

Attention Holdup^{*}

Francesco Fabbri[†]

September 29, 2026

Abstract

A novel holdup problem emerges when a consumer invests in learning about product quality set by a producer. Flexible information acquisition, modeled using rational inattention, mitigates holdup 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, the consumer learns more precisely. However, more precise information makes downward quality deviations more profitable; to deter them, the consumer trades less frequently. When attention costs vanish, holdup becomes inevitable: the consumer never trades, and the market fails. A new refinement, preventing free information acquisition, uniquely selects this binary-quality equilibrium.

KEYWORDS: holdup problem, rational inattention, quality provision.

JEL CODES: 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 Nageeb Ali, Pierpaolo Battigalli, Roland Benabou, Alex Bloedel, Sylvain Chassang, Yeon-Koo Che, Roberto Corrao, Mark Dean, Mira Frick, Yannai Gonczarowski, Ryota Iijima, Navin Kartik, Giacomo Lanzani, Elliot Lipnowski, Qingmin Liu, Alessandro Lizzeri, Leonardo Maltoni, Paola MoscarIELLO, Jacopo Perego, Andrea Prat, Giulio Principi, Doron Ravid, Anne-Katrin Roesler, Fedor Sandomirskiy, Elia Sartori, Ilya Segal, Carol Shou, Steve Tadelis, João Thereze, Quitzé Valenzuela-Stookey, Joel Watson, Alexander Wolitzky, and Leeat Yariv for their helpful comments, and to Carlo Maria Cusumano and Ferdinand Pieroth for many conversations on related modeling issues in earlier joint work. I also thank participants at the 2025 BFI Conference on Rational Inattention, 2026 Canadian Economic Theory Conference, 27th ACM Conference on Economics and Computation, Bay Area Economic Theory Workshop, Bocconi University, Boston University, Columbia University, CSEF Capri in Theory, Duke University, Georgetown University, Harvard University, Kellogg School of Management, MIT, NYU, Princeton University, UC Berkeley, UCLA, UC San Diego, University of Naples Federico II, and WashU.

[†]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 assess 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 a bilateral trade relationship in which a producer makes a relation-specific investment to supply a consumer. Such relationships are a classical setting for holdup: once investment costs are sunk, the consumer may appropriate the resulting surplus, leaving the producer insufficiently compensated. Contracts can protect investment incentives by disciplining quality provision. However, some dimensions of quality are difficult to contract and observe. To discipline quality provision along these dimensions, the consumer may need to acquire information before deciding whether to purchase. Information acquisition is itself a costly investment, undertaken only if, on average, it compensates its cost by improving purchasing decisions. With little incentive to process information, product quality—being costly to improve and hard to assess—may decline. In particular, without processing information, quality may fall to a level that makes trade nonviable, causing this market to fail.

In the context of a quality provision game, we show that flexible information acquisition mitigates the holdup problem by allowing the consumer to acquire imperfect information. Indeed, information that is too imprecise does not sufficiently discipline quality provision, while perfect information is also problematic, since the producer can profitably deviate to the lowest acceptable quality and prevent the consumer from recovering the cost of the information investment. Flexible information therefore allows the consumer to choose exactly the degree of imprecision needed to provide incentives to invest in both quality and information. Crucially, these incentives depend on information processing costs. When processing costs decrease, the consumer learns quality more precisely, leading the producer to provide higher quality levels. We refer to this force as the *attention effect*. At the same time, more precise information exposes the consumer to the producer's opportunistic behavior by making deviations to lower quality more attractive. To deter such deviations, the consumer participates in the market less often, reducing trade frequency. We refer to this force as the *holdup effect*.

The interplay between the attention and holdup effects becomes particularly stark as information costs vanish. As information becomes nearly perfect, slightly lower qualities would still be accepted with high probability, giving the producer strong incentives to deviate downward. By the holdup effect, the consumer deters these deviations by trading less often. In the limit, trade vanishes, and the market fails. Thus, while information investments can be held up when costs are too high—causing market failure—the holdup problem also emerges at the opposite extreme, where information is effectively free. 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 the 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 reduction being zero when the consumer’s beliefs are deterministic. To exclude it, we introduce *Shannon best responses*, which complement credibility by requiring attention strategies to remain optimal under arbitrarily small positive information costs. Shannon best responses ensure that information costs are positive 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 a *no-learning* equilibrium in which the consumer believes quality is too low to pay attention and leaves the market without buying, and, as a consequence, the producer sets a low quality to save on production costs.

We highlight two key features of Shannon-Nash equilibria with positive quality provision: *inefficient trade* and *binary quality*. In particular, the producer randomizes between zero and a higher-than-price quality. We provide intuition for these results. Trade must be inefficient because the consumer cannot accept the producer’s offer for sure without accepting every offer

for sure. But, if the consumer always accepts, the producer optimally deviates to the lowest possible quality. To understand why quality provision is binary, first note that quality must be random, with realizations both below and above the price. When best responses are Shannon, attention is costly, and the consumer never pays attention if quality is deterministic. Without attention, however, the producer would again deviate to lower qualities. Finally, quality provision is binary because, to deter 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 mitigates the holdup problem in information investments 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, $\lambda > 0$, 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 attention 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. At $\lambda = \bar{\lambda}$, the market fails: the producer becomes 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 often to sustain higher-than-price quality. Therefore, the attention investment is no longer compensated.

As previewed, when the attention cost vanishes, $\lambda \rightarrow 0$, the market fails. To discourage deviations, the consumer resorts increasingly often to the no-trade outside option, until quality provision and trade vanish, as in the no-learning equilibrium. By contrast, when information is free, $\lambda = 0$, only one credible equilibrium exists, which is also Shannon-Nash: the *free perfect information* equilibrium—the consumer pays full attention for free, and the producer sets quality equal to the price. This discrepancy stems from the *discontinuity* of credible best responses at zero information costs. Recall that credibility requires the consumer to treat every quality as if it occurs in equilibrium. Therefore, when information is free, credibility requires the consumer to accept every quality above the price: purchasing at these qualities yields a gain, while any additional information is costless. With positive information costs, however, these gains must be weighed against the cost of acquiring information. In equilibrium, as information becomes cheaper, high quality becomes rarer, shrinking the gains from acquiring information. As a consequence, progressively cheaper information does not lead the consumer to accept every quality above the price, as credibility would require at $\lambda = 0$.

This result reinforces, through a different channel, the message of Ravid et al. (2022), who caution against approximating economic environments with small information costs by models with free information. In their setting, the argument rests on equilibrium selection: as information costs vanish, the market converges to the worst free-learning equilibrium. In our setting, where both quality provision and information acquisition are endogenous, the discrepancy

instead reflects a failure of prediction: when information is free, $\lambda = 0$, the unique Shannon-Nash equilibrium predicts the free perfect information outcome; as information costs vanish, $\lambda \rightarrow 0$, the Shannon-Nash equilibrium instead converges to no quality provision and no trade, as in the no-learning equilibrium.

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.

So far, our analysis 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 above a certain threshold, $\bar{p}$. To build intuition, when $p \rightarrow \bar{p}$, we have that $\bar{\lambda} \rightarrow \lambda$, which implies that the producer becomes indifferent between providing zero quality or a quality equal to 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 with price, which implies that the producer would choose the lowest available price sustaining positive quality. Intuitively, low prices provide *commitment benefits* for the producer by changing the consumer's incentives to acquire information. By committing to a lower price, the producer makes attention more costly relative to the price, inducing less selective purchasing and more frequent trade. In particular, as $p \rightarrow \bar{p}$, the threshold $\bar{\lambda}$ approaches λ . As with higher attention costs, lower prices therefore induce equilibria with a higher average trade probability. At the same time, a lower price reduces the producer's revenue from each sale. We show that the benefits of higher trade efficiency outweigh the revenue losses.

Next, we study what happens when production costs vary. We show that profits are higher when costs are higher. As in the price comparison, the result reflects a commitment benefit: high production costs induce lower high-quality levels and allow the consumer to trade more often, increasing trade efficiency. In particular, we show that the positive effects of higher production costs—a lower high-quality level and a higher probability of trade—always outweigh the negative effect of making production more expensive.

We conclude by discussing additional results and alternative assumptions. First, we show that *average quality* increases with attention and production costs, even though the high-quality level decreases, and is non-monotone in price. Then, we show that the binary-quality structure of our equilibrium extends to other information environments satisfying *convex-concave* information—a generalization of additive normal noise. Market failure as information costs vanish is not an artifact of entropic costs. We establish this result in a setting where the consumer observes quality with *additive normal noise* and chooses its precision, the inverse of the variance, paying a cost proportional to it. However, market failure is not the unique limiting prediction: another equilibrium sustains positive trade in the limit. Finally, we extend the analysis to *convex production costs* and show that, conditional on binary-quality Shannon-Nash equilibria, the attention and holdup effects and the limiting market failure result continue to hold. Supplemental Appendix A discusses recoverable production costs, equilibrium purification, consumer heterogeneity, and equilibria without credibility.

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 as it is already explored by Kessler (1998). Roesler and Szentes (2017) explore the role of imperfect information in a setting in which the seller observes the buyer’s signal structure before setting the price. In our framework, the producer’s quality provision needs to incentivize consumer attention; when learning is free, the consumer cannot credibly induce quality above the price. Furthermore, when attention costs vanish, the equilibrium outcome approximates the no-learning equilibrium, which implies zero surplus for both consumer and producer. This result is reminiscent of Ravid et al. (2022), who show that to sustain an equilibrium when information costs are positive, the buyer optimally ignores a large amount of information, selecting the worst free-learning equilibrium when costs vanish. Our result differs in two ways. First, as information costs vanish, our model always leads to complete market failure, with no quality provision and no trade. Second, the consumer distinguishes equilibrium qualities with increasing precision, but high quality becomes increasingly rare because quality provision is endogenous in our model. Thus, the market fails despite near-perfect information, and the limiting outcome lies outside the set of free-learning Shannon-Nash equilibria.

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 et al., 2024). Inefficiencies arise since the producer, lacking 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 equilibria with positive quality provision.

We connect to the classical literature on the holdup problem by interpreting information acquisition as a relation-specific and non-contractible investment (Groux, 1984; Grossman and Hart, 1986; Aghion et al., 1994; Anderlini and Felli, 2006).¹ 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 net 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) and Krähmer (2026) in which the consumer’s investment results in a distribution of valuations for the good. For this reason, as in their case, the consumer has information rents and positive utility.

Our findings relate to the literature studying how costly information acquisition about exogenous traits shapes market outcomes (Matějka, 2016; Martin, 2017; Boyacı and Akçay, 2018; Yang, 2020; Mensch, 2022; Chatterjee et al., 2025).² Mensch et al. (2026) study a model of monopolistic screening and show that quality can be below the efficient level for all types. In our setting, quality is stochastic and may be higher than the free-information 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 environment features moral hazard in product quality, linking us to the moral hazard literature (Mirrlees, 1976; Holmström, 1979; Georgiadis et al., 2024). In particular, one can view our model as an endogenous market for lemons (Akerlof, 1970). In Akerlof’s original market for lemons, individual qualities are exogenous. In our model, by contrast, the quality distribution is endogenously determined by the producer’s moral-hazard choice, while the consumer’s beliefs are endogenously shaped by information acquisition. Similarly to us, Georgiadis and Szentes (2020) study 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, lacking commitment power, acquires it only through imperfect information.

¹See Che and Sákovic (2008) for a detailed literature review. See also Brennan (2007), who shows how costly contract reading can generate a holdup problem and prevent agreement.

²See 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 is in Supplemental Appendix B.

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 *no-learning* 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 no-learning 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. The reason is that imperfect information provides commitment power to the consumer. When information is imperfect, the consumer's purchasing decision does not respond perfectly to realized quality. This imperfect responsiveness disciplines the producer, who cannot

profitably deviate to lower quality levels without affecting the probability of trade.

To illustrate the role of imperfect information, suppose the information technology reveals only whether quality is above or below a fixed *detection 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 deliver sharp predictions 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, a range of quality levels 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.³
- (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.⁴

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

Throughout the analysis, we assume that c is *linear* in quality and normalize the cost of producing zero quality to zero, that is, $c(q) = \alpha \cdot q$, where $\alpha \in (0, 1)$. Constant marginal production costs simplify the study of the producer's incentives while preserving the key trade-off: higher quality increases trade probability but also production costs. Specifically, linear costs ensure that Shannon-Nash equilibria with positive quality provision involve only two quality levels, ruling out more complex distributions that may arise under convex costs.⁵ Nevertheless, the attention

³We 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.

⁴Restricting the consumer's strategy space to attention strategies is without loss under entropy-based information costs. See Matějka and McKay (2015) and Ravid (2019).

⁵We discuss this argument in more detail when we sketch the proofs of our results in Section 3.3.

and holdup effects, and market failure as information costs vanish, do not require linear costs: as we show in Section 4, these results continue to hold under convex production costs when we restrict attention to binary-quality Shannon-Nash equilibria.

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.

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 Supplemental Appendix A, 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, cannot be made indifferent across different qualities when the consumer acquires information, and, as a consequence, the consumer has no incentives to do so.

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, $H(x) = -x \log(x) - (1-x) \log(1-x)$.⁶ 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 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 $q \in Q$.⁷

⁶We follow the usual convention, that is, $H(0) = H(1) = 0$.

⁷All the formal results maintain $\lambda > 0$, while the free-information case, $\lambda = 0$, is discussed separately.

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 *no-learning* 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.⁸

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.

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 a best response to β . Therefore, every quality $q \neq \gamma \cdot p$

⁸This 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.

occurs with zero probability, and, by equation (1), the consumer utility is unaffected by zero-probability events. Since β is a best response to qualities that occur with positive probability, it is a best response to σ . Notice that γ is set such that 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 in equilibrium.

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 may become suboptimal when the producer randomizes over every quality. 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 $\tilde{\sigma}_n \rightarrow \sigma$ strongly,⁹
- β is a best response to $\tilde{\sigma}_n$.

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

Credibility is similar in spirit to Selten’s (1975) trembling-hand perfection, as both require robustness to small mistakes in opponents’ behavior. Under credibility, the perturbations of the producer’s strategy may vary across the off-equilibrium quality levels being tested.

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).

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 rationally inattentive consumer follows the logit formula (2) μ -almost surely. Specifically, when choosing the probability of accepting an offer of

⁹A 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 $A \subseteq Y$.

quality $q \in \text{supp } \mu$, the consumer weighs the gains from trade, $q - p$, against the attention costs, λ , and adjusts by the *average trade probability*, $\pi := \mathbb{E}_\mu[\beta]$.¹⁰

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.¹¹

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 zero 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 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.¹²

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

¹⁰In 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.

¹¹Credible best responses need not remain credible as information costs vanish. To see this, consider $\mu = \delta_0$ and $\beta = 0$, as in the no-learning equilibrium. For every $\lambda > 0$, β is a best response to μ and satisfies equation (2) everywhere with $\pi = 0$; hence, by Lemma 1, it is credible. At $\lambda = 0$, however, it is not credible anymore. Indeed, consider any $\tilde{q} > p$ and any perturbed belief assigning positive probability to $\tilde{q}$. Changing the purchase probability at $\tilde{q}$ from zero to one strictly increases expected utility without incurring any information cost. Thus, $\beta = 0$ is not a best response to any such perturbed belief, showing that the credible-best-response correspondence does not have a closed graph at zero information costs.

¹²See Proposition 13 and Corollary 3 in Supplemental Appendix C for a formal treatment.

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*

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

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.¹³

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.*

¹³An 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 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: the equilibrium strategy is robust to arbitrarily small fixed costs since, as we show in Proposition 3, the consumer has positive utility in equilibrium.

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. Thus, our robustness criterion ensures that the consumer pays attention only when there is something to learn, $|\text{supp}\mu| > 1$.

