

Surplus Squeeze and Informational Hold-up*

Peter Achim[†] Willy Lefez[‡]

July 11, 2025

Abstract

We study a static bilateral trade setting with moral hazard, where a seller privately chooses quality and a buyer may pay to verify it. We show that buyer-side information acquisition can lead to informational hold-up through a mechanism we call surplus squeezing: precise verification enables the seller to extract all buyer surplus, deterring inspection and causing trade to unravel. When verification is noisy, uncertainty preserves buyer surplus and sustains trade. Our framework highlights how strategic responses to learning can distort investment incentives, offering a new perspective on the limits of information precision in mitigating moral hazard.

1 Introduction

Markets plagued by hidden actions often lead to inefficiencies: sellers cut corners, agents shirk, quality deteriorates, and buyers lose trust. In classical moral hazard settings, two broad remedies have been proposed. One relies on contractual commitment. If parties agree on payment schemes that are contingent on observable outcome, they can shape the incentives of the informed party even when the latter's actions aren't directly observable. The other approach relies on reputation: when the interaction is infinitely repeated, the informed party has the option to build and maintain reputation which can incentivize it to exert costly actions. Yet many important transactions happen only once, in environments where commitments are impossible and there is no future to reward or punish. How, then,

*We are extremely grateful for helpful comments from Kemal Özbek, Helmut Bester, Roland Strausz, and participants at the Humboldt microeconomics colloquium. Peter Achim gratefully acknowledges financial support from the Leverhulme Trust through Research Fellowship RF-2025-484\7. Willy Lefez gratefully acknowledges support from the DFG-funded Collaborative Research Center SFB/TRR 190 "Rationality and Competition" and from the European Research Council under grant agreement ERC-PRIVDIMA.

[†]Dept. of Economics, University of York. Email: peter.achim@york.ac.uk

[‡]Dept. of Economics, Humboldt University. Email: willy.lefez@hu-berlin.de

does real-world trade survive in high-stakes, one-shot exchanges—from purchasing a used car to underwriting a loan—when neither contracting nor repetition can guarantee honest effort?

A natural answer, familiar to anyone who has ever paid for an inspection or certification, is that the buyer can choose to learn more. Instead of relying on contracts or reputations, she pays an upfront cost—whether a surveyor’s fee, an audit charge, or a certification cost—to acquire information about the seller’s hidden action or product quality. This costly-information-acquisition channel is widespread in practice: buyers routinely bear inspection fees in markets for homes, cars, and machinery; employers monitor and verify employee performance; financial institutions commission third-party due diligence. Yet, despite its ubiquity, economic theory has scarcely explored how endogenous, buyer-financed verification interacts with a seller’s hidden-action incentives. Does allowing the buyer to pay for perfect or highly precise information always improve market outcomes, as classical intuition would suggest? Or might the very precision of information that buyers seek undermine their incentive to invest—and with it, the entire market?

In this paper we uncover a striking and counterintuitive answer: more precise information can sometimes destroy trade rather than support it. The culprit is what we call “surplus squeezing.” When a buyer pays to verify quality with almost perfect accuracy, she commits an irreversible cost before the seller has revealed his hidden action. Once that cost is sunk, the seller is free to choose the lowest level of quality that still clears the buyer’s post-verification acceptance rule. In effect, the seller uses the buyer’s investment as a commitment device: he adjusts quality just high enough to trigger purchase, capturing nearly all of the surplus the buyer hoped to earn from learning. Anticipating this outcome, the buyer refuses to pay for verification—no inspection, no information, no trade. Market breakdown is the inevitable result.

To make this argument precise, we construct a simple one-period bilateral-trade model. A monopolistic seller first chooses a costly quality level for a single unit of a good. The buyer values the good equally to its quality and must decide whether to pay a fixed verification fee in order to observe a signal of quality before deciding whether to transact at a posted price. We compare two polar cases of the verification technology. In the perfect-verification regime, the buyer’s inspection reveals the seller’s true quality exactly. In the noisy-verification regime, she receives only a signal correlated with quality—noisy enough that the seller cannot fully anticipate the buyer’s posterior belief from her signal realization.

Under perfect verification, intuition might lead one to expect that the buyer’s ability to detect any deviation from promised quality would discipline the seller. Indeed, if that inspection were costless, the seller would choose a quality level equal to the price, restoring full efficiency. But once the buyer must pay any strictly positive cost, perfect precision backfires: the seller knows that after the buyer inspects, any quality at least as large as the price guarantees a sale. Consequently, he supplies exactly the price-equivalent quality and pockets the difference between value and cost. Since this leaves the buyer with zero surplus net of her inspection fee, she will not pay to inspect, depriving the market of any

information discipline. The only equilibrium is one in which no inspection occurs, the seller provides the lowest-cost quality—zero—and trade collapses.

The perfect-verification result reveals the extreme version of surplus squeezing. But real-world checks are rarely free or infinitely precise. We therefore analyze a more nuanced, noisy-verification regime, where the buyer’s inspection yields a conditionally normal signal with finite precision. Noise reintroduces uncertainty: if the seller lowers quality to save on cost, there remains a positive probability that the buyer’s signal will fall below her purchase threshold, causing the sale to fail. The seller’s expected profit function thus becomes strictly quasi-concave in quality, with a unique interior maximizer that balances marginal production cost against marginal sales probability. In equilibrium, the seller randomizes between zero quality and this interior optimum, while the buyer uses a simple cutoff rule to decide whether to buy based on the observed signal. Crucially, the buyer’s inspection yields positive expected net surplus—enough to cover the verification fee—so she continues to inspect and trade can occur.

Yet this equilibrium only exists for an intermediate range of signal precisions. If the inspection technology is too poor (near-zero precision), the buyer gains almost no information, so the seller again has no incentive to provide any positive quality. If the technology is too precise (near-perfect), surplus squeezing kicks back in, reducing the buyer’s net gain to zero and deterring inspections. Only in the “Goldilocks zone” of moderate precision does costly verification sustain positive-quality trade. We derive closed-form thresholds for the verification cost and precision that delineate this zone.

Our findings speak to at least three strands of literature. First, they extend classical moral hazard and hidden-action models by endogenizing the buyer’s information acquisition choice. While much of contract theory takes the informational environment as exogenous, we show that who pays for information—and how precise that information is—can fundamentally alter incentives. Second, we connect to costly-state-verification frameworks (Townsend, 1979; Border and Sobel, 1987), but with a twist: the buyer chooses whether to verify and the seller responds by choosing hidden action, rather than a principal verifying an agent’s cash flows. Third, our equilibrium construction resembles inspection games (e.g. Avenhaus et al., 2002), but with a continuous action space for the seller. This continuity allows the seller to finely calibrate quality to the buyer’s acceptance rule, giving rise to surplus squeezing—a phenomenon absent in binary-action inspection models.

Beyond its theoretical contribution, our analysis carries practical implications. In markets from real estate to online platforms, buyers often face a choice between free but noisy reviews and costly but comprehensive inspections. Our results caution that pushing for ever more precision—say, mandating exhaustive, standardized testing—can paradoxically undermine incentives to provide quality if buyers must foot the bill. Policy interventions that subsidize verification or share its cost, or design mechanisms that commit sellers to leave a non-negligible surplus for buyers, may be necessary to avoid collapse.

Our description of the informational hold-up problem is closely related to the classic

| | Context | Direct Effect | Indirect Effect |
|-----------------------|---|--|---------------------------|
| Hold-up Problem | Seller makes specific investment, buyer makes take it or leave it offer | Buyer sets price to extract all surplus | Insufficient investment |
| Diamond Paradox | Buyer searches a specific seller, a seller posts a price | Seller sets price to extract surplus difference between the value of the good and outside option | Insufficient search |
| Informational Hold-up | Buyer verifies quality, seller produces quality level | Seller chooses least-cost quality the buyer is willing to purchase post verification | Insufficient verification |

Table 1: Summary of Direct and Indirect Effects in Bargaining Scenarios

hold-up problem, due to Klein et al. (1978) and Grossman and Hart (1986), but has some conceptual differences. In the classic setting, one party (typically the seller) must make a relationship-specific investment before trade, such as building a specialized input or tailoring production to a buyer’s needs. After the investment is sunk, the counterparty (typically the buyer) has the ability to make a take-it-or-leave-it offer, capturing the surplus generated by that investment. The direct effect is appropriation: the investor cannot recover the full return on her investment. The indirect effect is underinvestment: anticipating this ex post extraction, the seller scales back or avoids investing altogether, leading to inefficient outcomes. Our mechanism is structurally similar: the buyer makes a costly, relationship-specific investment not in physical capital, but in information, by paying to verify the seller’s hidden quality. Once that investment is sunk, the seller adjusts his behavior to appropriate the resulting informational surplus by producing just enough quality to pass the buyer’s acceptance threshold. The buyer, foreseeing this, opts not to inspect in the first place, and trade collapses. The difference lies in what is invested and who controls the surplus: in our case, the party who makes the investment (the buyer) also learns something, but this informational gain is nullified by the seller’s strategic response (see Table 1.)

