Indeed, by Proposition 1, $\pi \in (0, 1)$ can occur only if the producer randomizes. Finally, the result on binary quality follows since, to prevent producer deviations, the isoprofit curve, which is linear, must lie above the consumer's best response, which is s-shaped in equilibrium. Binary quality is common to other models of information acquisition where best responses are s-shaped.¹³

Main result. Theorem 1 advances this discussion by identifying the form that trade inefficiency and random quality provision take when the consumer is rational inattentive.

THEOREM 1. *Fix (α, p) . There is a threshold $\bar{\lambda} > 0$ such that a unique Shannon-Nash equilibrium with positive quality provision (σ, β) exists if and only if $\lambda < \bar{\lambda}$. In this equilibrium:*

- *The higher-than-price quality $q^*(\lambda) > p$ is decreasing in the attention cost;*
- *The average trade probability $\pi^*(\lambda) := \mathbb{E}_\sigma[\beta] \in (0, 1)$ is increasing in the attention cost.*

Furthermore, as $\lambda \rightarrow \bar{\lambda}$, $q^(\lambda) \rightarrow p$ and $\sigma(0) \rightarrow 0$. As $\lambda \rightarrow 0$, $q^*(\lambda) \rightarrow p/\alpha$, but the market fails: $\sigma(0) \rightarrow 1$ and $\pi^*(\lambda) \rightarrow 0$.*

By establishing the existence of a unique Shannon-Nash equilibrium with positive quality provision, Theorem 1 shows that flexible information acquisition can partially compensate for the lack of contractual commitment. Specifically, an equilibrium exists if and only if the attention cost is not too high. When $\lambda > \bar{\lambda}$, the information investment is held up, as the producer finds it optimal to deviate and set quality to zero. When

¹³For instance, see Section 4 for a model where the consumer learns the producer's quality by acquiring convex-concave information.

the attention cost vanishes, $\lambda \rightarrow 0$, market also fails: the Shannon-Nash equilibrium converges to the equilibrium with pessimistic expectations as the market is composed of lemons only. The interaction of two forces drives these results—the *learning effect* and the *participation effect*, which imply, respectively, that the high-quality level and the average trade probability decrease and increase in the attention cost.

The *learning effect* causes the high-quality level to decrease in the attention cost. Intuitively, as the attention cost λ decreases, the consumer *learns* the producer’s offer more precisely and makes fewer mistakes. Specifically, the probability of buying a high quality product, $\beta(q^*)$, increases, while buying a zero quality one, $\beta(0)$, decreases. This reduces the profits from producing zero quality, $p \cdot \beta(0)$, as the producer benefits less frequently from the consumer erroneously purchasing a zero quality product. But the producer has to be indifferent in equilibrium, which implies that the profits from producing the high quality must also decrease. Since $\beta(q^*)$ increases, the high quality has to increase until the profits of the two qualities coincide, leading to the result.

The *participation effect* implies that the average trade probability increases in the attention cost. By the learning effect, when λ decreases, the high-quality level increases, say from \hat{q} to q^* , to equate profits between low and high qualities. However, the indifference between 0 and q^* does not guarantee an equilibrium: the producer may still find it optimal to deviate to other quality levels. For instance, if a more attentive consumer traded on average more frequently, that is, π increases as λ decreases, then the probability of buying would increase for all quality levels. In particular, it would increase for the previous high quality, \hat{q} , as well. But this is not possible: if $\beta(\hat{q})$ increases, the producer would obtain higher profits by deviating to \hat{q} . Therefore, a more attentive consumer *participates* in the market less often to prevent the producer from deviating to off-equilibrium quality levels. Anticipating this, the producer renders the market less appealing by reducing, when necessary, the frequency of high quality.

The interplay between the learning and the participation effect severely affects the market when the attention cost vanishes. In this scenario, the consumer’s mistakes frequency also vanishes due to the learning effect, $\beta(q^*) \rightarrow 1$ and $\beta(0) \rightarrow 0$. Therefore, to sustain trade, the producer needs to provide the maximum feasible quality, $q^* \rightarrow p/\alpha$. However, as high quality progressively converges to its maximum level, the participation effect makes the consumer trade extremely rarely to prevent the producer from deviating to lower qualities close to q^* , which would be accepted with a high probability. In the limit, the market is composed of lemons only, as the consumer’s choice never to trade




Figure 1: The high quality, trade probability, and frequency of low quality in the unique Shannon-Nash equilibrium with positive quality, plotted in terms of the attention cost. Parameters: $\alpha = 0.1$, $p = 10$.

forces the high-quality frequency to converge to zero.

We contrast this limiting result with the free-learning benchmark. When attention is free, that is, $\lambda = 0$, there are two Pareto-ranked Shannon-Nash equilibria: the pessimistic expectations equilibrium and a perfect information equilibrium—the consumer pays full attention for free, and the producer sets quality equal to the price.¹⁴ By Theorem 1, in the limit of $\lambda \rightarrow 0$, the pessimistic expectations equilibrium is selected, the worst of the two.

The above conclusion is reminiscent of a result in Ravid et al. (2022), which also states that, when information costs vanish, equilibria converge to the worst free-learning equilibrium, albeit for different reasons. In Ravid et al., a buyer learns about her exogenous valuation before deciding whether to accept a take-it-or-leave-it offer made by a seller. In this setting, to sustain an equilibrium when information costs are positive, the buyer ignores a large amount of information, selecting the worst-free learning equilibrium when costs vanish. However, our consumer does not ignore information; on the contrary, as $\lambda \rightarrow 0$, the consumer distinguishes equilibrium qualities with greater precision, $\beta(q^*) \rightarrow 1$ and $\beta(0) \rightarrow 0$. The difference is that, in our framework, uncertainty is endogenous, and the consumer acquires imperfect information to discourage the producer from providing off-equilibrium qualities while paying attention to equilibrium ones. Nevertheless, when $\lambda \rightarrow 0$, the market fails since quality provision is insufficient to sustain trade.

¹⁴By Lemma 1, when $\lambda = 0$, credible best responses imply that the consumer always buys a quality higher than the price and never a quality lower than the price.

As the attention cost approaches the threshold, $\lambda \rightarrow \bar{\lambda}$, the quality provision approximates that of perfect information, $q^* \rightarrow p$ and $\sigma(0) \rightarrow 0$. Inattention generates mistakes, which increase profits and lower the high-quality level. Ultimately, when the attention cost reaches $\lambda = \bar{\lambda}$, the producer is indifferent between 0 and p , and the market fails since the consumer no longer has any incentive to trade. In particular, for $\lambda \geq \bar{\lambda}$, investing in information no longer prevents the holdup problem. The reason is that the consumer makes mistakes too frequently— $\beta(0)$ is too high—to prevent the producer’s deviations.

Figure 1 provides a graphical representation of Theorem 1. The graphs describe how the high quality and the average trade probability vary in the attention cost, illustrating the learning effect (on the left) and participation effect (on the right).

Producer and consumer surplus. We now discuss how the producer profits and the consumer utility vary as a function of the attention cost.

PROPOSITION 3. *Fix (α, p) . In the unique Shannon-Nash equilibrium with positive quality provision, the producer profits are positive, increasing in λ , and approach zero as $\lambda \rightarrow 0$. The consumer utility is positive and approaches zero as $\lambda \rightarrow 0$ and $\lambda \rightarrow \bar{\lambda}$.*

Figure 2 provides a graphical illustration of Proposition 3 by plotting the producer profits and the consumer utility as a function of unit attention cost.

As the unit of attention cost increases, the profits increase as the producer benefits from both the learning and the participation effect. The learning effect implies a lower high-quality level, while the participation effect a higher trade probability. When $\lambda \rightarrow 0$, the average trade probability $\pi \rightarrow 0$ as well, and the producer obtains zero profits.

The learning and participation effects pull consumer utility in opposite directions. As consumer attention increases, the learning effect raises the high-quality level, but the participation effect makes the market less appealing by reducing trade frequency. Both extremes of attention cost ($\lambda \rightarrow 0$ or $\lambda \rightarrow \bar{\lambda}$) lead to poor consumer outcomes: in the former case, trade probability drops to zero, while in the latter, product quality declines to match the price. As a result, utility converges to zero at both extremes.

Utility is positive as the consumer extracts information rents by paying attention. The argument mirrors that of Ravid (2020). In the Shannon-Nash equilibrium, the consumer’s attention strategy β is non-constant σ -almost surely. By known results, this implies that β belongs to a region where the mutual information is strictly convex, meaning that the consumer’s objective is strictly concave. Therefore, β is strictly optimal: the consumer

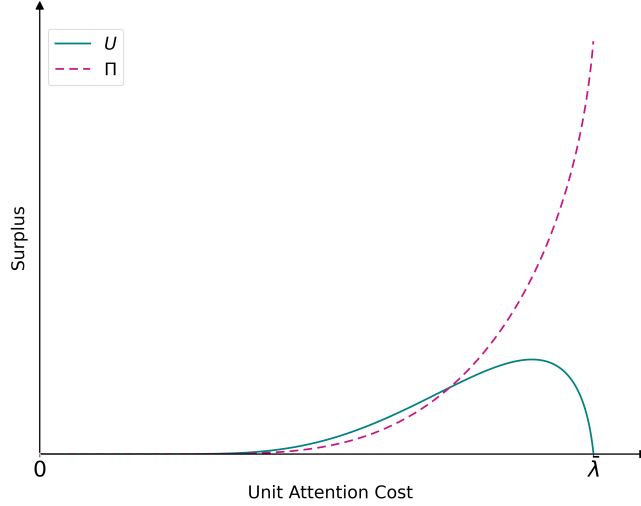


Figure 2: The producer profits and consumer utility in the unique Shannon-Nash equilibrium with positive provision provision, plotted in terms of the attention cost. Parameters: $\alpha = 0.1$, $p = 10$.

strictly prefers β over the option of never purchasing, which yields zero utility.

Proposition 3 has relevant social welfare implications. If the consumer could *manipulate* her attention level, she would prefer to be more inattentive for values of λ close to zero and less inattentive for values of λ close to $\bar{\lambda}$. The producer, instead, would always prefer to face a less attentive consumer. As a consequence, decreasing the attention level is Pareto-efficient when the consumer is particularly attentive, while increasing the attention level can only benefit the consumer.

Attention threshold. We conclude this section by discussing how the threshold $\bar{\lambda}$ responds to changes in the production costs and the price.

PROPOSITION 4. *The threshold $\bar{\lambda} := \bar{\lambda}(\alpha, p) > 0$ is decreasing in α , and increasing in p .*

By Theorem 1, when the attention cost reaches the threshold, $\lambda = \bar{\lambda}$, the producer is indifferent between offering zero quality or a quality equal to the price. If production costs, α , increase, the producer will no longer remain indifferent and strictly prefer to offer zero quality. In this scenario, at the current λ , the consumer is making mistakes too frequently to discourage the producer from deviating to lower the quality level. Consequently, the attention threshold decreases. A similar logic applies when the price decreases: the producer will strictly prefer to offer zero quality, as the reduced price incentivizes this deviation. As a result, the consumer attention threshold decreases.

3.1 Price setting

In this section, we fix the production and attention costs to explore the market outcome as the price varies. This analysis facilitates the understanding of a game where a player—such as the producer or the consumer—or a bargaining protocol sets the price before the quality provision game takes place. The chosen price determines a subgame, which we assume selects the unique Shannon-Nash equilibrium with positive quality (Theorem 1).

THEOREM 2. *Fix (α, λ) . The unique Shannon-Nash equilibrium with positive quality provision exists if and only if $p > \bar{p}$, where $\bar{p} := \bar{p}(\alpha, \lambda) > 0$ is a threshold increasing in both arguments. The producer profits are positive, decreasing in p , and approach zero as $p \rightarrow \infty$. The consumer utility is positive and approaches zero as $p \rightarrow \bar{p}$ and $p \rightarrow \infty$.*

The price threshold \bar{p} represents the minimum price at which the market can operate. It follows a logic analogous to the attention threshold, $\bar{\lambda}$, in Theorem 1. For a given (α, λ) , when $p = \bar{p}$, the producer is indifferent between offering zero quality or a quality equal to the price. If $p < \bar{p}$, the production and information costs become too high to sustain positive quality provision: the producer finds it optimal to set quality to zero, causing the equilibrium to break down. For a similar reason, the threshold \bar{p} is increasing in both α and λ . As either parameter increases, the relative gains from trade diminish compared to the associated costs. Therefore, the lowest price at which trade occurs increases.

The incentives in the quality provision games determine the producer and consumer's preferences over prices. Following Theorem 2, if the consumer could set the price at the beginning of the quality provision game, she would choose an intermediate price level. Instead, the producer would select the lowest possible price compatible with the threshold \bar{p} . To build intuition, low prices offer *commitment benefits* for the producer by increasing the average trade probability and strengthening the incentives to provide the high-quality product more frequently. As the price p approaches \bar{p} , the attention threshold $\bar{\lambda}$ converges to λ , implying that the high-quality level q^* is close to p . In this scenario, the frequency of high quality converges to 1.¹⁵ This implies that, when all prices are available, there is no producer-optimal price: the producer sets prices arbitrarily close to the threshold \bar{p} . However, at \bar{p} , the Shannon-Nash equilibrium with positive quality provision fails to exist since $q^* = p$, excluding the consumer from the market.

How consumer utility varies as the price varies builds on the preceding discussion.

¹⁵See the proof of Theorem 2 in Appendix B for a formal analysis of the relations between π^* , q^* , and p .




Figure 3: The producer profits and consumer utility in the unique Shannon-Nash equilibrium with positive quality provision, plotted in terms of the price. Parameters: $\alpha = 0.1$, $\lambda = 10$.

When $p \rightarrow \bar{p}$, the high-quality level q^* converges to the price, leading to utility converging to 0. When $p \rightarrow \infty$, by Proposition 4, the parameter λ becomes smaller relative to the increasing threshold $\bar{\lambda}$. This mirrors the scenario where $\lambda \rightarrow 0$, leading to the outcome $\pi^* \rightarrow 0$, which again drives utility to 0. The argument establishing that the consumer utility is positive in equilibrium follows the same reasoning as in Proposition 3.

The analysis of producer profits is more subtle. On the one hand, an increase in price raises the producer's gains from trade. On the other, higher prices select equilibria with higher high-quality levels and lower average trade probability, reducing profits. Theorem 2 demonstrates that the combined negative effects of higher high quality and lower trade probability always outweigh the positive effect of increased gains from trade: profits are decreasing in the price. Furthermore, the negative effects become extreme at high prices: as $p \rightarrow \infty$, profits converge to 0, driven by $q^* \rightarrow \infty$ and $\pi^* \rightarrow 0$.

The result that low prices provide commitment benefits for the producer contrasts with a finding in Wolitzky (2023). In Wolitzky, information is free for the consumer, exogenous, and satisfies specific regularity requirements. The producer chooses price and quality simultaneously; the equilibria in which price is chosen before quality are ruled out to avoid the multiplicity that arises when off-equilibrium prices induce negative beliefs about quality. In this setting, there is a unique equilibrium with positive quality provision in which the producer sets $p = q$, and the consumer decides whether to purchase based

on the signal realization. In this equilibrium, a high price, as opposed to a low price, has commitment benefits for the producer, as it improves the likelihood of trade by inducing to provide high quality. Two are the main differences with our model. First, robustness to non-degenerate entropic costs implies that information processing is never free and, as a result, deterministic equilibria cannot be Shannon-Nash. Second, by disentangling the choice of quality from the price, the producer's objective is not strictly concave in the quality level, allowing for randomized equilibria that enable monitoring incentives.

Figure 3 graphically illustrates Theorem 2 by plotting the producer profits and the consumer utility in terms of prices.

3.2 Cost-reducing investments

We investigate how variations in production costs affect the market outcome while holding both the price and attention cost fixed. This analysis can be framed as the producer making an observable investment before the quality-provision game takes place to determine his production costs. The resulting costs induce a subgame that selects the unique Shannon-Nash equilibrium of Theorem 1.

PROPOSITION 5. Fix (λ, p) . The unique Shannon-Nash equilibrium with positive quality provision exists if and only if $\alpha < \bar{\alpha} < 1$, where $\bar{\alpha} := \bar{\alpha}(\lambda, p) > 0$ is a threshold decreasing in λ , and increasing in p . Profits are positive, increasing in α , and approach zero as $\alpha \rightarrow 0$. The consumer utility is positive and approaches zero as $\alpha \rightarrow 0$ and $\alpha \rightarrow \bar{\alpha}$.

The threshold $\bar{\alpha}$ determines the maximum production costs at which the market operates. Its logic aligns to the previous thresholds, $\bar{\lambda}$ and \bar{p} . When $\alpha = \bar{\alpha}$, the producer is indifferent between producing a quality of zero or equal to the price. If λ increases or p decreases, producing a low quality becomes relatively more profitable: the producer is no longer indifferent when $\alpha = \bar{\alpha}$ and deviates to zero quality. As a result, trade occurs only at lower production costs.

The incentives in the quality provision game shape the producer's preferences over cost-reducing investments. Again, when $\alpha \rightarrow \bar{\alpha}$, $\bar{\lambda} \rightarrow \lambda$, and by Theorem 1, the high quality converges to the price as its frequency approaches 1. In contrast, when $\alpha \rightarrow 0$, the parameter λ becomes relative small compared to $\bar{\lambda}$, mirroring the scenario where $\lambda \rightarrow 0$. In this case, $q^* \rightarrow \infty$ and $\pi^* \rightarrow 0$.¹⁶ High production costs, and therefore modest

¹⁶See the proof of Proposition 5 in Appendix B for a formal analysis of the relations between π^* , q^* , α .