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 these 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 depends on the reference measure. Shannon best responses instead preserve the logit formula by selecting credible best responses (see Lemma 1).¹⁴

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 et al. (2024). They study, respectively, monopolistic and competitive pricing when the consumer is rationally inattentive and simultaneously learns about the price

¹⁴See Proposition 7 in Denti et al. (2022) for further details, and Supplemental Appendix A for the properties of best responses that are not Shannon, but only robust to non-degenerate entropic costs.

and quality of the good, which is exogenous in their settings. When quality is deterministic, each of these games has 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, whenever prices are non-constant, 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 attention and holdup effect, and the limiting cases of attention costs. Finally, we study how surplus varies as the attention cost varies. Later in Section 4, we discuss how the average quality varies with attention costs, production costs, and price.

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 $q^* > p$.*

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.¹⁵

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

¹⁵Random quality provision need not rely on deliberate randomization by the producer. In Supplemental Appendix A, we purify the Shannon-Nash equilibrium through small, privately observed production cost differences, with quality chosen deterministically conditional on those differences.

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.¹⁶

Main result. Theorem 1 advances this discussion by identifying the form that trade inefficiency and random quality provision take when the consumer is rationally 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 high quality $q^*(\lambda) > p$ is strictly decreasing in λ ; (attention effect)*
- *The average trade probability $\pi^*(\lambda) \in (0, 1)$ is strictly increasing in λ . (holdup effect)*

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 mitigates the holdup problem in information investment and sustains positive quality. Specifically, such an equilibrium exists if and only if the attention cost is below $\bar{\lambda}$. When $\lambda > \bar{\lambda}$, information investment is held up, as the producer finds it optimal to deviate and set quality to zero. When the attention cost vanishes, $\lambda \rightarrow 0$, the market also fails: quality provision and trade vanish, yielding the same outcome as the no-learning equilibrium, with a market composed of lemons only. The interaction of two forces drives these results—the *attention effect* and the *holdup effect*, which imply, respectively, that the high-quality level and the average trade probability decrease and increase in the attention cost.

The *attention 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

¹⁶For instance, see Section 4 for a model where the consumer learns the producer's quality by acquiring convex-concave information, a framework which includes additive normal noise.

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, implying the result.

The *holdup effect* implies that the average trade probability increases in the attention cost. By the attention 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. To see why, suppose that a more attentive consumer traded on average at least as frequently, that is, π does not decrease as λ decreases. By equation (2), then the probability of buying would increase for all quality levels above the price. 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 $\hat{q}$, while equilibrium profits $p\beta(0)$ fall as λ decreases. 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 attention and the holdup effect severely affects the market when the attention cost vanishes. In this scenario, the consumer frequency of mistakes also vanishes due to the attention 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 holdup 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 never trades forcing the high-quality frequency to converge to zero. Intuitively, as $\lambda \rightarrow 0$, consumer information becomes arbitrarily precise, but the market outcome approximates no-learning.

We contrast this limiting result with the *free-learning* benchmark. When attention is free, $\lambda = 0$, there is a unique credible equilibrium, which is also Shannon-Nash: the *free perfect information* equilibrium, in which the producer sets quality equal to the price and the consumer buys with certainty. To see this, recall that credibility requires attention strategies to remain optimal under perturbations assigning positive probability to every quality. Since information is free, optimality under these perturbations, together with equilibrium conditions, implies $\beta(q) = 1$ for $q \geq p$ and $\beta(q) = 0$ for $q < p$. However, Theorem 1 shows that, as $\lambda \rightarrow 0$, the Shannon-Nash equilibrium converges to no quality provision and no trade, as in the no-learning equilibrium. This discrepancy reflects the *discontinuity* of credible best responses at zero information costs. Along the equilibrium sequence, the consumer distinguishes equilibrium qualities with increasing precision, $\beta(q^*) \rightarrow 1$ and $\beta(0) \rightarrow 0$, but high quality becomes increasingly rare, $\sigma(0) \rightarrow 1$. This decline shrinks the gains from acquiring information, so cheaper information does not lead the




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: $\alpha = 0.1$, $p = 10$.

consumer to accept every off-equilibrium quality above the price, even under the credibility perturbations that assign these qualities positive probability. Indeed, for every $q \in (p, p/\alpha)$, $\beta(q) \rightarrow 0$, while credibility at $\lambda = 0$ requires $\beta(q) = 1$.¹⁷

The market also fails when the attention cost reaches the opposite threshold, $\lambda = \bar{\lambda}$. As $\lambda \rightarrow \bar{\lambda}$, we have that $q^* \rightarrow p$ and $\sigma(0) \rightarrow 0$, and the market approaches the artificial credible equilibrium described in Section 2.2, where the producer extracts the entire surplus.¹⁸ In our model, higher inattention generates more frequent mistakes, which increase profits and reduce the high-quality level. Along this sequence, as formalized by Proposition 3 below, producer profits progressively increase while consumer utility approaches zero. Ultimately, when the attention cost $\lambda = \bar{\lambda}$, under the limiting attention strategy, the producer is indifferent between providing quality 0 and p , but the market fails since the consumer no longer has any incentive to pay attention. Only the no-learning Shannon-Nash equilibrium remains, and profits drop to zero. For $\lambda \geq \bar{\lambda}$, investing in information no longer mitigates the holdup problem. The reason is that the consumer makes mistakes too frequently— $\beta(0)$ is too high—to sustain higher-than-price quality.

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 attention effect (on the left) and holdup effect (on the right).

¹⁷While the free-learning benchmark fails to approximate Shannon-Nash equilibria with vanishing information costs, the limit still selects one of the worst free-learning Nash equilibria, without the credibility refinement, where both players obtain zero surplus.

¹⁸This limiting profile is also a Nash equilibrium of the free-information game without credibility, in which the consumer buys sufficiently often to sustain quality equal to the price.




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

Producer and consumer surplus. Our discussion so far connects naturally to how producer profits and consumer utility vary with 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 $\lambda \rightarrow \bar{\lambda}$.*

As the unit of attention cost increases, the profits increase as the producer benefits from both the attention and the holdup effect. The attention effect implies a lower high-quality level, while the holdup 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 attention and holdup effects pull consumer utility in opposite directions. As consumer attention increases, the attention effect benefits the consumer by reducing mistakes, while the holdup 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 strictly prefers β over the option of never purchasing, which yields zero utility.

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.

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 a Pareto improvement 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, $\lambda = \bar{\lambda}$, under the limiting attention strategy, the producer is indifferent between providing 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 $p = \bar{p}$, under the limiting attention strategy, the producer is indifferent between providing 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 making attention more costly relative to the price, they induce less selective purchasing, allowing the producer to sell 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 1.¹⁹ 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. 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; a refinement requires equilibria to survive when price is chosen before quality, ruling out those

¹⁹The proof of Theorem 2 in Supplemental Appendix C establishes that π^* decreases in p , while q^* increases.



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: $\alpha = 0.1$, $\lambda = 10$.

sustained by adverse beliefs following off-equilibrium prices. 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 strengthens the incentive to make a sale and therefore 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 with positive quality provision 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.




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: $\lambda = 10$, $p = 10$.

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 with the previous thresholds, $\bar{\lambda}$ and $\bar{p}$. When $\alpha = \bar{\alpha}$, under the limiting attention strategy, the producer is indifferent between providing zero quality or a quality equal to the price. If $\alpha > \bar{\alpha}$, production costs are too high at the given price and attention cost to sustain positive quality provision. For a similar reason, the threshold decreases when λ increases or p decreases. Either change breaks the indifference at $\alpha = \bar{\alpha}$, making deviations to zero quality more attractive: the former makes purchasing less selective and the latter makes purchasing zero quality a less costly mistake for the consumer.

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 relatively small compared to $\bar{\lambda}$, mirroring the scenario where $\lambda \rightarrow 0$. In this case, $q^* \rightarrow \infty$ and $\pi^* \rightarrow 0$.²⁰ High production costs, and therefore modest cost-reducing investments, serve as a *commitment device* by limiting the high-quality level. This avoids the combination of extreme quality and infrequent trade that emerges when production becomes arbitrarily cheap.

²⁰The proof of Proposition 5 in Supplemental Appendix C establishes that π^* increases in α , while q^* decreases.

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—overcome 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 production costs, illustrating Proposition 5.

3.3 Outline of the main arguments

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

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 (2) 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 $\text{supp } \sigma = \{0, q^*\}$, $q^* > p$, $\sigma(0) := \sigma_0 \in (0, 1)$, β satisfies equation (2) for every $q \in Q$ with $\mathbb{E}_\sigma[\beta] := \pi^* \in (0, 1)$, and the following conditions hold:*

$$\begin{aligned}
 (i) \quad & p \cdot \frac{\pi^* \cdot e^{\frac{q^*-p}{\lambda}}}{\pi^* \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi^*} - \alpha \cdot q^* = p \cdot \frac{\pi^* \cdot e^{\frac{-p}{\lambda}}}{\pi^* \cdot e^{\frac{-p}{\lambda}} + 1 - \pi^*} & (\text{indifference}) \\
 (ii) \quad & \frac{\pi^* \cdot e^{\frac{q^*-p}{\lambda}}}{\pi^* \cdot e^{\frac{q^*-p}{\lambda}} + 1 - \pi^*} = \frac{1}{2} + \frac{1}{2} \cdot \sqrt{1 - 4 \cdot \frac{\alpha \lambda}{p}} & (\text{tangency}) \\
 (iii) \quad & \pi^* = \frac{(1 - \sigma_0) \cdot (1 - e^{\frac{q^*-p}{\lambda}}) - \sigma_0 \cdot (e^{\frac{-p}{\lambda}} - 1)}{(e^{\frac{q^*-p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)} & (\text{consumer optimal})
 \end{aligned}$$

The profile (σ, β) is a Shannon-Nash equilibrium with positive quality provision if and only if it satisfies binary-quality (by Lemma 2), β is optimal (by Lemma 1), 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).²¹

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.

²¹See Proposition 1 in the Supplemental Appendix of Matějka and McKay (2015).

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 Supplemental Appendix C.

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^* 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 condition fails.

Attention and holdup 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 attention effect, while the tangency condition, preventing deviation to off-equilibrium qualities, to the holdup 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 $q^* \in (p, p/\alpha)$, 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.

Average quality. We complement our analysis by discussing how the average quality varies with attention costs, production costs, and prices. We present these comparisons separately because our previous results, and their interpretation, do not rely on them. In particular, we show that the average quality increases with both attention costs and production costs, even though the high-quality level decreases. The relation with prices is instead non-monotone.

PROPOSITION 6. *In the unique Shannon-Nash equilibrium with positive quality provision, the average quality $\mathbb{E}_\sigma[q] = \sigma(q^*)q^*$ satisfies the following properties, holding the other parameters fixed:*

- *It is strictly increasing in $\lambda \in (0, \bar{\lambda})$, with $\mathbb{E}_\sigma[q] \rightarrow 0$, as $\lambda \rightarrow 0$, and $\mathbb{E}_\sigma[q] \rightarrow p$, as $\lambda \rightarrow \bar{\lambda}$;*
- *It is strictly increasing in $\alpha \in (0, \bar{\alpha})$, with $\mathbb{E}_\sigma[q] \rightarrow 0$, as $\alpha \rightarrow 0$, and $\mathbb{E}_\sigma[q] \rightarrow p$, as $\alpha \rightarrow \bar{\alpha}$;*
- *It is non-monotone in the price $p \in (\bar{p}, \infty)$. In particular,*

$$\lim_{p \rightarrow \bar{p}} \mathbb{E}_\sigma[q] = \bar{p}, \quad \lim_{p \rightarrow \bar{p}} \frac{d}{dp} \mathbb{E}_\sigma[q] = 1, \quad \lim_{p \rightarrow \infty} \mathbb{E}_\sigma[q] = 0.$$

As a result, the average quality is maximized at an interior price.

Higher attention costs and higher production costs reduce the high-quality level but increase its frequency. The latter effect dominates, raising average quality. Higher prices instead raise the high-quality level and reduce its frequency. Close to $\bar{p}$, average quality increases with price; as prices become arbitrarily high, high quality becomes sufficiently rare that average quality approaches zero.

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 $s_1 > s_2, q_1 > q_2$, $f(s_2|q_2) \cdot f(s_1|q_1) > f(s_1|q_2) \cdot f(s_2|q_1)$;

(iv) For interior signals, the following holds. 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 $(q_s, p/\alpha]$.

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 identify consumer strategies that differ only 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 $\epsilon \sim N(0, 1/\tau)$ is a zero-mean, independent, normally distributed random variable with variance $1/\tau > 0$, and κ is proportional to τ . Furthermore, the property (iv) is satisfied whenever ϵ is drawn from a distribution unimodal at zero.²²

PROPOSITION 7. *If there exists a Nash equilibrium with positive quality provision when information acquisition satisfies convex-concave information, then it 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.

Gaussian information. A natural question is whether market failure with vanishing information costs can arise in other information environments, or is an artifact of entropic costs or, more generally, prior-dependent costs. We provide a positive answer to this question. When the consumer observes quality with additive normal noise and chooses its precision, the inverse of the variance, paying a cost proportional to it, we construct a Nash equilibrium in which quality provision and trade vanish as information costs approach zero, as in the no-learning equilibrium. However, another equilibrium also exists: high-quality approaches the price, its frequency approaches one, and trade remains positive. Consumer utility vanishes along both families. In the former, the whole market shuts down; in the latter, the producer captures all surplus. Market

²²A 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 Supplemental Appendix D for a formal argument.

failure for vanishing information costs is not an artifact of entropic costs, although different models of flexible information acquisition may allow for broader predictions.

Formally, the consumer chooses a precision $\tau > 0$, subject to the cost $\lambda \cdot \tau$, and observes $s = q + \epsilon$, where $\epsilon \sim N(0, 1/\tau)$ is independent of quality. After observing the signal, the consumer decides whether to buy. Alternatively, the consumer can acquire no information at no cost and buy following the prior. The rest of the game is unchanged.

PROPOSITION 8. *Fix $(\alpha, p) \in (0, 1) \times (0, \infty)$. There exists $\bar{\lambda}_G > 0$ such that, for every $\lambda \in (0, \bar{\lambda}_G)$, there are at least two distinct binary-quality Nash equilibria with positive quality provision. These equilibria, $\{(\sigma_i, \beta_i)\}_{i \in \{1, 2\}}$, have precision $\tau_i > 0$, $\text{supp} \sigma_i = \{0, q_i\}$ with $p < q_i < p/\alpha$, average trade probability π_i , and satisfy the following:*

- As $\lambda \rightarrow 0$, $\tau_1 \rightarrow \infty$ and $q_1 \rightarrow p/\alpha$, but the market fails: $\sigma_1(0) \rightarrow 1$ and $\pi_1 \rightarrow 0$. Furthermore, $\beta_1(0) \rightarrow 0$ and $\beta_1(q_1) \rightarrow 1$.
- As $\lambda \rightarrow 0$, $\tau_2 \rightarrow \tilde{\tau} \in (0, \infty)$, $q_2 \rightarrow p$, and $\sigma_2(0) \rightarrow 0$, while $\pi_2 \rightarrow \tilde{\pi} \in (\alpha, 1)$. The constants $\tilde{\tau}$ and $\tilde{\pi}$ depend only on (α, p) and α , respectively.

Producer profits and consumer utility are positive in both equilibria for every $\lambda \in (0, \bar{\lambda}_G)$. As $\lambda \rightarrow 0$, consumer utility vanishes in both, while producer profits converge to zero in the first and to $p(\tilde{\pi} - \alpha) > 0$ in the second.