The logic also parallels the Diamond search paradox, but again with a conceptual shift. In Diamond (1971), the buyer invests by paying a cost to search for sellers. Upon meeting one, the seller posts a price that extracts the surplus relative to the buyer’s outside option. The direct effect is that the seller captures the value created by the buyer’s search; the indirect effect is that the seller are able to extract all surplus in equilibrium. Our informational hold-up shares this structure, with verification replacing search. The buyer inspects to distinguish high from low quality, but once the seller understands the inspection standard, he tailors quality to barely meet it, again extracting the surplus that the buyer hoped to secure through information. But unlike the Diamond setting, where appropriation leaves trade intact but inefficient, informational hold-up can destroy trade altogether: precise verification invites complete surplus squeezing, deterring inspection

and unraveling the very mechanism that could discipline hidden action. This highlights a key novelty of our mechanism: while both classic hold-up and the Diamond paradox concern the division of surplus, our setting introduces a case where strategic appropriation eliminates the surplus itself, unless verification is kept imperfect enough to preserve a rent for the buyer.

In the remainder of the paper, we first set out the model and describe the two verification regimes. We then prove the impossibility of trade under costly perfect verification, before turning to the noisy-verification case, where we characterize the unique trading equilibrium and derive the conditions under which it exists. We conclude by discussing extensions—nonlinear quality costs, private verification fees, and endogenous choice of inspection precision—and outline directions for future work on dynamic information acquisition and mechanism design under endogenous verification. By highlighting how strategic responses to buyer learning can distort the very incentives that learning is meant to shore up, our paper offers a cautionary lesson: in the battle against moral hazard, information is not always power.

Literature Review Our paper extends the classic hold-up problem in incomplete-contract settings, where one party must make a relationship-specific investment but fears ex-post exploitation by the counter-party later on. This fear of opportunistic renegotiation leads to underinvestment and possible trade breakdown. Our observation that this logic extends to settings where the investment relates to information production, with small differences. the buyer must spend costly resources to learn about her surplus, but the surplus is controlled by the seller who will choose the quality to be as small as possible while still acceptable to the buyer, and thereby squeeze out all surplus for the buyer. The core issue in both settings is that the party which invests does not control the division of surplus.

There is a growing literature that studies the role of buyer learning in bilateral trade environments.¹ Roesler and Szentes (2017) study a static monopoly pricing environment and identify the buyer’s optimal information structure. Ravid et al. (2022) build on this by letting the buyer choose flexibly how much to learn while holding quality fixed; they demonstrate that when the cost per unit of information falls toward zero, equilibrium outcomes converge to the worst “no-learning” outcome under a trembling-hand refinement, thereby uncovering a strategic free-rider problem in buyer learning.

Ravid (2020) examines a bilateral trade game with one seller and one buyer in which the seller learns the good’s quality and then posts a take-it-or-leave-it price, and the buyer pays an entropy-based cost to receive information about price and quality; he finds that

¹Our model is also related to the literature on buyer learning through intermediaries (Armstrong and Zhou, 2011; Inderst and Ottaviani, 2012; Teh and Wright, 2022; Evans and Park, 2024; Achim et al., 2025). These papers examine how buyers rely on third parties—such as platforms, advisors, or experts—to acquire or interpret information about products. While conceptually related to our setting, this literature emphasizes delegated information acquisition, where the intermediary filters or distorts what the buyer learns.

when costs run high the buyer declines to learn and trade fails, while when costs fall trade occurs but the seller distorts prices to capture the resulting surplus. The key difference lies in what each model makes endogenous and how the buyer learns. Ravid treats the good’s quality and the buyer’s valuation as fixed by nature, and the buyer chooses a continuous “attention level” to learn about the seller’s price-quality offer at an entropy cost. In our paper the seller instead sets the buyer’s value by choosing quality under a cost, and the buyer faces a binary decision—to pay a fixed fee for perfect verification or remain completely uninformed—so that perfect information itself destroys the seller’s incentive to provide quality.

Fabbri (2025), in independent work, introduces an “attention hold-up” mechanism that is similar in spirit to ours: a buyer learns about seller-chosen quality, and in the limit—as the cost of information vanishes—trade breaks down. Fabbri models learning as a flexible choice over information structures with entropy-based costs. In contrast, we adopt a simpler verification framework, closer to classical costly state verification: the buyer either pays a fixed cost for a (possibly noisy) signal or remains uninformed. This approach is mathematically tractable and makes the core logic of informational hold-up fully transparent: when buyer learning is too precise, the seller uses it to squeeze the buyer’s surplus to save on cost at fixed price; when it’s too noisy, the buyer learns too little to justify verifying. Only at intermediate precision does verification preserve enough surplus to support trade. Our results further show that the flexibility in the learning process is not crucial. Even with a fixed-precision signal, costly learning alone suffices to trigger an informational hold-up, and not just in the limit, but robustly at interior parameter values.

While our model shares features with the costly state verification literature, it applies to a setting of hidden action rather than hidden state. Classical models of moral hazard, such as Holmström (1979), study static principal–agent problems in which incentives for effort are provided through output-contingent contracts. In this literature, the informational environment is typically taken as given, and it does not consider incentives for information acquisition explicitly. In dynamic environments, moral hazard is often addressed through monitoring, where the principal gradually learns about the agent’s action from noisy signals over time (e.g., Mailath and Morris (2002); Fudenberg and Levine (1994)). By contrast, verification refers to ex-post access to a signal about a persistent outcome or state that results from a hidden action, a distinction emphasized in the costly state verification literature Townsend (1979); Border and Sobel (1987). Our model follows this tradition: the buyer does not observe the seller’s chosen quality but can verify it by paying a fixed cost. Although the underlying friction arises from hidden action, the structure of learning is closer to verification than to monitoring. Related models that combine hidden action with the possibility of verification include Strausz (2005), who studies optimal certification under unverifiable effort, and Achim and Knoepfle (2024, 2025), who explore incentive provision when the principal can verify whether a hidden action was taken.

2 Model

Setup. We explore a market in which a single risk-neutral seller offers a good to a single risk-neutral buyer. We assume the price for the good $p > 0$ is exogenously given. Prior to meeting the buyer, the seller chooses the good's quality $q \in [0, \infty)$ which affects the buyer's value from owning the good. Producing quality q entails a linear cost

$$C(q) = cq$$

Before purchasing the good, the buyer has the option to verify quality at a cost $k > 0$. The verification process is assumed to be exogenous, and produces a signal $s \in \mathbb{R}$ drawn from a distribution $\nu(\cdot|q)$ that depends on true quality. The observed signal induces a posterior mean regarding quality, denoted by $\tilde{q}(s)$.

Timing.

1. The seller chooses a quality level q .
2. The buyer decides whether to verify quality ($a = 1$) or not ($a = 0$).
3. If the buyer verifies quality ($a = 1$), she observes a signal s .
4. The buyer decides whether to purchase the good ($b = 1$) or not ($b = 0$).

Payoffs. We assume buyer's value from owning the good is equal to its quality. The payoff for the buyer if the good's quality is q , is therefore equal to $b(q - p) - ak$. The profit for the seller's profit is given by $bp - cq$.

Strategies. A strategy for the seller is represented by a probability distribution F on \mathbb{R} . We denote expectations with respect to F by $\mathbb{E}_F[\cdot]$ and define $\mu_0 = \mathbb{E}_F[q]$ to be the prior mean over quality with respect to F . A strategy for the buyer consists of a verification decision, and for each possible observation following the verification stage (include the empty signal if the buyer chooses not to verify) a purchasing decision. We let $\alpha \in [0, 1]$ denote the probability that the buyer verifies quality. For each signal s , we denote by $\beta(s) \in [0, 1]$ the probability that the buyer purchases the good after observing s . The buyer's purchase decision if she does not verify quality is denoted by $\beta(\emptyset)$.

Expected payoffs. Suppose the seller's choice of quality is given by the distribution function F , and the buyer's strategy is (α, β) . The buyer's expected payoff is then given by

$$U(\alpha, \beta, F) = \alpha \left[\int_q \left(\int_s \beta(s)(q - p) d\nu(s | q) \right) dF(q) - k \right] + (1 - \alpha)\beta(\emptyset) \left(\int_q q dF(q) - p \right)$$

The seller's profit is derived from the revenue generated when the buyer purchases, net of the cost of providing the chosen level of quality. Specifically, whenever the buyer decides

to purchase, the seller receive the price p but incurs a cost cq to deliver the chosen quality q . Hence, the seller's expected profit is

$$V(\alpha, \beta, F) = \int_q \left[\left(\alpha \int_s \beta(s) d\nu(s | q) + (1 - \alpha)\beta(\emptyset) \right) p - cq \right] dF(q).$$

Equilibrium. We consider (weak) perfect Bayesian equilibrium (PBE) of this environment which we define formally as follows.

Definition 1 (Perfect Bayesian Equilibrium). *A Perfect Bayesian Equilibrium is a tuple $(F^*, (\alpha^*, \beta^*), \tilde{q}(s))$ with the following properties:*

1. *Seller's Optimality: Given the price p and the buyer's strategy (α^*, β^*) , the seller's strategy F^* maximizes her expected profit,*

$$F^* \in \arg \max_F V(\alpha^*, \beta^*, F),$$

2. *Buyer's Optimality: Given the seller's strategy F^* and posterior mean $\tilde{q}(s)$, the buyer's purchase and verification decisions are optimal:*

$$\alpha^* \in \arg \max_{\alpha} U(\alpha, \beta^*, F^*), \quad \beta^*(s) \in \arg \max_b b \cdot (\tilde{q}(s) - p)$$

3. *Consistency of Beliefs: The posterior mean $\tilde{q}(s)$ is consistent with Bayes' rule, whenever applicable.*

The central question which we aim to explore in this paper is whether the buyer's ability to verify quality can result in gain from trade being realized in equilibrium. To this end, it is useful to briefly recognize, if verification were perfect and costless ($k = 0$), meaning that the buyer can directly observe the seller's chosen quality, then in equilibrium, the seller would find it optimal to produce $q = p$. Any higher quality would increase the cost for the seller without an increasing in revenue. Any lower quality $q < p$ would be immediately detected and rejected by the buyer.