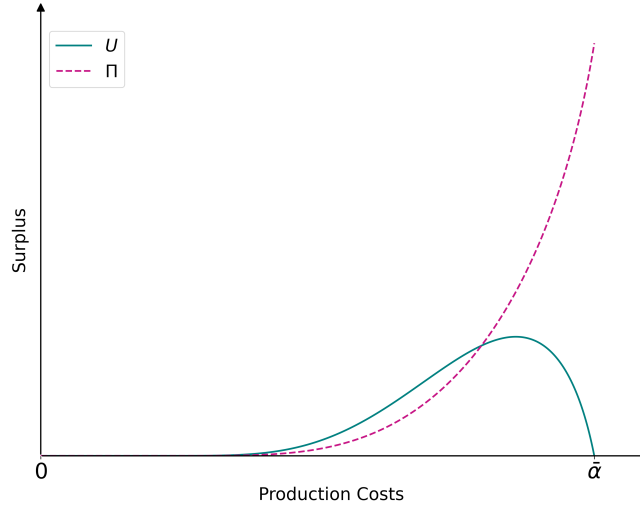


Figure 4: The producer profits and consumer utility in the unique Shannon-Nash equilibrium with positive quality provision, plotted in terms of the production costs. Parameters: $\lambda = 10$, $p = 10$.

cost-reducing investments, serve as a *commitment device* for the producer by providing an incentive to produce lower high-quality levels more frequently.

Proposition 5 establishes that the producer always benefits from committing to high production costs. That is, the positive effects of a higher production cost—a lower high-quality level and a higher probability of trade—overcomes the negative effect of making production more expensive. In particular, to sustain trade when production costs approach zero, the producer is forced to produce extreme high-quality levels. As a result, high-quality frequency vanishes, and profits converge to zero as the consumer never participates in the market.

The results on consumer utility follow from this discussion. Utility converges to zero at both extremes of production costs ($\alpha \rightarrow 0$ and $\alpha \rightarrow \bar{\alpha}$) since the market becomes unappealing for the consumer: either the frequency of high quality approaches zero or the high quality converges to the price. Consumer utility is positive by Proposition 3.

Figure 4 plots the producer profits and the consumer utility in terms of the information costs, illustrating Proposition 5.

3.3 Outline of the main arguments

We present a concise summary of the main arguments in the proofs of Proposition 2 and Theorem 1, with the detailed proofs provided in Appendix B.

Binary quality. We start by discussing the structure of the producer equilibrium strategy.

LEMMA 2. *If (σ, β) is a Shannon-Nash equilibrium with positive quality provision, then $\text{supp}\sigma = \{0, q^*\}$ where $q^* > p$.*

The intuition for this result proceeds through the following two steps.

Step 1.—By Lemma 1, the consumer’s best response β follows the adjusted multinomial logit formula (1) for every quality level. Each attention strategy is determined by the average trade probability, π . Depending on π , the shape of β over the positive reals can vary: it is either strictly concave when π is sufficiently high, or *s-shaped* for lower values of π —that is, strictly convex for low qualities and strictly concave for high qualities. When β is strictly concave, no Shannon-Nash equilibrium with positive quality provision can be sustained. In this case, the producer faces a strictly concave demand, which leads to a unique producer-optimal quality level. However, by Proposition 1, $\pi \in \{0, 1\}$ when σ is deterministic, meaning β cannot be strictly concave. Therefore, in equilibrium, β takes an s-shaped form.

Step 2.—Given linear production costs, the producer isoprofit curve is linear in conditional trade probabilities. For a profit level $\bar{\Pi} \geq 0$,

$$\bar{\Pi} = p \cdot \beta(q) - \alpha \cdot q \implies \beta(q) = q \cdot \frac{\alpha}{p} + \frac{\bar{\Pi}}{p}$$

To prevent the producer from deviating to other quality levels, the isoprofit curve must lie above the demand. If the demand exceeds the isoprofit curve at any quality, then the producer obtains higher profits by deviating to this quality as the intercept of the isoprofit line, and hence profits, would increase. By combining the s-shaped form of the consumer’s best response with the linear form of the isoprofit curve, it follows that these curves intersect in two points: zero and the tangent point in the concave region of the attention strategy. Therefore, the producer’s strategy has support over exactly 2 qualities.

Equilibrium behavior. We are finally ready to characterize Shannon-Nash equilibria with positive quality provision.

LEMMA 3. *Fix (α, λ, p) . The profile (σ, β) is a Shannon-Nash equilibrium with positive quality provision if and only if $\text{supp}\sigma = \{0, q^*\}$, $q^* > p$, $\sigma(0) := \sigma_0 \in (0, 1)$, $\mathbb{E}_\sigma[\beta] := \pi^* \in (0, 1)$, and the following conditions hold:*

$$(i) \quad p \cdot \frac{\pi^* \cdot e^{\frac{q^*-p}{\lambda}}}{\pi^* \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi^*} - \alpha \cdot q^* = p \cdot \frac{\pi^* \cdot e^{-\frac{p}{\lambda}}}{\pi^* \cdot e^{-\frac{p}{\lambda}} + 1 - \pi^*} \quad (\text{indifference})$$

$$(ii) \quad \frac{\pi^* \cdot e^{\frac{q^*-p}{\lambda}}}{\pi^* \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi^*} = \frac{1}{2} + \frac{1}{2} \cdot \sqrt{1 - 4 \cdot \frac{\alpha \lambda}{p}} \quad (\text{tangency})$$

$$(iii) \quad \pi^* = \frac{(1 - \sigma_0) \cdot (1 - e^{\frac{q^*-p}{\lambda}}) - \sigma_0 \cdot (e^{-\frac{p}{\lambda}} - 1)}{(e^{\frac{q^*-p}{\lambda}} - 1) \cdot (e^{-\frac{p}{\lambda}} - 1)} \quad (\text{consumer optimal})$$

The profile (σ, β) is a Shannon-Nash equilibrium with positive quality provision if and only if it satisfies binary-quality (by Lemma 2), inefficient trade—if $\pi^* = 1$ the producer deviates to zero quality—together with three conditions. Conditions (i) and (ii) discipline producer's behavior. Condition (i) imposes indifference between producing qualities equal to q^* and 0; condition (ii) requires that q^* is the point in which the isoprofit line is tangent to the concave part of the demand. Condition (iii) ensures that π^* is a best response to the equilibrium strategy of the seller. This condition is peculiar to the binary nature of the problem and is due to Matějka and McKay (2015).¹⁷

The necessity of these conditions follows by Lemma 2. Since Shannon-Nash equilibria are binary-quality, condition (i) holds as the producer must be indifferent between the two quality levels. Furthermore, condition (ii) holds as well, as the producer would otherwise find it optimal to deviate to other qualities. Finally, the consumer's best response must be interior, which is equivalent to condition (iii). Sufficiency is immediate.

We proceed by combining conditions (i) and (ii) to show they admit a unique solution. This solution is not available in closed form, and the complete argument is technically involved; we defer it to Appendix B.

LEMMA 4. *Fix (α, p) . There exists a unique pair (π^*, q^*) with $\pi^* \in (0, 1)$ and $q^* > p$ satisfying equations (i) and (ii) in Lemma 3 if and only if $\lambda < \bar{\lambda}$.*

For any such pair (π^*, q^*) , we establish that there exists a unique $\sigma(0) \in (0, 1)$ such that condition (iii) holds. Consequently, by Lemma 3, a unique Shannon-Nash equilibrium with positive quality provision exists if and only if $\lambda < \bar{\lambda}$. Notice that, when $\lambda \geq \bar{\lambda}$, the pair (π^*, q^*) ceases to exist since, as we discuss below, q^* is decreasing in λ and $q^* \leq p$.

The threshold $\bar{\lambda}$ is obtained when the producer is indifferent between 0 and p . To allow for trade when q^* is close to p , $\sigma(0)$ must approach zero and, as a consequence, π^*

¹⁷See Proposition 1 in the Supplemental Appendix of Matějka and McKay (2015).

converges to $\beta(q^*)$. The threshold $\bar{\lambda}$ is the unique value of $\lambda > 0$ that satisfies the equation obtained by combining conditions (i) and (ii) when $q^* = p$ and $\pi^* = \beta(q^*)$.

One final comment is in order. The tangency condition, obtained by first-order methods, requires $\beta(q^*)$ to equal a value that does not depend on q^* , but only on the parameter (α, λ, p) . Such value is well-defined only if, for fixed (α, p) , $\lambda < \lambda^e := p/(4\alpha)$. We show that $\bar{\lambda} < \lambda^e$: the market ceases to exist as $q^* \rightarrow p$ before the tangency conditions fails.

Learning and participation effect. The analysis of these effects follows from the property of the unique pair that, by Lemma 4, solves conditions (i) and (ii) when $\lambda < \bar{\lambda}$. In particular, one could associate the indifference condition, requiring the producer to set the high quality to be indifferent in equilibrium, to the learning effect, while the tangency condition, preventing deviation to off-equilibrium qualities, to the participation effect.

First, we show that q^* is differentiable by applying the implicit function theorem. As a consequence, π^* is differentiable as well. The limiting result $q \rightarrow p/\alpha$ when $\lambda \rightarrow 0$ is established by standard arguments. To prove that q^* is monotone, we first show that for every $\lambda < \bar{\lambda}$, there is a unique $q^* > p$ satisfying equations (i) and (ii), and viceversa: for every $q^* > p$, there exists a unique $\lambda < \bar{\lambda}$ which satisfies equations (i) and (ii). This establishes a one-to-one relationship between q^* and λ ; monotonicity follows by the limiting results on q^* . A similar approach shows that $\pi^* \rightarrow 0$ as $\lambda \rightarrow 0$ and the monotonicity of π^* .

4 Discussion

We conclude by further discussing the model's assumptions and results.

Beyond moral hazard. We drop the assumption of sunk production costs by assuming that costs are *recoverable*: the producer incurs production costs only if trade occurs. As a result, the producer no longer faces a moral hazard problem in quality provision. Formally, given a consumer's attention strategy $\beta: Q \rightarrow [0, 1]$, the producer sets product quality $\sigma \in \Delta(Q)$ to maximize

$$\mathcal{P}(\sigma, \beta) = \int \beta(q) \cdot (p - c(q)) d\sigma(q),$$

where $c(q) = \alpha \cdot q$ for $\alpha \in (0, 1)$. The rest of the problem is unchanged.

When costs are recoverable instead of sunk, it is harder for the consumer to make

the producer indifferent across different quality levels. Intuitively, with sunk costs, the producer faces no losses in expectations, while with recoverable costs no losses almost surely. The following proposition conforms with this intuition.

PROPOSITION 6. *When production costs are recoverable, there is no Shannon-Nash equilibrium with positive quality provision.*

As a consequence, the only Shannon-Nash equilibrium is the pessimistic expectations equilibrium, in which there is no positive quality and no trade. Flexible information does not mitigate the holdup problem when costs are recoverable: no information incentivizes monitoring by inducing the producer to randomize across qualities.

The argument proceeds as follows. Since the consumer's best response is credible, it satisfies the logit formula (2) for every quality level. Because this demand is downward sloping, we calculate its inverse in terms of the quality level. By rewriting the producer's objective using the inverse demand, we show that this objective is strictly concave, and therefore, the producer has a unique best reply. By Proposition 1, when $|\text{supp } \sigma| = 1$, we have $\pi \in \{0, 1\}$, which does not allow for any positive quality provision in equilibrium.

Finally, when costs are recoverable, there is a credible equilibrium with deterministic and positive quality, which follows the same logic of the artificial equilibrium described in Section 2.2. In this equilibrium, the quality equals the price, and the consumer buys with an intermediate probability set at a specific level to prevent the producer's deviations. This attention strategy is optimal only due to degenerate beliefs which imply zero entropic costs: if costs were minimal, the consumer would prefer not to pay attention. This equilibrium exists since the framework with recoverable costs is strategically equivalent to Ravid's (2020) model in the case of deterministic quality.

Convex-concave information. The binary nature of the unique Shannon-Nash equilibrium with positive quality provision relies on the s-shaped form of the consumer's best responses. We test this claim in a different environment with general information costs in which the consumer can acquire information with a convex-concave shape.

We say that information acquisition satisfies *convex-concave information* when, prior to her purchasing decision, the consumer observes a signal $s \in S \subseteq \mathbb{R}$ by selecting a signal distribution conditional on each quality level $q \in Q$, $F(s|q)$. Assume $S = [\underline{s}, \bar{s}]$, where $\underline{s} \in [-\infty, 0)$ and $\bar{s} \in (0, -\infty]$. The signal distribution F satisfies the following conditions:

- (i) F has a density f ;
- (ii) The density f is full-support, $f(s|q) > 0$ for every (s, q) ;
- (iii) The density f satisfies the *strict monotone likelihood ratio property* (MLRP), that is, for every $s_1 > s_2, q_1 > q_2, f(s_2|q_2) \cdot f(s_1|q_1) > f(s_2|q_1) \cdot f(s_1|q_2)$;
- (iv) For every $s \leq 0, F(s|\cdot)$ is strictly convex. For every $s > 0$, one of the following holds:
 - (a) $F(s|\cdot)$ is strictly concave; (b) there is a $q_s \in [0, p/\alpha], F(s|\cdot)$ is strictly concave on $[0, q_s)$, and strictly convex on $(q_s, q/\alpha]$.

The consumer has access to all distributions satisfying the conditions above by incurring a cost $\kappa : F \mapsto \mathbb{R}$ which satisfies $\kappa(F) > 0$. After the signal realizes, the consumer decides whether to buy, $b : S \rightarrow [0, 1]$. By condition (i), we focus on consumer strategies that differ on a measure-zero set of signals. Furthermore, the consumer can gather no information at no costs and buy based on her beliefs. The rest of the game is unchanged.

Normal noise satisfies all four conditions above. In this case, the consumer observes $s = q + \epsilon$, where $\epsilon \sim N(0, \sigma^2)$ is a zero-mean, independent, normally distributed random variable with variance $\sigma^2 > 0$, and κ is proportional to $1/\sigma^2$. Furthermore, the property (iv) is satisfied whenever ϵ is drawn from a distribution unimodal at zero.¹⁸

PROPOSITION 7. *If there exists a Nash equilibrium with positive quality provision when information acquisition satisfies convex-concave information, then is binary-quality: the producer randomizes between 0 and $q^* > p$.*

The intuition for the result is as follows. The assumptions of full-support and strict MLRP, by Proposition 1 in Milgrom (1981), imply that the consumer's best response follows a threshold strategy: the consumer buys whenever the signal exceeds a threshold. Condition (iv) ensures that the probability of purchasing given any threshold is either strictly concave, strictly convex, or convex-concave. In the former case, there is no equilibrium with positive quality as the producer has no incentives to randomize since he is facing a strictly concave objective. Following an argument analogous to the one outlined for Lemma 2, in the latter two cases, the producer is indifferent between, at most, two quality levels, proving the proposition.

¹⁸A probability distribution H is *unimodal* at $m \in \mathbb{R}$ if it is strictly convex for all $x < m$, and strictly concave for all $x > m$. See Appendix C for a formal argument.

Consumer heterogeneity. Acquiring imperfect information mitigates the holdup problem by providing commitment power to the consumer. We now investigate whether a similar effect occurs when the consumer has some private information. In particular, we consider consumer private types that affect preference intensity for product quality. To isolate the effect of private information, we assume that monitoring generates no commitment benefits. All consumer types have access to the same perfect monitoring technology: they incur the same fixed cost to learn about quality perfectly.

The consumer's strategy space is defined as follows. The consumer either never buys: $\beta = 0$; always buys: $\beta = 1$; or monitors perfectly with some probability: β always buys a quality higher than the price, may buy a quality equal to the price, and never buys a quality lower than the price.¹⁹

A consumer type is drawn at the beginning of the game, the consumer observes her type, while the producer does not. Let $\Theta \subset [0, \infty)$ be a finite set of types distributed according to a common prior $\rho \in \Delta(\Theta)$. For every producer's strategy $\sigma \in \Delta(Q)$, a consumer of type $\theta \in \Theta$ selects β to maximize

$$U_{\theta, \kappa}(\sigma, \beta) = \int \beta(q) \cdot (\theta \cdot q - p) d\sigma(q) - \kappa(\beta),$$

where the monitoring cost $\kappa(\beta) = \kappa \cdot \mathbb{1}_{\{\beta \neq 0, \beta \neq 1\}}$ with $\kappa > 0$, meaning that each consumer type incurs the same fixed monitoring cost. As we discuss below, the case with infinite consumer types is analogous.

Again, there is no equilibrium except the one with pessimistic expectations. In every equilibrium with positive quality provision, some consumer type has to monitor with positive probability. To ensure monitoring, the producer randomizes across different qualities. Let $\underline{\theta}$ be the lowest consumer type that monitors in equilibrium. Notice that $\underline{\theta} > 0$: a consumer that has no utility for the product never monitors. Since monitoring is costly, type $\underline{\theta}$ monitors only if the producer provides qualities higher than $p/\underline{\theta}$ with positive probability. Given price p , the quality threshold $p/\underline{\theta}$ is exactly what makes type $\underline{\theta}$ indifferent between buying and not buying. Therefore, every type $\theta > \underline{\theta}$ strictly prefers to buy when they observe a quality equal to $p/\underline{\theta}$. But this implies that the producer can deviate by slightly reducing the quality level of every quality strictly higher than $p/\underline{\theta}$ while still assigning to this set of qualities the same probability. Since these qualities are produced with positive probability, the producer saves on production costs. Furthermore,

¹⁹A formal definition of this strategy space is in Appendix A.

every consumer type that monitors in equilibrium, including $\underline{\theta}$, keeps buying the product upon observing a quality higher than $p/\underline{\theta}$, leaving the gains from trade unaltered.

A similar argument applies when consumer types are infinite. The only difference is that now there may not exist the minimum consumer type that monitors in equilibrium. Therefore, we need to consider $\inf \theta$, the infimum types of consumer that monitors. Notice that $\inf \theta > 0$; if $\inf \theta = 0$, then consumer types close to 0 would strictly prefer to monitor over not buy, but the limit type $\theta = 0$ strictly prefers not to buy.

Dropping credibility. We conclude our analysis by investigating the properties of best responses that are not Shannon, but only robust to non-degenerate entropic costs.

