UNIVERSITÀ COMMERCIALE "LUIGI BOCCONI"
PHD SCHOOL

PhD program in: Economics and Finance
Cycle: 35th
Disciplinary Field (code): SECS-P/01

Essays in Microeconomic Theory

Advisor: Fernando VEGA-REDONDO

PhD Thesis by
Francesco MASSAZZA
ID number: 3112135

Year: 2025

Abstract

This dissertation consists of three essays in microeconomic theory. The first chapter, co-authored with Daria Stepushina, develops a buyer–seller model in which the seller may voluntarily disclose partially verifiable evidence of product quality and the buyer can pay for a potentially noisy authentication test. Each technology's precision ranges from uninformative to perfectly accurate. We fully characterize all pure- and mixed-strategy equilibria under natural refinements on off-path beliefs and seller payoffs. Our analysis reveals that (i) full separation is attainable only when seller evidence is sufficiently precise relative to the quality gap or with a perfect (and free) authentication: no positive-cost authentication, however accurate, can substitute for weak seller disclosure; (ii) authentication can never fully eliminate fraud, yet perfectly accurate tests minimize it; and (iii) when disclosed evidence is sufficiently uninformative, forgoing authentication yields higher welfare than a perfect (but costly) authentication, despite high fraud levels. These findings deliver clear prescriptions for designing certification regimes. The second chapter analyzes Bayesian persuasion in a network context, focusing on heterogeneous receivers and correlated messages. The sender aims to maximize the adoption of her preferred action by skeptical receivers without dissuading believers. This chapter introduces a novel definition for soft news strategies in networks and presents results on their characterization. The third chapter investigates the influence of pre-electoral polls on electoral outcomes. It considers strategic behavior among voters and pollsters and addresses empirical discrepancies between poll predictions and actual results. By proposing a theoretical model, this chapter aims to explain the mechanisms behind bandwagon and underdog effects, study competition among pollsters with misaligned interests, and explore the potential for polls to influence electoral outcomes.

Contents

Acknowledgements

I am beyond grateful to my advisor, Fernando Vega-Redondo, for his continued support throughout my PhD. He believed in me and helped me through the hardest times of this endeavor. I will never be able to thank him enough for his invaluable guidance, patience, and encouragement.

I deeply thank my committee members, Nenad Kos and Tangren Feng, for their availability and advice; the Bocconi PhD school and the Fondazione Romeo ed Enrica Invernizzi for the infrastructure and financial support provided over the last five years; the program directors, Max Croce and Marco Ottaviani; and the administrative assistants, Angela, Silvia, and Giulia, for their dedication to the program.

I would like to acknowledge all the Bocconi faculty I was lucky to collaborate with during these years: Alfredo Di Tillio, Guido Tabellini, Fausto Panunzi, Tito Boeri, Thomas Le Barbanchon, Serena Negrelli, Salvo Nunnari, Francesco Scervini, Claudio Brenna, Fernando Vega-Redondo, Nenad Kos, and the amazing staff members: Marika, Alessandra, Erika, Giovanna, and Rebecca, for their help and support with all the teaching activities.

I am truly indebted to Amalia, Diego, Inês, Jack, and Zheng: this journey would not have been possible without your friendship, constant support, and enlightening discussions. You are not only great minds but also amazing people. Thanks to Jeffrey, Lorenzo, and Piero for the hours we spent together in the office. Thanks to all my PhD colleagues, particularly Giuliano, Jacopo, Goonj, Jaime, Pietro, Melika, Daria, Francesco, Giacomo, the La Strada-Bar Bruto crew, and the I Giovani di Ieri teammates. If you are not on the list, you should be.

Finally, I am extremely grateful to my family for their unconditional support and encouragement, to my friends for the laughter and good memories, and to Cate, who is probably the reason why I embarked on this journey and certainly the reason why I was able to persevere through the storms. I dedicate this dissertation to them.

Introduction

This dissertation comprises three essays in microeconomic theory. While each essay explores distinct environments, they collectively investigate how information provision influences key economic outcomes. The first chapter, co-authored with Daria Stepushina, examines equilibrium dynamics in a seller-buyer scenario where private information about product quality can be disclosed and acquired by the respective party. The second chapter analyzes Bayesian persuasion in a network context, focusing on heterogeneous receivers and correlated messages, with the sender aiming to maximize adoption of her preferred action by sceptical receivers while not dissuading believers. The third chapter proposes to investigate the influence of pre-electoral polls on electoral outcomes, considering strategic behavior among voters and pollsters, and addressing empirical discrepancies between poll predictions and actual results.