In contrast, with costly verification, it is easy to see that there is always an equilibrium where trade fails. If the buyer never verifies ($\alpha = 0$) and never purchases ($\beta(\emptyset) = 0$), then the seller optimally chooses $q = 0$. In this scenario, producing any positive quality would be strictly worse for the seller, since there won't be a sale. On the other hand, anticipating zero quality, the buyer optimally refuses to verify or to buy. The question

However, when verification is noisy or costly, the threat of learning may be insufficient to discipline the seller, and trade may still fail. Indeed, imperfect verification cannot sustain trade unless it is sufficiently precise or cheap. In Lemma 1, we show that there is a threshold precision τ (or a threshold cost k) below which no equilibrium with trade can exist. For example, suppose verification is mandatory and costless but imprecise. If the seller were to produce $q = p$ in a candidate equilibrium, then the buyer's posterior $\mu_F(s)$ would always equal p , regardless of the observed signal s . In that case, no information ever arrives, so a seller who secretly deviates to $q = 0$ could still generate signals that

lead the buyer to purchase, making the deviation profitable. Therefore, noisy verification must be sufficiently precise (or verification must be sufficiently cheap) to restore trade. In subsequent sections, we characterize exactly how large τ (or how small k) must be for a trading equilibrium to exist.

3 Perfect Verification

We first turn to the benchmark of perfect verification. Under perfect verification, the buyer observes with perfect accuracy the quality level q that the seller has chosen. Formally, this means that for all s in the support of $\nu(s|q)$ we have $\tilde{q}(s) = q$. In this setting, verification functions like the costly monitoring technologies familiar from repeated-game theory, and not like a Bayesian signal in the classic Blackwell sense: once the buyer verifies, her posterior belief places probability one on the seller's true quality, irrespective of any prior distribution or previous beliefs.

At first glance, this strong form of learning seems to grant the buyer maximal power to discipline the seller, since any deviation from a promised quality would be detected with certainty. Indeed, it is natural to think that more precise information should improve market outcomes. In standard adverse-selection or moral hazard settings, richer signals about hidden types or unobserved actions enable contracts or strategies to be more finely conditioned on realized outcomes, and thereby reduces inefficiencies compared with the coarse or noisy observations that typically prevail. By analogy, perfect verification allows the buyer to condition trade (and any penalties) directly on the observed q , and one would expect that this at least alleviates the incentive problems for the seller.

Our main result exposes a striking counterpoint: whenever verification carries even the slightest positive cost, no equilibrium with trade exists under perfect verification. Formally, under the assumption of perfect but costly verification (i.e., $k > 0$), there is a unique equilibrium outcome in which the seller produces zero quality and the buyer neither verifies nor buys. In other words, the very pledge of perfect information intensifies the familiar hold-up problem to the point that exchange unravels entirely. This result is stated more formally in the following theorem.

Theorem 1. *Under perfect verification, there is no trade in equilibrium.*

Why does trade collapse so completely? The heart of the problem lies in the interplay between the sunk cost of verification and the seller's opportunistic incentives. First, imagine an equilibrium in which the buyer never verifies. In such an equilibrium, the buyer's decision to purchase depends solely on her prior expectation of quality, q^e , relative to the price p : she will buy so long as $q^e \geq p$. But if the buyer does not verify, the seller can profitably deviate by producing zero quality, since this eliminates the cost without affecting revenue. Since verification never occurs, the buyer cannot detect or punish this deviation, and so the seller strictly prefers to satisfy no-quality at zero cost rather than deliver any positive quality. Anticipating this, the buyer cannot rationally choose to purchase absent verification, so no-verification equilibria with trade cannot persist.

It is natural to wonder whether trade may perhaps be sustained by mixed strategies, where seller mixes over different quality levels and the buyer verifies only occasionally. The answer is no. To see why, suppose the buyer and the seller were to randomize. Verification makes sense for the buyer precisely only if his expected surplus from trade, net of the cost, is nonnegative. Because verification is perfect, after paying k the buyer learns the exact realization of q and trades whenever $q \geq p$. Suppose the distribution over qualities chosen by the seller are represented by the CDF F . Then the buyer's expected gain from verifying is

$$\int_p^{\bar{q}} (q - p) dF(q) - k.$$

But note that whenever the buyer verifies, the seller *knows* that any quality $q \geq p$ will lead to trade. Among all such qualities $q > p$, the seller thus can always raise profit by lowering quality to $q' = (q + p)/2$, because this reduces her cost while still guaranteeing sale. In other words, once the buyer has sunk the cost of verification, the seller can “rip off” the buyer by delivering quality barely above p , capturing nearly all the surplus above p . In equilibrium, therefore, the support of the seller's random strategy cannot contain any quality level that exceeds price p , which in turn means that the buyer's gain from verification cannot be positive.

We thus find that despite the fact that perfect verification is the most informative possible signal and would restore full efficiency if costless, any strictly positive verification cost magnifies the hold-up problem rather than resolving it. The sink of the buyer's monitoring expenditure leaves her vulnerable: once the cost is sunk, the seller can guarantee trade by delivering only an infinitesimal improvement over the price, capturing nearly the entire surplus and leaving the buyer with nothing net of cost. Anticipating this, the buyer refrains from verifying, trade fails, and quality production collapses. This negative result underlines a crucial lesson: more precise information does not necessarily ameliorate incentive problems when its acquisition itself can be exploited opportunistically. Only in the knife-edge case of $k = 0$ does perfect verification restore the efficient outcome $q = p$; the slightest cost suffices to unravel all trade.

4 Noisy Verification

We now turn to the case of noisy verification, where the buyer does not learn the seller's chosen quality exactly, but instead observes a signal that is informative yet imprecise. Specifically, after the seller selects a quality level q , the buyer, if he chooses to verify, observes a signal $s = q + \epsilon$, where ϵ is normally distributed noise with mean zero and precision ρ . The higher the precision, the more tightly concentrated the signal is around the true quality. Denote by $\Phi(\cdot)$ the CDF of the standard normal distribution, and $\phi = \Phi'$ is the density. The CDF of the noise term ϵ is denoted by $H(\epsilon) = \Phi(\rho\epsilon)$ and its density is $h(\epsilon) = \rho\phi(\rho\epsilon)$.

Given a prior belief about how the seller chooses quality, the buyer updates her belief upon seeing a signal using Bayes' rule. Since the noise distribution is smooth and satisfies a monotone likelihood ratio property, the posterior distribution over quality after observing s will shift in a disciplined way: higher signals make higher quality more likely in a first-order stochastic dominance sense. As a result, the buyer's posterior expectation of quality will be a strictly increasing function of s . This means the buyer's optimal decision rule can be represented by a simple threshold: she will buy if the signal is high enough to convince her that the expected quality exceeds the price, and reject otherwise.

This learning environment introduces a crucial wedge between what the buyer sees and what the seller knows. Because the buyer's signal is noisy, the seller cannot perfectly anticipate what the buyer will infer about quality. This uncertainty plays a key role in limiting the seller's ability to precisely tailor quality in response to the buyer's inspection threshold. In contrast to the perfect verification case, where the seller could fully exploit the buyer's learning, the presence of noise makes such appropriation risky. Lowering quality too much increases the chance that the signal will fall short, causing the buyer to walk away. As we will see, this basic feature of noisy learning is what ultimately makes trade sustainable in equilibrium.

4.1 The need for random quality

Before discussing the equilibrium construction, it is crucial to understand when and why the buyer finds it worthwhile to pay the verification cost k . Intuitively, paying to verify is only valuable if the inspection can alter her purchase decision. If inspection produces no actionable information, the buyer would refuse to inspect and trade collapses under moral hazard.

Concretely, let the seller's strategy be a distribution F over quality levels $q \geq 0$, with prior mean

$$\tilde{q} = \int_0^\infty q dF(q).$$

The distribution F is a pure strategy for the seller if it places all mass at a single point \tilde{q} . In that case, no matter what noisy signal $s = q + \epsilon$ the buyer observes, her posterior collapses to \tilde{q} . The buyer learns nothing new: her post-verification expectation of quality remains exactly \tilde{q} , so her decision to trade—buy if $\tilde{q} \geq p$, reject otherwise—cannot be improved by inspection. In that scenario, the buyer's payoff with verification

$$\max\{q_s - p, 0\} = \max\{\tilde{q} - p, 0\}$$

matches her payoff without inspection, $\max\{\tilde{q} - p, 0\}$. Therefore the the *value of verification*, given by

$$y_F = \int_0^\infty \int_{-\infty}^\infty \max\{q_s - p, 0\} h(s - q) ds dF(q) - \max\{\tilde{q} - p, 0\}$$

vanishes. Since the buyer would pay $k > 0$ for zero gain, the buyer never verifies.