COROLLARY 1. *The following holds:*

- (i) *Let $\mu \in \Delta(Q)$ with $|\text{supp}\mu| > 1$. Then, β is a best response to μ that is robust to non-degenerate entropic costs if and only if it is a best response to μ .*
- (ii) *Let $\mu \in \Delta(Q)$ with $|\text{supp}\mu| = 1$. Then, β is a best response to μ that is robust to non-degenerate entropic costs if and only if it is a constant best response to μ .*

As in Proposition 1, robustness to non-degenerate entropic costs has no bite when quality is stochastic. When quality is deterministic, this assumption removes all incentives to pay attention: if quality is higher or lower than the price, the consumer always or never buys, respectively setting $\beta = 1$ or $\beta = 0$; if quality is equal to the price, the consumer sets a constant $\beta \in [0, 1]$ to avoid incurring $\lambda \cdot \varepsilon > 0$ information costs. Notice that this behavior is consistent with optimality since the logit formula (2) restricts consumer's behavior only for qualities that occur almost surely.

By Corollary 1 point (ii), robustness to non-degenerate entropic costs suffices to exclude the equilibria with non-credible threats discussed in Section 2.1. However, other more sophisticated forms of non-credible threats arise when the best responses are not Shannon. For instance, let $\tilde{q} \in (p, p/\alpha)$ be such that there is $\pi \in (0, 1)$ that makes the producer indifferent between providing 0 and \tilde{q} .²⁰ When best responses are credible, we must ensure that \tilde{q} satisfies a tangency condition as in Lemma 3 to prevent the producer from deviating to other quality levels. This is not necessary when credibility is dropped as the consumer can set $\beta(q) = 0$ for all $q \notin \{0, \tilde{q}\}$, preventing any deviation. Differently from the non-credible threats previously discussed, the consumer incurs positive information

²⁰For all $\lambda < \bar{\lambda}$, the existence of at least one such $\tilde{q} \in (p, p/\alpha)$ follows by Theorem 1.

costs while setting $\pi \in (0, 1)$. Nevertheless, one can show that, for every $\tilde{q} \in (p, p/\alpha)$ for which such a π exists, there is a producer randomization over 0 and \tilde{q} that incentivizes the consumer to set π at the equilibrium level.

We say that the profile (σ, β) is an *equilibrium with non-degenerate entropic costs* if it is a Nash equilibrium and β is robust to non-degenerate entropic costs.

COROLLARY 2. *Fix (α, λ, p) . The profile (σ, β) is an equilibrium with non-degenerate entropic costs and positive quality provision if and only if $|\text{supp}\sigma| \in \{2, 3\}$, $\sigma(0) := \sigma_0 \in (0, 1)$, $\mathbb{E}_\sigma[\beta] := \pi^* \in (0, 1)$, and the following conditions hold:*

(i) *For every $q \in \text{supp}\sigma$, $\Pi(q, \beta) = k$ for some $k \geq 0$; (indifference)*

(ii) *$\pi^* \in \arg\max \sum_q \sigma(q) \cdot (\lambda \cdot \ln(\pi \cdot e^{\frac{q-p}{\lambda}} + 1) - \pi)$ (consumer optimal)*

Sufficiency follows from the previous discussion. If the producer is indifferent across two or three quality levels, then the consumer can prevent deviation by never purchasing off-equilibrium qualities at no additional attention costs. The consumer optimal condition follows by Lemma 2 in Matějka and McKay (2015).

To understand why, in equilibrium, the producer randomizes across at most three quality levels, consider again the argument outlined in Section 3.3. Without credibility, the consumer's best response satisfies equation (2) only for qualities that occur in equilibrium. Therefore, the linear isoprofit curve does not need to lie above the consumer demand if this were defined everywhere but can intersect it. When π is sufficiently high and β defined everywhere is strictly concave, there are at most two intersection points; when π is lower and β is s-shaped, there are at most three intersection points.

Appendix

The appendix is organized as follows. Appendix A examines the holdup problem discussed in Section 1.2. Appendix B provides formal proofs of our main results from Section 2.2 and Section 3. Finally, Appendix C presents the results related to Section 4.

A Costly monitoring

This section considers different versions of the quality provision game described in Section 2. They match the models discussed in Section 1.2.

Perfect monitoring. The consumer has access to a restricted set of attention strategies,

$$\mathcal{B}^{pm} := \Delta\left(\{0\} \cup \{1\} \cup \{\beta_\gamma = \mathbb{1}_{\{q>p\}} + \gamma \cdot \mathbb{1}_{\{q=p\}} : \gamma \in [0, 1]\}\right).$$

That is, the consumer never buys, always buys, or buys any quality strictly higher than the price with probability one and a quality equal to the price with probability $\gamma \in [0, 1]$. Furthermore, we allow the consumer to randomize across these strategies. Notice that this is equivalent to a *perfect monitoring technology* that detects every quality perfectly.

For every producer's strategy $\sigma \in \Delta(Q)$, the consumer selects $\beta \in \mathcal{B}^{pm}$ to maximize

$$U_\kappa(\sigma, \beta) = \int \beta(q) \cdot (q - p) d\sigma(q) - \kappa(\beta),$$

where the monitoring cost $\kappa(\beta) = \kappa \cdot \mathbb{1}_{\{\beta \neq 0, \beta \neq 1\}}$ with $\kappa > 0$, meaning that the consumer incurs a fixed monitoring cost. The rest of the game is unchanged.

The strategy profile $(\sigma, \beta) = (\mathbb{1}_{\{q=0\}}, 0)$ is a Nash equilibrium. If the consumer never buys, the unique best response of the producer is never to produce a positive quality. If the producer does not produce a positive quality, the unique best response of the consumer is never to buy. This is the pessimistic expectations equilibrium.

To see that no other equilibrium exists, consider the following steps.

First, if any equilibrium with $\sigma \neq \mathbb{1}_{\{q=0\}}$, the consumer has to monitor with positive probability. Assume $\beta \in \Delta(\{0, 1\})$. For every $\sigma \neq \mathbb{1}_{\{q=0\}}$, the producer best responds by setting quality equal to zero,

$$\Pi(\mathbb{1}_{\{q=0\}}, \beta) = p \cdot \beta > p \cdot \beta - \alpha \cdot \mathbb{E}_\sigma[q] = \Pi(\sigma, \beta).$$

As a consequence, the producer randomizes across different qualities. Let $\sigma = \mathbb{1}_{\{q=\tilde{q}\}}$ for some $\tilde{q} \in [0, p/\alpha]$. If $\tilde{q} \in [0, p)$, then the consumer best respond by setting $\beta = 0$; if $\tilde{q} \in (p, p/\alpha]$, then the consumer sets $\beta = 1$; $\tilde{q} = p$, then the consumer is indifferent between playing 0 and 1 and never monitors.

In any equilibrium, $\text{supp } \sigma \cap (p, p/\alpha] \neq \emptyset$. For every equilibrium with $\text{supp } \sigma \subseteq [0, p]$ in which the consumer has to monitor with positive probability, it is optimal to deviate and never buy, $U_\kappa < 0 = U_\kappa(\sigma, 0)$.

Assume (σ, β) is an equilibrium with $\sigma \neq \mathbb{1}_{\{q=0\}}$. Then, it has to satisfy the properties described above, namely, the consumer monitors with positive probability, the producer randomizes, and $\text{supp } \sigma \cap (p, p/\alpha] \neq \emptyset$. Let $\hat{q} = \mathbb{E}_\sigma[q] > 0$. Consider a strategy $\tilde{\sigma}$ such that $\sigma(A) = \tilde{\sigma}(A)$ for every Borel measurable $A \subseteq [0, p]$, and $\tilde{q} = \mathbb{E}_{\tilde{\sigma}}[q] = \hat{q} - \epsilon$ for some $\epsilon > 0$. We have that,

$$\begin{aligned}
\Pi(\sigma, \beta) &= p \cdot \int \beta(q) d\sigma(q) - \alpha \bar{q} = p \cdot \left(\int_{[0,p)} \beta(q) d\sigma(q) + \beta(p)\sigma(p) + \int_{(p,p/\alpha]} \beta(q) d\sigma(q) \right) - \alpha \bar{q} \\
&= p \cdot \left(\int_{[0,p)} \beta(q) d\tilde{\sigma}(q) + \beta(p)\tilde{\sigma}(p) + \int_{(p,p/\alpha]} \beta(q) d\sigma(q) \right) - \alpha \bar{q} \\
&= p \cdot \left(\int_{[0,p)} \beta(q) d\tilde{\sigma}(q) + \beta(p)\tilde{\sigma}(p) + \bar{\beta} \int_{(p,p/\alpha]} d\sigma(q) \right) - \alpha \bar{q} \\
&= p \cdot \left(\int_{[0,p)} \beta(q) d\tilde{\sigma}(q) + \beta(p)\tilde{\sigma}(p) + \bar{\beta} \int_{(p,p/\alpha]} d\tilde{\sigma}(q) \right) - \alpha \bar{q} \\
&< p \cdot \left(\int_{[0,p)} \beta(q) d\tilde{\sigma}(q) + \beta(p)\tilde{\sigma}(p) + \int_{(p,p/\alpha]} \beta(q) d\tilde{\sigma}(q) \right) - \alpha \bar{q} = \Pi(\tilde{\sigma}, \beta),
\end{aligned}$$

where the first, second, and last equalities holds by definition, the third equality holds since $\sigma(q) = \tilde{\sigma}(q)$ for every $q \leq p$, the fourth equality holds since every $\beta \in \mathcal{B}^{pm}$ is constant on $(p, p/\alpha]$, which we set equal to $\bar{\beta}$ in this case, the fifth equality holds since $\sigma((p, p/\alpha]) = \tilde{\sigma}((p, p/\alpha])$, and the inequality holds since $\tilde{q} < \hat{q}$.

Imperfect monitoring. Assume now that the consumer has access to a different set of attention strategies. For a fixed $\bar{q} \in (p, p/\alpha]$, the consumer action space is

$$\mathcal{B}^{\bar{q}} := \Delta \left(\{0\} \cup \{1\} \cup \{\beta_{\bar{q}} = \mathbb{1}_{\{q \geq \bar{q}\}}\} \right).$$

That is, the consumer never buys, always buys, or buys with probability one whenever the quality is higher than or equal to the threshold $\bar{q} \in (p, p/\alpha]$. We allow the consumer

to randomize over these strategies. Notice that this represents an *imperfect monitoring technology* that detects whether the quality is above or below the threshold \bar{q} . To simplify, we implicitly assume that the consumer does not buy when she receives the signal that quality is below the threshold \bar{q} ; removing this assumption does not alter our equilibrium analysis. The utility U_κ and the rest of the problem is unchanged.

The pessimistic expectations equilibrium remains a valid equilibrium. The argument is analogous to the previous one and is thus omitted.

In this case, two equilibria (σ_1, β_1) and (σ_2, β_2) , different than pessimistic expectations, exist if $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$. These equilibria are defined as:

$$\beta_1 = \frac{\bar{\alpha}}{p} \cdot \beta_{\bar{q}} + \left(1 - \frac{\bar{\alpha}}{p}\right) \cdot 0 \quad \text{and} \quad \sigma_1 = \frac{\kappa}{\bar{q} - p} \cdot \bar{q} + \left(1 - \frac{\kappa}{\bar{q} - p}\right) \cdot 0$$

$$\beta_2 = \frac{\bar{\alpha}}{p} \cdot \beta_{\bar{q}} + \left(1 - \frac{\bar{\alpha}}{p}\right) \cdot 1 \quad \text{and} \quad \sigma_2 = \frac{p - \kappa}{p} \cdot \bar{q} + \left(1 - \frac{p - \kappa}{p}\right) \cdot 0,$$

where $\bar{\alpha} = \alpha \cdot \bar{q}$, which, by hypothesis, satisfies $\bar{\alpha} \leq p$. Furthermore, $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$ ensures $\kappa/(\bar{q} - p) \in (0, 1)$ and $(p - \kappa)/p \in (0, 1)$.

In both equilibria, the consumer monitors with probability α/p . In β_1 , the consumer never buys when she does not monitor, while in β_2 , she always buys when she does not monitor. The producer is randomizing between \bar{q} and 0: in σ_1 the weight associated with \bar{q} is $\kappa/(\bar{q} - p)$, while in σ_2 is $(p - \kappa)/p$. Notice that, since $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$,

$$1 > \frac{p - \kappa}{p} > \frac{p}{\bar{q}} > \frac{\kappa}{\bar{q} - p} > 0.$$

Finally, for $\kappa \rightarrow 0$, $\sigma_1(\bar{q}) \rightarrow 0$, while $\sigma_2(\bar{q}) \rightarrow 1$.

We check that (σ_1, β_1) is an equilibrium. First, notice that the producer is not willing to deviate to other qualities. Under β_1 , any $q > \bar{q}$ induces the same trade probability as \bar{q} , but entails higher production costs; any $q \in (0, \bar{q})$ induces the same trade probability as zero quality, but higher costs than zero quality. Furthermore, the producer is indifferent between producing 0 and \bar{q} ,

$$p \cdot \beta_1(\bar{q}) - \bar{\alpha} = p \cdot \beta_1(0) = 0 \iff \beta_1(\bar{q}) = \frac{\bar{\alpha}}{p}.$$

The consumer is indifferent between $\beta_{\bar{q}}$ and 0, which strictly prefers over 1. To see this,

$$\sigma_1(\bar{q}) \cdot (\bar{q} - p) - \kappa = 0 > \sigma_1(\bar{q}) \cdot \bar{q} - p \iff \sigma_1(\bar{q}) = \frac{\kappa}{\bar{q} - p} < \frac{p}{\bar{q}}.$$

Therefore, (σ_1, β_1) is an equilibrium.

We now show that (σ_2, β_2) is an equilibrium. The argument establishing that the producer is not willing to deviate to other quality levels is analogous to the one above and thus omitted. The producer is indifferent between producing 0 and \bar{q} ,

$$p \cdot \beta_2(\bar{q}) - \bar{\alpha} = p \cdot \beta_2(0) = p \cdot \left(1 - \frac{\bar{\alpha}}{p}\right) \iff \beta_2(\bar{q}) = 1.$$

The consumer is indifferent between $\beta_{\bar{q}}$ and 1, which strictly prefers over 0. To see this,

$$\sigma_2(\bar{q}) \cdot (\bar{q} - p) - \kappa = \sigma_2(\bar{q}) \cdot \bar{q} - p > 0 \iff \sigma_2(\bar{q}) = \frac{p - \kappa}{p} > \frac{p}{\bar{q}}.$$

Therefore, (σ_2, β_2) is an equilibrium.

For every $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$, (σ_1, β_1) and (σ_2, β_2) are no longer equilibria. However, if $\kappa = \frac{p}{\bar{q}}(\bar{q} - p)$, they remain equilibria as the above argument still applies. In particular,

$$1 > \frac{p - \kappa}{p} = \frac{p}{\bar{q}} = \frac{\kappa}{\bar{q} - p} > 0.$$

Due to this indifference it emerges a class of equilibria that mixes the previous two: the producer plays $\sigma_1 = \sigma_2$; the consumer plays, with probability $\bar{\alpha}/p$, $\beta_{\bar{q}}$, and, with probability $1 - \bar{\alpha}/p$, any randomization between 0 and 1.

Flexible imperfect monitoring. Finally, assume that the consumer has access to the following attention strategies

$$\bar{\mathcal{B}} := \Delta \left(\{0\} \cup \{1\} \cup \{ \beta_{\bar{q}} = \mathbb{1}_{\{q \geq \bar{q}\}} : \bar{q} \in (p, p/\alpha] \} \right).$$

That is, the consumer never buys, always buys, or buys with probability one whenever the quality is higher than or equal to some threshold $\bar{q} \in (p, p/\alpha]$. We allow the consumer to randomize over these strategies. The utility U_κ and the rest of the problem is unchanged.

For every \bar{q} with $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$, the strategy profile (σ_i, β_i) , for $i \in \{1, 2\}$, defined as above with respect to \bar{q} is an equilibrium. Since the strategy space of the producer is

the same, σ_i is a best response to β_i . At the same time, the consumer does not want to deviate to any other detection threshold: playing β_q with $q < \bar{q}$ yields the same utility as $\beta_{\bar{q}}$, while playing one with $q > \bar{q}$ yields $-\kappa$. It follows that β_i is a best response to σ_i .

B Rational inattention

Shannon-Nash equilibrium. We start by showing the results of Section 2.2. The following proposition discusses the artificial equilibrium with zero information costs.

PROPOSITION 8. *Fix (α, p) . There is a credible equilibrium with deterministic and positive quality provision if and only if $\lambda \leq \bar{\lambda}$, where $\bar{\lambda} > 0$. In this equilibrium, the seller provides quality $\sigma = \mathbb{1}_{\{q=p\}}$, the consumer accepts with probability $\pi = \frac{1}{2} + \frac{1}{2} \cdot \sqrt{1 - 4 \cdot \alpha \lambda / p} \in (0, 1)$. In particular, the consumer incurs no information costs, $I(\sigma, \beta) = 0$.*

Proof. We show that $\sigma = \delta_p$ must hold in any equilibrium with deterministic and positive quality provision. Indeed, if $\sigma = \delta_{\tilde{q}}$ with $\tilde{q} < p$, then $\pi = 0$ is the unique best response. Similarly, if $\sigma = \delta_{\tilde{q}}$ with $\tilde{q} > p$, then $\pi = 1$ is the unique best response. The former case yields a deterministic equilibrium with non-positive quality, while the latter case does not lead to any equilibrium as the seller profitably deviates to $\sigma = \delta_0$ when $\pi = 1$.

Let $\sigma = \delta_p$. In this case, the consumer is indifferent to any $\pi \in [0, 1]$ and, in equilibrium, suitably set $\pi = \pi^* \in (0, 1)$ such that the producer finds it optimal to provide a quality level equal to the price. To see this, notice that $I(\mu, \beta) = 0$ whenever $|\text{supp } \mu| = 1$ for $\mu \in \Delta(Q)$.