In the first equilibrium, the consumer distinguishes qualities with increasing precision, but high quality becomes increasingly rare. As in Theorem 1, the market fails. In the second equilibrium, instead, high quality becomes increasingly common and approaches the price, shrinking the consumer's gains from acquiring information. Along this equilibrium, precision remains bounded even as its cost vanishes. Although the producer captures all surplus in the limit, this equilibrium does not converge to the free-perfect-information equilibrium with certain trade: quality approaches the price, but trade remains imperfect.

There is a parallel between the detection-threshold case discussed in Section 1.2, credible equilibria under entropic costs, and equilibria with Gaussian information. Each setting admits a no-learning equilibrium and additional equilibria with contrasting limiting behavior: one approaches market failure, while another sustains positive trade as information costs vanish. Under entropic costs, the latter is the artificial credible equilibrium described in Section 2.2, which approaches the free-perfect-information equilibrium as $\lambda \rightarrow 0$. The limits differ across settings. With a fixed detection threshold above the price, quality remains above the price and consumer utility converges to a strictly positive value along the positive limiting equilibrium. With Gaussian information, quality instead approaches the price, utility vanishes, and trade remains below one.

Convex production costs. We now extend our results to the case of convex production costs. Specifically, we say that production costs are *convex* if $c : [0, \infty) \rightarrow [0, \infty)$ is twice continuously differentiable, strictly increasing, and convex, with $c(0) = 0$. We further assume $c(p) < p$, and let $\bar{q} := c^{-1}(p)$ be the maximum feasible quality given the fixed price p . As in the linear case, it is without loss to restrict to qualities in $[0, \bar{q}]$, since higher qualities yield negative profits. The rest of the game is unchanged.

Convex costs make equilibrium quality provision more involved. Recall the argument outlined in step 2 after Lemma 2: with linear costs, the isoprofit curve is also linear and must lie above the s-shaped demand. With convex costs, instead, the isoprofit curve is convex; $\beta(q) = (c(q) + \bar{\Pi})/p$. This allows for knife-edge cases where the isoprofit curve touches the demand at a low quality, but then coincides with it over an interval in its convex region, and finally touches it again at a higher quality. Therefore, the same argument no longer restricts equilibrium quality to be binary.

To avoid this issue, we focus on binary-quality Shannon-Nash equilibria with positive quality provision, without establishing their existence or uniqueness. Furthermore, to carry out the comparative statics, we assume that the two support qualities and the average trade probability are continuously differentiable in the parameter being varied.

PROPOSITION 9. *Fix convex production costs, c , and the price, p . Consider binary-quality Shannon-Nash equilibria with positive quality provision indexed by λ , and let $\text{supp } \sigma = \{q_l(\lambda), q_h(\lambda)\}$, where $q_l(\lambda) < p < q_h(\lambda)$. In these equilibria:*

- *The higher-than-price quality $q_h(\lambda) > p$ is strictly decreasing in the attention cost;*
- *The average trade probability $\pi^*(\lambda) \in (0, 1)$ is strictly increasing in the attention cost.*

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

The attention and holdup effects play the same role as in Theorem 1. Lower attention costs induce higher high-quality levels, but also reduce the average trade probability needed to discourage producer deviations. As attention costs vanish, high quality approaches its maximum level, but it becomes increasingly rare, while the lower quality converges to zero. Thus, the market fails.

Notice that binary quality does not require $q_l = 0$. To see this, consider the case $c(q) = q^2$. Positive quality provision, $\pi^* \in (0, 1)$, and equation (2) imply

$$\left. \frac{d}{dq} \Pi(q, \beta) \right|_{q=0} = \frac{p}{\lambda} \beta(0)(1 - \beta(0)) > 0,$$

meaning that the producer prefers a small positive quality to zero. Proposition 9 allows both equilibrium qualities to be positive and establishes that $q_l \rightarrow 0$ as $\lambda \rightarrow 0$.

PROPOSITION 10. *Fix convex production costs, c . In binary-quality Shannon-Nash equilibria with positive quality provision, producer profits and consumer utility are positive. Profits are strictly increasing in λ , holding p fixed, and strictly decreasing in p , holding λ fixed. Furthermore, if $c(q) = \alpha \cdot \tilde{c}(q)$ for some $\alpha > 0$, profits are strictly increasing in α , holding $(\tilde{c}, p, \lambda)$ fixed. For fixed p , profits and consumer utility approach zero as $\lambda \rightarrow 0$.*

As in the linear cost model, higher attention costs and lower prices provide commitment benefits that increase profits. For the consumer, the attention and holdup effects continue to work in opposite directions: greater precision reduces mistakes but lowers trade frequency. The limiting result implies, also in this more general convex setting, that positive attention costs benefit the whole market.

Notice that the comparison for production costs requires specifying how the cost function changes. As in the linear case, a higher α raises costs proportionally at every positive quality. With convex costs, a non-proportional increase can affect the costs of the two equilibrium qualities differently, and its effect on profits is less immediate.

References

- AGHION, P., M. DEWATRIPONT, AND P. REY (1994): “Renegotiation design with unverifiable information,” *Econometrica*, 257–282.
- AKERLOF, G. A. (1970): “The Market for “Lemons”: Quality uncertainty and the market mechanism,” *Quarterly Journal of Economics*, 84, 488–500.
- ANDERLINI, L. AND L. FELLI (2006): “Transaction costs and the robustness of the Coase theorem,” *The Economic Journal*, 116, 223–245.
- 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.
- BRENNAN, J. R. (2007): “Contracting with Reading Costs and Renegotiation Costs,” Ph.D. thesis, University of California, San Diego.

- CHATTERJEE, K., M. DONG, AND T. HOSHINO (2025): “Bargaining and information acquisition,” *American Economic Journal: Microeconomics*, 17, 75–102.
- CHE, Y.-K. AND J. SÁKOVICS (2008): “Hold-up problem,” in *The New Palgrave Dictionary of Economics*, Springer, 1–7.
- 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*, 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.
- KRÄHMER, D. (2026): “The hold-up problem with flexible unobservable investments,” *Journal of Economic Theory*, 235, 106205.

- 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., D. RAVID, AND J. THEREZE (2026): “Monopoly, quality, and costly information,” *Working Paper*.
- MILGROM, P. R. (1981): “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 380–391.
- MIRRELEES, 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.
- 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.

Supplemental Appendix

The supplemental appendix is organized as follows. Supplemental Appendix A discusses further results about the model's assumptions, while Supplemental Appendix B examines the holdup problem discussed in Section 1.2. Supplemental Appendix C provides formal proofs of our main results from Section 2.2 and Section 3. Finally, Supplemental Appendix D presents the proofs of the results in Section 4 and Supplemental Appendix A.

A Further discussion

We continue and conclude our discussion related to the model's assumptions and results started in Section 4.

Recoverable production costs. We drop the assumption of sunk production costs by assuming that costs are *recoverable*: the producer incurs production costs only if trade occurs. 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 expectation, while with recoverable costs no losses almost surely. The following proposition conforms with this intuition.

PROPOSITION 11. *When production costs are recoverable, there is no Shannon-Nash equilibrium where trade occurs with positive probability.*

Flexible information does not mitigate the holdup problem when costs are recoverable: no Shannon-Nash equilibrium involves positive trade. The producer may provide positive quality, or randomize across different qualities, as long as the resulting distribution is too unattractive for the consumer to acquire information or engage in trade. Producing qualities that are never sold is not an issue for the producer, since production costs are recovered when no trade occurs.

The argument proceeds as follows. Suppose there is an equilibrium with $\pi \in (0, 1)$. Since the consumer's best response is credible, it satisfies the logit formula (2) for every quality level. Because this demand is upward 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 $\pi \in \{0, 1\}$, a contradiction. Finally, the case $\pi \in \{0, 1\}$ does not allow for positive trade probability in equilibrium.

Notice that, 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.²³ 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.

Purification. Another natural question is whether the unique Shannon-Nash mixed equilibrium with positive quality provision (Theorem 1) can be purified through small production cost differences, privately observed by the producer, with deterministic qualities provided conditional on those differences. The answer is again affirmative.

Formally, we define a *perturbed game* as follows. Suppose that the producer's marginal production cost is $\alpha + \epsilon\theta$, where $\epsilon \in (0, \min\{\alpha, 1 - \alpha\})$, and θ is uniformly distributed on $\Theta := [-1, 1]$ and privately observed by the producer before choosing quality. The consumer knows this distribution but does not observe θ ; the rest of the game is unchanged.

Let $\check{\sigma} : \Theta \rightarrow \Delta([0, \infty))$ be a type-dependent producer's strategy. With a slight abuse of notation, denote by $\sigma \in \Delta([0, \infty))$ the quality distribution induced by $\check{\sigma}$ and the distribution of producer types. A strategy profile $(\check{\sigma}, \beta)$ is a *Bayes-Nash equilibrium* (BNE) if $\beta : [0, \infty) \rightarrow [0, 1]$ is a best response to σ , and $\check{\sigma}(\theta)$ is a best response to β for every $\theta \in \Theta$.

PROPOSITION 12. *Fix (α, p) , and $\lambda \in (0, \bar{\lambda})$. Let (σ^*, β^*) be the Shannon-Nash equilibrium with positive quality provision, $\text{supp } \sigma^* = \{0, q^*\}$, of Theorem 1. There exists $\bar{\epsilon} > 0$ such that, for every $\epsilon \in (0, \bar{\epsilon})$, the perturbed game admits a BNE, $(\check{\sigma}_\epsilon, \beta_\epsilon)$, in which $\check{\sigma}_\epsilon(\theta)$ assigns probability one to a quality $q_\epsilon(\theta) \in \{0\} \cup (p, p/\alpha)$ for every θ , and β_ϵ is a Shannon best response. In these equilibria:*

- *There is a threshold $\theta_\epsilon \in (-1, 1)$ such that $q_\epsilon(\theta) > p$ for $\theta < \theta_\epsilon$, and $q_\epsilon(\theta) = 0$ for $\theta > \theta_\epsilon$; the threshold type is indifferent and can choose either quality;*
- *As $\epsilon \rightarrow 0$, we have that $\theta_\epsilon \rightarrow 2\sigma^*(q^*) - 1$ and*

$$\sup_{\theta < \theta_\epsilon} |q_\epsilon(\theta) - q^*| \rightarrow 0, \quad \sigma_\epsilon \implies \sigma^*, \quad \sup_{q \geq 0} |\beta_\epsilon(q) - \beta^*(q)| \rightarrow 0,$$

²³Following the same calculations as in the proof of Proposition 13, this credible equilibrium exists for $\lambda < \frac{p(1-\alpha)}{\alpha}$ with $\pi = 1 - \frac{\alpha\lambda}{p(1-\alpha)}$. At the threshold, $\lambda = \frac{p(1-\alpha)}{\alpha}$, it becomes a Shannon-Nash equilibrium: the consumer never buys, $\beta = 0$; the producer sets $q = p$ since production costs are fully recovered.

where $\Rightarrow$ denotes weak convergence.²⁴ Therefore, average quality, average trade probability, expected profits, and consumer utility converge to their values under (σ^*, β^*) .

The proposition purifies the producer's mixed strategy: uncertainty about quality arises from private cost differences and every producer type makes a deterministic choice. Types with lower costs choose a positive quality, higher than the price, while types with higher costs produce zero quality. As the cost differences vanish, the market converges to the Shannon-Nash equilibrium with positive quality provision: positive qualities converge to its high-quality level and the fraction of types choosing them converges to its equilibrium probability.

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: upon monitoring, the consumer always buys when the valuation exceeds the price, may buy when the valuation equals the price, and never buys when the valuation is below the price.²⁵

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 no-learning. 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

²⁴A sequence $(\nu_n)_{n \in \mathbb{N}}$ of Borel probability measures over $Y \subseteq \mathbb{R}$ converges weakly to a Borel probability measure ν if $\int f(q) d\nu_n(q) \rightarrow \int f(q) d\nu(q)$ for every bounded continuous function $f : Y \rightarrow \mathbb{R}$.

²⁵A formal definition of this strategy space is in Supplemental Appendix B.

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, 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 revenues from trade unaltered.

A similar argument applies when consumer types are infinite. The only difference is that now there may not exist a lowest consumer type that monitors in equilibrium. Therefore, we need to consider $\underline{\theta}$, the infimum of the types that monitor with positive probability. For every such type θ , monitoring is weakly preferred to not buying, so $\int \max\{\theta q - p, 0\} d\sigma(q) \geq \kappa$. Since quality is bounded, taking a sequence of these types converging to $\underline{\theta}$ gives $\int \max\{\underline{\theta} q - p, 0\} d\sigma(q) \geq \kappa > 0$. This implies that $\underline{\theta} > 0$ and that the producer provides qualities strictly higher than $p/\underline{\theta}$ with positive probability. Therefore, the same deviation described above is profitable.

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 either a best response to μ with $\mathbb{E}_\mu[\beta] \in (0, 1)$ or a constant 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 .*

Similarly to Proposition 1, robustness to non-degenerate entropic costs has a limited bite when quality is stochastic: it excludes best responses that are constant μ -almost surely, i.e., $\mathbb{E}_\mu[\beta] \in \{0, 1\}$, but are not constant everywhere. 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 the 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 $\tilde{q}$.²⁶ 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 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) . Consider a profile (σ, β) such that $\beta(q) = 0$ for every $q \notin \text{supp } \sigma$. This profile is an equilibrium with non-degenerate entropic costs and positive quality provision if and only if $|\text{supp } \sigma| \in \{2, 3\}$, β satisfies equation (2) on $\text{supp } \sigma$ with $\pi = \pi^* \in (0, 1)$, and the following conditions hold:*

- (i) *For every $q \in \text{supp } \sigma$, $\Pi(q, \beta) = k$ for some $k \geq 0$; (indifference)*
- (ii) $\pi^* \in \arg \max_{\pi \in [0, 1]} \sum_q \sigma(q) \cdot (\lambda \cdot \ln(\pi \cdot e^{\frac{q-p}{\lambda}} + 1 - \pi))$ (consumer optimal)

First, notice that the condition $\beta(q) = 0$ for every $q \notin \text{supp } \sigma$ is immaterial for equilibrium outcomes: setting off-path purchase probabilities to zero preserves the consumer's utility and cannot make producer deviations more profitable. Under this restriction, condition (i) ensures that σ is a best response, since off-path qualities yield non-positive profits. Condition (ii), together with the logit formula, ensures consumer optimality by Lemma 2 in Matějka and McKay (2015). By Corollary 1, $\pi^* \in (0, 1)$ implies robustness.

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.

²⁶For all $\lambda < \bar{\lambda}$, the existence of at least one such $\tilde{q} \in (p, p/\alpha)$ follows by Theorem 1.

B 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 for pure strategies is $\kappa(\beta) = \kappa \cdot \mathbb{1}_{\{\beta \neq 0, \beta \neq 1\}}$, with $\kappa > 0$. For mixed strategies, costs are evaluated in expectation: monitoring with probability b incurs expected cost $b \cdot \kappa$. 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 no-learning equilibrium.

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

First, if any equilibrium with $\sigma \neq \mathbb{1}_{\{q=0\}}$ exists, 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, in any such equilibrium, 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 responds by setting $\beta = 0$; if $\tilde{q} \in (p, p/\alpha]$, then the consumer sets $\beta = 1$; if $\tilde{q} = p$, then the consumer is indifferent between playing 0 and 1 and never monitors.

In any such 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,

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

where the first, second, and last equalities hold 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 $\hat{\beta}$ 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 are unchanged.

The no-learning 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 no-learning, 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 $\bar{\alpha}/p$. 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, and strictly prefers both to 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, and strictly prefers both to 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.