A similar dead end arises if F only supports qualities on one side of the price. Suppose that every possible q exceeds p . Then after verification the posterior mean q_s always exceeds p , so $\max\{q_s - p, 0\} = q_s - p$. But averaging q_s against the signal density simply reproduces the prior mean \tilde{q} , so again

$$y_F = \int_0^\infty \int_{-\infty}^\infty (q_s - p) h(s - q) ds dF(q) - (\tilde{q} - p) = \tilde{q} - p - (\tilde{q} - p) = 0.$$

There is no decision-relevant uncertainty, and verification remains futile. Likewise, if F places all weight below p , then $q_s < p$ for all s , the buyer never trades, and verification is again valueless.

The upshot is that verification can be valuable and thus trade possible only if the seller's quality distribution F genuinely straddles the buyer's reservation price. Put differently, there must be some chance of a "bad" quality that inspection might avoid and some chance of a "good" quality that inspection might reveal. Absent that two-sided uncertainty, the buyer's posterior never crosses the price threshold, and $\max\{q_s - p, 0\}$ is either identically zero or identical to its no-verification counterpart.

Formally, one shows by Jensen's inequality (together with the piecewise-linearity of $\max\{\cdot, 0\}$) that the "value of verification"

$$y_F = \int_0^\infty \int_{-\infty}^\infty \max\{q_s - p, 0\} h(s - q) ds dF(q) - \max\{\tilde{q} - p, 0\}$$

is strictly positive only if F mixes across points both below and above p . This observation underpins the following key lemma:

Lemma 1. *There is no equilibrium with trade in which the seller uses a pure strategy. In a trade equilibrium, the seller strategy supports qualities strictly below and above the price.*

In other words, Lemma 1 tells us that for costly verification to be part of an equilibrium, the seller must leave genuine uncertainty about quality. Having established this, we are now ready to characterize how the buyer's optimal threshold rule and the seller's profit-maximizing mix of quality levels together generate a nondegenerate trading equilibrium.

4.2 The purchasing decision of the buyer

Having established that the seller's strategy F established must be mixed and support quality realizations both below and above the price p , we now characterize the buyer's optimal response after verification. Specifically, we study the buyer's purchasing decision once she receives a signal

$$s = q + \epsilon,$$

where ϵ is the noise term drawn from a distribution with density $h(\epsilon)$. Equivalently, the likelihood of observing signal s when the true quality is q is given by $h(s - q)$, since

$\epsilon = s - q$. This kernel $h(s - q)$ captures the accuracy of the verification technology: a sharply peaked h (high precision) means s is typically very close to q , whereas a flatter h (low precision) indicates a very noisy signal. Upon receiving signal s , the buyer forms a posterior belief about the underlying quality.

Since the signal is drawn from a distribution centered at q , and the seller's strategy is given by F , Bayes' rule yields the posterior expectation of quality:

$$q_s = \mathbb{E}[q \mid s] = \frac{\int_0^\infty q h(s - q) dF(q)}{\int_0^\infty h(s - q) dF(q)}.$$

Given a posted price p , the buyer will choose to purchase if and only if her expected value exceeds the price, that is, if $q_s \geq p$. Her decision rule, then, can be represented as a function from signal realizations to binary actions: buy or not buy.

Now, because the likelihood function $h(s - q)$ satisfies the Monotone Likelihood Ratio Property (which holds because the PDF of the normal distribution is log-concave), the posterior expectation q_s is strictly increasing in the observed signal s . Higher signals make higher qualities more likely in the Bayesian sense, and thus push the posterior mean upward. This structural monotonicity implies that the buyer's purchasing decision reduces to a simple threshold rule: there exists a unique cutoff signal \bar{s} such that the buyer purchases the good if and only if $s \geq \bar{s}$. At the margin, the signal \bar{s} is defined by the indifference condition $q_{\bar{s}} = p$. This yields the following lemma.

Lemma 2. *The buyer's optimal purchasing strategy is a threshold rule. There exists a unique signal threshold \bar{s} such that the buyer purchases if and only if $s \geq \bar{s}$. The threshold is characterized by the condition:*

$$\mathbb{E}[q \mid \bar{s}] = p.$$

The fact that a unique threshold characterizes the purchasing rule directly follows from the monotone likelihood property of the signal distribution making q_s monotone in s . What is less immediate is that the threshold exists whenever F supports qualities below and above p . Such property does not hold for all noise structures, but the normal noise structure with long ranging tails ensures q_s spans the supports of F , thereby ensuring $q_s = p$ for some s .

Conditional on verification at a produced quality q , the probability of purchase under the optimal purchase rule, is, therefore, $1 - H(\bar{s} - q)$. The probability that the buyer purchases is elastic in the seller choice of quality q , thereby introducing a trade-off akin to monopoly pricing.

4.3 Quality distribution in equilibria with trade

Having pinned down the buyer's threshold rule $s \geq \bar{s}$, we now turn to the seller's problem: choosing a quality level q to maximize expected profit, taking the buyer's verification probability α and her cutoff \bar{s} as given. When the buyer verifies (which occurs with

probability α), the buyer purchases if and only if her noisy signal exceeds \bar{s} . Since $\Pr(s \geq \bar{s} | q) = 1 - H(\bar{s} - q)$, the seller's expected revenue from a unit of quality q is

$$\alpha [1 - H(\bar{s} - q)] p$$

conditional on the buyer verifying quality. If the buyer does not verify (which happens with probability $1 - \alpha$), the buyer either buys or not based solely on her prior mean \tilde{q} : the buyer purchases if $\tilde{q} \geq p$, yielding an additional contribution of

$$(1 - \alpha) \mathbf{1}_{\{\tilde{q} \geq p\}} p.$$

Subtracting the linear production cost cq gives the seller's overall expected payoff function:

$$V(q) = \alpha [1 - H(\bar{s} - q)] p + (1 - \alpha) \mathbf{1}_{\{\tilde{q} \geq p\}} p - cq.$$

This objective embodies a classic trade-off much like a monopolist facing a downward-sloping demand curve. Raising quality from q to $q + dq$ increases the probability of sale under verification by $\alpha h(\bar{s} - q) dq$, yielding marginal revenue $\alpha h(\bar{s} - q) p$. But that same increment in quality costs the seller $c dq$. Balancing these two forces yields the first-order condition for an interior optimum:

$$\alpha h(\bar{s} - q) p = c.$$

Because $h(\cdot)$ is the PDF of a normal distribution, the left-hand side has a familiar bell-shape, and thus there are at most two solutions to this equation, depending on where the bell-curve intersects with the horizontal line c . Moreover, a solution corresponds to a maximum of the seller's profit function only if the marginal revenue $\alpha h(\bar{s} - q) p$ cuts the constant marginal cost c from above. We denote this value by q^* . This means that the profit function for the seller has at most one interior maximum. The only other possible maximum lies at the lower boundary $q = 0$. From Lemma 1, we know that the seller must randomize between values above and below p , and therefore, if there exists an equilibrium with trade, it must be the case that $q^* > p$, and the seller randomizes between q^* and $q = 0$.

Putting these pieces together, one obtains the following sharp characterization of all equilibria with trade.

Proposition 1. *In every equilibrium with trade, the seller's quality distribution places positive mass only on two points: $q = 0$ and a unique strictly positive quality $q^* > p$ satisfying*

$$\alpha h(\bar{s} - q^*) p = c.$$

Knowing that the equilibrium distribution has binary support, we can write down the buyer's posterior in closed form and derive her precise cutoff rule. Suppose the seller randomizes between producing no quality $q = 0$ with probability $1 - \mu$ and the high quality $q = q^*$ with probability μ . Upon observing the noisy signal

$$s = q + \epsilon, \quad \epsilon \sim N\left(0, \frac{1}{\rho^2}\right),$$

the buyer's entire posterior mass collapses onto the two atoms $q = 0$ or $q = q^*$. By Bayes' theorem, the probability that the seller chose the high quality q^* given signal s is

$$\Pr(q = q^* | s) = \frac{\mu h(s - q^*)}{\mu h(s - q^*) + (1 - \mu) h(s)},$$

where $h(x) = \rho \phi(\rho x)$ is the normal density with precision ρ . Since the buyer's posterior mean is simply this probability times q^* , we obtain

$$q_s = \Pr(q = q^* | s) q^* = \frac{\mu \rho \phi(\rho(s - q^*))}{\mu \rho \phi(\rho(s - q^*)) + (1 - \mu) \rho \phi(\rho s)} q^*.$$

The buyer then purchases if and only if $q_s \geq p$. Equating $q_s = p$ and solving for the cutoff signal $s = \bar{s}$ involves canceling the $\rho \phi$ terms, isolating the exponential, taking logarithms, and rearranging. One then obtains the optimal purchase threshold for the buyer's signal in closed form:

$$\bar{s} = \frac{q^*}{2} - \frac{1}{2\rho^2 q^*} \ln \left[\frac{1 - \mu}{\mu} \left(\frac{q^*}{p} - 1 \right) \right].$$

It is straightforward to show that the cutoff \bar{s} rises when the seller's mixing weight μ on the high-quality atom or the high-quality level q^* itself increases: in both cases the buyer's prior expectation of good quality is stronger, so he demands a correspondingly higher signal to be sure of trading. Similarly, \bar{s} rises when the price p increases, since a more expensive purchase requires additional reassurance from the signal to justify buying.