The producer solves

$$\max_q \beta(q) \cdot p - \alpha \cdot q \quad (3)$$

where, by Lemma 1, β is determined by equation (2). In particular, fixing any $\pi \in (0, 1)$ in equation (2), we have that β is twice continuously differentiable, increasing, strictly convex for all q such that $\beta(q) < 1/2$, and strictly concave for all q such that $\beta(q) > 1/2$. These facts follow from simple calculations:

$$\frac{d}{dq} \beta(q) = \frac{1}{\lambda} \beta(q) (1 - \beta(q)), \quad (4)$$

which is always positive if $\pi \in (0, 1)$, and

$$\frac{d^2}{dq^2} \beta(q) = \frac{1}{\lambda^2} \beta(q) (1 - \beta(q)) (1 - 2\beta(q)),$$

which is positive and negative according to the conditions specified above.

Plugging in equation (4) in the first order conditions of problem (3), we obtain

$$\frac{d}{dq}\beta(q) \cdot p - \alpha = 0 \implies \beta(q)(1 - \beta(q)) = \frac{\alpha \cdot \lambda}{p}.$$

This equation admits two solutions for $\beta(q)$, corresponding to the local minimum and the local maximum. Since we are interested in the latter, which is obtained for values of q such that $\beta(q) > 1/2$, the relevant solution is

$$\beta(q) = \frac{1}{2} + \frac{\sqrt{1 - 4 \cdot \alpha \lambda / p}}{2}. \quad (5)$$

Therefore, in order for $\sigma = \delta_p$ to be optimal, the consumer sets $\pi^* = \beta(p) = 1/2 + 1/2 \cdot \sqrt{1 - 4 \cdot \alpha \lambda / p}$. Let β^* the resulting attention strategy defined by equation (2) when $\pi = \pi^*$. In this equilibrium, the producer's profits are

$$\Pi(\delta_p, \beta^*) = \frac{1 + \sqrt{1 - 4 \cdot \alpha \lambda / p}}{2} \cdot p - \alpha \cdot p \quad (6)$$

which are positive if and only if $\beta(p) \geq \alpha$. Furthermore, this equilibrium exists as long as $\beta(p)$ is well-defined, i.e., for all α, λ, p such that $\alpha \lambda / p < 1/4$. Rewriting this condition in terms of the parameter λ we obtain $\lambda \in (0, \lambda^e)$ where $\lambda^e := p/(4\alpha)$.

We need to ensure there are no possible deviations for the producer. Notice that, if one of such deviation exists, it must be in the convex region of β , i.e., the set of all $q \geq 0$ such that $\beta(q) < 1/2$, since π^* is set to make $\sigma = \delta_p$ optimal in the concave region. By convexity, if a deviation exists in the convex region, it must be $q = 0$. Therefore, an equilibrium with deterministic quality provision exists for all $\lambda > 0$ such that

$$\Pi(\delta_p, \beta^*) = K(\lambda) \cdot p - \alpha \cdot p \geq p \cdot \frac{K(\lambda) \cdot e^{-\frac{p}{\lambda}}}{K(\lambda) \cdot e^{-\frac{p}{\lambda}} + 1 - K(\lambda)} = p \cdot \beta^*(0) = \Pi(\delta_0, \beta^*),$$

where $K(\lambda) := 1/2 + 1/2 \cdot \sqrt{1 - 4 \cdot \alpha \lambda / p}$. Since $\Pi(\sigma_0, \beta^*) \geq 0$, then the producer's profits as displayed in equation (6) are positive. Let

$$\alpha(\lambda) := K(\lambda) - \frac{K(\lambda) \cdot e^{-\frac{p}{\lambda}}}{K(\lambda) \cdot e^{-\frac{p}{\lambda}} + 1 - K(\lambda)} = K(\lambda) - \frac{1}{1 + \frac{1 - K(\lambda)}{K(\lambda)} e^{\frac{p}{\lambda}}}. \quad (7)$$

We show that there exists a unique $\bar{\lambda} \in (0, \lambda^e)$ such that $\alpha(\bar{\lambda}) = \alpha$. First notice that, $\lim_{\lambda \rightarrow 0} \alpha(\lambda) = 1 > \alpha$. Furthermore,

$$\lim_{\lambda \rightarrow \lambda^e} \alpha(\lambda) = 1/2 - \frac{1}{1 + e^{4\alpha}} < \alpha,$$

where the inequality holds for all $\alpha > 0$. We now show $\bar{\lambda}$ is unique. Notice that $\frac{d}{d\lambda} K(\lambda) = \frac{-\alpha/p}{\sqrt{1-4\alpha\lambda/p}} < 0$, which implies that $\frac{d}{d\lambda} \frac{1-K(\lambda)}{K(\lambda)} > 0$. In particular,

$$\frac{d}{d\lambda} \frac{1-K(\lambda)}{K(\lambda)} = \frac{d}{d\lambda} \left(\frac{1}{K(\lambda)} - 1 \right) = \frac{\alpha/p}{K(\lambda)^2 \cdot \sqrt{1-4\alpha\lambda/p}},$$

which is finite for low values of λ , i.e., there exists $\epsilon > 0$, such that $\frac{d}{d\lambda} \frac{1-K(\lambda)}{K(\lambda)} < +\infty$ for all $\lambda \in (0, \epsilon)$. Therefore,

$$\frac{d}{d\lambda} \frac{1-K(\lambda)}{K(\lambda)} e^{\frac{p}{\lambda}} = e^{\frac{p}{\lambda}} \left(\frac{d}{d\lambda} \frac{1-K(\lambda)}{K(\lambda)} + \frac{1-K(\lambda)}{K(\lambda)} \cdot \left(\frac{-p}{\lambda^2} \right) \right).$$

By inspection, this derivative is negative for small values of $\lambda > 0$, and possibly positive for larger values of λ . As a result, $\alpha(\lambda)$ is decreasing for small values of λ , and possibly increasing for larger values. By combining this observation with $\lim_{\lambda \rightarrow 0} \alpha(\lambda) - \alpha > 0$ and $\lim_{\lambda \rightarrow \lambda^e} \alpha(\lambda) - \alpha < 0$, we obtain there is at most one λ such that $\alpha(\bar{\lambda}) = \alpha$.

Finally, the same argument establishes that $\Pi(\delta_p, \beta^*) \geq \Pi(\delta_0, \beta^*)$ for every $\lambda \in (0, \bar{\lambda}]$, concluding the proof. ■

PROOF OF PROPOSITION 1. We show points (i) and (ii) separately.

(i) “ \implies ” This direction is trivial.

“ \impliedby ” Let β be a credible best response to some $\mu \in \Delta(Q)$ with $|\text{supp}\mu| > 1$. We want to show that β is a Shannon best response, i.e., there exists $\epsilon > 0$ such that $\beta \in \text{argmax} \int \beta(q)(q-p)dq - \lambda \cdot I_\epsilon(\beta, \mu)$.

We distinguish two cases. First, let $\pi := \mathbb{E}_\mu[\beta] \in \{0, 1\}$. Then, β is also a Shannon best response since β is constant everywhere and $I_\epsilon(\beta, \mu) = I(\beta, \mu) = 0$.

Now, assume $\pi \notin \{0, 1\}$. Since $|\text{supp}\mu| > 1$ and β is credible, equation (2) holds, which implies that β is non-constant μ -almost surely. As a result $I(\beta, \mu) > 0$. Therefore, by

setting $\varepsilon := I(\beta, \mu)$, we have that, for every attention strategy $\tilde{\beta}$,

$$\int \tilde{\beta}(q)(q-p) d\sigma(q) - \lambda \cdot I_\varepsilon(\tilde{\beta}, \mu) \leq U(\tilde{\beta}) \leq U(\beta) = \int \beta(q)(q-p) d\sigma(q) - \lambda \cdot I_\varepsilon(\beta, \mu),$$

where the first inequality holds since $I_\varepsilon(\tilde{\beta}, \mu) \geq I(\tilde{\beta}, \mu)$, the second since β is a best response, and the last equality since $I(\beta, \mu) = I_\varepsilon(\beta, \mu)$.

(ii) “ \Leftarrow ” Let $|\text{supp}\mu| = 1$ for some $\mu \in \Delta(Q)$. First notice that, if $\text{supp}\mu \neq p$, then any credible best response to μ is constant on $[0, p/\alpha]$ implying is also robust to non-degenerate entropic costs. If $\text{supp}\mu = p$, then $\tilde{\beta}$ with $\mathbb{E}_\mu[\tilde{\beta}] \in \{0, 1\}$ is constant on $[0, p/\alpha]$ and hence robust to non-degenerate entropic costs.

“ \Rightarrow ” Let $|\text{supp}\mu| = 1$ for some $\mu \in \Delta(Q)$. We show that any credible best response $\tilde{\beta}$ with $\mathbb{E}_\mu[\tilde{\beta}] \in (0, 1)$ is not robust to non-degenerate entropic costs. Since $\tilde{\beta}$ is credible, by Lemma 1, $\tilde{\beta}$ is not constant on $[0, p/\alpha]$. Furthermore, since $\tilde{\beta}$ is a best response, by the previous direction, it must hold that $\text{supp}\mu = p$. Finally, for every $\varepsilon > 0$, $U(\beta) = 0 > -\lambda \cdot I_\varepsilon(\tilde{\beta}, \mu) = -\lambda \cdot \varepsilon$, where the strategy β satisfies $\mathbb{E}_\mu[\beta] = 1$. Therefore, $\tilde{\beta}$ has to satisfy $\mathbb{E}_\mu[\tilde{\beta}] \in \{0, 1\}$ as desired. ■

COROLLARY 3. *The credible equilibrium with deterministic and positive quality provision described in Proposition 8 is not Shannon-Nash.*

Proof. The statement follows immediately from point (ii) of Proposition 1, since, in any Shannon-Nash equilibrium, $\pi \in \{0, 1\}$ whenever $\sigma = \delta_p$. ■

Main result. The proof of Theorem 1 proceeds as follows. First, we show Proposition 2, establishing that any Shannon-Nash equilibrium with positive quality provision is binary-quality, i.e., has to satisfy $|\text{supp}\sigma| = 2$. We then construct one binary-quality equilibrium and show no other binary-quality equilibria exist, implying uniqueness. Finally, we prove that the properties listed in Theorem 1.

PROOF OF PROPOSITION 2. Consider the following lemma.

LEMMA 5. *In any credible equilibrium (σ, β) with $|\text{supp}\sigma| > 1$, $\pi := \mathbb{E}_\sigma[\beta] \in (0, 1)$.*

Proof. If $\pi = 0$, then the producer’s profits are negative, which implies that setting quality equal to zero is a profitable deviation; if $\pi = 1$, then the consumer always accept, and again setting quality equal to zero is a profitable deviation. □

LEMMA 6. *In any Shannon-Nash equilibrium with positive quality provision, the producer randomizes between qualities below and above the price:*

$$\text{supp } \sigma \cap [0, p) \neq \emptyset \quad \text{and} \quad \text{supp } \sigma \cap (p, p/\alpha] \neq \emptyset.$$

Proof. Let (σ, β) be a Shannon-Nash equilibrium with positive quality provision. Assume $|\text{supp } \sigma| = 1$. Then by Proposition 1 point (ii), $\pi \in \{0, 1\}$. By credibility, $\beta = 1$ everywhere, but, in this case, it is optimal for the producer to set quality to zero. The statement follows by Lemma 5. \square

Recall that, by Lemma 1, any credible best response β satisfies equation (2). Fixing any $\pi \in (0, 1)$ in equation (2), we have that $\beta(q)$ has range in $[\beta(0), 1)$, is strictly increasing, strictly convex for all q such that $\beta(q) < 1/2$, and strictly concave for all q such that $\beta(q) > 1/2$. See the proof of Proposition 8 for an argument establishing these facts.

LEMMA 7. *Let $\lambda \in (0, \lambda^e)$ and β be an attention strategy defined by equation (2) for some $\pi \in (0, 1)$. If σ is a best response to β with $|\text{supp } \sigma| > 1$, then $\text{supp } \sigma = \{0, q^*\}$, where $q^* > 0$ is determined by equation (5), i.e.,*

$$\beta(q^*) = \frac{1}{2} + \frac{\sqrt{1 - 4 \cdot \alpha \lambda / p}}{2}.$$

Proof. Fix $\pi \in (0, 1)$, and let β be the corresponding attention strategy defined by equation (2). Notice that, if σ is a best response to β with $|\text{supp } \sigma| > 1$, then β is not strictly concave over the interval $[0, p/\alpha]$. By the properties of β stated above this holds if and only if β is convex for low values of q and concave otherwise, which occurs if and only if $\beta(0) < 1/2$.

We first show that $|\text{supp } \sigma| \in \{2, 3\}$. If σ is a best response to β , then, for every $q \in \text{supp } \sigma$, we have that

$$q \in \text{argmax } \beta(q) \cdot p - \alpha \cdot q,$$

which implies that

$$\beta(q) \cdot p - \alpha \cdot q = \bar{\Pi}, \tag{8}$$

for some profit level $\bar{\Pi} \geq 0$. Notice that $\bar{\Pi} \geq 0$ since $\Pi(\delta_0) = 0$. By equation (8), the producer's isoprofit curves in terms of $\beta(q)$ are linear in q for any profit level $\bar{\Pi}$, i.e.,

$\beta(q) = (\alpha/p) \cdot q + \bar{\Pi}/p$. Thus, for every $q \in \text{supp } \sigma$, we have that

$$h(q) := \beta(q) - \frac{\alpha}{p} \cdot q - \frac{\bar{\Pi}}{p} = 0. \quad (9)$$

Since β is convex for low values of q and concave otherwise, we have that h is convex and concave for the same values. This argument implies that h has at most three zeros for any $\bar{\Pi} \geq 0$. Therefore, if $|\text{supp } \sigma| > 1$, then $|\text{supp } \sigma| \in \{2, 3\}$.

We show that, if $|\text{supp } \sigma| > 1$, then $|\text{supp } \sigma| = 2$. By contradiction, assume $|\text{supp } \sigma| = 3$. This implies that $h(\tilde{q}) > 0$ for some $\tilde{q} > 0$. To see this, notice that, if $h \leq 0$ and h is first convex and then concave, then h has at most two zeros, i.e., a local maximum in the concave region of β and zero. But if $h(\tilde{q}) > 0$ for some $\tilde{q} > 0$, then $\delta_{\tilde{q}}$ is a profitable deviation for the producer. Let $\bar{\Pi}$ be the profit level associated with σ . We have that,

$$h(\tilde{q}) = \beta(\tilde{q}) - \frac{\alpha}{p} \cdot \tilde{q} - \frac{\bar{\Pi}}{p} > 0 \implies \bar{\Pi} < \beta(\tilde{q}) \cdot p - \alpha \cdot \tilde{q} = \Pi(\delta_{\tilde{q}}). \quad (10)$$

We now show that, for every $\bar{\Pi} \geq 0$, there exists a unique $q^* > 0$ satisfying $\frac{d}{dq} h(q^*) = 0$ and $\beta(q^*) > 1/2$. By equation (5), if such a q^* exists, then it is determined by

$$\beta(q^*) = \frac{1}{2} + \frac{\sqrt{1 - 4 \cdot \alpha \lambda / p}}{2} =: K(\lambda).$$

For every $(\alpha, p) \in (0, 1) \times \mathbb{R}_+$, if $\lambda < \lambda^e := p/(4\alpha)$, $K(\lambda)$ is well-defined and has range in $(1/2, 1)$. We have that $\beta(0) < 1/2$. Furthermore, by equation (2), $\lim_{q \rightarrow \infty} \beta(q) = 1$. Since β is continuous, by the intermediate value theorem, we have that there exists $q > 0$ such that $\beta(q) = K(\lambda)$ for some $\lambda < \lambda^e$. Uniqueness follows since β is strictly increasing.

We conclude by showing that $\text{supp } \sigma = \{0, q^*\}$. Equation (10) implies that, if $h(q) > 0$ for some $q \in Q$, then δ_q is a profitable deviation. But this implies that, if σ is a best response to β with $|\text{supp } \sigma| > 1$, then $h(q) \leq 0$ for every $q \in Q$. Furthermore, since h has at most two zeros when $h \leq 0$, if $|\text{supp } \sigma| > 1$ and $\text{supp } \sigma \neq \{0, q^*\}$, then either $h(0) < 0$ or $h(q^*) < 0$. In both of these cases, the function h has at most one zero, leading to a contradiction. \square

COROLLARY 4. *If (σ, β) is a Shannon-Nash equilibrium with positive quality provision, then $\text{supp } \sigma = \{0, q^*\}$ where $q^* > 0$ is defined by equation (5).*

This concludes the proof of Proposition 2. \blacksquare

PROOF OF THEOREM 1. In order for the producer to be indifferent across two or more qualities, it is necessary that β is not strictly concave over the interval $[0, p/\alpha]$. Given the shape of β , this is achieved if and only if $\beta(0) < 1/2$. Therefore, the consumer sets π s.t.

$$\beta(0) = \frac{\pi \cdot e^{-\frac{p}{\lambda}}}{\pi \cdot e^{-\frac{p}{\lambda}} + 1 - \pi} < \frac{1}{2} \iff \pi < \frac{1}{e^{-\frac{p}{\lambda}} + 1} =: \pi_\lambda.$$

For every $(\alpha, p) \in (0, 1) \times \mathbb{R}_+$, $\lambda < \lambda^e$, if $\pi < \pi_\lambda$, and thus $\beta(0) < 1/2$, the proof above establishes the existence of $q^* > 0$ that satisfies equation (5).