Notice that, for $\kappa < \frac{p}{\bar{q}}(\bar{q} - p)$, the two equilibria described above are the only equilibria with positive quality provision when $\bar{q} < p/\alpha$. When $\bar{q} = p/\alpha$, instead, there is an interval of equilibria with these two equilibria as its endpoints. Within this interval, the consumer always monitors, while the producer provides $\bar{q}$ with any probability in $[\frac{\kappa}{\bar{q}-p}, \frac{p-\kappa}{p}]$ and sets quality to zero otherwise. Producer profits are zero, while consumer utility is strictly positive except at $\sigma(\bar{q}) = \frac{\kappa}{\bar{q}-p}$, where it equals zero.

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,

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

This indifference gives rise to 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 are 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 .

C 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 13. *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 sets $\pi = \pi^* \in (0, 1)$ 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^* be 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 non-negative if and only if $\beta(p) \geq \alpha$. Furthermore, this equilibrium requires that $\beta(p)$ is well-defined and strictly larger than $1/2$, i.e., $\alpha\lambda/p < 1/4$. 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 such a 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 $\Pi(\delta_0, \beta^*) > 0$, 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 $\bar{\lambda} \in (0, \lambda^e)$ 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. By equation (13) and Lemma 9, $\alpha(\lambda) = \alpha$ if and only if $q^*(\lambda) = p$. Since $q^*(\lambda)$ is strictly decreasing by Lemma 10, there is at most one such λ .

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 $\varepsilon > 0$ such that $\beta \in \arg\max \int \beta(q)(q - p)d\mu(q) - \lambda \cdot I_\varepsilon(\mu, \beta)$.

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 $I_\varepsilon(\mu, \beta) = I(\mu, \beta) = 0$.

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 $I(\mu, \beta) > 0$. Therefore, by setting

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

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

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

(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 that it 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 $\varepsilon > 0$, $U(\beta) = 0 > -\lambda \cdot I_\varepsilon(\mu, \tilde{\beta}) = -\lambda \cdot \varepsilon$, where the strategy β satisfies $\mathbb{E}_\mu[\beta] = 1$. Therefore, any Shannon best response $\tilde{\beta}$ 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 13 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 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 accepts, 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, β is constant 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 13 for an argument establishing these facts.

Notice that positive quality provision requires $\lambda < \lambda^e := p/(4\alpha)$; otherwise, producer profits are strictly decreasing in quality.

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 neither strictly concave nor strictly convex 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 implies $\beta(0) < 1/2$.

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

$$q \in \arg\max \beta(q) \cdot p - \alpha \cdot q,$$

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

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

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. ■

PROOF OF THEOREM 1. In order for two or more qualities to maximize profits, 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 a 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,

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

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 $\lambda \in (0, \lambda^e)$. We now show that, for every $\lambda \in (0, \lambda^e)$, 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

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

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

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

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,

$$1 - \frac{2e^{\frac{q}{\lambda}} (1 - K(\lambda))}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q}{\lambda}}} \Big|_{q=0} > 0$$

$$\lim_{q \rightarrow +\infty} 1 - \frac{2e^{\frac{q}{\lambda}} (1 - K(\lambda))}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q}{\lambda}}} < 0$$

$$\frac{d}{dq} \left(1 - \frac{2e^{\frac{q}{\lambda}} (1 - K(\lambda))}{K(\lambda) + (1 - K(\lambda)) \cdot e^{\frac{q}{\lambda}}} \right) < 0$$

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, strictly 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). Fix $\tilde{q} \in (0, p/\alpha)$. Since $K(\lambda) \rightarrow 1$ and $(1 - K(\lambda))e^{\frac{\tilde{q}}{\lambda}} \rightarrow +\infty$, we have $\lim_{\lambda \rightarrow 0} f(\lambda, \tilde{q}) = \alpha \tilde{q} - p < 0$. By the proof of Lemma 9, $f(\lambda, q) < 0$ for $0 < q < q^*(\lambda)$ and $f(\lambda, q) > 0$ for $q > q^*(\lambda)$. Therefore, for sufficiently small λ , $q^*(\lambda) > \tilde{q}$. Furthermore, equation (13) implies $\alpha q^*(\lambda) = p(K(\lambda) - \beta(0)) < p$. Since $\tilde{q}$ can be arbitrarily close to p/α , it follows that $\lim_{\lambda \rightarrow 0} q^*(\lambda) = p/\alpha$. The same equation then implies that $\lim_{\lambda \rightarrow 0} \beta(0) = 0$.

We show that, for every $q \in (0, p/\alpha)$, 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 . Since $q^*(\lambda)$ is continuous, this implies that it is strictly monotone. Together with $q^*(\lambda) < p/\alpha$ and $q^*(\lambda) \rightarrow p/\alpha$ as $\lambda \rightarrow 0$, this implies that $q^*(\lambda)$ is decreasing.

Fix $\bar{q} \in (0, p/\alpha)$. 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 $\bar{q} \in (0, p/\alpha)$, $f(\lambda, \bar{q})$ 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 $\bar{q} \in (0, p/\alpha)$, we have that

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

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

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

Combining these expressions and using $K(\lambda)(1-K(\lambda)) = \alpha\lambda/p$, we obtain

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

At any λ with $f(\lambda, \bar{q}) = 0$, $\bar{q}$ is the unique positive solution identified in Lemma 9. By the proof therein, $\frac{d}{dq} f(\lambda, \bar{q}) > 0$. Since $\frac{d}{d\lambda} K(\lambda) < 0$, both terms above are positive. Thus, at every solution of $f(\lambda, \bar{q}) = 0$, the function crosses from negative to positive, implying that $\hat{\lambda}$ is unique.

Finally, the derivatives of f with respect to λ and q , holding the other argument fixed, are both strictly positive when evaluated at the positive solution. Therefore, by the implicit function theorem, we have that $\frac{d}{d\lambda} q^*(\lambda) < 0$. $\square$

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). Notice that equation (15) depends on p and λ only through their ratio, λ/p . Therefore, by uniqueness of the threshold, $\bar{\lambda}(\alpha, p) = p \bar{\lambda}(\alpha, 1)$. Since $\bar{\lambda}(\alpha, 1) > 0$, $\bar{\lambda}(\alpha, p)$ is strictly increasing in p .

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). Now let, for fixed λ ,

$$\hat{f}(\alpha) := \hat{K}(\alpha) - \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{p}{\lambda}}}, \quad (16)$$

where $\hat{K}(\alpha) := 1/2 + 1/2\sqrt{1 - 4\alpha\lambda/p}$. Differentiating with respect to α , holding λ fixed, gives

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

At $\lambda = \bar{\lambda}$, $q^*(\bar{\lambda}) = p$. By the proof of Lemma 9,

$$\frac{d}{dq} f(\bar{\lambda}, p) = \alpha \left(1 - \frac{e^{\frac{p}{\bar{\lambda}}}}{(K(\bar{\lambda}) + (1 - K(\bar{\lambda}))e^{\frac{p}{\bar{\lambda}}})^2} \right) > 0.$$

Since $\frac{d}{d\alpha} \hat{K}(\alpha) < 0$, this implies that $\frac{d}{d\alpha} \hat{f}(\alpha) < 0$ at the threshold. Furthermore, $f(\lambda, p) = p(\alpha - \hat{f}(\alpha))$, where $\hat{f}$ also depends on the fixed λ . Thus, holding λ fixed,

$$\frac{d}{d\alpha} f(\lambda, p) = p \left(1 - \frac{d}{d\alpha} \hat{f}(\alpha) \right) > 0$$

at $\lambda = \bar{\lambda}$. By the proof of Lemma 10, the derivative of $f(\lambda, p)$ with respect to λ , holding α fixed, is also strictly positive at the threshold. Applying the implicit function theorem to $f(\bar{\lambda}, p) = 0$, with p fixed, therefore gives

$$\frac{d\bar{\lambda}}{d\alpha} = -\frac{\frac{d}{d\alpha} f(\bar{\lambda}, p)}{\frac{d}{d\lambda} f(\bar{\lambda}, p)} < 0. \quad \square$$

LEMMA 13. For every $\lambda \in (0, \bar{\lambda})$, let $\pi(\lambda)$ be the solution $\pi \in (0, \pi_\lambda)$ induced by $q^*(\lambda) > 0$. Then, $\pi(\lambda)$ is continuously differentiable, strictly 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 induced $\pi \in (0, 1)$. We now show that, for every $\pi \in (0, 1)$, there is at most one $\lambda \in (0, \bar{\lambda})$ inducing that value. This establishes a one-to-one relationship between π and λ that, together with continuity and the fact that $\lim_{\lambda \rightarrow 0} \pi(\lambda) = 0$, shows that $\pi(\lambda)$ is strictly increasing.

Fix $\bar{\pi} \in (0, 1)$. By rearranging equation (11), we obtain

$$z(\lambda) := q^*(\lambda) - p - \lambda \cdot \ln\left(\frac{1 - \bar{\pi}}{\bar{\pi}}\right) - \lambda \cdot \ln\left(\frac{K(\lambda)}{1 - K(\lambda)}\right) = 0. \quad (17)$$

We show that equation (17) has at most one solution $\lambda \in (0, \bar{\lambda})$. By Lemma 10, $\frac{d}{d\lambda} q^*(\lambda) < 0$. Furthermore,

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

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

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

where the last equality follows by definition of $K(\lambda)$. At any solution of $z(\lambda) = 0$, equation (17) and the preceding calculations imply

$$\frac{d}{d\lambda} z(\lambda) = \frac{d}{d\lambda} q^*(\lambda) - \frac{q^*(\lambda) - p}{\lambda} - \frac{p}{\alpha} \frac{d}{d\lambda} K(\lambda).$$

By the derivative identity established in the proof of Lemma 10 and the implicit function theorem,

$$\frac{d}{d\lambda} q^*(\lambda) - \frac{p}{\alpha} \frac{d}{d\lambda} K(\lambda) = - \frac{\alpha q^*(\lambda) e^{\frac{q^*(\lambda)}{\lambda}}}{\lambda (K(\lambda) + (1 - K(\lambda)) e^{\frac{q^*(\lambda)}{\lambda}})^2 \frac{d}{dq} f(\lambda, q^*(\lambda))} < 0.$$

Since $q^*(\lambda) > p$, it follows that $\frac{d}{d\lambda} z(\lambda) < 0$ at every solution of $z(\lambda) = 0$. Thus, at every solution, the function crosses from positive to negative, implying that there is at most one such solution. $\square$

COROLLARY 5. *For every $\lambda \in (0, \bar{\lambda})$, let $\pi(\lambda) \in (0, K(\bar{\lambda}))$ be the unique value identified in Lemma 9. 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 $\pi(\lambda)$ through equation (2) satisfies $\beta \in \arg \max U(\sigma, \beta)$.*

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 a 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 $\{\text{buy, not buy}\}$ 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

$$C(\sigma_0) := \frac{(1 - \sigma_0) \cdot (1 - e^{\frac{q^* - p}{\lambda}}) - \sigma_0 \cdot (e^{\frac{-p}{\lambda}} - 1)}{(e^{\frac{q^* - p}{\lambda}} - 1) \cdot (e^{\frac{-p}{\lambda}} - 1)} \in (0, 1). \quad (18)$$

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

$$C(0) = \frac{1}{1 - e^{\frac{-p}{\lambda}}} > 1$$

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

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 strictly 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 (18).

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 $\pi(\lambda) \in (0, K(\bar{\lambda}))$ be the consumer's equilibrium strategy. By inverting equation (18), 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 β', β'' such that $\mu(\{q : \beta'(q) \neq \beta''(q)\}) > 0$ and β' is non-constant μ -almost surely, we have that, for every $\gamma \in (0, 1)$,

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

Let β^0 be 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. ■

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

PROOF OF THEOREM 2. Fix $(\alpha, \lambda) \in (0, 1) \times \mathbb{R}_+$, and let $p^e := 4\alpha\lambda$ and $\tilde{K}(p) := \frac{1}{2} + \frac{1}{2}\sqrt{1 - \frac{4\alpha\lambda}{p}}$. 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 $p \leq \bar{p}$, $\lambda \geq \bar{\lambda}$. Therefore, the unique Shannon-Nash equilibrium with positive quality provision exists if and only if $p > \bar{p}$. 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

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

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*

$$p \cdot \frac{\tilde{K}(p)}{\tilde{K}(p) + (1 - \tilde{K}(p)) \cdot e^{\frac{q^*}{\lambda}}} + \alpha \cdot q^* - p \cdot \tilde{K}(p) = 0. \quad (19)$$

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

Proof. A unique and positive solution of equation (19) exists by Lemma 9 since $\lambda < \bar{\lambda}$ if and only if $p > \bar{p}$. Similarly, $q^*(p) > p$ is continuously differentiable by the argument in Lemma 10. At $p = \bar{p}$, the unique positive solution of equation (19) is $q = \bar{p}$, by the definition of $\bar{p}$. The same implicit function theorem argument gives continuity at this point, implying that $q^*(p) \rightarrow \bar{p}$ as $p \rightarrow \bar{p}$.

We now establish the limit as $p \rightarrow \infty$. Since $q^*(p) > p$, we have

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

Using $\tilde{K}(p)(1 - \tilde{K}(p)) = \alpha\lambda/p$, we obtain

$$\frac{(1 - \tilde{K}(p))e^{p/\lambda}}{p} = \frac{\alpha\lambda e^{p/\lambda}}{p^2 \tilde{K}(p)} \rightarrow \infty.$$

Therefore, $\tilde{h}(p) \rightarrow 0$, and hence $p\beta(0) \rightarrow 0$. Equation (19) now gives

$$\alpha q^*(p) - p = p(\tilde{K}(p) - 1) - p\beta(0) = -\frac{\alpha\lambda}{\tilde{K}(p)} - p\beta(0) \rightarrow -\alpha\lambda,$$

since $\tilde{K}(p) \rightarrow 1$. In particular, $\beta(0) \rightarrow 0$ as well. □

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 $p \in (\bar{p}, \infty)$. □

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

$$\Pi(p) = p \cdot \tilde{K}(p) - \alpha \cdot q^*(p), \quad (20)$$

where $\tilde{K}(p) = \beta(q^*(p))$, and $q^*(p) > p$ is defined by equation (19). Furthermore, by Lemma 16,

$$\frac{d}{dp} \Pi(p) = \tilde{K}(p) + p \frac{d}{dp} \tilde{K}(p) - \alpha \frac{d}{dp} q^*(p). \quad (21)$$

We proceed by finding an expression for the derivative of $q^*(p)$ by differentiating both sides of equation (19) 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 (21), we obtain

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

Let

$$D(p) := (\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}})^2 - e^{\frac{q^*(p)}{\lambda}}.$$

By the proof of Lemma 9, the derivative of the left-hand side of equation (19) with respect to q , evaluated at $q^*(p)$, satisfies

$$\alpha \cdot \frac{D(p)}{(\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p)}{\lambda}})^2} > 0.$$

Therefore, $D(p) > 0$. By rearranging the expression above, we obtain

$$\frac{d}{dp} \Pi(p) = \frac{\tilde{K}(p)^2 (1 - e^{\frac{q^*(p)}{\lambda}})}{D(p)} < 0,$$

where the inequality follows since $q^*(p) > 0$ and $D(p) > 0$, 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, by Lemma 16, $q^*(p) - p \rightarrow 0$, and $0 \leq U(\sigma, \beta) \leq q^*(p) - p$. 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 (19) 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

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

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. Since $1 - \tilde{K}(p) > 0$, the last inequality is equivalent to $e^{\frac{q^*(p) - p}{\lambda}} > 1$, which holds since $q^*(p) > p$, 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

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

where the first inequality holds by $q^*(p) > p$, the second inequality follows by taking limits in $\sigma(q^*(p))\tilde{K}(p) < \pi(p)$ for $p > \hat{p}$, and the last equality follows since, by Lemma 16, $(q^*(p) - p)/p \rightarrow 1/\alpha - 1 > 0$. Therefore, there exist constants $k_1, k_2 > 0$ such that $k_1 p \leq q^*(p) - p \leq k_2 p$ for all sufficiently large p . Since $1 - \tilde{K}(p) = \alpha\lambda/(p\tilde{K}(p)) \geq \alpha\lambda/p$, we have

$$0 \leq \frac{\tilde{K}(p)(q^*(p) - p)}{\tilde{K}(p) + (1 - \tilde{K}(p))e^{\frac{q^*(p) - p}{\lambda}}} \leq \frac{k_2 p^2}{\alpha\lambda} e^{-\frac{k_1 p}{\lambda}} \rightarrow 0.$$

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 $\alpha^e := \frac{p}{4\lambda}$. We first show that there exists a unique $\alpha \in (0, \alpha^e)$, which we term $\bar{\alpha}$, such that

$$\hat{g}(\alpha) := \hat{K}(\alpha) - \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{p}{\lambda}}} - \alpha = 0, \quad (22)$$

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 $\hat{g}(\bar{\alpha}) = 0$, 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 $x := p/\lambda > 0$. By differentiating, we obtain

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

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 (22).