4.4 Surplus Squeezing and Moral Hazard

We now examine what happens when the buyer's verification becomes increasingly precise. Formally, this corresponds to the case where the signal precision parameter ρ becomes arbitrarily large, so that the noise in the buyer's signal $s = q + \epsilon$, where $\epsilon \sim N(0, 1/\rho^2)$, vanishes. In this high-precision limit, the signal becomes tightly concentrated around the true quality q , but crucially, it never becomes perfectly revealing. Unlike in the perfect verification case, posteriors under noisy verification always remain mixtures over the seller's original support. That is, the buyer never learns the exact quality.

Despite this structural difference, the limit of high precision leads to a similar strategic effect. As $\rho \rightarrow \infty$, the buyer's signal increasingly pins down the true quality, and the seller responds by selecting a quality q^* that makes the buyer just indifferent at the threshold signal \bar{s} . In formal terms, we show that the profit-maximizing q^* necessarily goes to \bar{s} as precision grows arbitrarily large. But at \bar{s} , the buyer is indifferent between purchasing or not. In other words, producing $q^* = \bar{s}$ at an arbitrarily large precision level secures trade (signal realizations are concentrated at \bar{s}) but provides no surplus to the buyer. Similar to the perfect verification case, high precision results in surplus squeezing. The produces a quality level that barely ensures trade post verification leaving the buyer with no surplus. With zero net surplus after paying verification costs, the buyer eventually refuses to verify, leaving the seller with no incentive to provide anything above the minimal quality, resulting in no trade.

Proposition 2 (Surplus squeeze with Extreme Precision). *If the signal’s precision ρ tends to infinity, there is no trade in equilibrium.*

Sketch of Reasoning. When $\rho \rightarrow \infty$, the buyer’s signal (if acquired) becomes nearly perfect. Since the seller already knows the true quality, he optimally sets quality and price to extract all surplus, leaving the buyer with no net gain after paying the verification cost k . Anticipating this, the buyer chooses not to verify. Without verification, the seller then chooses the lowest possible quality, leading to a breakdown of trade. \square

The other case of extreme precision is when the signal becomes arbitrarily imprecise. Intuitively, if the signal is too noisy, the buyer gains little useful information from verification, making the cost of verifying unjustifiable. In such a situation, the buyer optimally opts not to verify and, as a consequence, no trade occurs.

This intuition is formalized in the following lemma, which shows that when the buyer’s signal is too imprecise—every equilibrium results in no trade.

Proposition 3. *If the signal’s precision goes to 0, there is no trade in equilibrium.*

Lemma 3 and Proposition 2’s results underscore the necessity of having a sufficient private information on the buyer side, by having a not too precise nor too imprecise verification technology. A too precise signal leads to surplus squeezing, similar to perfect information, undermining verification incentives. In contrast, a too imprecise signal leads to severe moral hazard where the seller has no incentives to produce positive quality.

4.5 On the Existence of an Equilibrium with Trade

We now establish conditions under which an equilibrium with trade and verification exists. To focus the analysis, we restrict attention to pure verification strategies: that is, the buyer either always verifies or never does. While mixed strategies with $\alpha \in (0, 1)$ may exist for knife-edge values of the verification cost k , they are fragile and typically leave the buyer with no surplus. Instead, we construct an equilibrium in which the buyer verifies with probability one.

Our approach is to characterize the seller’s optimal quality choice q^* in this setting. Recall that in equilibrium the seller must mix between $q^* > 0$ and 0 to sustain the buyer’s incentive to verify. We begin by analyzing the seller’s first-order condition for q^* and show that such an interior solution exists under appropriate conditions on precision ρ , price p , and marginal cost c .

In particular, we derive the first-order condition from the seller’s profit maximization problem under verification. For any $q^* > p$, one can find a mixing probability μ that makes q^* optimal. The marginal profit condition pins down a specific value ψ_ρ through the relationship

$$\phi(\psi_\rho) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\psi_\rho^2}{2}\right) = \frac{c}{\rho p}.$$

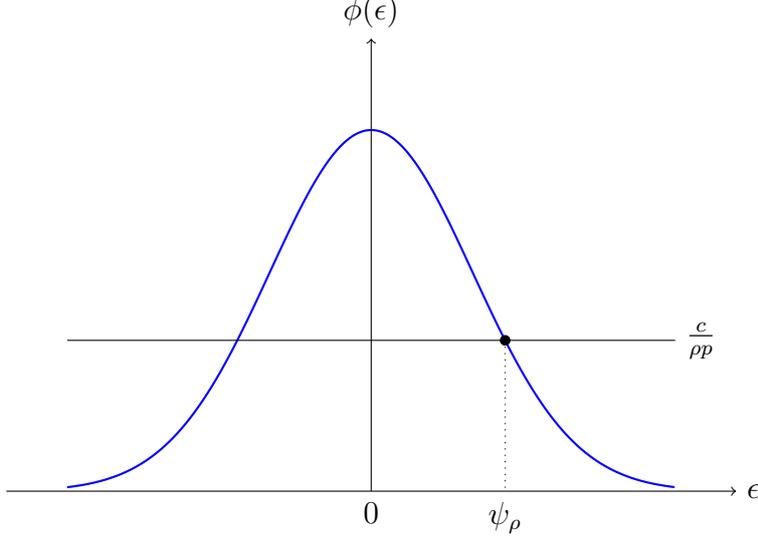


Figure 1: First-Order-Condition on q^*

Solving this equation gives

$$\psi_\rho = \sqrt{-2 \ln \left(\frac{c\sqrt{2\pi}}{\rho p} \right)}.$$

This expression is well-defined only when the argument of the logarithm is strictly less than one, which yields the necessary condition

$$\rho p \geq c\sqrt{2\pi}.$$

This requirement reflects a basic trade-off: the price–precision combination must be high enough to justify the marginal cost of increasing quality.

Figure 1 illustrates this condition graphically. The horizontal line at $c/(\rho p)$ intersects the standard normal density at two symmetric points. We focus on the positive root ψ_ρ , which identifies the seller’s optimal deviation margin.

Next, we relate this optimality condition to the buyer’s purchase decision. The buyer’s signal cutoff \bar{s} , and thus the likelihood of trade, depends on the seller’s mixing strategy. We define the normalized deviation

$$\psi(\mu, q^*) = \rho(\bar{s} - q^*) = -\frac{\rho q^*}{2} - \frac{1}{\rho q^*} \ln \left[\frac{1 - \mu}{\mu} \left(\frac{q^*}{p} - 1 \right) \right],$$

where \bar{s} is given explicitly by

$$\bar{s} = \frac{q^*}{2} - \frac{1}{\rho^2 q^*} \ln \left[\frac{1 - \mu}{\mu} \left(\frac{q^*}{p} - 1 \right) \right].$$

These equations jointly determine μ and q^* . For the seller’s strategy to be optimal, the normalized deviation $\psi(\mu, q^*)$ must satisfy the marginal condition:

$$\phi \left[\psi(\mu, q^*) \right] = \frac{c}{\rho p}. \quad (1)$$

This is a key implicit equation that governs the seller's incentives. For fixed values of $q^* > p$, we can find a corresponding μ such that this condition holds. That is, any high-quality level above the price can be made optimal by suitably choosing the mixing weight. We denote the solution to this equation $\psi_\rho = \psi(\mu(q^*), q^*)$.

To ensure the seller is willing to mix, we must verify that q^* is at least weakly more profitable than producing 0. This amounts to checking that the expected gain in trade probability outweighs the marginal cost:

$$\Phi(\psi_\rho + \rho p) - \Phi(\psi_\rho) > c.$$

This inequality ensures that the slope of the seller's payoff from raising quality is positive at $q = 0$ and reaches zero only at $q^* > 0$. Intuitively, there must be a large enough shift in the buyer's trade probability to compensate for the higher cost of quality.

Finally, we address the buyer's decision to verify. If the verification cost k is not too large, the buyer finds it worthwhile to pay for information, given the seller's strategic mixing. In that case, the game admits a verification equilibrium in which the buyer inspects with certainty and the seller mixes between 0 and q^* .

Theorem 2 (Existence of a Verification Equilibrium). *Let p be fixed and suppose*

$$\rho p \geq c\sqrt{2\pi} \quad \text{and} \quad \Phi(\psi_\rho + \rho p) - \Phi(\psi_\rho) > c.$$

Then there is a range $k \in (0, \bar{k}]$ such that there exists a unique verification equilibrium at price p .

This result shows that trade can be sustained in equilibrium with costly and noisy verification, provided the parameters are aligned to support both seller incentives and buyer willingness to acquire information. The equilibrium balances the moral hazard problem through endogenous uncertainty: the seller randomizes to make verification valuable, and the buyer verifies to discipline the seller.

5 Discussion

The central idea behind the informational hold-up problem is what we call a “surplus squeeze”: the buyer pays to verify, but the seller adjusts quality downward to exactly meet the buyer's minimum standard. As a result, verification gains no surplus to the buyer. Anticipating this, the buyer stops verifying altogether. The seller, knowing that the buyer will neither verify nor buy, provides no quality and trade collapses. Importantly, this logic is not the result of the particular feature of our model, but arises purely from the interaction of a sunk cost and informational precision.