Chapter 1 develops a unified framework that bridges the seller-driven disclosure literature (Grossman, 1981, Milgrom, 1981, Dye 1985) and the buyer-driven certification literature (Matthews and Postlewaite 1985, Bester and Ritzberger 2001, Stahl and Strausz 2017) by allowing both parties to choose information technologies of varying precision. A privately informed seller simultaneously selects a price and a level of partially verifiable evidence, generating a noisy public signal of quality; a buyer then observes price, evidence, and its realization, and decides whether to purchase outright, pay for a costly authentication test, or walk away. By letting evidence precision and authentication accuracy range from uninformative to perfect, we characterize every separating, pooling, and mixed-strategy equilibrium under an intuitive-criterion refinement and a seller-payoff maximization refinements, derive sharp existence conditions, and compare each regime's impact on trade volume, welfare, and fraud incidence. We show that full efficiency arises only when seller evidence is sufficiently precise, that perfect (but costly) authentication can never fully substitute for weak disclosure, and that—paradoxically—foregoing authentication can sometimes raise total welfare when disclosure is uninformative. These findings elucidate the fundamental trade-offs between voluntary disclosure and authentication, and yield clear policy prescriptions for designing verification mechanisms in quality-uncertain markets.

In Chapter 2, I consider a model of Bayesian persuasion with one sender and multiple heterogeneous receivers connected in a network. To persuade the receivers to take her favored action, the sender commits to a signaling policy that sends private correlated messages to receivers, who observe all realizations in their neighborhood. The receivers differ in their priors and are divided into believers and skeptics. I analyze the sender's optimal strategies when her objective is to maximize the expected probability that skeptics take her favored action while never dissuading the believers. In this context, I consider soft news strategies (Innocenti, 2021). I provide a novel definition for soft news strategies in networks and present preliminary results about the characterization of the set of all such strategies.

In Chapter 3, presented as a research proposal, I address the question, "Can electoral polls affect voter behavior and, ultimately, election outcomes?" The bandwagon and underdog effects are two debated and complementary explanations for why electoral poll predictions might significantly differ from actual electoral results (e.g., Trump's election in 2016, and the 2022 US midterm results that pointed toward a close election). I propose a theoretical model to explain the mechanisms behind these effects, study competition among pollsters with misaligned interests, and explore the potential for polls to influence electoral outcomes. Finally, I intend to test the theoretical predictions with empirical data to provide a comprehensive understanding of the impact of opinion polls on elections.

Chapter 1

Trust but Verify: Quality Evidence and Authentication

1.1 Introduction

Asymmetric information about product quality generates market inefficiencies, providing sellers with an incentive to disclose information and buyers with a motive to acquire it. High-quality sellers seek to credibly convey their product's attributes, while consumers invest in trustworthy signals to avoid overpaying for lower-quality goods. To mitigate these frictions, sellers may voluntarily disclose evidence—lab reports, provenance certificates, or third-party attestations—although such signals are often only partially verifiable. Symmetrically, buyers can pay for third-party authentication—paid inspection services or platform-provided guarantees—that yield valuable, but possibly noisy, information.

Second-hand e-commerce platforms clearly illustrate these trade-offs. They offer sellers a broader customer base, and the ability to resell used products, while buyers benefit from lower search costs, access to more listings, and competitive pricing. However, the inability to physically inspect used goods beforehand exacerbates information asymmetries. To overcome this, sellers can post photos, ratings, or test reports that vary in credibility, while buyers may subscribe to paid authentication services—ranging from AI checks on item images to in-person laboratory tests—that come with their own fees and error rates.

Despite the prevalence of disclosure and authentication in markets with opaque product quality, key questions remain about how the verifiability of seller-provided evidence and the precision of authentication jointly shape overall relevant economic outcome and, ultimately, affect trade gains and fraud probability. The classic disclosure literature (Grossman 1981; Milgrom 1981; Dye 1985) and the buyer-signal literature (Matthews & Postlewaite 1985) typically analyze only one side of the market: either how sellers reveal information or how buyers pay for signals. Moreover, they often assume costless, perfectly precise verification, thereby sidestepping real-world frictions of noisy tests, limited verifiability, and endogenous pricing of authentication services. Those simplifications leave aside whether imperfect information technologies—available both to sellers and to buyers—can jointly mitigate market inefficiencies.[1]

To fill this gap, our model allows for both seller-provided evidence and buyer-purchased authentication, each of which can range from uninformative to perfectly precise.[2] Given the motivation of the paper, we evaluate equilibrium outcomes along two key dimensions. Total *welfare* is the combined surplus of the agents when trade occurs, and *fraud* probability is the likelihood that a low-quality good sells at a price above its full-information value. An equilibrium is fully efficient if it maximizes total welfare and ensures the probability of fraud is zero.