LEMMA 8. *Let $\lambda \in (0, \lambda^e)$ and β be an attention strategy defined by equation (2) for some $\pi \in (0, \pi_\lambda)$ such that $\Pi(\delta_0, \beta) = \Pi(\delta_{q^*}, \beta)$, where $q^* > 0$ is defined by equation (5). If $\text{supp } \sigma = \{0, q^*\}$, then σ is best response to β . Furthermore, if $\text{supp } \sigma \not\subseteq \{0, q^*\}$, then σ is not a best response to β .*

Proof. The first part of the statement follows since $q^* \in \arg \max \beta(q) \cdot p - \alpha \cdot q$. If $|\text{supp } \sigma| > 1$, the second part of the statement follows by Lemma 7; if $|\text{supp } \sigma| = 1$, it follows by strict concavity of β at q^* . \square

LEMMA 9. *For every $\lambda \in (0, \lambda^e)$, there exists a unique $\pi \in (0, \pi_\lambda)$ s.t. the attention strategy β induced by equation (2) satisfies $\Pi(\delta_0, \beta) = \Pi(\delta_{q^*}, \beta)$, where $q^* > 0$ satisfies equation (5).*

Proof. For $\lambda < \lambda^e$, we show the existence of a unique $\pi \in (0, 1)$ such that, for some $q^* > 0$,

$$(i) \quad p \cdot \frac{\pi \cdot e^{-\frac{p}{\lambda}}}{\pi \cdot e^{-\frac{p}{\lambda}} + 1 - \pi} = p \cdot \frac{\pi \cdot e^{\frac{q^*-p}{\lambda}}}{\pi \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi} - \alpha \cdot q^*$$

$$(ii) \quad \frac{\pi \cdot e^{\frac{q^*-p}{\lambda}}}{\pi \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi} = K(\lambda)$$

Notice that, for every $q^* > 0$, there exists a unique $\pi \in (0, 1)$ satisfying point (ii). We proceed by showing the existence of a unique $q^* > 0$ satisfying both points (i) and (ii), and later check that the corresponding π belongs to $(0, 1)$. In particular, if for some $q^* > 0$, there exists $\pi \in (0, 1)$ such that points (i) and (ii) hold, then the induced attention strategy β cannot be strictly concave over $[0, p/\alpha]$. As a consequence, $\pi \in (0, \pi_\lambda)$.

By point (ii), we express π as a function of λ and q^* ,

$$\pi = \frac{K(\lambda)}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q^*-p}{\lambda}}}. \quad (11)$$

After some calculations, we have that

$$\frac{1 - \pi}{\pi} = \frac{(1 - K(\lambda)) \cdot e^{\frac{q^* - p}{\lambda}}}{K(\lambda)}.$$

As a consequence, we can express $\beta(0)$ as a function of λ and q^* as well,

$$\beta(0) = \frac{\pi \cdot e^{\frac{-p}{\lambda}}}{\pi \cdot e^{\frac{-p}{\lambda}} + 1 - \pi} = \frac{K(\lambda)}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q^*}{\lambda}}}. \quad (12)$$

By combining equation (12) and point (i) above, we obtain an equation describing q^* as a function of λ ,

$$p \cdot \frac{K(\lambda)}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q^*}{\lambda}}} + \alpha \cdot q^* - p \cdot K(\lambda) = 0 \quad (13)$$

Notice that $q^* = 0$ is always a solution for every $\lambda > 0$. We now show that, for every $\lambda > 0$, there exists a unique $q^* > 0$ satisfying (13). The value of π is retrieved by equation (11).

Denote by $f(q)$ the left hand side of equation (13). We have that

$$\frac{d}{dq} f(q) = \alpha \cdot \frac{p e^{\frac{q}{\lambda}} K(\lambda) (1 - K(\lambda))}{\lambda (K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q}{\lambda}})^2} = \alpha \cdot \left(1 - \frac{e^{\frac{q}{\lambda}}}{(K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q}{\lambda}})^2} \right),$$

where the second equality follows since $K(\lambda)(1 - K(\lambda)) = \alpha \lambda / p$. Notice, $\frac{d}{dq} f(0) = 0$. Furthermore, by rearranging, we have that

$$\frac{d}{dq} f(q) < 0 \iff A(q) := e^{\frac{2q}{\lambda}} (1 - K(\lambda))^2 + e^{\frac{q}{\lambda}} (2K(\lambda)(1 - K(\lambda)) - 1) + K(\lambda) < 0.$$

It follows that

$$\frac{d}{dq} A(q) = \frac{2}{\lambda} \cdot e^{\frac{2q}{\lambda}} (1 - K(\lambda))^2 + \frac{1}{\lambda} \cdot e^{\frac{q}{\lambda}} (2K(\lambda)(1 - K(\lambda)) - 1),$$

which implies

$$\frac{d}{dq} A(0) = \frac{1}{\lambda} \cdot (1 - 2K(\lambda)) < 0,$$

where the inequality follows since $K(\lambda) > 1/2$. Since $A(0) = 0$, this implies that $f(q) < 0$

for some $q > 0$ in a neighborhood of zero. Furthermore, after some calculations,

$$\frac{d^2}{dq^2}f(q) = -\frac{\alpha e^{\frac{q}{\lambda}}}{\lambda(K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q}{\lambda}})^2} \cdot \left(1 - \frac{2e^{\frac{q}{\lambda}}(1-K(\lambda))}{K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q}{\lambda}}}\right).$$

Since $K(\lambda) \in (1/2, 1)$, it follows that $\frac{d^2}{dq^2}f(q) < 0$ for small values of $q > 0$ and positive otherwise. Indeed,

$$\begin{aligned} 1 - \frac{2e^{\frac{q}{\lambda}}(1-K(\lambda))}{K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q}{\lambda}}} \Big|_{q=0} &> 0 \\ \lim_{q \rightarrow +\infty} 1 - \frac{2e^{\frac{q}{\lambda}}(1-K(\lambda))}{K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q}{\lambda}}} &< 0 \\ \frac{d}{dq} \left(1 - \frac{2e^{\frac{q}{\lambda}}(1-K(\lambda))}{K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q}{\lambda}}} \right) &< 0 \end{aligned}$$

To summarize, $f(q)$ behaves as follows: $f(0) = 0$, $\frac{d}{dq}f(0) = 0$, $f(q) < 0$ for some $q > 0$ in a neighborhood of zero, f is concave for small values of q and then convex, and $\lim_{q \rightarrow +\infty} f(q) = +\infty$. This implies that, for every $\lambda < \lambda^e$, there exists a unique $q^* > 0$ such that $f(q^*) = 0$.

For every $\lambda < \lambda^e$, let $q^* > 0$ satisfy equation (13). By equation (11), we retrieve $\pi \in (0, 1)$ as follows

$$\pi = \frac{K(\lambda)}{K(\lambda) + (1-K(\lambda)) \cdot e^{\frac{q^*-p}{\lambda}}}.$$

Since q^* is unique, π is unique as well. Furthermore, $\pi > 0$, and $\pi < 1$ since $(1-K(\lambda)) \cdot e^{\frac{q^*-p}{\lambda}} > 0$. \square

LEMMA 10. *For every $\lambda < \lambda^e$, let $q^*(\lambda)$ be the solution $q^* > 0$ satisfying equation (13) expressed as a function of λ . We have that $q^*(\lambda)$ is continuously differentiable, decreasing, and satisfies $\lim_{\lambda \rightarrow 0} q(\lambda) = p/\alpha$.*

Proof. For every $\lambda < \lambda^e$, the differentiability of $q^*(\lambda)$ follows by the implicit function theorem. Indeed, by Lemma 9, there is a point (λ, q^*) on the curve defined by equation (13). That is, $f(\lambda, q^*) = 0$ for some $\lambda > 0$, where, with a slight abuse of notation, f denotes the left hand side of equation (13) as a function of λ and q . Furthermore, by the proof

of Lemma 9, $\frac{d}{dq}f(\lambda, q^*) > 0$. This implies that in a neighborhood of (λ, q^*) we can write $q^* = q^*(\lambda)$ for a continuously differentiable $q^*(\lambda)$.

We now investigate how $q^*(\lambda)$ behaves when $\lambda \rightarrow 0$. We established that for every $\lambda < \lambda^e$, there exists a unique $q^* > 0$ solving equation (13). Let $\bar{q} = \lim_{\lambda \rightarrow 0} q^*(\lambda)$. For $\bar{q} \in (p, p/\alpha)$, $\lim_{\lambda \rightarrow 0} K(\lambda) = 1$, $\lim_{\lambda \rightarrow 0} (1 - K(\lambda)) \cdot e^{\frac{\bar{q}}{\lambda}} = +\infty$, and $\lim_{(\lambda, q) \rightarrow (0, \bar{q})} f(\lambda, q) = p - \alpha \cdot \bar{q}$. Therefore, it must be $\lim_{\lambda \rightarrow 0} q^*(\lambda) = p/\alpha$. Notice that this argument implies that $\lim_{\lambda \rightarrow 0} \beta(0) = 0$.

We show that, for every $q \in (p, p/\alpha)$, there exists a unique $\lambda < \lambda^e$ such that $f(\lambda, q) = 0$. By Lemma 9, this establishes a one-to-one relationship between q and λ . Together with the fact that $q^*(\lambda) \rightarrow p/\alpha$ as $\lambda \rightarrow 0$, it implies that $q^*(\lambda)$ is decreasing.

Fix $\bar{q} \in (p, p/\alpha)$. By equation (13), we have that $\lim_{\lambda \rightarrow 0} f(\lambda, \bar{q}) = \alpha \cdot \bar{q} - p < 0$, and

$$\begin{aligned} \lim_{\lambda \rightarrow \lambda^e} f(\lambda, \bar{q}) &= \frac{p}{1 + e^{4\alpha \cdot \frac{\bar{q}}{p}}} + \alpha \cdot \bar{q} - 1/2 \cdot p \\ &= p \cdot \left(\frac{1}{1 + e^{4\alpha \cdot \frac{\bar{q}}{p}}} + \alpha \cdot \frac{\bar{q}}{p} - 1/2 \right) = p \cdot \left(\frac{1}{1 + e^{4\gamma}} + \gamma - 1/2 \right) > 0, \end{aligned}$$

where $\gamma := \alpha \cdot \frac{\bar{q}}{p} > 0$, and the last step is a known inequality.

By the intermediate value theorem, since, for every $\bar{q} \in (p, p/\alpha)$, $f(\lambda, \bar{q})$ is continuous in λ , we have that there exists $\hat{\lambda} \in (0, \lambda^e)$ such that $f(\hat{\lambda}, \bar{q}) = 0$.

We now show that such a $\hat{\lambda}$ is unique. Fix $\bar{q} \in (p, p/\alpha)$, we have that

$$\frac{d}{d\lambda} f(\lambda, \bar{q}) = \frac{d}{d\lambda} \left(\frac{p}{1 + e^{\frac{\bar{q}}{\lambda}}} + \alpha \cdot \bar{q} - 1/2 \cdot p \right) = p \cdot \frac{d}{d\lambda} \left(\frac{1}{1 + \frac{1-K(\lambda)}{K(\lambda)} e^{\frac{\bar{q}}{\lambda}}} \right) - p \cdot \frac{d}{d\lambda} K(\lambda).$$

Notice that $-\frac{d}{d\lambda} K(\lambda) > 0$. We focus on the denominator of the first term. After some calculations, we have that

$$\frac{d}{d\lambda} \left(1 + \frac{1-K(\lambda)}{K(\lambda)} e^{\frac{\bar{q}}{\lambda}} \right) = \frac{d}{d\lambda} \left(\frac{1}{K(\lambda)} - 1 \right) e^{\frac{\bar{q}}{\lambda}} = \frac{e^{\frac{\bar{q}}{\lambda}}}{K(\lambda)} \cdot \left(\frac{\alpha}{pK(\lambda)\sqrt{1-4\frac{\alpha\lambda}{p}}} - \frac{\bar{q}(1-K(\lambda))}{\lambda^2} \right)$$

By inspection, the sign of the derivative is determined by

$$\left(\frac{\alpha}{pK(\lambda)\sqrt{1-4\frac{\alpha\lambda}{p}}} - \frac{\bar{q}(1-K(\lambda))}{\lambda^2} \right) > 0 \iff B(\lambda) := \left(\frac{1}{\sqrt{1-4\frac{\alpha\lambda}{p}}} - \frac{\bar{q}}{\lambda} \right) > 0,$$

where the implication follows since $K(\lambda) > 0$, and $K(\lambda)(1 - K(\lambda)) = \alpha\lambda/p$. Notice that B is negative for small values of λ , i.e., there exists $\epsilon > 0$, such that $B(\lambda) < 0$ for all $\lambda \in (0, \epsilon)$, and increasing in λ . This implies that, for small values of λ , $\frac{d}{d\lambda}f(\lambda, \bar{q})$ is increasing, while it may be decreasing for higher values. Combining this observation with the fact that $f(\lambda, \bar{q})$ is negative for $\lambda \rightarrow 0$ and positive at $\lambda = \lambda^e$, we conclude that there exists a unique $\hat{\lambda}$ such that $f(\hat{\lambda}, \bar{q}) = 0$. \square

LEMMA 11. *There exists $\hat{\lambda} \in (0, \lambda^e)$, such that, for all $\lambda \in (0, \hat{\lambda})$, $q^*(\lambda) > p$. Furthermore, $\hat{\lambda} = \bar{\lambda}$.*

Proof. The existence of $\hat{\lambda}$ follows by Lemma 10. We show that $\hat{\lambda} = \bar{\lambda}$. We have that $\lim_{q \rightarrow p} f(\hat{\lambda}, q) = 0$ if and only if

$$\frac{K(\hat{\lambda})}{K(\hat{\lambda}) + (1 - K(\hat{\lambda})) \cdot e^{\frac{p}{\hat{\lambda}}}} = K(\hat{\lambda}) - \alpha. \quad (14)$$

The left hand side of equation (14) is positive, which implies that $\bar{\lambda}$ satisfies $K(\bar{\lambda}) > \alpha$, hence the producer's profits are positive. Notice that equation (7) is equivalent to equation (14) determining $\bar{\lambda}$, hence $\hat{\lambda} = \bar{\lambda}$. \square

LEMMA 12. *The threshold $\bar{\lambda} := \bar{\lambda}(\alpha, p) > 0$ is decreasing in the first argument, and increasing in the second.*

Proof. Recall that $\bar{\lambda}$ is determined by equation (7), that is,

$$K(\bar{\lambda}) - \frac{K(\bar{\lambda})}{K(\bar{\lambda}) + (1 - K(\bar{\lambda})) \cdot e^{\frac{p}{\bar{\lambda}}}} = \alpha. \quad (15)$$

Fix $\alpha \in (0, 1)$. Furthermore, by Lemma 11, for every $p \in \mathbb{R}_+$, there is a unique $\bar{\lambda}$ satisfying equation (15). Let $p^e = 4\alpha\lambda$. We show that, every $\lambda \in (0, \lambda^e)$, there exists a unique $p > p^e$ satisfying

$$\tilde{f}(p) := \tilde{K}(p) - \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{p}{\bar{\lambda}}}} = \alpha. \quad (16)$$

where $\tilde{K}(p) := 1/2 + 1/2\sqrt{1 - 4\alpha\lambda/p}$. We have that $\lim_{p \rightarrow \infty} \tilde{f}(p) = 1 > \alpha$, and

$$\lim_{p \rightarrow p^e} \tilde{f}(p) = \frac{1}{2} - \frac{1}{1 + e^{4\alpha}} < \alpha,$$

where the inequality holds for every $\alpha > 0$. This shows there exists a $p > p^e$ satisfying equation (16). We now show that such a p is unique. Notice that $\tilde{K}(p)$ is increasing and, in particular,

$$\frac{d}{dp} \tilde{K}(p) = \frac{\alpha \lambda}{p^2 \sqrt{1 - 4\alpha \lambda / p}} = \frac{\tilde{K}(p)(1 - \tilde{K}(p))}{p \cdot (2\tilde{K}(p) - 1)}$$

where the last equality holds since $\alpha \lambda / p = \tilde{K}(p)(1 - \tilde{K}(p))$. Since $\frac{d}{dp} \frac{1 - \tilde{K}(p)}{K(p)} = \frac{-\frac{d}{dp} \tilde{K}(p)}{\tilde{K}(p)^2}$, after some calculations, we obtain

$$\frac{d}{dp} \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} e^{\frac{p}{\lambda}} = \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} e^{\frac{p}{\lambda}} \left(\frac{1}{\lambda} - \frac{1}{p \cdot (2\tilde{K}(p) - 1)} \right).$$

Notice that such a derivative is negative for values of p in a neighborhood of p^e and increasing for larger values. As a result, \tilde{f} is possibly decreasing for values of p in a neighborhood of p^e and increasing for larger values as well. Therefore, there exists a unique p satisfying $\tilde{f}(p) = \alpha$, establishing a one-to-one relationship between λ and p . Now, notice that equation (15) depends on p and λ only through their ratio, λ/p . Therefore, $\bar{\lambda}(p)$ is increasing.