Noise in the buyer's signal limits the seller's ability to surplus squeeze. When the buyer's signal is imprecise, there is greater variation in the buyer's posterior belief regarding quality, so that the seller's knowledge of true quality is less correlated with the buyer's posterior. Therefore, if the seller chooses a quality level close to the price, there

larger probability that the seller will receive an unfavorable signal and refuse to buy post-verification. As a result, it is optimal for the seller to leave a larger surplus to the buyer, which in turn makes it valuable for the buyer to verify.

Relation to Inspection Games This insight links closely to a classic game-theoretic literature on inspection games (see Avenhaus et al., 2002, for a survey). These are models where one party (an inspector or auditor) decides whether to monitor, and another party (typically an agent or regulated firm) decides whether to comply. In equilibrium, the inspector randomizes inspections, and the agent randomizes compliance, so that each keeps the other indifferent. Importantly, as inspection becomes cheaper or more accurate, compliance becomes more common—more transparency leads to better behavior.

Our model can be seen as a natural extension of this idea, but with a crucial twist. The “agent” in our case (the seller) does not choose between two fixed actions (comply or cheat), but selects a level of quality along a continuum. This means the seller can finely tune their behavior to the buyer’s standard. And when verification is perfect, the seller will tune precisely to the buyer’s threshold. But this ends up “squeezing out” all gains for the buyer. Unlike inspection games, where partial compliance is sustained through mutual randomization, our model collapses to a unique pure-strategy equilibrium—no verification, no quality, no trade. The introduction of a continuous action space for one side fundamentally alters the equilibrium structure.

Why Private Information Alone Doesn’t Help One might think that introducing private information might help improve matters. For example, consider the case in which the buyer’s verification cost k is private information. This means the seller only knows that k is drawn from distribution G , but she doesn’t know exactly how much the buyer has to pay for verification. Could this information asymmetry restore trade? Unfortunately, it doesn’t help. Even when the seller doesn’t know precisely who the buyer is, he knows that buyers with very high costs won’t verify. So he can safely lower quality for that part of the population. Those buyers, getting squeezed, drop out. But then slightly lower-cost buyers anticipate that the seller is lowering quality in response, and they stop verifying too. This logic continues all the way up the distribution—trade “unravels from the bottom.”

This is a familiar phenomenon in models with adverse selection: marginal types drop out, making things worse for others. The key point here is that private information about the buyer’s own verification cost does not protect the buyer from being squeezed, because the seller tailors quality based on average behavior. As a result, there is no gain from adding private cost heterogeneity. What turns out to matter is uncertainty that the seller has about the buyer’s signal or payoff—something that prevents him from tailoring quality perfectly to extract surplus.

Robustness to Noise Specifications While our baseline model uses a Gaussian error in verification, the main insights don’t depend on that specific choice. What matters is

that the noise smooths the link between true quality and what the buyer sees. As long as the noise distribution has the property that better quality makes higher signals more likely (technically, a monotone likelihood ratio), the results hold. The seller's expected profit remains hump-shaped, and we get a clean two-point quality support in equilibrium. Perfect information kills trade by enabling surplus squeeze while some noise helps restore it.

Random Perfect Verification and Commitment Suppose the buyer can commit to verifying the good with some probability, and the verification is perfect when it occurs. One might expect this to improve outcomes, but it doesn't. The core issue is that the seller chooses quality based on the buyer's posterior belief—which becomes fully informative if verification occurs. Even if the buyer only verifies with some probability, the seller knows what the posterior will be in that branch and anticipates the buyer's action accordingly. This leads to the same surplus squeeze as in the full verification case.

What does help is commitment to not buying following a bad signal. This kind of commitment changes the seller's incentives directly: if low quality leads to a guaranteed lost sale, the seller may prefer to provide high quality to avoid that outcome. In this sense, the logic mirrors that of ultimatum games, where the responder's ability to reject low offers, even at a cost, disciplines the proposer's behavior.

Cost of Quality and Generalizations We assume throughout that the seller's cost of quality is linear, so the marginal cost is constant. We make this assumption for tractability, but it is not essential for our results. If one assumes instead that the cost of quality is convex, then this complicates the analysis in a number of ways. Most importantly, the seller's optimality condition does not necessarily have a unique maximum, but may have multiple, or even a continuum of maxima. It then becomes more difficult to establish equilibrium existence, but one can generalize our results and finding provided that the marginal cost is well-behaved (in particular, the high-order derivatives of the marginal cost should be positive and bounded). In this case, our main findings go through: as long as the verification signal is noisy, the seller randomizes between zero and some (not necessarily unique) optimum. Convex costs may shift the optimal quality point but don't undo the mechanism.

Endogenous Choice of Verification Technology Finally, imagine a setting where the buyer can choose how to verify, either paying a small cost for noisy inspection or a higher cost for perfect inspection. At first glance, the perfect signal seems better. But in equilibrium, the buyer understands that perfect inspection will lead the seller to perfectly tailor quality to her cutoff, squeezing out all her surplus. The buyer ends up preferring noisy inspection, because the uncertainty leaves enough slack to prevent the seller from extracting everything. In this sense, the buyer deliberately chooses to remain partially ignorant in order to elicit better behavior. This result is perhaps the most striking. It shows that more information is not always better—even from the buyer's perspective.

When the seller can respond opportunistically, transparency backfires. It is precisely the uncertainty—the noise—that allows some rent to be preserved, and hence some quality to be provided. This echoes broader themes in the literature on relational contracts and information design: sometimes, it’s optimal not to know everything.

One could also consider settings in which the buyer learns flexibly by choosing an information structure rather than a fixed verification technology. This raises deeper questions about how information is acquired—whether through signal design, noise control, or attention allocation. The approach closest to ours is the noise-precision model of Achim and Ozbek (2025), in which buyers endogenously choose the noise level of their signal.

6 Conclusion

This paper examines how costly buyer learning interacts with moral hazard in a static bilateral trade setting without commitment. We show that while verification can alleviate incentive problems, overly precise information induces surplus squeezing: the seller exploits the buyer’s learning by adjusting quality to extract surplus, ultimately discouraging inspection and causing trade to break down. By contrast, noisy verification preserves some uncertainty and sustains buyer surplus, enabling trade.

Our model captures an informational form of hold-up, where the uninformed party’s investment in learning is strategically undermined by the informed party’s response. While our framework abstracts from dynamic considerations and richer contractual instruments, it offers a transparent benchmark for understanding how information precision shapes incentives in bilateral exchange. Future research could explore these dynamics in repeated settings, allow for richer learning technologies or intermediary structures, and examine welfare implications when verification precision itself is endogenous.

References

- Achim, P. and J. Knoepfle (2024). Relational enforcement. *Theoretical Economics* 19(2), 823–863.
- Achim, P. and J. Knoepfle (2025). The tension between trust and oversight in long-term relationships. arXiv preprint arXiv:2504.02696.
- Achim, P. and K. Ozbek (2025). Flexible learning via noise reduction. arXiv preprint arXiv:2503.20741.
- Achim, P., L. Yang, and B. Li (2025). Endorsements and referrals: Product recommendations in bilateral trade. Available at SSRN 5168225.
- Armstrong, M. and J. Zhou (2011). Paying for prominence. *The Economic Journal* 121(556), F368–F395.

- Avenhaus, R., B. Von Stengel, and S. Zamir (2002). Inspection games. *Handbook of game theory with economic applications* 3, 1947–1987.
- Border, K. C. and J. Sobel (1987). Samurai accountant: A theory of auditing and plunder. *The Review of economic studies* 54(4), 525–540.
- Diamond, P. A. (1971). A model of price adjustment. *Journal of economic theory* 3(2), 156–168.
- Evans, R. and I.-U. Park (2024). Selling information for bilateral trade. Working Paper.
- Fabbri, F. (2025). Attention holdup. Working Paper.
- Fudenberg, D. and D. K. Levine (1994). Efficiency and observability with long-run and short-run players. *Journal of Economic Theory* 62(1), 103–135.
- 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(4), 691–719.
- Holmström, B. (1979). Moral hazard and observability. *The Bell Journal of Economics* 10(1), 74–91.
- Inderst, R. and M. Ottaviani (2012). Competition through commissions and kickbacks. *American Economic Review* 102(2), 780–809.
- Klein, B., R. A. Crawford, and A. A. Alchian (1978). Vertical integration, appropriable rents, and the competitive contracting process. *Journal of Law and Economics* 21(2), 297–326.
- Mailath, G. J. and S. Morris (2002). Repeated games with almost-public monitoring. *Journal of Economic theory* 102(1), 189–228.
- Ravid, D. (2020). Ultimatum bargaining with rational inattention. *American Economic Review* 110(9), 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(2), 346–387.
- Roesler, A.-K. and B. Szentes (2017). Buyer-optimal learning and monopoly pricing. *American Economic Review* 107(7), 2072–2080.
- Strausz, R. (2005). Honest certification and the threat of capture. *International Journal of Industrial Organization* 23(1-2), 45–62.
- Teh, T.-H. and J. Wright (2022). Intermediation and steering: Competition in prices and commissions. *American Economic Journal: Microeconomics* 14(2), 281–321.
- Townsend, R. M. (1979). Optimal contracts and competitive markets with costly state verification. *Journal of Economic Theory* 21(2), 265–293.