We now show that $\bar{\alpha}$ is unique. Using $\hat{K}(\alpha)(1 - \hat{K}(\alpha)) = \frac{\alpha\lambda}{p}$, we can rewrite

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

Since $\frac{d}{d\alpha} \hat{K}(\alpha) < 0$ and $e^{\frac{p}{\lambda}} > 1$, the denominator in parentheses is strictly increasing in α . Thus,

the expression in parentheses is strictly decreasing. Since $\alpha > 0$, equation $\hat{g}(\alpha) = 0$ has at most one solution in $(0, \alpha^e)$. Therefore, $\bar{\alpha}$ is unique.

LEMMA 20. *In the unique Shannon-Nash equilibrium with positive quality provision, $\pi(\alpha)$ is strictly 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 $\mathcal{P}_\alpha$ 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

$$p\hat{K}(\alpha) - \alpha \left(p + \lambda \ln \left(\frac{\hat{K}(\alpha)}{1 - \hat{K}(\alpha)} \frac{1 - \pi(\alpha)}{\pi(\alpha)} \right) \right) = p \frac{\pi(\alpha) e^{\frac{-p}{\lambda}}}{\pi(\alpha) e^{\frac{-p}{\lambda}} + 1 - \pi(\alpha)} \quad (23)$$

Fix $\hat{\pi} \in \mathcal{P}_\alpha$. We show that there exists a unique α satisfying equation (23). Notice that in equilibrium $\hat{K}(\alpha) > \hat{\pi}$, and therefore we restrict attention 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

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

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

By the same implicit function theorem argument as in Lemma 10, $q^*(\alpha)$ is continuously differentiable, and equation (11) therefore implies that $\pi(\alpha)$ is continuous. The previous argument establishes a one-to-one relationship between π and α , that, combined with continuity and the fact that $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$ shows $\pi(\alpha)$ is strictly increasing.

To see that $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$, proceed by contradiction. Assume there exist $\varepsilon > 0$ and a sequence $\alpha_n \rightarrow 0$ such that $\pi(\alpha_n) \geq \varepsilon$. Since $\pi(\alpha_n) < \pi_\lambda < 1$, the sequence $\ln(\frac{1 - \pi(\alpha_n)}{\pi(\alpha_n)})$ is bounded. Furthermore, $\hat{K}(\alpha_n) \rightarrow 1$ and $\alpha_n \ln(\frac{\hat{K}(\alpha_n)}{1 - \hat{K}(\alpha_n)}) \rightarrow 0$. Therefore, the left hand side of equation (23) converges to p . However, $\pi(\alpha_n) < \pi_\lambda$ implies that the right hand side is less than $\frac{p}{2}$, a contradiction. Hence, $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$. $\square$

LEMMA 21. *The producer profits are positive, increasing in α , and converge 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 $\beta(0) > 0$; $\Pi(\alpha)$ is strictly increasing since it is strictly increasing in $\pi(\alpha)$, which is strictly increasing by Lemma 20, and $\lim_{\alpha \rightarrow 0} \Pi(\alpha) = 0$ since $\lim_{\alpha \rightarrow 0} \pi(\alpha) = 0$ by Lemma 20. 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

$$\alpha \cdot q^* = p \cdot \hat{K}(\alpha) - p \cdot \frac{\hat{K}(\alpha)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha)) \cdot e^{\frac{q^*}{\lambda}}}. \quad (24)$$

The solution of equation (24) 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 (24) 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

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

Since $q^*(\alpha) > p$, we have

$$\hat{E}(\alpha) = \frac{e^{\frac{q^*(\alpha)}{\lambda}} - e^{\frac{q^*(\alpha)-p}{\lambda}}}{e^{\frac{q^*(\alpha)}{\lambda}} - 1} < 1.$$

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 $\alpha \in (0, \bar{\alpha})$. The consumer utility gains from trade when quality equals $q^*(\alpha)$, which we

label $\hat{U}(q^*(\alpha))$, satisfy

$$0 \leq \hat{U}(q^*(\alpha)) = \sigma(q^*(\alpha))\hat{K}(\alpha)(q^*(\alpha) - p) < \frac{\hat{K}(\alpha)(q^*(\alpha) - p)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha) - p}{\lambda}}}.$$

Since $\alpha q^*(\alpha) \rightarrow p$, there exist constants $k_1, k_2 > 0$ such that $\frac{k_1}{\alpha} \leq q^*(\alpha) - p \leq \frac{k_2}{\alpha}$ for sufficiently small α . Using $\hat{K}(\alpha)(1 - \hat{K}(\alpha)) = \alpha\lambda/p$, we obtain

$$0 \leq \frac{\hat{K}(\alpha)(q^*(\alpha) - p)}{\hat{K}(\alpha) + (1 - \hat{K}(\alpha))e^{\frac{q^*(\alpha) - p}{\lambda}}} \leq \frac{pk_2}{\lambda\alpha^2}e^{-\frac{k_1}{\alpha\lambda}} \rightarrow 0.$$

This concludes the proof since $0 \leq U(\sigma, \beta) \leq \hat{U}(q^*(\alpha))$. □

This concludes the proof of Proposition 5. ■

D Extensions and robustness

This section presents formal arguments establishing the results discussed in Section 4 and Supplemental Appendix A.

Average quality. We show Proposition 6 below.

PROOF OF PROPOSITION 6. Producer indifference and the definition of π^* imply

$$\alpha q^* = p(\beta(q^*) - \beta(0)), \quad \pi^* = \beta(0) + \sigma(q^*)(\beta(q^*) - \beta(0)).$$

Consequently,

$$\mathbb{E}_\sigma[q] = \sigma(q^*)q^* = \frac{p}{\alpha}(\pi^* - \beta(0)). \tag{25}$$

The producer's first-order condition gives

$$\frac{p}{\lambda}\beta(q^*)(1 - \beta(q^*)) = \alpha.$$

By the argument in the proof of Proposition 9, applied to $c(q) = \alpha q$, $q_l = 0$, and $q_h = q^*$, we have $\beta(q^*) + \beta(0) < 1$. Consequently,

$$\beta(q^*)(1 - \beta(q^*)) - \beta(0)(1 - \beta(0)) > 0.$$

For the attention cost, fix (α, p) . Differentiating the logit formula along the equilibrium,

implies

$$\begin{aligned}\frac{d}{d\lambda}\beta(0) &= \beta(0)(1-\beta(0))\left(\frac{d}{d\lambda}\ln\left(\frac{\pi^*}{1-\pi^*}\right) + \frac{p}{\lambda^2}\right), \\ \frac{d}{d\lambda}\beta(q^*) &= \beta(q^*)(1-\beta(q^*))\left(\frac{d}{d\lambda}\ln\left(\frac{\pi^*}{1-\pi^*}\right) + \frac{1}{\lambda}\frac{d}{d\lambda}q^* - \frac{q^*-p}{\lambda^2}\right).\end{aligned}$$

Then, we differentiate producer indifference and use the first-order condition to simplify. We obtain

$$\begin{aligned}\frac{d}{d\lambda}\ln\left(\frac{\pi^*}{1-\pi^*}\right) &= \frac{\beta(q^*)(1-\beta(q^*))(q^*-p) + \beta(0)(1-\beta(0))p}{\lambda^2(\beta(q^*)(1-\beta(q^*)) - \beta(0)(1-\beta(0)))}, \\ \frac{d}{d\lambda}\beta(0) &= \frac{\beta(q^*)(1-\beta(q^*))\beta(0)(1-\beta(0))q^*}{\lambda^2(\beta(q^*)(1-\beta(q^*)) - \beta(0)(1-\beta(0)))}.\end{aligned}$$

Since

$$\frac{d}{d\lambda}\pi^* = \pi^*(1-\pi^*)\frac{d}{d\lambda}\ln\left(\frac{\pi^*}{1-\pi^*}\right),$$

it follows that

$$\begin{aligned}\frac{d}{d\lambda}(\pi^* - \beta(0)) &= \frac{\pi^*(1-\pi^*)\beta(q^*)(1-\beta(q^*))q^*}{\lambda^2(\beta(q^*)(1-\beta(q^*)) - \beta(0)(1-\beta(0)))} \\ &\quad \cdot \left(1 - \frac{\beta(0)(1-\beta(0))}{\pi^*(1-\pi^*)} - \frac{p}{q^*}\left(1 - \frac{\beta(0)(1-\beta(0))}{\beta(q^*)(1-\beta(q^*))}\right)\right).\end{aligned}\tag{26}$$

The term in parentheses is strictly positive. To see this, the logit formula implies

$$\ln\left(\frac{\pi^*}{1-\pi^*}\right) = \left(1 - \frac{p}{q^*}\right)\ln\left(\frac{\beta(0)}{1-\beta(0)}\right) + \frac{p}{q^*}\ln\left(\frac{\beta(q^*)}{1-\beta(q^*)}\right).$$

Since $q^* > p$, this is a strict convex combination of two distinct log odds. The function $u \mapsto 2 + e^u + e^{-u}$ is strictly convex, and, for every $v \in (0, 1)$,

$$2 + e^{\ln\left(\frac{v}{1-v}\right)} + e^{-\ln\left(\frac{v}{1-v}\right)} = \frac{1}{v(1-v)}.$$

Therefore,

$$\frac{1}{\pi^*(1-\pi^*)} < \left(1 - \frac{p}{q^*}\right)\frac{1}{\beta(0)(1-\beta(0))} + \frac{p}{q^*}\frac{1}{\beta(q^*)(1-\beta(q^*))}.$$

Multiplying by $\beta(0)(1-\beta(0))$ and rearranging gives the required inequality. Equations (25) and (26) imply that average quality is strictly increasing in λ .

Finally, Theorem 1 gives $\sigma(q^*) \rightarrow 0$ and $q^* \rightarrow p/\alpha$ as $\lambda \rightarrow 0$, while $\sigma(q^*) \rightarrow 1$ and $q^* \rightarrow p$ as $\lambda \rightarrow \bar{\lambda}$. Consequently, $\mathbb{E}_\sigma[q] \rightarrow 0$ as $\lambda \rightarrow 0$, and $\mathbb{E}_\sigma[q] \rightarrow p$ as $\lambda \rightarrow \bar{\lambda}$.

For the production cost, fix (λ, p) and vary α within the positive-quality equilibrium region.

Differentiating producer indifference gives

$$p\left(\frac{d}{d\alpha}\beta(q^*) - \frac{d}{d\alpha}\beta(0)\right) = q^* + \alpha\frac{d}{d\alpha}q^*.$$

The logit formula implies

$$\begin{aligned}\frac{d}{d\alpha}\beta(q^*) &= \beta(q^*)(1 - \beta(q^*))\left(\frac{d}{d\alpha}\ln\left(\frac{\pi^*}{1 - \pi^*}\right) + \frac{1}{\lambda}\frac{d}{d\alpha}q^*\right), \\ \frac{d}{d\alpha}\beta(0) &= \beta(0)(1 - \beta(0))\frac{d}{d\alpha}\ln\left(\frac{\pi^*}{1 - \pi^*}\right).\end{aligned}$$

Using the producer's first-order condition, we can simplify, and

$$\frac{d}{d\alpha}\ln\left(\frac{\pi^*}{1 - \pi^*}\right) = \frac{q^*}{p(\beta(q^*)(1 - \beta(q^*)) - \beta(0)(1 - \beta(0)))}.$$

Therefore,

$$\begin{aligned}\frac{d}{d\alpha}\mathbb{E}_\sigma[q] &= -\frac{\mathbb{E}_\sigma[q]}{\alpha} + \frac{p}{\alpha}(\pi^*(1 - \pi^*) - \beta(0)(1 - \beta(0)))\frac{d}{d\alpha}\ln\left(\frac{\pi^*}{1 - \pi^*}\right) \\ &= \frac{q^*}{\alpha}\left(\frac{\pi^*(1 - \pi^*) - \beta(0)(1 - \beta(0))}{\beta(q^*)(1 - \beta(q^*)) - \beta(0)(1 - \beta(0))} - \sigma(q^*)\right) \\ &= \frac{\mathbb{E}_\sigma[q]}{\alpha}\frac{\beta(q^*) - \pi^*}{1 - \beta(q^*) - \beta(0)} > 0.\end{aligned}$$

The last equality uses

$$\pi^*(1 - \pi^*) - \beta(0)(1 - \beta(0)) = (\pi^* - \beta(0))(1 - \pi^* - \beta(0)), \quad \sigma(q^*) = \frac{\pi^* - \beta(0)}{\beta(q^*) - \beta(0)}.$$

The strict inequality follows since $0 < \sigma(q^*) < 1$, so $\beta(0) < \pi^* < \beta(q^*)$.

We now establish the limits as production costs vary. As $\alpha \rightarrow \bar{\alpha}$, the equilibrium limits established in the production-cost comparative statics give $q^* \rightarrow p$ and $\sigma(q^*) \rightarrow 1$. Hence, $\mathbb{E}_\sigma[q] \rightarrow p$ as $\alpha \rightarrow \bar{\alpha}$. As $\alpha \rightarrow 0$, we have $\alpha q^* \rightarrow p$. Thus $q^* \geq \frac{p}{2\alpha}$ for sufficiently small α . The first-order condition also gives

$$1 - \beta(q^*) = \frac{\alpha\lambda}{p\beta(q^*)} \geq \frac{\alpha\lambda}{p}.$$

Consequently, the logit formula implies

$$0 < \pi^* = \frac{\beta(q^*)}{\beta(q^*) + (1 - \beta(q^*))e^{\frac{q^* - p}{\lambda}}} \leq \frac{p}{\alpha\lambda}e^{\frac{p - q^*}{\lambda}} \leq \frac{p}{\alpha\lambda}e^{\frac{p}{\lambda}}e^{\frac{-p}{2\alpha\lambda}}.$$

By equation (25),

$$0 < \mathbb{E}_\sigma[q] \leq \frac{p}{\alpha} \pi^* \leq \frac{p^2}{\alpha^2 \lambda} e^{\frac{p}{\lambda}} e^{\frac{-p}{2\alpha\lambda}} \rightarrow 0,$$

since exponential decay dominates the factor $1/\alpha^2$.