Fix $p \in \mathbb{R}_+$. By Lemma 11, for every $\alpha \in (0, 1)$, there is a unique λ , which we call $\bar{\lambda}$, satisfying equation (15). Furthermore, by the proof of Proposition 8, we have that for every $\lambda \in (0, \bar{\lambda}]$, $\frac{d}{d\lambda} \alpha(\bar{\lambda}) < 0$. Formally,

$$\begin{aligned} \frac{d}{d\lambda} \alpha(\lambda) &= \frac{d}{d\lambda} \left(K(\lambda) - \frac{1}{1 + \frac{1 - K(\lambda)}{K(\lambda)} e^{\frac{p}{\lambda}}} \right) \\ &= \frac{d}{d\lambda} K(\lambda) - \frac{\frac{d}{d\lambda} K(\lambda) e^{\frac{p}{\lambda}}}{(K(\lambda) + (1 - K(\lambda)) e^{\frac{p}{\lambda}})^2} + \frac{e^{\frac{p}{\lambda}} \alpha / \lambda}{(K(\lambda) + (1 - K(\lambda)) e^{\frac{p}{\lambda}})^2}, \end{aligned}$$

where the second equality holds since $K(\lambda)(1 - K(\lambda)) = \alpha \lambda / p$, and

$$\begin{aligned} \frac{d}{d\lambda} \alpha(\lambda) < 0 &\implies \frac{d}{d\lambda} K(\lambda) \cdot \left(1 - \frac{e^{\frac{p}{\lambda}}}{(K(\lambda) + (1 - K(\lambda)) e^{\frac{p}{\lambda}})^2} \right) < 0 \\ &\implies \left(1 - \frac{e^{\frac{p}{\lambda}}}{(K(\lambda) + (1 - K(\lambda)) e^{\frac{p}{\lambda}})^2} \right) > 0, \end{aligned}$$

where the second implication holds since $\frac{d}{d\lambda}K(\lambda) < 0$. Now, let

$$\hat{f}(\alpha) := \hat{K}(\alpha) - \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{p}{\lambda}}} = \alpha, \quad (17)$$

where $\hat{K}(\alpha) := 1/2 + 1/2\sqrt{1 - 4\alpha\lambda/p}$. We have that,

$$\frac{d}{d\alpha}\hat{f}(\alpha) = \frac{d}{d\alpha}\hat{K}(\alpha) \cdot \left(1 - \frac{e^{\frac{p}{\lambda}}}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}})^2}\right).$$

By the previous argument, for every $\lambda \in (0, \bar{\lambda}]$, $\frac{d}{d\alpha}\hat{f}(\alpha) < 0$, since $\frac{d}{d\alpha}\hat{K}(\alpha) < 0$ and $(1 - e^{\frac{p}{\lambda}}/(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}})^2) > 0$. Therefore, if α increases, the right hand side of equation (17) increases, but the left hand side, namely $\hat{f}(\alpha)$, decreases for every $\lambda \in (0, \bar{\lambda}]$. Since $\frac{d}{d\lambda}\alpha(\lambda) < 0$, this implies that $\bar{\lambda}$ decreases. \square

LEMMA 13. *For every $\lambda \in (0, \bar{\lambda})$, let $\pi(\lambda)$ be the solution $\pi \in (0, \pi_\lambda)$ induced by $q^*(\lambda) > 0$. We have that $\pi(\lambda)$ is continuously differentiable, increasing, and satisfies $\lim_{\lambda \rightarrow 0} \pi(\lambda) = 0$.*

Proof. Recall that π is retrieved by equation (11) by plugging in q^* solving equation (13). By Lemma 10, $q^*(\lambda)$ is continuously differentiable, which implies that $\pi(\lambda)$ is continuously differentiable as well. Furthermore, again by Lemma 10, $\lim_{\lambda \rightarrow 0} q^*(\lambda) = p/\alpha$, implying that

$$\lim_{\lambda \rightarrow 0} \pi(\lambda) = \lim_{\lambda \rightarrow 0} \frac{K(\lambda)}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q^*(\lambda) - p}{\lambda}}} = 0,$$

since $(1 - K(\lambda)) \cdot e^{\frac{q^*(\lambda) - p}{\lambda}} \rightarrow +\infty$ as $\lambda \rightarrow 0$. By Lemma 11, $q^*(\lambda) \rightarrow p$ as $\lambda \rightarrow \bar{\lambda}$, and

$$\lim_{\lambda \rightarrow \bar{\lambda}} \pi(\lambda) = \lim_{\lambda \rightarrow \bar{\lambda}} \frac{K(\lambda)}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q^*(\lambda) - p}{\lambda}}} = K(\bar{\lambda}).$$

Notice that $K(\bar{\lambda}) \leq \pi_{\bar{\lambda}}$ by construction.

By combining Lemma 10 and Lemma 11, for every $(\alpha, p) \in (0, 1) \times \mathbb{R}_+$, $\lambda \in (0, \bar{\lambda})$, there exists a unique $q^* \in (p, p/\alpha)$, and a unique $\pi \in (0, K(\bar{\lambda}))$. We now show that, for every $\pi \in (0, K(\bar{\lambda}))$, $\lambda \in (0, \bar{\lambda})$ is also uniquely determined. This establishes a one-to-one relationship between π and λ that, joint with the fact that $\lim_{\lambda \rightarrow 0} \pi(\lambda) = 0$, shows that $\pi(\lambda)$ is increasing.

Fix $\bar{\pi} \in (0, K(\bar{\lambda}))$. By rearranging equation (11), we obtain

$$z(\lambda) := q^* - p - \lambda \cdot \ln\left(\frac{1 - \bar{\pi}}{\bar{\pi}}\right) - \lambda \cdot \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) = 0. \quad (18)$$

For $\lambda \rightarrow 0$, we have that $K(\lambda) \rightarrow 1$, $-\lambda \cdot \ln\frac{K(\lambda)}{1-K(\lambda)} \rightarrow 0$, $q^* \rightarrow p/\alpha$, hence $z(\lambda) \rightarrow p/\alpha - p > 0$. For $\lambda \rightarrow \bar{\lambda}$, we have that $q^* \rightarrow p$. Furthermore, for every $\lambda \in (0, \bar{\lambda})$, $K(\lambda) > K(\bar{\lambda}) > \bar{\pi}$, which implies $1 - K(\lambda) < 1 - \bar{\pi}$, and therefore

$$\ln\left(\frac{(1 - \bar{\pi}) \cdot K(\lambda)}{\bar{\pi} \cdot (1 - K(\lambda))}\right) > 0.$$

This shows that, for $\lim_{\lambda \rightarrow \bar{\lambda}} z(\lambda) < 0$.

By the intermediate value theorem, for every $\bar{\pi} \in (0, K(\bar{\lambda}))$, there is a $\lambda \in (0, \bar{\lambda})$ satisfying equation (18). We further show such a λ is unique. By Lemma 10, $\frac{d}{d\lambda} q^*(\lambda) < 0$, and clearly

$$\frac{d}{d\lambda} - \lambda \cdot \ln\left(\frac{1 - \bar{\pi}}{\bar{\pi}}\right) < 0.$$

After some calculations,

$$\frac{d}{d\lambda} \frac{K(\lambda)}{1 - K(\lambda)} = \frac{\frac{d}{d\lambda} K(\lambda)}{1 - K(\lambda)} + \frac{K(\lambda) \frac{d}{d\lambda} K(\lambda)}{(1 - K(\lambda))^2} = \frac{\frac{d}{d\lambda} K(\lambda)}{(1 - K(\lambda))^2},$$

and

$$\begin{aligned} \frac{d}{d\lambda} \lambda \cdot \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) &= \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) + \lambda \cdot \left(\frac{1 - K(\lambda)}{K(\lambda)}\right) \frac{\frac{d}{d\lambda} K(\lambda)}{(1 - K(\lambda))^2} = \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) + \frac{p}{\alpha} \frac{d}{d\lambda} K(\lambda) \\ &= \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) - \frac{1}{\sqrt{1 - 4\frac{\alpha\lambda}{p}}} := D(\lambda), \end{aligned}$$

where the last inequality follows by definition of $K(\lambda)$. By inspection, $D(\lambda)$ is positive for small values of λ , i.e., there exists $\epsilon > 0$ such that $D(\lambda) > 0$ for all $\lambda \in (0, \epsilon)$, and possibly positive for larger values. Clearly, the opposite holds for $-D(\lambda)$. This argument implies that the function z defined by equation (18) is always decreasing for lower values of λ , and possibly increasing for larger values. Combining this observation with the fact that $z(\lambda)$ is positive for $\lambda \rightarrow 0$ and negative when $\lambda \rightarrow \bar{\lambda}$, we conclude that there exists a unique λ such that $z(\lambda) = 0$. \square

COROLLARY 5. For every $\lambda \in (0, \bar{\lambda})$, there exists a unique $\pi(\lambda) \in (0, K(\bar{\lambda}))$ such that, if σ is a best response to the attention strategy $\beta(\lambda)$ defined by $\pi(\lambda)$ by equation (2), then $\text{supp}\sigma \subseteq \{0, q^*(\lambda)\}$.

LEMMA 14. For every $\lambda \in (0, \bar{\lambda})$, there exists a unique $\sigma \in \Delta(\{0, q^*(\lambda)\})$ such that the induced attention strategy β defined by equation (2) satisfies $\beta \in \text{argmax} U(\sigma, \beta)$.

Proof. For every $\lambda \in (0, \bar{\lambda})$, $\pi \in (0, 1)$, we show that there exists a $\sigma \in \Delta(\{0, q^*(\lambda)\})$ such that π is best response to σ . The statement follows by setting $\pi = \pi(\lambda)$.

This consumer problem can be mapped to a decision problem where the action set is {buy, not buy} and the state space is $\{0, q^*(\lambda)\}$. In particular, this decision problem is equivalent to Problem 1 in Matějka and McKay (2015) featuring a binary action set, with one safe action, and a binary state space. Therefore, for any $\sigma \in \Delta(\{0, q^*(\lambda)\})$, the optimal π is determined by

$$\pi = \max \left\{ 0, \min \left\{ 1, \frac{(1 - \sigma_0) \cdot (1 - e^{\frac{q^* - p}{\lambda}}) - \sigma_0 \cdot (e^{\frac{-p}{\lambda}} - 1)}{(e^{\frac{q^* - p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)} \right\} \right\},$$

where $\sigma_0 := \sigma(0) \in (0, 1)$. This solution is interior if and only if

$$C(\sigma_0) := \frac{(1 - \sigma_0) \cdot (1 - e^{\frac{q^* - p}{\lambda}}) - \sigma_0 \cdot (e^{\frac{-p}{\lambda}} - 1)}{(e^{\frac{q^* - p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)} \in (0, 1). \quad (19)$$

It is immediate to check that $C(\sigma_0)$ is decreasing in σ_0 . Indeed, since $q^* > p$,

$$\frac{d}{d\sigma_0} C(\sigma_0) = -\frac{e^{-\frac{p}{\lambda}} - e^{\frac{q^* - p}{\lambda}}}{(e^{\frac{q^* - p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)} < 0.$$

Furthermore, for every $\lambda \in (0, \bar{\lambda})$, we have that

$$\begin{aligned} C(0) &= \frac{1}{1 - e^{\frac{-p}{\lambda}}} > 1 \\ C(1) &= \frac{-1}{e^{\frac{q^* - p}{\lambda}} - 1} < 0, \end{aligned}$$

where again the second inequality follows since $q^* > p$.

By the intermediate value theorem, for every $\pi \in (0, 1)$, there exists $\sigma_0^* \in (0, 1)$ such

that $C(\sigma_0^*) = \pi$. Furthermore, since C is monotone in σ_0 , such a σ_0^* is unique. \square

REMARK 1. Notice that a trade probability-quality pair (π, q) uniquely pins down the binary-quality Shannon-Nash equilibrium we are interested in, including the producer's strategy σ . Indeed, by definition, $\pi = \sigma(0)\beta(0) + (1 - \sigma(0))\beta(q)$, where $\beta(q)$ is determined by the function K through parameters (α, λ, p) , and $\beta(0)$ is determined in equation (12) by the function K through parameters (α, λ, p) and by q . Therefore, an alternative approach to show Lemma 14 is to invert this equation, which allows us to rewrite $\sigma(0)$ in terms of π , $\beta(0)$ and $\beta(q)$, and check that such a σ_0 satisfies equation (19).

COROLLARY 6. *For every $\lambda \in (0, \bar{\lambda})$, there exists a unique Shannon-Nash equilibrium (σ, β) with positive quality provision.*

LEMMA 15. *For every $\lambda \in (0, \bar{\lambda})$, let $\sigma(\lambda)$ be the producer's strategy, identified by Lemma 14, which is part of the unique Shannon-Nash equilibrium with positive quality provision. We have that $\lim_{\lambda \rightarrow 0} \sigma_0(\lambda) = 1$ and $\lim_{\lambda \rightarrow \bar{\lambda}} \sigma_0(\lambda) = 0$, where $\sigma_0(\lambda) := \sigma(\lambda)(0)$.*

Proof. For every $\lambda \in (0, \bar{\lambda})$, let $\pi(\lambda) \in (0, \pi(\bar{\lambda}))$ be the consumer's equilibrium strategy. By inverting equation (19), we have that

$$\sigma_0(\lambda) = \frac{(1 - e^{\frac{q^*(\lambda)-p}{\lambda}}) - \pi(\lambda) \cdot (e^{\frac{q^*(\lambda)-p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)}{(-e^{\frac{q^*(\lambda)-p}{\lambda}} + e^{\frac{-p}{\lambda}})} = \frac{(e^{\frac{q^*(\lambda)-p}{\lambda}} - 1) \cdot (1 - \pi(\lambda) \cdot (1 - e^{\frac{-p}{\lambda}}))}{(e^{\frac{q^*(\lambda)-p}{\lambda}} - e^{\frac{-p}{\lambda}})}.$$

From this equation, it is easy to see that as $\lambda \rightarrow 0$, $\pi(\lambda) \rightarrow 0$, $q^*(\lambda) \rightarrow p/\alpha$, which implies that $\sigma_0(\lambda) \rightarrow 1$; as $\lambda \rightarrow \bar{\lambda}$, $\pi(\lambda) \rightarrow K(\bar{\lambda})$, $q^*(\lambda) \rightarrow p$, implying $\sigma_0(\lambda) \rightarrow 0$. \square

This concludes the proof of Theorem 1. \blacksquare

We relate the results in Section 3.3 to the ones in the proof of Proposition 2 and Theorem 1. Lemma 2 follows by Corollary 4; Lemma 3 follows by Lemma 5, Lemma 8, and the proof of Lemma 9; the proof of Lemma 4 by Lemma 9 and 11.

Comparative statics on the attention cost.

PROOF OF PROPOSITION 3. In the unique Shannon-Nash equilibrium with positive quality provision (σ, β) , the producer's profits as a function of $\lambda \in (0, \bar{\lambda})$ are given by

$$\Pi(\lambda) = p \cdot \beta(0) = p \cdot \frac{\pi(\lambda) \cdot e^{\frac{-p}{\lambda}}}{\pi(\lambda) \cdot e^{\frac{-p}{\lambda}} + 1 - \pi(\lambda)}.$$

Profits are positive since, by equation (12) and Lemma 10, $\beta(0) > 0$ for every $\lambda \in (0, \bar{\lambda})$. Furthermore, profits converge to zero as $\lambda \rightarrow 0$ since, by the proof of Lemma 10, $\beta(0) \rightarrow 0$. By Lemma 13, $\pi(\lambda)$ is increasing. Therefore, $\Pi(\lambda)$ is increasing as well.

We now show that, for every $\lambda \in (0, \bar{\lambda})$, the consumer utility is positive. By known properties, mutual information is strictly convex over the region of attention strategies that are non-constant almost surely (see Ravid, 2020). Formally, for any $\beta' \neq \beta''$ such that β' is non-constant μ -almost surely, we have that, for every $\gamma \in (0, 1)$,

$$I(\gamma \cdot \beta' + (1 - \gamma) \cdot \beta'', \mu) < \gamma \cdot I(\beta', \mu) + (1 - \gamma) \cdot I(\beta'', \mu).$$

Let β^0 the attention strategy recommending the consumer to never accept the producer's offer, i.e., $\mathbb{E}_\mu[\beta^0] = 0$, and consider $\text{co}\{\beta^0, \beta\}$, the convex hull containing all the convex combinations of β^0 and β . Any $\tilde{\beta}$ in the interior of $\text{co}\{\beta^0, \beta\}$ is non-constant μ -almost surely, implying that I is strictly convex over $\text{co}\{\beta^0, \beta\}$. As the consumer's gains from trade, i.e., the consumer's objective net of the information costs, are affine in the attention strategies, we have that the consumer objective as a whole is strictly concave over $\text{co}\{\beta^0, \beta\}$. Since β is a best response, this implies that $U(\beta) > U(\beta^0) = 0$.

Finally, $\lim_{\lambda \rightarrow 0} U(\sigma, \beta) = 0$ and $\lim_{\lambda \rightarrow \bar{\lambda}} U(\sigma, \beta) = 0$ follow by the definition of U and $\lim_{\lambda \rightarrow 0} \sigma_0(\lambda) = 1$ and $\lim_{\lambda \rightarrow \bar{\lambda}} q(\lambda) = p$, respectively. \blacksquare

Comparative statics on the price level. We show the result of Section 3.1.

PROOF OF THEOREM 2. Fix $(\alpha, \lambda) \in (0, 1) \times \mathbb{R}_+$. By the proof of Lemma 12, there exists $\bar{p} \in (p^e, \infty)$ such that for every $p > \bar{p}$, $\lambda > \bar{\lambda}$. Furthermore, for $p \leq \bar{p}$, $\lambda \leq \bar{\lambda}$. Therefore, the unique Shannon-Nash equilibrium with positive quality provision exists if and only $p > \bar{p}$. Now, the fact that $\bar{p} := \bar{p}(\alpha, \lambda)$ is increasing in the second argument follows by the fact that $\bar{\lambda}$ is increasing in p (see Lemma 12). To show that \bar{p} is increasing in the first argument, consider again the following equation

$$\tilde{f}(p) := \tilde{K}(p) - \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{p}{\bar{\lambda}}}} = \alpha.$$

By Lemma 12, as α increases, $\bar{\lambda}$ decreases. Since $\tilde{f}(p)$ depends on λ and p only through their ratio, λ/p , an increase in α corresponds to a decrease in the ratio λ/p . For λ fixed, the decrease in such a ratio occurs only if \bar{p} increases, concluding the argument.

LEMMA 16. For every $p < \bar{p}$, let $q^*(p)$ be the unique solution $q^* > p$ satisfying the following equation

$$p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*}{\lambda}}} + \alpha \cdot q^* - p \cdot \tilde{K}(p) = 0. \quad (20)$$

We have that $q^*(p)$ is continuously differentiable, and satisfies $\lim_{p \rightarrow \bar{p}} q(p) = \bar{p}$ and $\lim_{p \rightarrow \infty} \alpha q(p) - p = 0$.

Proof. A unique and positive solution of equation (20) exists by Lemma 9 since $\lambda < \bar{\lambda}$ if and only if $p > \bar{p}$. Similarly, $q^*(p) > p$ is continuously differentiable by Lemma 10. As $p \rightarrow \bar{p}$, we have that $\bar{\lambda} \rightarrow \lambda$, which, by Lemma 10, implies that $q \rightarrow \bar{p}$.

We now show that $\lim_{p \rightarrow \infty} \alpha q(p) - p = 0$. Since $q^* > p$, we have that

$$0 \leq p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*}{\lambda}}} \leq p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{p}{\lambda}}} := \tilde{h}(p),$$

for every $p > \bar{p}$. Since $\lim_{p \rightarrow \infty} (1 - \tilde{K}(p)) \cdot e^{\frac{p}{\lambda}} \cdot \frac{1}{p} = \infty$, we have that $\lim_{p \rightarrow \infty} \tilde{h}(p) = 0$. The statement follows by equation (20), $\tilde{K}(p) \rightarrow 1$ as $p \rightarrow \infty$, and $\tilde{K}(p) < 1$. Notice that this implies that $\beta(0) \rightarrow 0$ as well. \square

LEMMA 17. For every $p \in (\bar{p}, \infty)$, the producer profits are positive and converge to zero as $p \rightarrow \infty$.