A Proofs

Proof of Proposition 1. First, consider the case where the buyer does not verify in equilibrium. The buyer purchases if she believes the expected quality produced is greater than the price, $q^e \geq p$. In this case, producing $q = 0$ is a strictly profitable deviation for the seller, as the buyer still purchases (no verification), and so the seller makes the same revenue at 0 cost. Therefore, the buyer cannot purchase the seller's product in equilibrium if she does not verify. Consequently, there is no equilibrium with trade if the buyer does not verify.

Suppose instead the buyer verifies with a positive probability $\alpha > 0$ in equilibrium. Verification is a best response if the buyer believes the seller plays a strategy F such that:

$$\int_p (q - p) dF(q) \geq k$$

Therefore, F supports quality levels such that $q - p \geq k$ (otherwise the buyer cannot expect a surplus from trade of at least k). However, producing $q > p$ is strictly worse than producing a q' such that $q > q' > p$: It generates the same level of trade at a strictly lower cost. Therefore, the seller cannot produce any quality $q > p$ in equilibrium. In turn, the buyer's expected surplus from trade post-verification is 0, and the buyer does not verify in equilibrium.

The only equilibrium under perfect verification is no verification, no trade, and no quality production. \square

B Proofs

Proof of Lemma 1. If the seller produces quality q_0 with probability 1:

$$F(q) = \mathbb{1}_{\{q \geq q_0\}}(q),$$

then the buyer's posterior belief is $q_s = q_0$ for all signal realizations. The value of information boils down to:

$$y_F = \int_{\mathbb{R}} \max\{q_0 - p, 0\} h(s - q_0) ds - \max\{q_0 - p, 0\} = 0.$$

The buyer makes the same decision whether she verifies or not: the value of verification is 0, hence the buyer does not verify.

However, if the buyer does not verify, then no positive quality level can be sustained in equilibrium (see no verification benchmark discussion). Therefore the buyer expects $q_0 = 0$ and does not trade.

If the seller plays F that supports qualities only above price, then the buyer updates his expected belief but does not adapt his purchasing decision based on the signal. In this case, verification has no value. Indeed, if $q_s \geq p$ for all s then $\max\{q_s - p, 0\} = q_s - p$ for all s , thereby:

$$y_F = \int_{\mathbb{R}_+} \int_{\mathbb{R}} (q_s - p) h(s - q) ds dF(q) - (\tilde{q} - p) = \tilde{q} - p - (\tilde{q} - p) = 0$$

The buyer does not verify, leading to no trade.

Lastly if the seller strategy only supports quality below price, then there is no trade (nor verification). \square

Proof of Lemma 2. Without loss of generality, we shall assume that the support of F is compact. Since the noise density $h(\varepsilon)$ is log-concave, it satisfies the Monotone Likelihood-Ratio property (MLRP) in q . Hence the posterior mean $q_s = \mathbb{E}[q | s]$ is strictly increasing in s . Let

$$\bar{q} = \max(\text{supp } F), \quad \underline{q} = \min(\text{supp } F).$$

We now show that

$$\lim_{s \rightarrow +\infty} q_s = \bar{q}, \quad \lim_{s \rightarrow -\infty} q_s = \underline{q}.$$

Since h is a log-concave density (for example, Gaussian) centered at zero, we have

$$\frac{h(s - q)}{h(s - \bar{q})} = \exp\left\{-\frac{\bar{q}^2 - q^2 - 2s(\bar{q} - q)}{2\sigma^2}\right\} \rightarrow -\infty \quad \text{as } s \rightarrow +\infty.$$

Hence

$$\lim_{s \rightarrow +\infty} \frac{h(s - q)}{h(s - \bar{q})} = 0 \quad \text{for every } q < \bar{q}.$$

Thus the posterior mass concentrates arbitrarily close to \bar{q} and so $\lim_{s \rightarrow +\infty} \mathbb{P}(q < \bar{q} | s) = 0$, which implies

$$\lim_{s \rightarrow +\infty} q_s = \lim_{s \rightarrow +\infty} \mathbb{E}[q | s] = \bar{q}.$$

An identical argument applies on the left. Fix any $q > \underline{q}$. As $s \rightarrow -\infty$,

$$\frac{h(s - q)}{h(s - \underline{q})} = \exp\left\{-\frac{(s - q)^2 - (s - \underline{q})^2}{2\sigma^2}\right\} \rightarrow 0.$$

and hence the posterior mass collapses onto $q = \underline{q}$, and

$$\lim_{s \rightarrow -\infty} q_s = \lim_{s \rightarrow -\infty} \mathbb{E}[q | s] = \underline{q}.$$

Finally, note that if since q_s is strictly increasing in s by MLRP, and so for $\underline{q} < p < \bar{q}$, there is a unique \bar{s} such that

$$q_{\bar{s}} = \mathbb{E}[q | \bar{s}] = p.$$

If instead, $\underline{q} < \bar{q} \leq p$, then $q_s < p$ so that for all s , so the buyer never purchases. Similarly, if $\bar{q} > \underline{q} \geq p$, then $q_s > p$ for all s , so the buyer always purchases regardless of the signal. \square

Proof of Proposition 1. The seller's payoff is

$$v(q) = \alpha(1 - H(\bar{s} - q))p + (1 - \alpha)p - cq,$$

The seller's optimal choice of quality for fixed p^* corresponds to a maximum of $v(q)$. If $\alpha = 0$, then $v'(q) < 0$, so the optimal level of quality is $q = 0$, and hence $\tilde{q} = 0$. Thus, $\tilde{q} > 0$ implies $\alpha > 0$. We now show that $\tilde{q} > 0$ implies that $v(\cdot)$ must have two maxima. Consider the derivative of u :

$$v'(q) = \alpha h(\bar{s} - q)p - c$$

Note that $h(\cdot)$ is the density of the normal distribution, and as such is quasi-concave, which implies that the identity

$$v'(q) = \alpha h(\bar{s} - q)p - c = 0$$

has either one, two, or no solutions. Note that

$$\lim_{q \rightarrow 0} v'(q) = \lim_{q \rightarrow \infty} -c < 0.$$

If $v'(q) = 0$ has no solutions, then $v'(q) < 0$ for all q , so that $q = 0$ is the unique maximizer, so choosing any positive quality is not optimal for the seller.

Suppose $v'(q) = 0$ has two solutions. Denote these by q_0 and $q_1 > q_0$. Since h is quasi-concave, this is possible only if $v'(0) < 0$ and $\lim_{q \rightarrow \infty} v'(q) < 0$. Hence $v'(q)$ is increasing at $q = q_0$ and decreasing at $q = q_1$, that is, $v''(q_0) > 0$ and $v''(q_1) < 0$. This means that q_0 is a local minimum, and q_1 is a local maximum. So only q_1 can be a candidate for a global maximum.

We thus established that there are at most two candidates for an interior global maximum, $q = 0$ and $q = \max\{q | v'(q) = 0\} =: q^*$. If only one of these local maxima is a global maximum of u , then the seller would find it optimal to choose the maximizing quality level with probability 1, but then the buyer would not find it optimal to verify, contradicting the requirement that $\alpha > 0$.

It thus follows that if there exists an equilibrium with positive expected quality, $\tilde{q} = \mathbb{E}_F[q] > 0$, then the support of F must be $\{0, q^*\}$. \square

Proof of Proposition 2. We start by proving a useful lemma showing that, in equilibrium, μ is bounded away from 1, and that q^* is bounded away from p .

Lemma 3. *Consider a verification equilibrium with trade (q^*, μ) , (α, \bar{s}_μ) . For all precision level ρ , we have that:*

$$\begin{aligned} \frac{kc}{p(1-c)} &\leq \mu \leq 1 - \frac{k}{p} \\ \frac{k}{p}(p-k) + p &\leq q^* \leq \frac{p}{c} \end{aligned}$$

Proof. Consider the buyer's verification value:

$$y_F = \int_{\mathbb{R}_+} \int_{\mathbb{R}} \max\{q_s - p, 0\} h(s - q) ds dF(q) - \max\{\tilde{q} - p, 0\}.$$

When $\rho \rightarrow +\infty$ the density of the normal distribution $h(s - q)$ converges to the Dirac distribution δ_q , so that at infinite precision:

$$y_F^\infty = \int_{\mathbb{R}_+} \max\{q - p, 0\} dF(q) - \max\{\tilde{q} - p, 0\}.$$

Of course $y_F^\infty \geq y_F$ for all precision levels. Using the seller's strategy (q^*, μ) we have that:

$$y_F \leq \mu(q^* - p) - \max\{\mu q^* - p, 0\}$$

A necessary condition for the buyer to verify in equilibrium is that

$$\begin{aligned} \mu(q^* - p) - \max\{\mu q^* - p, 0\} - k \geq 0 &\iff \begin{cases} \mu(q^* - p) - \mu q^* + p - k \geq 0 & \text{and} \\ \mu(q^* - p) - k \geq 0 \end{cases} \\ &\iff \begin{cases} \mu \leq 1 - \frac{k}{p} & \text{and} \\ \mu \geq \frac{k}{q^* - p} \end{cases} \end{aligned}$$

Using the second inequality we have $q^* - p \geq \frac{k}{\mu} \geq \frac{k}{p}(p - k)$.