For the price, fix (α, λ) . Differentiating producer indifference and using the first-order condition implies

$$0 = \beta(q^*) - \beta(0) + p(\beta(q^*)(1 - \beta(q^*)) - \beta(0)(1 - \beta(0))) \left(\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) - \frac{1}{\lambda} \right).$$

Thus,

$$\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) = \frac{1}{\lambda} - \frac{1}{p(1 - \beta(q^*) - \beta(0))}.$$

Furthermore,

$$\begin{aligned} \frac{d}{dp} \pi^* &= \pi^*(1 - \pi^*) \frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right), \\ \frac{d}{dp} \beta(0) &= \beta(0)(1 - \beta(0)) \left(\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) - \frac{1}{\lambda} \right). \end{aligned}$$

Differentiating equation (25) and substituting these expressions yields

$$\begin{aligned} \frac{d}{dp} \mathbb{E}_\sigma[q] &= \frac{\pi^* - \beta(0)}{\alpha} + \frac{p\pi^*(1 - \pi^*)}{\alpha\lambda} - \frac{\pi^*(1 - \pi^*) - \beta(0)(1 - \beta(0))}{\alpha(1 - \beta(q^*) - \beta(0))} \\ &= \frac{p\pi^*(1 - \pi^*)}{\alpha\lambda} - \frac{(\pi^* - \beta(0))(\beta(q^*) - \pi^*)}{\alpha(1 - \beta(q^*) - \beta(0))}. \end{aligned}$$

As $p \rightarrow \bar{p}$, $q^* \rightarrow \bar{p}$ and $\sigma(q^*) \rightarrow 1$, so $\mathbb{E}_\sigma[q] \rightarrow \bar{p}$ and $\pi^* - \beta(q^*) \rightarrow 0$. As $p \rightarrow \bar{p}$, the purchase probabilities $\beta(q^*)$ and $\beta(0)$ converge to values strictly between zero and one. Since producer indifference and the first-order condition continue to hold at these limiting values, the argument in the proof of Proposition 9 also implies that their sum is strictly less than one. Therefore, the denominator $1 - \beta(q^*) - \beta(0)$ converges to a strictly positive value. Using $\frac{p}{\lambda} \beta(q^*)(1 - \beta(q^*)) = \alpha$ in the preceding expression, we obtain

$$\lim_{p \rightarrow \bar{p}} \mathbb{E}_\sigma[q] = \bar{p}, \quad \lim_{p \rightarrow \bar{p}} \frac{d}{dp} \mathbb{E}_\sigma[q] = 1.$$

In particular, average quality is strictly increasing for prices sufficiently close to $\bar{p}$.

To establish the limit at high prices, recall that $\beta(q^*) \rightarrow 1$, $p(1 - \beta(q^*)) \rightarrow \alpha\lambda$, and $\alpha q^* - p \rightarrow -\alpha\lambda$. For sufficiently large p , these limits imply

$$1 - \beta(q^*) \geq \frac{\alpha\lambda}{2p}, \quad q^* - p \geq \frac{p(1 - \alpha)}{\alpha} - 2\lambda.$$

Consequently, the equilibrium expression for π^* gives

$$0 < \pi^* = \frac{\beta(q^*)}{\beta(q^*) + (1 - \beta(q^*))e^{\frac{q^* - p}{\lambda}}} \leq \frac{2p}{\alpha\lambda} e^2 e^{\frac{-p(1-\alpha)}{\alpha\lambda}}.$$

Therefore,

$$0 < \mathbb{E}_\sigma[q] = \frac{p}{\alpha}(\pi^* - \beta(0)) \leq \frac{p}{\alpha}\pi^* \leq \frac{2p^2}{\alpha^2\lambda} e^2 e^{\frac{-p(1-\alpha)}{\alpha\lambda}} \rightarrow 0.$$

This limit, together with the strict increase near $\bar{p}$, rules out monotonicity.

Finally, average quality is continuous, increases near $\bar{p}$, and converges to zero as $p \rightarrow \infty$. Its maximum is therefore attained on a closed, bounded interval, with both endpoints yielding less than some interior price. By the extreme value theorem, average quality attains a global maximum at an interior price. ■

Convex-concave information. We start by showing that condition (iv) holds when the consumer observes signal $s = q + \epsilon$ where $\epsilon \sim H$ has a distribution unimodal at zero. 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 $0 < \hat{s} \leq p/\alpha$, 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 a 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 $s^* \leq 0$, 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 the function is either strictly convex over $[0, p/\alpha]$ or convex-concave. In the strictly convex case, $q^* = 0$. In the convex-concave case, by full support and strict MLRP, $1 - F(s^*|q)$ 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 $q^* > p$. ■

Gaussian information. We show Proposition 8 below.

PROOF OF PROPOSITION 8. We denote by Φ and ϕ the standard normal distribution function and density, respectively. We construct equilibria indexed by the precision τ , and then show that this construction implies the existence of two equilibria for every sufficiently small positive λ . We write q^* for the positive support quality and t for the signal cutoff above which the consumer buys. Throughout the proof, plain π denotes the mathematical constant, while π^* denotes the average trade probability.

If the consumer chooses precision $\tau > 0$ and buys if and only if the signal is above t , then $\beta(q) = \Phi((q - t)\sqrt{\tau})$. Producer indifference between 0 and q^* and the first-order condition at q^* can be written as

$$p \left(\Phi((q^* - t)\sqrt{\tau}) - \Phi(-t\sqrt{\tau}) \right) = \alpha q^*, \quad (27)$$

$$p\sqrt{\tau}\phi((q^* - t)\sqrt{\tau}) = \alpha. \quad (28)$$

We construct a solution with $t < q^* < 2t$. For every $\tau > (2\pi\alpha^2)/p^2$, substituting the normal density into equation (28) gives

$$\frac{p\sqrt{\tau}}{\sqrt{2\pi}} e^{-\frac{(q^* - t)^2\tau}{2}} = \alpha.$$

Taking logarithms and rearranging, we obtain

$$(q^* - t)^2\tau = \ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right) > 0.$$

Since $q^* > t$, by applying the positive square root, we obtain

$$q^*(\tau) = t(\tau) + \frac{1}{\sqrt{\tau}} \sqrt{\ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right)}. \quad (29)$$

Substituting into producer indifference and using the first-order condition, we obtain

$$\Phi\left(\sqrt{\ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right)}\right) - \Phi(-t\sqrt{\tau}) - \left(t\sqrt{\tau} + \sqrt{\ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right)}\right)\phi\left(\sqrt{\ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right)}\right) = 0. \quad (30)$$

For every $\tau > (2\pi\alpha^2)/p^2$, this equation has a unique solution $t(\tau)$ satisfying

$$t(\tau) > \frac{1}{\sqrt{\tau}} \sqrt{\ln\left(\frac{p^2\tau}{2\pi\alpha^2}\right)} := \underline{t}$$

To see this, notice that, at $t = \underline{t}$, the left hand side of (30) is strictly positive: for every $u > 0$,

$$2\Phi(u) - 1 - 2u\phi(u) > 0,$$

since this expression is zero at $u = 0$ and its derivative is $2u^2\phi(u) > 0$. Above this lower bound on t , differentiating the left hand side of equation (30) wrt t implies

$$\sqrt{\tau} \left(\phi(t\sqrt{\tau}) - \phi \left(\sqrt{\ln \left(\frac{p^2\tau}{2\pi\alpha^2} \right)} \right) \right) < 0.$$

The left hand side tends to minus infinity as $t \rightarrow \infty$, proving existence and uniqueness. In particular, equation (29) implies $0 < q^* - t < t$.

The implicit function theorem implies that $t(\tau)$ and $q^*(\tau)$ are continuously differentiable. Differentiating equation (30) with respect to τ implies

$$\frac{d}{d\tau} (t(\tau)\sqrt{\tau}) = \frac{q^*(\tau)}{2\sqrt{\tau} \left(1 - \frac{\phi(t(\tau)\sqrt{\tau})}{\phi((q^*(\tau) - t(\tau))\sqrt{\tau})} \right)} > 0.$$

Also, $(q^*(\tau) - t(\tau))\sqrt{\tau} = \sqrt{\ln \left(\frac{p^2\tau}{2\pi\alpha^2} \right)}$ is strictly increasing. Equation (27) therefore implies that $q^*(\tau)$ is strictly increasing.

As $\tau \rightarrow (2\pi\alpha^2)/p^2$, $(q^*(\tau) - t(\tau))\sqrt{\tau}$ approaches zero. Equation (30) implies that $t(\tau)\sqrt{\tau}$ remains bounded, since

$$q^*(\tau)\sqrt{\tau} \leq \frac{1}{\phi((q^*(\tau) - t(\tau))\sqrt{\tau})}.$$

Any limit point $\nu \geq 0$ of $t(\tau)\sqrt{\tau}$ must satisfy

$$\Phi(\nu) - \frac{1}{2} = \nu\phi(0),$$

which is possible only for $\nu = 0$, since $\phi(u) < \phi(0)$ for $u > 0$. Thus, $q^*(\tau) \rightarrow 0$. As $\tau \rightarrow \infty$, both $t(\tau)\sqrt{\tau}$ and $(q^*(\tau) - t(\tau))\sqrt{\tau}$ approach infinity, so equation (27) gives $q^*(\tau) \rightarrow p/\alpha$.

Consequently, there is a unique $\bar{\tau} > (2\pi\alpha^2)/p^2$ such that $q^*(\bar{\tau}) = p$. For every $\tau > \bar{\tau}$, we have $p < q^*(\tau) < p/\alpha$. Holding the consumer's strategy fixed, producer profits satisfy

$$\frac{d}{dq} \Pi(q, \beta) = p\sqrt{\tau}\phi((q - t)\sqrt{\tau}) - \alpha = \alpha \left(\frac{\phi((q - t)\sqrt{\tau})}{\phi((q^* - t)\sqrt{\tau})} - 1 \right).$$

This derivative is negative for $0 \leq q < 2t - q^*$, positive for $2t - q^* < q < q^*$, and negative for

$q > q^*$. Therefore, the only possible global maximizers are 0 and q^* . Producer indifference makes both optimal, so every randomization over these qualities is a best response. Profits equal $p\Phi(-t\sqrt{\tau}) > 0$.

Fix $\tau > \tilde{\tau}$. For notational convenience, we suppress the dependence of q^* and t on τ . Set the probability of providing q^* as

$$\sigma(q^*) = \frac{p\phi(t\sqrt{\tau})}{(q^* - p)\phi((q^* - t)\sqrt{\tau}) + p\phi(t\sqrt{\tau})} \in (0, 1), \quad \sigma(0) = 1 - \sigma(q^*). \quad (31)$$

Notice that this choice satisfies the Bayesian cutoff condition

$$\sigma(q^*)(q^* - p)\phi((q^* - t)\sqrt{\tau}) = (1 - \sigma(q^*))p\phi(t\sqrt{\tau}).$$

Since the normal likelihood ratio is strictly increasing in the signal, buying above t is optimal. Moreover,

$$\frac{\sigma(q^*)q^*}{p} = \frac{q^*\phi(t\sqrt{\tau})}{(q^* - p)\phi((q^* - t)\sqrt{\tau}) + p\phi(t\sqrt{\tau})} < 1,$$

because $q^* > p$ and $\phi(t\sqrt{\tau}) < \phi((q^* - t)\sqrt{\tau})$. The best choice without information is therefore never buying, which yields zero utility.

We next determine the cost parameter λ for which the constructed precision τ is optimal for the consumer. Hold q^* and $\sigma(q^*)$ fixed. At any alternative precision $u > 0$, the optimal signal cutoff is

$$\frac{q^*}{2} + \frac{1}{q^*u} \ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right).$$

The logarithm is strictly positive because $\sigma(q^*)q^* < p$. The consumer's expected utility before subtracting information costs, denoted by $G(u)$, is therefore

$$\begin{aligned} G(u) = & \sigma(q^*)(q^* - p)\Phi \left(\frac{q^*\sqrt{u}}{2} - \frac{1}{q^*\sqrt{u}} \ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right) \right) \\ & - (1 - \sigma(q^*))p\Phi \left(-\frac{q^*\sqrt{u}}{2} - \frac{1}{q^*\sqrt{u}} \ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right) \right). \end{aligned}$$

Set $G(0) = 0$, consistently with optimality without information. Differentiation implies

$$\frac{d}{du}G(u) = \frac{q^* \sqrt{\sigma(q^*)(q^* - p)(1 - \sigma(q^*))p}}{2\sqrt{2\pi u}} \exp \left(-\frac{(q^*)^2 u}{8} - \frac{1}{2(q^*)^2 u} \left(\ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right) \right)^2 \right) \quad (32)$$

The derivative of the logarithm of this expression has the same sign as

$$\frac{1}{2(q^*)^2} \left(\ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right) \right)^2 - \frac{u}{2} - \frac{(q^*)^2 u^2}{8}.$$

This expression is strictly decreasing in $u > 0$, positive near zero, and negative for large u . Hence, $\frac{d}{du} G(u)$ first strictly increases and then strictly decreases, and tends to zero at both endpoints.

At the constructed precision τ , the cutoff condition gives

$$\ln \left(\frac{(1 - \sigma(q^*))p}{\sigma(q^*)(q^* - p)} \right) = \frac{\tau}{2} (t^2 - (q^* - t)^2).$$

Consequently,

$$\left. \frac{d}{du} G(u) \right|_{u=\tau} = \frac{\sigma(q^*)(q^* - p)\phi((q^* - t)\sqrt{\tau})q^*}{2\sqrt{\tau}}, \quad \frac{\tau \left. \frac{d^2}{du^2} G(u) \right|_{u=\tau}}{\left. \frac{d}{du} G(u) \right|_{u=\tau}} = -\frac{1 + \tau t(q^* - t)}{2} < 0.$$

Choose the cost coefficient

$$\lambda(\tau) = \frac{\sigma(q^*)(q^* - p)\phi((q^* - t)\sqrt{\tau})q^*}{2\sqrt{\tau}} = \frac{\sigma(q^*)(q^* - p)}{2\tau} \left(\Phi((q^* - t)\sqrt{\tau}) - \Phi(-t\sqrt{\tau}) \right) > 0, \quad (33)$$

where the equality follows from producer indifference and the first-order condition. Thus, τ satisfies the first-order condition for precision and lies on the strictly decreasing part of $\frac{d}{du} G$. It is the unique positive local maximum of $G(u) - \lambda(\tau)u$. Since $G(u) \leq \sigma(q^*)(q^* - p)$, net utility tends to minus infinity as $u \rightarrow \infty$. The only possible global maximizers are therefore 0 and τ .

The constructed precision gives strictly positive utility. The cutoff condition and equation (33) imply

$$\begin{aligned} U &= \sigma(q^*)(q^* - p) \left(\frac{\Phi((q^* - t)\sqrt{\tau}) + \Phi(-t\sqrt{\tau})}{2} - \frac{\Phi(-t\sqrt{\tau})\phi((q^* - t)\sqrt{\tau})}{\phi(t\sqrt{\tau})} \right) \\ &= \sigma(q^*)(q^* - p)\phi((q^* - t)\sqrt{\tau}) \cdot \left(\frac{q^*\sqrt{\tau}}{2} - \frac{\Phi(-t\sqrt{\tau})}{\phi(t\sqrt{\tau})} \left(1 - \frac{\phi(t\sqrt{\tau})}{\phi((q^* - t)\sqrt{\tau})} \right) \right), \end{aligned} \quad (34)$$

where the second equality uses producer indifference and the first-order condition. The normal tail bound gives

$$\Phi(-t\sqrt{\tau}) = \int_{t\sqrt{\tau}}^{\infty} \phi(v) dv < \frac{1}{t\sqrt{\tau}} \int_{t\sqrt{\tau}}^{\infty} v\phi(v) dv = \frac{\phi(t\sqrt{\tau})}{t\sqrt{\tau}}.$$

Moreover, since $0 < q^* - t < t$,

$$1 - \frac{\phi(t\sqrt{\tau})}{\phi((q^* - t)\sqrt{\tau})} = 1 - e^{-\frac{\tau}{2}(t^2 - (q^* - t)^2)} < \frac{\tau}{2}(t^2 - (q^* - t)^2).$$

Therefore,

$$\frac{\Phi(-t\sqrt{\tau})}{\phi(t\sqrt{\tau})} \left(1 - \frac{\phi(t\sqrt{\tau})}{\phi((q^* - t)\sqrt{\tau})} \right) < \frac{\sqrt{\tau}(2t - q^*)q^*}{2t} < \frac{q^*\sqrt{\tau}}{2},$$

where the last inequality follows from $q^* > t$. Equation (34) implies $U > 0$. The consumer strictly prefers τ to no information, completing the verification of global optimality. Randomizing over experiments or purchasing rules cannot improve on this global optimum. We have constructed a binary-quality Nash equilibrium at $\lambda(\tau)$ for every $\tau > \tilde{\tau}$.