Proof. In equilibrium, the producer profits equal

$$\Pi = p \cdot \beta(0) > 0,$$

where the inequality holds since $\beta(0) > 0$ by Proposition 3 and the fact that, for every $p > \bar{p}$, we have that $\lambda < \bar{\lambda}$. The fact that profits converge to zero as $p \rightarrow \infty$ follows from the proof of Lemma 16, since by equation (12),

$$p \cdot \beta(0) = p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*}{\lambda}}}$$

for every $p \in (\bar{p}, \infty)$. \square

LEMMA 18. In the unique Shannon-Nash equilibrium with positive quality provision (σ, β) , the producer profits are decreasing in p .

Proof. The producer profits as a function of $p \in (\bar{p}, \infty)$, are given by

$$\Pi(p) = p \cdot \tilde{K}(p) - \alpha \cdot q^*(p), \quad (21)$$

where $\tilde{K}(p) = \beta(q^*(p))$, and $q^*(p) > p$ is defined by equation (20). Furthermore, by Lemma 16,

$$\frac{d}{dp} \Pi(p) = \tilde{K}(p) + p \frac{d}{dp} \tilde{K}(p) - \alpha \frac{d}{dp} q^*(p). \quad (22)$$

We proceed by finding an expression for the derivative of $q^*(p)$ by differentiating both sides of equation (20) by p . We obtain,

$$\begin{aligned} \frac{d}{dp} \left(p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p)}{\lambda}}} + \alpha \cdot q^*(p) - p \cdot \tilde{K}(p) \right) &= 0 \\ \alpha \frac{d}{dp} q^*(p) - \tilde{K}(p) - p \frac{d}{dp} \tilde{K}(p) + \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p)}{\lambda}}} - p \left(\frac{e^{\frac{q^*(p)}{\lambda}} \left(\frac{d}{dp} \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} + \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} \frac{d}{dp} \frac{q^*(p)}{\lambda} \right)}{\left(1 + \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} e^{\frac{q^*(p)}{\lambda}} \right)^2} \right) &= 0 \\ \alpha \frac{d}{dp} q^*(p) \left(1 - \frac{e^{\frac{q^*(p)}{\lambda}}}{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2} \right) & \\ = \tilde{K}(p) \left(1 - \frac{1}{\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}}} \right) + p \left(1 - \frac{e^{\frac{q^*(p)}{\lambda}}}{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2} \right) \frac{d}{dp} \tilde{K}(p) & \\ \alpha \frac{d}{dp} q^*(p) = p \frac{d}{dp} \tilde{K}(p) + \tilde{K}(p) \left(1 - \frac{1}{\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}}} \right) \left(1 - \frac{e^{\frac{q^*(p)}{\lambda}}}{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2} \right)^{-1} & \\ \alpha \frac{d}{dp} q^*(p) = p \frac{d}{dp} \tilde{K}(p) + \tilde{K}(p) \left(\frac{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2 - (\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})}{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2 - e^{\frac{q^*(p)}{\lambda}}} \right), & \end{aligned}$$

where the first passage follows by differentiating and rearranging, and the second passage by the facts that $\tilde{K}(p)(1 - \tilde{K}(p)) = \alpha\lambda/p$ and $\frac{d}{dp} \frac{1 - \tilde{K}(p)}{\tilde{K}(p)} = \frac{-\frac{d}{dp} \tilde{K}(p)}{\tilde{K}(p)^2}$.

By plugging the above expression for $\alpha \frac{d}{dp} q^*(p)$ into equation (22), we obtain

$$\frac{d}{dp} \Pi(p) = \tilde{K}(p) \left(1 - \frac{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2 - (\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})}{(\tilde{K}(p) + (1 - \tilde{K}(p)) e^{\frac{q^*(p)}{\lambda}})^2 - e^{\frac{q^*(p)}{\lambda}}} \right).$$

Notice that

$$\begin{aligned} \frac{d}{dp}\Pi(p) < 0 &\iff C(p) := \left(\frac{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}})^2 - (\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}}))}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}})^2 - e^{\frac{q^*(p)}{\lambda}}} \right) > 1 \\ &\iff D(p) := (\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}})^2 - e^{\frac{q^*(p)}{\lambda}} > 0 \end{aligned}$$

where the last implication follows since $\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}} > 1$ and $|C(p)| \geq 1$. By Lemma 17, we have that $\Pi(p) > 0$ for every $p \in (\bar{p}, \infty)$ and $\lim_{p \rightarrow \infty} \Pi(p) = 0$. This implies that there exists \hat{p} such that $\frac{d}{dp}\Pi(\hat{p}) < 0$, which is equivalent to $D(\hat{p}) > 0$. By contradiction, assume $D(p) < 0$ for some $p > \bar{p}$. By Lemma 16, we can apply the intermediate value theorem to obtain that $D(\tilde{p}) = 0$ for some $\tilde{p} > \bar{p}$. But this implies that $\frac{d}{dp}\Pi(\tilde{p})$ is not well-defined, contradicting Lemma 16 claiming that $q^*(p)$, and therefore $p \cdot \beta(0)$ by equation (12), is continuously differentiable for every $p > \bar{p}$. This argument shows that $D(p) > 0$ which is equivalent $\frac{d}{dp}\Pi(p) < 0$, concluding the proof. \square

LEMMA 19. *In the unique Shannon-Nash equilibrium with positive quality provision (σ, β) , the consumer utility is positive, and converges to zero as $p \rightarrow \bar{p}$ and $p \rightarrow \infty$.*

Proof. For every $p \in (\bar{p}, \infty)$, the consumer utility is positive by Proposition 3. To see that $\lim_{p \rightarrow \bar{p}} U(\sigma, \beta) = 0$ notice that, as $p \rightarrow \bar{p}$, $\bar{\lambda} \rightarrow \lambda$, which implies $q^* \rightarrow p$. To check that $\lim_{p \rightarrow \infty} U(\sigma, \beta) = 0$, recall that, by equation (11), $\pi(p)$ is defined as

$$\pi(p) = \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p) - p}{\lambda}}},$$

where $q^*(p) > p$ solves equation (20) for $p \in (\bar{p}, \infty)$. Furthermore, by definition,

$$\pi(p) = (1 - \sigma(q^*(p)))\beta(0) + \sigma(q^*(p))\beta(q^*(p)) \iff \sigma(q^*(p)) = \frac{\pi(p) - \beta(0)}{\beta(q^*(p)) - \beta(0)},$$

where $\sigma(q^*(p))$ denotes the probability that the producer plays the quality $q^*(p)$ when the price is p . By equation (12), and after some calculations, we obtain

$$\sigma(q^*(p)) = \frac{\frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p) - p}{\lambda}}} - \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p)}{\lambda}}}{\tilde{K}(p) - \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p)}{\lambda}}}},$$

$$\begin{aligned}
&= \frac{(1 - \tilde{K}(p))(e^{\frac{q^*(p)}{\lambda}} - e^{\frac{q^*(p)-p}{\lambda}})}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}} - 1)} \cdot \frac{1}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)-p}{\lambda}})} \\
&:= E(p) \cdot \frac{1}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)-p}{\lambda}})}
\end{aligned}$$

Notice that, there exists $\hat{p} > \bar{p}$ such that $E(p) < 1$ for all $p > \hat{p}$. To see this, notice that

$$E(p) < 1 \iff \frac{(1 - \tilde{K}(p))(e^{\frac{q^*(p)}{\lambda}} - e^{\frac{q^*(p)-p}{\lambda}})}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}} - 1)} < 1 \iff -(1 - \tilde{K}(p))e^{\frac{q^*(p)-p}{\lambda}} < \tilde{K}(p) - 1$$

where the second implication follows since $q^*(p) > p$, and $(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{p}{\lambda}} - 1) > 0$ for p sufficiently large. By Lemma 16, $\lim_{p \rightarrow \infty} \alpha q^*(p) - p = 0$, for $\alpha \in (0, 1)$, which implies that $-(1 - \tilde{K}(p))e^{\frac{q^*(p)-p}{\lambda}} \rightarrow -\infty$, proving the step. As a result,

$$\sigma(q^*(p))\tilde{K}(p) < \frac{\tilde{K}(p)}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)-p}{\lambda}})} = \pi(p)$$

for all $p > \hat{p}$. The consumer utility gains from trade when quality equals $q^*(p)$, which we label $\tilde{U}(q^*(p))$, converge to zero as $p \rightarrow \infty$ since

$$\begin{aligned}
0 &\leq \lim_{p \rightarrow \infty} \tilde{U}(q^*(p)) := \lim_{p \rightarrow \infty} \sigma(q^*(p))\beta(q^*(p))(q^*(p) - p) = \lim_{p \rightarrow \infty} \sigma(q^*(p))\tilde{K}(p)(q^*(p) - p) \\
&< \lim_{p \rightarrow \infty} \frac{\tilde{K}(p)(q^*(p) - p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*(p)-p}{\lambda}}} = 0
\end{aligned}$$

where the first inequality holds by $q^*(p) > p$, the second inequality since $\sigma(q^*(p))\tilde{K}(p) < \pi(p)$ for $p > \hat{p}$, and last equality since, by Lemma 16, $\lim_{p \rightarrow \infty} \alpha q^*(p) - p = 0$, which implies that $q^*(p)$ approximates p/α in the limit, and

$$\lim_{p \rightarrow \infty} \frac{p \cdot \tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{p}{\lambda}}} = 0.$$

This concludes the proof since $\tilde{U}(q^*(p)) > U(\sigma, \beta)$ for every $p > \bar{p}$. □

This concludes the proof of Theorem 2. ■

Comparative statics on production costs. We prove the result of Section 3.2.

PROOF OF PROPOSITION 5. Fix $(p, \lambda) \in \mathbb{R}_+^2$ and let $\alpha^e := p/4\lambda$. We first show that there exists a unique $\alpha \in (0, \alpha^e)$, which we term $\bar{\alpha}$, such that

$$\hat{g}(\alpha) := \hat{K}(\alpha) - \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{p}{\lambda}}} - \alpha = 0, \quad (23)$$

where $\hat{K}(\alpha) := 1/2 + 1/2\sqrt{1 - 4\alpha\lambda/p}$. The fact that $\bar{\alpha}(p, \lambda)$ is increasing in the first component and decreasing in the second follows from the proof of Theorem 2 and Lemma 12, respectively. Finally, the unique Shannon-Nash equilibrium with positive quality provision exists if and only if $\alpha \in (0, \bar{\alpha})$. Indeed, when $\alpha = \bar{\alpha}$, we have that $g(\bar{\alpha}) = 0$, and hence $\lambda = \bar{\lambda}$. By Lemma 12, $\bar{\lambda}$ is decreasing in α which implies that $\alpha > \bar{\alpha}$ if and only if $\lambda > \bar{\lambda}$.

We have that $\lim_{\alpha \rightarrow 0} \hat{g}(\alpha) = 0$ since $\lim_{\alpha \rightarrow 0} \hat{K}(\alpha) = 1$ and

$$\lim_{\alpha \rightarrow \alpha^e} \hat{g}(\alpha) = \frac{1}{2} - \frac{1}{1 + e^{\frac{p}{\lambda}}} - \frac{p}{4\lambda} = \frac{1}{2} - \frac{1}{1 + e^x} - \frac{x}{4} < 0,$$

where the last inequality holds for all $x > 0$. By taking first order conditions, we obtain

$$\begin{aligned} \frac{d}{d\alpha} \hat{g}(\alpha) &= \frac{d}{d\alpha} \hat{K}(\alpha) + \frac{e^{\frac{p}{\lambda}} \frac{d}{d\alpha} \hat{K}(\alpha)}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}})^2} - 1 \\ &= \frac{-\lambda}{p(2\hat{K}(\alpha) - 1)} \left(1 - \frac{e^{\frac{p}{\lambda}}}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}})^2} \right) - 1 := M(\alpha) \cdot N(\alpha) - 1, \end{aligned}$$

where the second equality follows since $\frac{d}{d\alpha} \hat{K}(\alpha) = \frac{-\lambda}{p(2\hat{K}(\alpha) - 1)}$. Therefore, we have that

$$\frac{d}{d\alpha} \hat{g}(0) = -\frac{\lambda}{p}(1 - e^{\frac{p}{\lambda}}) - 1 = -\frac{1}{x}(1 - e^x) - 1 > 0,$$

where the equality holds since $\hat{K}(0) = 1$ and the inequality holds for all $x > 0$. By continuity, this implies that there exists $\bar{\alpha} \in (0, \alpha^e)$ which solves equation (23). We now show that $\bar{\alpha}$ is unique. To this end,

$$\frac{d^2}{d\alpha^2} \hat{g}(\alpha) = M(\alpha) \frac{d}{d\alpha} N(\alpha) + N(\alpha) \frac{d}{d\alpha} M(\alpha).$$

Since $\frac{d}{d\alpha} \hat{K}(\alpha) < 0$ and $e^{\frac{p}{\lambda}} > 1$, we have that $\frac{d}{d\alpha} (\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}}) > 0$, which implies

that $\frac{d}{d\alpha}N(\alpha) > 0$. As a consequence $M(\alpha)\frac{d}{d\alpha}N(\alpha) < 0$ since $M(\alpha) < 0$. Furthermore, $\frac{d}{d\alpha}M(\alpha) < 0$. This implies that $N(\alpha)\frac{d}{d\alpha}M(\alpha) > 0$ if and only if $N(\alpha) < 0$. By inspection, $N(\alpha)$ is negative for values of α in a neighborhood of 0, and positive for larger values of α ,

$$\lim_{\alpha \rightarrow \alpha^e} N(\alpha) = \left(1 - \frac{4e^{\frac{p}{\lambda}}}{(1 + e^{\frac{p}{\lambda}})^2}\right) = \left(1 - \frac{4e^x}{(1 + e^x)^2}\right) > 0,$$

where the first equality holds since $\hat{K}(\alpha^e) = \frac{1}{2}$, and the second equality for every $x > 0$. To summarize, the function \hat{g} satisfies the following: $\hat{g}(0) = 0$, $\frac{d}{d\alpha}g(0) > 0$, $\hat{g}(\alpha^e) < 0$, \hat{g} is possibly convex for values of α in a neighborhood of 0, and always concave for larger values, i.e., if $\frac{d^2}{d\alpha^2}g(\tilde{\alpha}) < 0$ for some $\tilde{\alpha} \in (0, \alpha^e)$, then $\frac{d^2}{d\alpha^2}g(\alpha) < 0$ for every $\alpha > \tilde{\alpha}$. This implies that there exists at most one $\bar{\alpha} \in (0, \alpha^e)$ such that $g(\bar{\alpha}) = 0$.

LEMMA 20. *In the unique Shannon-Nash equilibrium with positive quality provision, $\pi(\alpha)$ is increasing and $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$.*

Proof. Fix (λ, p) . By the argument made above, for every $\alpha \in (0, \bar{\alpha})$, there exists a unique $\pi(\alpha)$ that sustains the Shannon-Nash equilibrium. Let \mathcal{P}_α be the set of all these $\pi(\alpha)$, and recall that every $\pi \in \mathcal{P}_\alpha$ satisfies $\pi < \frac{1}{e^{\frac{p}{\lambda}} + 1} := \pi_\lambda$.

By combining equations (11) and (13), we obtain that π as a function of α has to satisfy

$$p\hat{K}(\alpha) - \alpha \left(p + \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \pi(\alpha)}{\pi(\alpha)} \right) \right) = p \frac{\pi(\alpha) e^{-\frac{p}{\lambda}}}{\pi(\alpha) e^{-\frac{p}{\lambda}} + 1 - \pi(\alpha)} \quad (24)$$

Fix $\hat{\pi} \in \mathcal{P}_\alpha$. We show that there exists a unique α satisfying equation (24). Notice that in equilibrium $\hat{K}(\alpha) > \hat{\pi}$, and therefore we restrict on this region without loss. Consider

$$p\hat{K}(\alpha) - \alpha \left(p + \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \hat{\pi}}{\hat{\pi}} \right) \right) = p \frac{\hat{\pi} e^{-\frac{p}{\lambda}}}{\hat{\pi} e^{-\frac{p}{\lambda}} + 1 - \hat{\pi}}$$

Notice that the right hand side is constant in α . Furthermore, the left hand side is decreasing since

$$\begin{aligned} & \frac{d}{d\alpha} \left(p\hat{K}(\alpha) - \alpha \left(p + \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \hat{\pi}}{\hat{\pi}} \right) \right) \right) \\ &= p \frac{d}{d\alpha} \hat{K}(\alpha) - p\hat{K}(\alpha) - p - \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \hat{\pi}}{\hat{\pi}} \right) - \alpha \lambda \frac{\frac{d}{d\alpha} \hat{K}(\alpha)}{\hat{K}(\alpha)(1 - \hat{K}(\alpha))} \end{aligned}$$

$$= p \frac{d}{d\alpha} \hat{K}(\alpha) - p \hat{K}(\alpha) - p - \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \hat{\pi}}{\hat{\pi}} \right) - p \frac{d}{d\alpha} \hat{K}(\alpha) < 0,$$

where the last equality follows since $\hat{K}(\alpha)(1 - \hat{K}(\alpha)) = \alpha \lambda / p$ and the inequality since $\hat{K}(\alpha) > \hat{\pi}$.