Now consider the seller's profit as a function of the quality q^* . We have that

$$v(q^*) \leq p - cq^*$$

For q^* to be a best response a necessary condition is that $v(q^*) \geq 0$, so we must have that:

$$p - cq^* \geq 0 \iff q^* \leq \frac{p}{c}$$

Plugging it in the previous lower bound on μ we have $\mu \geq \frac{kc}{p(1-c)}$. \square

We now prove the statement of proposition 2. The first-order condition characterizing q^* is

$$\alpha \rho \phi\left(\rho(\bar{s} - q^*)\right) p = c,$$

where $\phi(z) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z^2}{2}\right)$ is the standard normal density and \bar{s} is the buyer's cutoff signal. Isolating $(\bar{s} - q^*)$ using the first-order condition we obtain

$$\frac{(\bar{s} - q^*)^2}{2} = \frac{1}{\rho^2} \ln\left(\frac{\rho}{\sqrt{2\pi}} \frac{p}{c}\right).$$

Consider a sequence of equilibria indexed by ρ_n $(q^*(\rho_n), \mu(\rho_n)), (\alpha(\rho_n), \bar{s}_\mu(\rho_n))$ where $n \in \mathbb{N}$, and such that $\lim_{n \rightarrow \infty} \rho_n = \infty$. At all ρ_n :

$$\frac{(\bar{s}_\mu(\rho_n) - q^*(\rho_n))^2}{2} = \frac{1}{\rho_n^2} \ln\left(\frac{\rho_n}{\sqrt{2\pi}} \frac{p}{c}\right).$$

Therefore:

$$\lim_{n \rightarrow \infty} \frac{(\bar{s}_\mu(\rho_n) - q^*(\rho_n))^2}{2} = 0 \iff \lim_{n \rightarrow \infty} (\bar{s}_\mu(\rho_n) - q^*(\rho_n)) = 0.$$

Similar to perfect verification, with almost perfect verification producing $q' = \bar{s} + \epsilon$ secures almost all trades for minimal costs. At the limit, the best quality q^* equals the threshold.

Now consider the equation characterizing the threshold \bar{s}_μ :

$$\bar{s}_\mu = \frac{q^*}{2} - \frac{1}{2\rho^2 q^*} \ln \left[\frac{1 - \mu \left(\frac{q^*}{p} - 1 \right)}{\mu} \right].$$

Because $q^*(\rho_n)$ is uniformly bounded above $p + \frac{k}{p}(p - k)$ and below $\frac{p}{c}$, and that $\mu(\rho_n)$ is uniformly bounded away from 0 and 1, then:

$$\lim_{n \rightarrow \infty} \frac{1}{2\rho^2 q^*(\rho_n)} \ln \left[\frac{1 - \mu(\rho_n) \left(\frac{q^*(\rho_n)}{p} - 1 \right)}{\mu(\rho_n)} \right] = 0.$$

Therefore:

$$\lim_{n \rightarrow \infty} \left(s_\mu(\rho_n) - \frac{q^*(\rho_n)}{2} \right) = 0.$$

Using the Algebraic Limit Theorem we have that:

$$\begin{aligned} \lim_{n \rightarrow \infty} \left(s_\mu(\rho_n) - q^*(\rho_n) - s_\mu(\rho_n) + \frac{q^*(\rho_n)}{2} \right) &= 0 \\ \implies \lim_{n \rightarrow \infty} q^*(\rho_n) &= 0 \end{aligned}$$

But $q^*(\rho^n)$ is uniformly lower bounded by $\frac{k}{p}(p - k) + p$, a contradiction. \square

Proof of Proposition 3. The first-order condition is:

$$\frac{\rho}{\sqrt{2\pi}} \exp \left(\frac{-\rho^2(\bar{s} - q^*)^2}{2} \right) = c.$$

So that when the precision goes to 0 the marginal revenue goes to 0: When the buyer has an arbitrarily imprecise signal, increasing or reducing the quality produced has almost no effect on sales.

Rewriting the FOC as in proposition 2 we have that

$$\frac{(\bar{s} - q^*)^2}{2} = \frac{1}{\rho^2} \ln \left(\frac{\rho}{\sqrt{2\pi} c} \right),$$

so that

$$\lim_{\rho \rightarrow 0} \frac{(\bar{s} - q^*)^2}{2} = -\infty.$$

As marginal revenue goes to 0, the seller's incentive is to lower quality (reducing production costs without losing sales), down to $q^* = 0$ where there is no interior solution anymore. \square

Proof of Theorem 2. We call ψ_ρ the solution to (1):

$$\phi(\psi_\rho) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\psi_\rho^2}{2}\right) = \frac{c}{\rho p}.$$

Solving for the maximum ψ_ρ yields

$$\psi_\rho = \sqrt{-2 \ln\left(\frac{c\sqrt{2\pi}}{\rho p}\right)},$$

In order for (1) to hold, we must have

$$\rho p \geq c\sqrt{2\pi}$$

We first solve for $\psi(\mu, q^*) = \psi_\rho$ by picking a suitable $\mu(q^*)$. Recall that

$$\psi(\mu, q^*) = -\frac{\rho q^*}{2} - \frac{1}{\rho q^*} \ln \left[\frac{1-\mu}{\mu} \left(\frac{q^*}{p} - 1 \right) \right].$$

In equilibrium q satisfies the first-order condition, we need:

$$\phi(\psi(\mu, q^*)) = \frac{c}{\rho p}$$

$\psi(\mu, q^*)$ is defined and continuous in $\mu \in (0, 1)$ whenever $q^* > p$. Also, $\lim_{\mu \rightarrow 0} \psi(\mu, q^*) = -\infty$ and $\lim_{\mu \rightarrow 1} \psi(\mu, q^*) = +\infty$. Moreover, ψ is continuous, therefore for all $q^* > p$, there exists a $\mu(q^*)$ such that $\psi(\mu(q^*), q^*) = \psi_\rho$ that is such that q^* satisfies the first order condition.

Recall that accounting for buyer's verification incentives, we have (lemma 4) that $\frac{kc}{p(1-c)} \leq \mu \leq 1 - \frac{k}{p}$. So the μ satisfying the FOC for $\phi(\mu, q^*)$ might be outside these bounds. However, we can pick low enough k to ensure the desired μ is feasible.

Therefore, for all $q^* > p$ there exists a $k > 0$ and a $\mu(q^*)$ such that q^* indeed solves the FOC.

Second, we find a q^* such that at $\psi(\mu(q^*), q^*) = \psi_\rho$ the seller is indifferent between 0 and q^* . The probability of observing a signal above the approval threshold is $\Phi(\rho(\bar{s} - q^*)) = \Phi(\psi_\rho)$. When $q = 0$, then the probability of observing a signal above the approval threshold is $\Phi(\rho\bar{s}) = \Phi(\psi_\rho + \rho q^*)$ (where we use that $\rho(\bar{s} - q^*) = \psi_\rho$ to write $\rho\bar{s} = \psi_\rho + \rho q^*$). Hence, indifference requires

$$(1 - \Phi(\psi_\rho))p - c q^* = (1 - \Phi(\psi_\rho + \rho q^*))p. \quad (2)$$

Define the function

$$f(q^*) := (\Phi(\psi_\rho + \rho q^*) - \Phi(\psi_\rho))p - c q^*.$$

Then Equation (2) holds if and only if $f(q^*) = 0$.

First, $f(q^*)$ is strictly decreasing if and only if

$$\phi(\psi_\rho + \rho q^*) \leq \frac{c}{\rho p}.$$

Using the formula for the density of the normal and the value of ψ_ρ we have that:

$$\begin{aligned}\phi(\psi_\rho + \rho q^*) &= \frac{c}{\rho p} \exp\left\{-\frac{1}{2}(\psi_\rho^2 + 2\psi_\rho \rho q^* + (\rho q^*)^2)\right\} \\ &= \frac{c}{\rho p} \exp\left\{-\frac{1}{2}(2\psi_\rho \rho q^* + (\rho q^*)^2)\right\}.\end{aligned}$$

Therefore, f is strictly decreasing in q^* . We now use the Intermediate Value Theorem to show that there exists a (unique) finite $q^* \in (p, \frac{p}{c})$ such that $f(q^*) = 0$. Consider

$$\lim_{q^* \downarrow p} f(q^*) = (\Phi(\psi_\rho + \rho p) - \Phi(\psi_\rho))p - cp.$$

By hypothesis, we have

$$\Phi(\psi_\rho + \rho p) - \Phi(\psi_\rho) > c,$$

and thus $\lim_{q^* \downarrow p} f(q^*) > 0$. On the other hand, consider $q^* = \frac{p}{c}$:

$$f\left(\frac{p}{c}\right) = (1 - \Phi(\psi_\rho))p - p < 0.$$

To break even the seller must trade with probability 1 at $q^* = \frac{p}{c}$. But, $\Phi(\psi_\rho) > 0$, the buyer may draw signals below the threshold and therefore $\frac{p}{c}$ yields negative profit. The Intermediate Value Theorem implies that there exists a $q^* \in (p, \frac{p}{c})$ such that

$$f(q^*) = 0.$$

From lemma 4 we have: $\frac{k}{p}(p-k) + p \leq q^*$. So we can make any $q^* > p$ feasible by picking a low enough $k > 0$. Therefore, for all q^* satisfying indifference condition there is a k and a μ such that (μ, q^*) and $(\alpha = 1, \bar{s}_\mu)$ constitutes a perfect bayesian equilibrium of the game. \square