Now, let $\tau \rightarrow \infty$. We have $q^*(\tau) \rightarrow p/\alpha$ and, by equation (29),

$$q^*(\tau) - t(\tau) = \frac{1}{\sqrt{\tau}} \sqrt{\ln \left(\frac{p^2 \tau}{2\pi \alpha^2} \right)} \rightarrow 0.$$

Thus, $t(\tau) \rightarrow p/\alpha$, and

$$\frac{\phi(t(\tau)\sqrt{\tau})}{\phi((q^*(\tau) - t(\tau))\sqrt{\tau})} = e^{-\frac{\tau}{2}(t(\tau)^2 - (q^*(\tau) - t(\tau))^2)} \rightarrow 0.$$

Equation (31) gives $\sigma(q^*) \rightarrow 0$, while

$$\beta(0) = \Phi(-t\sqrt{\tau}) \rightarrow 0, \quad \beta(q^*) = \Phi((q^* - t)\sqrt{\tau}) \rightarrow 1.$$

Consequently,

$$\pi^* = \sigma(q^*)\beta(q^*) + (1 - \sigma(q^*))\beta(0) \rightarrow 0, \quad \Pi = p\Phi(-t\sqrt{\tau}) \rightarrow 0, \quad 0 < U \leq \sigma(q^*)(q^* - p) \rightarrow 0.$$

Equation (33) implies $\lambda(\tau) \rightarrow 0$. This gives the first limiting equilibrium in the statement.

Next, let $\tau \rightarrow \tilde{\tau}$ from above. By definition, $q^*(\tau) \rightarrow p$. Equation (31) gives $\sigma(q^*) \rightarrow 1$, and hence

$$\pi^* \rightarrow \tilde{\pi} := \Phi((p - t(\tilde{\tau}))\sqrt{\tilde{\tau}}).$$

Producer indifference at the limit implies

$$\Phi(-t(\tilde{\tau})\sqrt{\tilde{\tau}}) = \tilde{\pi} - \alpha > 0,$$

so $\tilde{\pi} \in (\alpha, 1)$. The defining equations for the limiting cutoff and precision are

$$\Phi\left(\sqrt{\ln\left(\frac{p^2\tilde{\tau}}{2\pi\alpha^2}\right)}\right) - \Phi\left(\sqrt{\ln\left(\frac{p^2\tilde{\tau}}{2\pi\alpha^2}\right)} - p\sqrt{\tilde{\tau}}\right) = \alpha,$$

$$t(\tilde{\tau})\sqrt{\tilde{\tau}} = p\sqrt{\tilde{\tau}} - \sqrt{\ln\left(\frac{p^2\tilde{\tau}}{2\pi\alpha^2}\right)}.$$

These equations depend on p and $\tilde{\tau}$ only through $p\sqrt{\tilde{\tau}}$. Uniqueness of the constructed solution therefore implies that $p\sqrt{\tilde{\tau}}$, and hence $\tilde{\pi}$, depend only on α , while $\tilde{\tau}$ depends only on (α, p) . Moreover,

$$\Pi \rightarrow p(\tilde{\pi} - \alpha) > 0, \quad 0 < U \leq \sigma(q^*)(q^* - p) \rightarrow 0.$$

Equation (33) gives $\lambda(\tau) \rightarrow 0$ because $q^* - p \rightarrow 0$. This is the second limiting equilibrium.

Finally, fix any $\tau_0 > \tilde{\tau}$ and set $\tilde{\lambda}_G := \lambda(\tau_0) > 0$. The function $\lambda(\tau)$ is continuous and strictly positive on $(\tilde{\tau}, \infty)$, and tends to zero at both endpoints. For every $\lambda \in (0, \tilde{\lambda}_G)$, the intermediate value theorem gives a solution of $\lambda(\tau) = \lambda$ in each of $(\tilde{\tau}, \tau_0)$ and (τ_0, ∞) . Select $\tau_2(\lambda)$ in the former and $\tau_1(\lambda)$ in the latter, and denote the corresponding equilibria by (σ_i, β_i) , their positive qualities by q_i , and their average trade probabilities by π_i , for $i \in \{1, 2\}$. The equilibria are distinct because $q^*(\tau)$ is strictly increasing.

As $\lambda \rightarrow 0$, any such selections satisfy $\tau_2(\lambda) \rightarrow \tilde{\tau}$ and $\tau_1(\lambda) \rightarrow \infty$. Otherwise, a subsequence would remain in a compact subset of $(\tilde{\tau}, \infty)$, on which continuity and strict positivity give a positive lower bound for $\lambda(\tau)$, a contradiction. The limits proved above therefore apply to the selected equilibria for every sufficiently small positive cost coefficient. This proves the proposition. $\blacksquare$

Convex production costs. We prove Propositions 9 and 10.

PROOF OF PROPOSITION 9. We first establish some properties of any binary-quality Shannon-Nash equilibrium with positive quality provision. By Proposition 1 and credibility, $\pi^* \in (0, 1)$. Moreover, the two equilibrium qualities satisfy $q_l < p < q_h$: otherwise, the consumer would optimally never buy or always buy without acquiring information, and the producer would deviate to zero quality. By equation (2), $0 < \beta(q_l) < \beta(q_h) < 1$ and

$$\frac{q_h - q_l}{\lambda} = \ln\left(\frac{\beta(q_h)}{1 - \beta(q_h)}\right) - \ln\left(\frac{\beta(q_l)}{1 - \beta(q_l)}\right) = \int_{\beta(q_l)}^{\beta(q_h)} \frac{1}{t(1-t)} dt. \quad (35)$$

The producer obtains positive profits, since providing quality zero yields $p\beta(0) > 0$. Therefore,

$q_h < \bar{q}$, and producer indifference and tangency at q_h imply

$$p(\beta(q_h) - \beta(q_l)) = c(q_h) - c(q_l), \quad (36)$$

$$\frac{p}{\lambda} \beta(q_h)(1 - \beta(q_h)) = \frac{d}{dq} c(q_h). \quad (37)$$

If $q_l > 0$, the corresponding first-order condition is $\frac{p}{\lambda} \beta(q_l)(1 - \beta(q_l)) = \frac{d}{dq} c(q_l)$.

We show that $\beta(q_h) + \beta(q_l) < 1$, which implies

$$\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)) = (\beta(q_h) - \beta(q_l))(1 - \beta(q_h) - \beta(q_l)) > 0.$$

By convexity of c , we have $c(q_h) - c(q_l) \leq \frac{d}{dq} c(q_h)(q_h - q_l)$, that together with equations (36) and (37) imply

$$\beta(q_h) - \beta(q_l) \leq \beta(q_h)(1 - \beta(q_h)) \frac{q_h - q_l}{\lambda}.$$

If $\beta(q_h) + \beta(q_l) \geq 1$, then $\beta(q_l) \geq 1 - \beta(q_h)$, and $t(1 - t) > \beta(q_h)(1 - \beta(q_h))$ for every $t \in (\beta(q_l), \beta(q_h))$. Equation (35) would then imply

$$\beta(q_h)(1 - \beta(q_h)) \frac{q_h - q_l}{\lambda} = \int_{\beta(q_l)}^{\beta(q_h)} \frac{\beta(q_h)(1 - \beta(q_h))}{t(1 - t)} dt < \beta(q_h) - \beta(q_l),$$

a contradiction. Hence, $\beta(q_h) + \beta(q_l) < 1$.

Differentiating equation (36), and using producer optimality, gives

$$\frac{d}{d\lambda} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) = \frac{\beta(q_h)(1 - \beta(q_h))(q_h - p) + \beta(q_l)(1 - \beta(q_l))(p - q_l)}{\lambda^2(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)))} > 0. \quad (38)$$

Indeed, the terms involving the derivatives of q_h and q_l cancel by the first-order conditions. If $q_l = 0$, then $\frac{d}{d\lambda} q_l = 0$, since q_l is differentiable and non-negative, so the same calculation applies. Equation (38) implies that π^* is strictly increasing in λ .

To show the attention effect, differentiate equation (37) and use equation (38). We obtain

$$\begin{aligned} & \left(\frac{p\beta(q_h)(1 - \beta(q_h))}{\lambda^2} (1 - 2\beta(q_h)) - \frac{d^2}{dq^2} c(q_h) \right) \frac{d}{d\lambda} q_h \\ &= \frac{p\beta(q_h)(1 - \beta(q_h))}{\lambda^2} \left(1 - \frac{(1 - 2\beta(q_h))\beta(q_l)(1 - \beta(q_l))(q_h - q_l)}{\lambda(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)))} \right). \end{aligned} \quad (39)$$

The right-hand side is positive. This is immediate if $\beta(q_h) \geq \frac{1}{2}$. If $\beta(q_h) < \frac{1}{2}$, then $t(1 - t) >$

$\beta(q_l)(1 - \beta(q_l))$ for $t \in (\beta(q_l), \beta(q_h))$, and equation (35) implies

$$\frac{\beta(q_l)(1 - \beta(q_l))(q_h - q_l)}{\lambda} < \beta(q_h) - \beta(q_l).$$

Since

$$\begin{aligned} (1 - 2\beta(q_h))(\beta(q_h) - \beta(q_l)) &< (\beta(q_h) - \beta(q_l))(1 - \beta(q_h) - \beta(q_l)) \\ &= \beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)), \end{aligned}$$

the right-hand side of equation (39) is again positive. The coefficient of $\frac{d}{d\lambda} q_h$ in equation (39) is the second derivative of the producer's objective at q_h , which is non-positive by optimality. Equation (39) implies that this coefficient cannot be zero. Therefore, $\frac{d}{d\lambda} q_h < 0$.

Consider any sequence of binary-quality equilibria with $\lambda \rightarrow 0$. Since $q_h \in (p, \bar{q})$ and c is twice continuously differentiable, equation (37) implies $\beta(q_h)(1 - \beta(q_h)) \rightarrow 0$. Using the same equation, the second-order condition at q_h can be written as

$$\frac{1}{\lambda} \frac{d}{dq} c(q_h)(1 - 2\beta(q_h)) - \frac{d^2}{dq^2} c(q_h) \leq 0.$$

Since $\frac{d}{dq} c(q_h) \geq \frac{d}{dq} c(p) > 0$ and $\frac{d^2}{dq^2} c$ is bounded on $[0, \bar{q}]$, this inequality rules out $\beta(q_h) \rightarrow 0$. Together with $\beta(q_h)(1 - \beta(q_h)) \rightarrow 0$, it implies $\beta(q_h) \rightarrow 1$. Since $\beta(q_h) + \beta(q_l) < 1$, we also have $\beta(q_l) \rightarrow 0$.

We show that $q_l \rightarrow 0$. Otherwise, there would exist an $\varepsilon > 0$ and a subsequence along which $q_l \geq \varepsilon$. Along this subsequence, q_l is interior and its first and second-order conditions imply

$$\frac{1}{\lambda} \frac{d}{dq} c(q_l)(1 - 2\beta(q_l)) - \frac{d^2}{dq^2} c(q_l) \leq 0.$$

But $\frac{d}{dq} c(q_l) \geq \frac{d}{dq} c(\varepsilon) > 0$ and $\beta(q_l) \rightarrow 0$, contradicting this inequality for sufficiently small λ . Thus, $q_l \rightarrow 0$. Together with the limits previously established, equation (36) then implies $c(q_h) \rightarrow p$, and therefore $q_h \rightarrow \bar{q}$.

By equation (2),

$$\pi^* = \frac{\beta(q_h)}{\beta(q_h) + (1 - \beta(q_h))e^{\frac{q_h - p}{\lambda}}}.$$

Since $\bar{q} > p$, there exists a $\delta > 0$ such that $q_h - p \geq \delta$ for sufficiently small λ . Furthermore, equation (37) implies

$$1 - \beta(q_h) = \frac{\lambda}{p} \frac{1}{\beta(q_h)} \frac{d}{dq} c(q_h) \geq \frac{\lambda}{p} \frac{d}{dq} c(p).$$

Consequently,

$$0 < \pi^* \leq \frac{p}{\lambda} \left(\frac{d}{dq} c(p) \right)^{-1} e^{\frac{-\delta}{\lambda}} \rightarrow 0.$$

Finally, $\pi^* = \sigma(q_l)\beta(q_l) + (1 - \sigma(q_l))\beta(q_h)$ implies

$$1 - \sigma(q_l) = \frac{\pi^* - \beta(q_l)}{\beta(q_h) - \beta(q_l)} \rightarrow 0,$$

concluding the proof. ■

PROOF OF PROPOSITION 10. Producer profits satisfy

$$\Pi = p\beta(q_l) - c(q_l) = p\beta(q_h) - c(q_h) > 0,$$

where positivity follows since $p\beta(0) > 0$. For the attention cost, equation (38) gives

$$\begin{aligned} \frac{d}{d\lambda} \Pi &= p\beta(q_l)(1 - \beta(q_l)) \left(\frac{d}{d\lambda} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) - \frac{q_l - p}{\lambda^2} \right) \\ &= \frac{p\beta(q_h)(1 - \beta(q_h))\beta(q_l)(1 - \beta(q_l))(q_h - q_l)}{\lambda^2(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)))} > 0. \end{aligned}$$

For the price, fix (c, λ) and differentiate equation (36). The terms involving changes in the support qualities again vanish by optimality, yielding

$$\beta(q_h) - \beta(q_l) + p(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l))) \cdot \left(\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) - \frac{1}{\lambda} \right) = 0.$$

Hence,

$$\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) = \frac{1}{\lambda} - \frac{\beta(q_h) - \beta(q_l)}{p(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l)))},$$

and

$$\begin{aligned} \frac{d}{dp} \Pi &= \beta(q_l) + p\beta(q_l)(1 - \beta(q_l)) \left(\frac{d}{dp} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) - \frac{1}{\lambda} \right) = \beta(q_l) - \frac{\beta(q_l)(1 - \beta(q_l))}{1 - \beta(q_h) - \beta(q_l)} \\ &= -\frac{\beta(q_l)\beta(q_h)}{1 - \beta(q_h) - \beta(q_l)} < 0. \end{aligned}$$

For the production costs, fix $(\tilde{c}, p, \lambda)$. Differentiating producer indifference gives

$$p(\beta(q_h)(1 - \beta(q_h)) - \beta(q_l)(1 - \beta(q_l))) \frac{d}{d\alpha} \ln \left(\frac{\pi^*}{1 - \pi^*} \right) = \tilde{c}(q_h) - \tilde{c}(q_l).$$

Consequently,

$$\begin{aligned}\frac{d}{d\alpha}\Pi &= p\beta(q_l)(1-\beta(q_l))\frac{d}{d\alpha}\ln\left(\frac{\pi^*}{1-\pi^*}\right) - \tilde{c}(q_l) = \frac{\beta(q_l)(1-\beta(q_l))(\tilde{c}(q_h) - \tilde{c}(q_l))}{\beta(q_h)(1-\beta(q_h)) - \beta(q_l)(1-\beta(q_l))} - \tilde{c}(q_l) \\ &= \frac{1}{\alpha}\left(\frac{p\beta(q_l)(1-\beta(q_l))}{1-\beta(q_h)-\beta(q_l)} - \alpha\tilde{c}(q_l)\right) = \frac{1}{\alpha}\left(\Pi + \frac{p\beta(q_l)\beta(q_h)}{1-\beta(q_h)-\beta(q_l)}\right) > 0,\end{aligned}$$

where the third equality uses $\alpha(\tilde{c}(q_h) - \tilde{c}(q_l)) = p(\beta(q_h) - \beta(q_l))$.