The previous argument establishes a one-to-one relationship between π and α , that combined with the fact that $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$ shows that $\pi(\alpha)$ is increasing. To see that $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$, take the limit for $\alpha \rightarrow 0$ in equation (24), to obtain

$$\begin{aligned} \lim_{\alpha \rightarrow 0} \text{LHS}(\alpha) &= p - \lambda \lim_{\alpha \rightarrow 0} \alpha \cdot \lim_{\alpha \rightarrow 0} \ln \left(\frac{1 - \pi(\alpha)}{\pi(\alpha)} \right) \\ \lim_{\alpha \rightarrow 0} \text{RHS}(\alpha) &= p \cdot \lim_{\alpha \rightarrow 0} \frac{\pi(\alpha) e^{-\frac{p}{\lambda}}}{\pi(\alpha) e^{-\frac{p}{\lambda}} + 1 - \pi(\alpha)} < p \end{aligned}$$

If $\lim_{\alpha \rightarrow 0} \pi(\alpha) \neq 0$, then $\lim_{\alpha \rightarrow 0} \text{LHS}(\alpha) = p$, contradicting the existence of $\pi(\alpha)$ satisfying equation (24) for every $\alpha \in (0, \bar{\alpha})$. \square

LEMMA 21. *The producer profits are positive, increasing in α , and converges to zero as $\alpha \rightarrow 0$.*

Proof. The producer profits as a function of $\alpha \in (0, \bar{\alpha})$ can be written as

$$\Pi(\alpha) = p \cdot \beta(0) = p \cdot \frac{1}{1 + \frac{1 - \pi(\alpha)}{\pi(\alpha)} e^{\frac{p}{\lambda}}}.$$

It follows that: $\Pi(\alpha) > 0$ since $\beta(0) > 0$; $\frac{d}{d\alpha} \Pi(\alpha) > 0$ since $\frac{d}{d\alpha} \pi(\alpha) > 0$ by Lemma 24, and $\lim_{\alpha \rightarrow 0} \Pi(\alpha) = 0$ since $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$ by Lemma 24. Notice that this shows that $\lim_{\alpha \rightarrow 0} \beta(0) = 0$. \square

LEMMA 22. *The consumer utility is positive, and converges to zero as $\alpha \rightarrow 0$ and $\alpha \rightarrow \bar{\alpha}$.*

Proof. For every $\alpha \in (0, \bar{\alpha})$, the consumer utility is positive by Proposition 3. To see that $\lim_{\alpha \rightarrow \bar{\alpha}} U(\sigma, \beta) = 0$ notice that, as $\alpha \rightarrow \bar{\alpha}$, $\bar{\lambda} \rightarrow \lambda$, which implies $q^* \rightarrow p$. We now show that $\lim_{\alpha \rightarrow 0} U(\sigma, \beta) = 0$. Let $q^*(\alpha)$ be the unique solution $q^* > p$ satisfying the following equation

$$\alpha \cdot q^* = p \cdot \hat{K}(\alpha) - p \cdot \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{q^*}{\lambda}}}. \quad (25)$$

The solution of equation (25) exists by Lemma 9 since $\lambda < \bar{\lambda}$ if and only if $\alpha < \bar{\alpha}$. Similarly, $q^*(\alpha) > p$ is continuously differentiable by Lemma 10, and $\lim_{\alpha \rightarrow 0} \alpha q^*(\alpha) = p$. This

latter point follows by equation (25) since $\lim_{\alpha \rightarrow 0} \beta(0) = 0$ by the proof of Lemma 21. By mimicking the steps of the proof of Lemma 19, we have that

$$\begin{aligned} \sigma(q^*(\alpha))\hat{K}(\alpha) &= \frac{(1 - \hat{K}(\alpha))(e^{\frac{q^*(\alpha)}{\lambda}} - e^{\frac{q^*(\alpha)-p}{\lambda}})}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)}{\lambda}} - 1)} \cdot \frac{\hat{K}(\alpha)}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)-p}{\lambda}})} \\ &:= \hat{E}(\alpha) \cdot \frac{\hat{K}(\alpha)}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)-p}{\lambda}})} \end{aligned}$$

There exists $\hat{\alpha} < \bar{\alpha}$ such that $\hat{E}(\alpha) < 1$ for all $\alpha < \hat{\alpha}$. To see this, notice that

$$\hat{E}(\alpha) < 1 \iff \frac{(1 - \hat{K}(\alpha))(e^{\frac{q^*(\alpha)}{\lambda}} - e^{\frac{q^*(\alpha)-p}{\lambda}})}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)}{\lambda}} - 1)} < 1 \iff -(1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)-p}{\lambda}} < \hat{K}(\alpha) - 1$$

where the second implication follows since $q^*(\alpha) > p$, and $(\hat{K}(\alpha) + (1 - \hat{K}(\hat{\alpha}))e^{\frac{p}{\lambda}} - 1) > 0$. By the argument above, $\lim_{\alpha \rightarrow 0} \alpha q^*(\alpha) = p$, which implies that $-(1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)-p}{\lambda}} \rightarrow -\infty$, proving the step. As a result,

$$\sigma(q^*(\alpha))\hat{K}(\alpha) < \frac{\hat{K}(\alpha)}{(\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha)-p}{\lambda}})} = \pi(\alpha)$$

for all $\alpha < \hat{\alpha}$. The consumer utility gains from trade when quality equals $q^*(\alpha)$, which we label $\hat{U}(q^*(\alpha))$, converge to zero as $\alpha \rightarrow 0$ since

$$\begin{aligned} 0 &\leq \lim_{\alpha \rightarrow 0} \hat{U}(q^*(\alpha)) := \lim_{\alpha \rightarrow 0} \sigma(q^*(\alpha))\beta(q^*(\alpha))(q^*(\alpha) - p) = \lim_{\alpha \rightarrow 0} \sigma(q^*(\alpha))\hat{K}(\alpha)(q^*(\alpha) - p) \\ &< \lim_{\alpha \rightarrow 0} \frac{\hat{K}(\alpha)(q^*(\alpha) - p)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{q^*(\alpha)-p}{\lambda}}} = 0 \end{aligned}$$

where the first inequality holds by $q^*(\alpha) > p$, the second inequality since $\sigma(q^*(\alpha))\hat{K}(\alpha) < \pi(\alpha)$ for $\alpha < \hat{\alpha}$, and last equality since $\lim_{\alpha \rightarrow 0} \alpha q^*(\alpha) = p$, which implies that $q^*(\alpha)$ approximates p/α in the limit, and

$$\lim_{\alpha \rightarrow 0} \frac{(\frac{1}{\alpha} - 1) \cdot \hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{(\frac{1}{\alpha} - 1)}} = 0.$$

This concludes the proof since $\hat{U}(q^*(\alpha)) > U(\sigma, \beta)$ for every $\alpha < \bar{\alpha}$. □

This concludes the proof of Proposition 5. ■

C Extensions and robustness

This section presents formal arguments establishing the results discussed in Section 4.

Beyond moral hazard. We start by showing Proposition 6.

PROOF OF PROPOSITION 6. Since the consumer's best response is credible, by Lemma 1, it satisfies equation (2). Fixing $\pi \in (0, 1)$ —notice that for $\pi \in \{0, 1\}$ there cannot be a Shannon-Nash equilibrium with positive quality—and rearranging equation (2), we obtain

$$q(\beta) = p + \lambda \cdot \ln\left(\frac{1-\pi}{\pi} \frac{\beta}{1-\beta}\right), \quad (26)$$

where $\beta \in (0, 1)$ is the probability of accepting the offer (p, q) . We rewrite the producer problem as follows:

$$\max_{\beta \in (0, \bar{\beta})} \beta \cdot \left(p - \alpha \cdot \left(p + \lambda \cdot \ln\left(\frac{1-\pi}{\pi} \frac{\beta}{1-\beta}\right) \right) \right) = \beta \cdot p \cdot (1 - \alpha) - \alpha \cdot \lambda \cdot \beta \cdot \ln\left(\frac{1-\pi}{\pi} \frac{\beta}{1-\beta}\right)$$

where $\bar{\beta} = \beta(p/\alpha)$. This objective is strictly concave since $f(x) = x \cdot \ln(\frac{x}{1-x})$, for $x \in (0, 1)$, is strictly convex. As a result, there is a unique β solving the problem and a unique q by equation (26). By Proposition 3, when $|\text{supp } \sigma = 1|$, Shannon best responses satisfy $\pi \in \{0, 1\}$, concluding the proof. ■

The argument establishing the existence of the artificial equilibrium in this setting is analogous to the proof of Proposition 2 in Ravid (2020) and thus omitted.

Convex-concave information. We start by showing that condition (iv) holds when the consumer observes signal $s = q + \epsilon$ where $\epsilon \sim H$ is a unimodal at zero random variable. For every $\hat{s} \in S$, we have that

$$F(\hat{s}|q) = F(s \leq \hat{s}|q) = H(q + \epsilon \leq \hat{s}) = H(\hat{s} - q).$$

If $\hat{s} \leq 0$, then H is strictly convex, and F is strictly convex as well. If $\hat{s} > p/\alpha$, then H is strictly concave for all $q \in [0, p/\alpha]$. As a result, F is strictly concave. If $\hat{s} \leq p/\alpha$, then let $q_{\hat{s}} = \hat{s}$. For all $q < q_{\hat{s}}$, H is strictly concave, while for all $q > q_{\hat{s}}$, H is strictly convex. The same properties hold for F .

PROOF OF PROPOSITION 7. Let (σ, F, b) be Nash equilibrium with positive quality provision. Since $\kappa(F) > 0$, $|\text{supp } \sigma| > 1$. By Proposition 1 in Milgrom (1981), since F satisfies full-

support and strict MLRP, and σ is not degenerate, then the consumer follows a cut-off strategy: for all $s < s^*$, $b(s) = 0$; for all $s \geq s^*$, $b(s) = 1$. As we are considering strategies that may differ on measure-zero sets, $b(s^*)$ is immaterial by condition (i). Furthermore, $s^* \in (\underline{s}, \bar{s})$ to ensure that the consumer trades with probability in $(0, 1)$.

The probability that the consumer buys when quality is $q \in Q$ is $1 - F(s^*|q)$. By condition (iv), if $s^* < 0$, then $1 - F(s^*|q)$ is strictly concave for all $q \in [0, p/\alpha]$. In this case, the producer's objective is strictly concave, meaning he has a unique and deterministic best response. If $s^* > 0$, then function is either strictly convex over $[0, p/\alpha]$ or convex-concave. Since the producer isoprofit curve in terms of conditional trade probability is increasing in q , then $1 - F(s^*|q)$ cannot decrease in q . Since the isoprofit curve is linear, following the same argument as Lemma 7, we have that the isoprofit curve has to lie above the curve defined by $1 - F(s^*|q)$, which implies that there are exactly two points of contact. Therefore, in equilibrium the producer randomizes between 0 and $q^* > p$. Notice that, when $1 - F(s^*|q)$ is strictly convex, $q^* = p/\alpha$. ■

Dropping credibility. We start by proving Corollary 1.

PROOF OF COROLLARY 1. The proof of point (i) is analogous to the argument establishing point (i) in Proposition 1 and is thus omitted. For point (ii), notice that, if β is a constant best response, then is also robust to non-degenerate entropic costs since the property imposes no additional requirements. If β is a best response that is robust to non-degenerate entropic costs when $|\text{supp } \sigma| = 1$, then β is constant since $\lambda \cdot \varepsilon > 0$. ■

We now show that, for every $\tilde{q} \in (p, p/\alpha)$ such that there exists $\pi \in (0, 1)$ with $\Pi(0, \beta) = \Pi(\tilde{q}, \beta)$, then there exists $\sigma(0) \in (0, 1)$ such that the consumer finds it optimal to set π . Since the consumer faces a binary decision problem, by Problem 1 in Matějka and McKay (2015), we know that the optimal π is determined by

$$\pi = \max \left\{ 0, \min \left\{ 1, \frac{(1 - \sigma_0) \cdot (1 - e^{-\frac{\tilde{q}-p}{\lambda}}) - \sigma_0 \cdot (e^{-\frac{p}{\lambda}} - 1)}{(e^{-\frac{\tilde{q}-p}{\lambda}} - 1) \cdot (e^{-\frac{p}{\lambda}} - 1)} \right\} \right\},$$

where $\sigma_0 := \sigma(0) \in (0, 1)$. This solution is interior if and only if

$$C(\sigma_0) = \frac{(1 - \sigma_0) \cdot (1 - e^{-\frac{\tilde{q}-p}{\lambda}}) - \sigma_0 \cdot (e^{-\frac{p}{\lambda}} - 1)}{(e^{-\frac{\tilde{q}-p}{\lambda}} - 1) \cdot (e^{-\frac{p}{\lambda}} - 1)} \in (0, 1).$$

As in the proof of Lemma 14, $C(\sigma_0)$ is continuous and decreasing in σ_0 . Furthermore, C

is positive when $\sigma_0 = 0$ and negative when $\sigma_0 = 1$. As a result, for every $\pi \in (0, 1)$, by the intermediate value theorem, there is a $\sigma_0 \in (0, 1)$ such $C(\sigma_0) = \pi$.

PROOF OF COROLLARY 2. The sufficiency part of the statement follows by the definition of Nash equilibrium and the characterization of robustness to non-degenerate entropic costs of Corollary 1 point (i). Necessity follows for the proof of Theorem 1. In particular, the fact that in any equilibrium with non-degenerate entropic costs the producer randomizes between at most three qualities follows the same argument of Lemma 7 establishing $|\text{supp } \sigma| \in \{2, 3\}$. ■

References

AGHION, P., M. DEWATRIPONT, AND P. REY (1994): “Renegotiation design with unverifiable information,” *Econometrica: Journal of the Econometric Society*, 257–282.

BATTIGALLI, P. AND G. MAGGI (2002): “Rigidity, discretion, and the costs of writing contracts,” *American Economic Review*, 92, 798–817.

BOYACI, T. AND Y. AKÇAY (2018): “Pricing when customers have limited attention,” *Management Science*, 64, 2995–3014.

CONDORELLI, D. AND B. SZENTES (2020): “Information design in the holdup problem,” *Journal of Political Economy*, 128, 681–709.

CSISZÁR, I. (1974): “On an extremum problem of information theory,” *Studia Scientiarum Mathematicarum Hungarica*, 9, 57–61.

CUSUMANO, C. M., F. FABBRI, AND F. PIEROTH (2024): “Competing to commit: Markets with rational inattention,” *American Economic Review*, 114, 285–306.

DENTI, T., M. MARINACCI, AND L. MONTRUCCHIO (2020): “A note on rational inattention and rate distortion theory,” *Decisions in Economics and Finance*, 43, 75–89.

DENTI, T., M. MARINACCI, AND A. RUSTICHINI (2022): “Experimental cost of information,” *American Economic Review*, 112, 3106–3123.

GEORGIADIS, G., D. RAVID, AND B. SZENTES (2024): “Flexible moral hazard problems,” *Econometrica*, 92, 387–409.

- GEORGIADIS, G. AND B. SZENTES (2020): “Optimal monitoring design,” *Econometrica*, 88, 2075–2107.
- GROSSMAN, S. J. AND O. D. HART (1986): “The costs and benefits of ownership: A theory of vertical and lateral integration,” *Journal of political economy*, 94, 691–719.
- GROUT, P. A. (1984): “Investment and wages in the absence of binding contracts: A Nash bargaining approach,” *Econometrica: Journal of the Econometric Society*, 449–460.
- GUL, F. (2001): “Unobservable investment and the hold-up problem,” *Econometrica*, 69, 343–376.
- HOLMSTRÖM, B. (1979): “Moral hazard and observability,” *The Bell journal of economics*, 74–91.
- KESSLER, A. S. (1998): “The value of ignorance,” *The RAND Journal of Economics*, 339–354.
- LAU, S. (2008): “Information and bargaining in the hold-up problem,” *The RAND Journal of Economics*, 39, 266–282.
- MAĆKOWIAK, B., F. MATĚJKA, AND M. WIEDERHOLT (2023): “Rational inattention: A review,” *Journal of Economic Literature*, 61, 226–273.
- MARTIN, D. (2017): “Strategic pricing with rational inattention to quality,” *Games and Economic Behavior*, 104, 131–145.
- MATĚJKA, F. (2016): “Rationally inattentive seller: Sales and discrete pricing,” *The Review of Economic Studies*, 83, 1125–1155.
- MATĚJKA, F. AND A. MCKAY (2015): “Rational inattention to discrete choices: A new foundation for the multinomial logit model,” *American Economic Review*, 105, 272–298.
- MENSCH, J. (2022): “Screening inattentive buyers,” *American Economic Review*, 112, 1949–1984.
- MENSCH, J. AND D. RAVID (2024): “Monopoly, product quality, and flexible learning,” *arXiv preprint arXiv:2202.09985*.
- MILGROM, P. R. (1981): “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 380–391.

- MIRRLLEES, J. A. (1976): "The optimal structure of incentives and authority within an organization," *The Bell Journal of Economics*, 105–131.
- RAVID, D. (2019): "Bargaining with rational inattention," *Available at SSRN 2957890*.
- (2020): "Ultimatum bargaining with rational inattention," *American Economic Review*, 110, 2948–2963.
- RAVID, D., A.-K. ROESLER, AND B. SZENTES (2022): "Learning before trading: on the inefficiency of ignoring free information," *Journal of Political Economy*, 130, 346–387.
- ROESLER, A.-K. AND B. SZENTES (2017): "Buyer-optimal learning and monopoly pricing," *American Economic Review*, 107, 2072–2080.
- SEGAL, I. (1999): "Complexity and renegotiation: A foundation for incomplete contracts," *The Review of Economic Studies*, 66, 57–82.
- SELTEN, R. (1975): "Reexamination of the perfectness concept for equilibrium points in extensive games," *International Journal of Game Theory*, 4, 25–55.
- SIMS, C. A. (2003): "Implications of rational inattention," *Journal of monetary Economics*, 50, 665–690.
- THEREZE, J. (2024): "Screening costly information," *Working Paper*.
- TIROLE, J. (2009): "Cognition and incomplete contracts," *American Economic Review*, 99, 265–294.
- WOLITZKY, A. (2023): "Unobserved-offers bargaining," *American Economic Review*, 113, 136–173.
- YANG, M. (2015): "Coordination with flexible information acquisition," *Journal of Economic Theory*, 158, 721–738.
- (2020): "Optimality of debt under flexible information acquisition," *The Review of Economic Studies*, 87, 487–536.