Consumer utility is positive by the same argument as in Proposition 3. For fixed p , Proposition 9 gives $\beta(q_l) \rightarrow 0$, $q_l \rightarrow 0$, and $\sigma(q_h) \rightarrow 0$ as $\lambda \rightarrow 0$. Therefore,

$$0 < \Pi = p\beta(q_l) - c(q_l) \leq p\beta(q_l) \rightarrow 0,$$

and, since $q_h \leq \bar{q}$,

$$0 < U(\sigma, \beta) \leq \sigma(q_h)(\bar{q} - p) \rightarrow 0.$$

This concludes the proof. ■

Recoverable production costs. We start by showing Proposition 11.

PROOF OF PROPOSITION 11. 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 trade—and rearranging equation (2), we obtain

$$q(\beta) = p + \lambda \cdot \ln\left(\frac{1-\pi}{\pi} \frac{\beta}{1-\beta}\right), \tag{40}$$

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

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

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 (40). By Proposition 1, when $|\text{supp } \sigma| = 1$, Shannon best responses satisfy $\pi \in \{0, 1\}$. ■

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.

Purification. We prove Proposition 12 below.

PROOF OF PROPOSITION 12. Fix (α, p, λ) as in the statement, and let $\pi^* := \mathbb{E}_{\sigma^*}[\beta^*]$. By the argument in the proof of Proposition 9, applied to linear production costs, we have $\beta^*(0) + \beta^*(q^*) < 1$. Consequently, $(\beta^*(q^*) - \beta^*(0))(1 - \beta^*(q^*) - \beta^*(0)) > 0$.

We first characterize producer optimality against the logit attention strategy of equation (2) for π close to π^* . For a type $\theta \in [-1, 1]$, profits have derivative

$$\frac{d}{dq} (p\beta(q) - (\alpha + \epsilon\theta)q) = \frac{p}{\lambda} \frac{\pi(1-\pi)e^{\frac{q-p}{\lambda}}}{\left(1 - \pi + \pi e^{\frac{q-p}{\lambda}}\right)^2} - (\alpha + \epsilon\theta).$$

Let $\hat{q}(\theta, \pi, \epsilon)$ denote the quality at which this derivative changes from positive to negative:

$$\hat{q}(\theta, \pi, \epsilon) = p + \lambda \ln \left(\frac{1 + \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}}}{1 - \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}}} \frac{1 - \pi}{\pi} \right). \quad (41)$$

At this quality,

$$\beta(\hat{q}(\theta, \pi, \epsilon)) = \frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}}.$$

Since $4\alpha\lambda < p$, this quality is continuously differentiable in (π, ϵ) near $(\pi^*, 0)$, and $\hat{q}(\theta, \pi^*, 0) = q^*$ for every θ . For (π, ϵ) sufficiently close to $(\pi^*, 0)$, the following hold for every $\theta \in [-1, 1]$:

$$\frac{\pi e^{-\frac{p}{\lambda}}}{1 - \pi + \pi e^{-\frac{p}{\lambda}}} < \frac{1}{2} - \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}}, \quad p < \hat{q}(\theta, \pi, \epsilon) < p/\alpha.$$

The first inequality follows from $\beta^*(0) < 1 - \beta^*(q^*)$, and the second from $p < q^* < p/\alpha$. As quality increases from zero, the profit derivative above is first negative, then positive, and then negative, changing sign when the purchase probability equals

$$\frac{1}{2} - \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}} \quad \text{and} \quad \frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}},$$

respectively. Since profits tend to minus infinity as $q \rightarrow \infty$, the only possible global maximizers are 0 and $\hat{q}(\theta, \pi, \epsilon)$.

For fixed π and $\epsilon > 0$, the profit difference between these qualities, 0 and $\hat{q}(\theta, \pi, \epsilon)$, is strictly decreasing in type. Indeed, the first-order condition at $\hat{q}(\theta, \pi, \epsilon)$ implies

$$\frac{d}{d\theta} \left[p \left(\frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon\theta)\lambda}{p}} - \frac{\pi e^{-\frac{p}{\lambda}}}{1 - \pi + \pi e^{-\frac{p}{\lambda}}} \right) - (\alpha + \epsilon\theta) \hat{q}(\theta, \pi, \epsilon) \right] = -\epsilon \hat{q}(\theta, \pi, \epsilon) < 0.$$

Thus, if one type is indifferent, all lower types strictly prefer the positive quality and all higher types strictly prefer zero.

We construct an equilibrium by solving for the average trade probability π and the threshold type $\hat{\theta}$. Consider the two equations

$$0 = p \left(\frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon \hat{\theta})\lambda}{p}} - \frac{\pi e^{-\frac{p}{\lambda}}}{1 - \pi + \pi e^{-\frac{p}{\lambda}}} \right) - (\alpha + \epsilon \hat{\theta}) \hat{q}(\hat{\theta}, \pi, \epsilon), \quad (42)$$

$$0 = \pi - \frac{\pi e^{-\frac{p}{\lambda}}}{1 - \pi + \pi e^{-\frac{p}{\lambda}}} - \frac{1}{2} \int_{-1}^{\hat{\theta}} \left(\frac{1}{2} + \frac{1}{2} \sqrt{1 - \frac{4(\alpha + \epsilon \theta)\lambda}{p}} - \frac{\pi e^{-\frac{p}{\lambda}}}{1 - \pi + \pi e^{-\frac{p}{\lambda}}} \right) d\theta. \quad (43)$$

The first makes the threshold type indifferent. The second equates π to the average purchase probability induced by the proposed quality choices. Both right-hand sides are continuously differentiable in an open neighborhood of $(\pi, \hat{\theta}, \epsilon) = (\pi^*, 2\sigma^*(q^*) - 1, 0)$, and both are zero at this point.

At $(\pi^*, 2\sigma^*(q^*) - 1, 0)$, the derivatives of the right-hand side of equation (42) with respect to π and $\hat{\theta}$ are, respectively,

$$\frac{p(\beta^*(q^*)(1 - \beta^*(q^*)) - \beta^*(0)(1 - \beta^*(0)))}{\pi^*(1 - \pi^*)} > 0, \quad 0.$$

The first expression follows by differentiating profits while holding quality fixed, using the producer's first-order condition at the positive quality. The derivative of the right-hand side of equation (43) with respect to $\hat{\theta}$ is

$$-\frac{\beta^*(q^*) - \beta^*(0)}{2} < 0.$$

Consequently, the determinant of the Jacobian of the two right-hand sides with respect to $(\pi, \hat{\theta})$ is

$$-\frac{p(\beta^*(q^*)(1 - \beta^*(q^*)) - \beta^*(0)(1 - \beta^*(0))) (\beta^*(q^*) - \beta^*(0))}{2\pi^*(1 - \pi^*)} < 0.$$

The implicit function theorem therefore gives continuously differentiable solutions π_ϵ and θ_ϵ for all sufficiently small ϵ , with $\pi_\epsilon \rightarrow \pi^*$, and $\theta_\epsilon \rightarrow 2\sigma^*(q^*) - 1$. Choose $\bar{\epsilon} \in (0, \min\{\alpha, 1 - \alpha\})$ sufficiently small that, for every $\epsilon \in (0, \bar{\epsilon})$, these solutions satisfy $\pi_\epsilon \in (0, 1)$ and $\theta_\epsilon \in (-1, 1)$, and the preceding best response characterization holds for every type.

Define

$$\beta_\epsilon(q) = \frac{\pi_\epsilon e^{\frac{q-p}{\lambda}}}{1 - \pi_\epsilon + \pi_\epsilon e^{\frac{q-p}{\lambda}}}, \quad q_\epsilon(\theta) = \begin{cases} \hat{q}(\theta, \pi_\epsilon, \epsilon), & \theta < \theta_\epsilon, \\ 0, & \theta \geq \theta_\epsilon. \end{cases}$$

Let $\check{\sigma}_\epsilon(\theta)$ assign probability one to $q_\epsilon(\theta)$, and let σ_ϵ denote the induced quality distribution.

Equation (42) and the strict decrease of the profit difference in type imply that these choices are global best responses for every type. The threshold type is indifferent and could instead choose the positive quality; this has no effect on the induced distribution because the type has probability zero.

We verify the optimality of the average trade probability π_ϵ . For $r \in [0, 1]$, the function

$$r \rightarrow \lambda \int \ln \left(1 - r + r e^{\frac{q-p}{\lambda}} \right) d\sigma_\epsilon(q),$$

is concave. Its derivative at $r = \pi_\epsilon$ is

$$\lambda \int \frac{e^{\frac{q-p}{\lambda}} - 1}{1 - \pi_\epsilon + \pi_\epsilon e^{\frac{q-p}{\lambda}}} d\sigma_\epsilon(q) = \frac{\lambda}{\pi_\epsilon(1 - \pi_\epsilon)} \left(\int \beta_\epsilon(q) d\sigma_\epsilon(q) - \pi_\epsilon \right) = 0,$$

where the last equality follows from equation (43). Thus π_ϵ maximizes this function.

To verify optimality over attention strategies, let γ be any alternative attention strategy with average trade probability $r \in (0, 1)$. Consumer utility satisfies

$$\begin{aligned} & \int \left[\gamma(q)(q-p) - \lambda \left(\gamma(q) \ln \left(\frac{\gamma(q)}{r} \right) + (1 - \gamma(q)) \ln \left(\frac{1 - \gamma(q)}{1 - r} \right) \right) \right] d\sigma_\epsilon(q) \\ & \leq \lambda \int \ln \left(1 - r + r e^{\frac{q-p}{\lambda}} \right) d\sigma_\epsilon(q) \leq \lambda \int \ln \left(1 - \pi_\epsilon + \pi_\epsilon e^{\frac{q-p}{\lambda}} \right) d\sigma_\epsilon(q). \end{aligned}$$

The first inequality bounds the utility of any alternative strategy γ by maximizing the expression in brackets separately at each quality. The maximizing trade probabilities follow the logit formula with parameter r . The second inequality follows because π_ϵ maximizes the resulting bound over $r \in [0, 1]$. Never buying and always buying correspond to the endpoints $r = 0$ and $r = 1$, and satisfy the same bound. Finally, β_ϵ follows the logit formula with parameter π_ϵ and has average trade probability π_ϵ , so it attains this bound. No alternative strategy therefore gives higher utility, and β_ϵ is a best response.

Furthermore, β_ϵ satisfies equation (2) at every quality and is therefore credible by Lemma 1. The induced quality distribution gives positive probability to zero and to qualities above p , and $\pi_\epsilon \in (0, 1)$. By Proposition 1, β_ϵ is a Shannon best response. Thus $(\check{\sigma}_\epsilon, \beta_\epsilon)$ is a BNE with the required properties.

As $\epsilon \rightarrow 0$, uniform continuity of the positive candidate quality near the original equilibrium, together with $\pi_\epsilon \rightarrow \pi^*$ and $\sup_{\theta \in [-1, 1]} |\epsilon \theta| \rightarrow 0$, gives

$$\sup_{\theta \in [-1, 1]} \left| p + \lambda \ln \left(\frac{1 + \sqrt{1 - \frac{4(\alpha + \epsilon \theta)\lambda}{p}}}{1 - \sqrt{1 - \frac{4(\alpha + \epsilon \theta)\lambda}{p}}} \frac{1 - \pi_\epsilon}{\pi_\epsilon} \right) - q^* \right| \rightarrow 0.$$

In particular, $\sup_{\theta < \theta_\epsilon} |q_\epsilon(\theta) - q^*| \rightarrow 0$. For every $\theta \neq 2\sigma^*(q^*) - 1$, the actual quality converges to q^* if $\theta < 2\sigma^*(q^*) - 1$, and to zero if $\theta > 2\sigma^*(q^*) - 1$. The fraction of types choosing positive quality satisfies $(1 + \theta_\epsilon)/2 \rightarrow \sigma^*(q^*)$. For every bounded continuous function f , dominated convergence therefore gives

$$\int f(q) d\sigma_\epsilon(q) = \frac{1}{2} \int_{-1}^1 f(q_\epsilon(\theta)) d\theta \rightarrow \sigma^*(q^*) f(q^*) + (1 - \sigma^*(q^*)) f(0).$$

Thus $\sigma_\epsilon \Rightarrow \sigma^*$. Moreover,

$$\sup_{q \geq 0} |\beta_\epsilon(q) - \beta^*(q)| \leq \frac{1}{4} \left| \ln \left(\frac{\pi_\epsilon}{1 - \pi_\epsilon} \right) - \ln \left(\frac{\pi^*}{1 - \pi^*} \right) \right| \rightarrow 0,$$

since the derivative of the logistic function is at most $\frac{1}{4}$.

All chosen qualities lie in $[0, p/\alpha]$ for sufficiently small ϵ . Together with uniform convergence of β_ϵ and continuity of β^* , the preceding convergence of quality choices allows us to apply dominated convergence to average quality, average trade probability, and expected producer profits.

To see why utility converges, notice that mutual information is

$$\begin{aligned} I(\sigma_\epsilon, \beta_\epsilon) = & -\pi_\epsilon \ln \pi_\epsilon - (1 - \pi_\epsilon) \ln(1 - \pi_\epsilon) + \frac{1}{2} \int_{-1}^1 [\beta_\epsilon(q_\epsilon(\theta)) \ln \beta_\epsilon(q_\epsilon(\theta)) \\ & + (1 - \beta_\epsilon(q_\epsilon(\theta))) \ln(1 - \beta_\epsilon(q_\epsilon(\theta)))] d\theta. \end{aligned}$$

The functions $v \ln v$ and $(1 - v) \ln(1 - v)$, extended continuously to the endpoints, are bounded on $[0, 1]$. Dominated convergence implies that this expression approaches $I(\sigma^*, \beta^*)$. Expected consumption surplus converges by the same argument, so consumer utility also converges. This proves the proposition. $\blacksquare$

Dropping credibility. We start by proving Corollary 1.

PROOF OF COROLLARY 1. For point (i), if $\mathbb{E}_\mu[\beta] \in (0, 1)$, the logit formula holds μ -almost surely and implies $I(\mu, \beta) > 0$. Robustness follows by the argument in Proposition 1. If $\mathbb{E}_\mu[\beta] \in \{0, 1\}$, mutual information is zero. A constant best response is robust, while a non-constant one is not as it incurs positive information costs. For point (ii), if β is a constant best response, then it 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 $|\text{supp } \mu| = 1$, then β is constant since $\lambda \cdot \epsilon > 0$. $\blacksquare$

We now show that, for every $\tilde{q} \in (p, p/\alpha)$ such that there exists $\pi \in (0, 1)$ with $\Pi(0, \beta) = \Pi(\tilde{q}, \beta)$,

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

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

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

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

As in the proof of Lemma 14, $C(\sigma_0)$ is continuous and decreasing in σ_0 . Furthermore, C is strictly greater than one 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 that $C(\sigma_0) = \pi$.

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). For necessity, producer optimality implies condition (i), while consumer optimality implies condition (ii) and equation (2) σ -almost surely. By Corollary 1, positive quality provision implies $|\text{supp } \sigma| > 1$ and $\pi^* \in (0, 1)$. Finally, condition (i) and equation (2) σ -almost surely imply that σ assigns probability one to qualities at which the logit curve intersects a linear isoprofit curve. Since the logit curve is strictly concave, strictly convex, or strictly convex and then strictly concave, by adapting the argument in Lemma 7, there are at most three intersections. Therefore, $|\text{supp } \sigma| \in \{2, 3\}$. ■