The equilibrium analysis yields three principal findings. First, market efficiency—via full revelation—is attainable only when the evidence signal is sufficiently precise relative to the quality gap, independently of authentication. Indeed, authentication cannot by itself foster full separation because it only becomes valuable when sellers pool their types. Buyers will pay for costly authentication only if equilibrium pooling leaves them uncertain about quality. Were sellers fully separated, authentication would serve no purpose and its fee would never be paid—so any positive demand for quality testing simply signals that pooling persists, meaning authentication cannot by itself enforce separation. Conversely, when evidence is sufficiently precise it sustains separation by making misrepresentation

---

[1] As a recent exception in a signaling framework, see Bester et al. (2021) on a competitive labor market in which workers signal ability through education and firms audit to learn about their productivity.

[2] By allowing each technology's accuracy to range from totally uninformative to perfectly precise, the model nests and generalizes previous analyses that consider only one side or treat the two sides separately (e.g. Stahl and Strausz, 2017).

too costly, which in turn deters untruthful pricing.

Second, when evidence precision is too low for separation, authentication accuracy tends to reduce fraud probability, but cannot restore full efficiency. Authentication only partially disciplines sellers when it is purchased with positive probability, and higher accuracy strengthens that effect. However, it never fully eliminates pooling: precisely because some sellers pool, buyers find authentication worthwhile. Even perfect testing cannot change this, as the resulting equilibrium still ties welfare gains to reductions in fraud. By contrast, when no authentication is available, welfare is maximized by expanding trade—even if that means more fraud—since total surplus increases with every transaction regardless of its fraudulent content.

Finally, and perhaps most remarkably, whenever evidence is sufficiently imprecise, foregoing authentication can yield higher welfare than perfect authentication. In particular, if the evidence signal is completely uninformative, any form of authentication is strictly detrimental to welfare. This last result reveals a strong complementarity between disclosure and authentication: in the absence of any credible, verifiable evidence, even flawless authentication lowers welfare relative to the no-authentication case—although skipping authentication altogether comes at the cost of a very high fraud probability.

When evidence becomes less precise, both the authentication and no-authentication regimes experience rising fraud alongside higher total welfare, but for distinct reasons. With authentication in place, lower precision makes it easier for low-quality goods to pass the test and spurs more buyers to authenticate, so the incentives to cheat and to reveal truthfully partially offset one another. In contrast, when no authentication is available, neither buyer behavior nor seller strategy adjusts as precision falls, so every additional act of fraud simply generates more trades and thus more surplus. Consequently, beyond a critical level of imprecision, the no-authentication regime yields strictly greater welfare despite—or indeed because of—its higher fraud incidence.

Crucially, these conclusions rest only on the assumption that perfect authentication requires paying a strictly positive price. In particular, they hold even when the actual cost of providing authentication is zero; any positive provision cost of authentication would only widen the welfare advantage of foregoing authentication when evidence is uninformative.

Formally, these results derive from the analysis of an augmented one-seller, one-buyer adverse-selection framework that incorporates both seller disclosure and buyer authentication. The seller, privately informed of her good's quality (high or low), simultaneously chooses a price and an evidence level (high or low). The buyer observes the price, the evidence level, and the resulting evidence signal, and chooses to either buy outright, pay the authentication fee to obtain a potentially noisy verification (and buy only if positive), or walk away. We abstract from any production costs associated with the good, the evidence disclosure process, and the authentication service.

Our framework, like many signaling models, admits multiple equilibria. To focus on the outcomes that rest on reasonable beliefs, we apply Cho and Kreps's intuitive criterion (as extended by Bester and Ritzberger) and further restrict attention to equilibria that are not strictly worse for both seller types. Under these refinements, every surviving equilibrium features the high-quality seller choosing the more precise evidence level, thereby guaranteeing her a payoff above that of the low-quality seller.

In pure-strategy equilibria, full separation arises only when the seller's evidence is sufficiently precise; otherwise, both types pool at a single price and evidence level. When authentication is perfectly accurate, buyers never purchase it in these pooling outcomes, anticipating that low-quality sellers could never safely overprice, buyers optimally save the fee. By contrast, if authentication is both noisy and inexpensive, pooling can coexist with buyers always paying for testing, which tempers but does not eliminate fraud.

Between these extremes lie mixed-strategy or "semi-separating" equilibria, which appear whenever evidence precision is too weak to sustain full separation yet too strong to enforce pure pooling. In these cases the high-type seller always signals at the top tier, while the low-type seller randomizes between mimicking and conceding her type just enough to keep the buyer indifferent. The buyer, unable to fully infer quality, likewise randomizes her decision to test.

A comparison of welfare and fraud across these equilibria yields a clear hierarchy: separating outcomes achieve the first-best, pooling without authentication maximizes trade volume (and hence welfare when separation is impossible) but also fraud, and pooling with testing lies in between. Mixed equilibria perform strictly worse than pooling with authentication on both dimensions, though when evidence is very uninformative, forgoing testing can paradoxically raise welfare even though it allows more fraud. If authentication were costless, perfect testing would restore full efficiency; but under positive fees, it can sometimes reduce social surplus by suppressing fraud-generating trades.

Related literature. Extensive research has been done on information acquisition by uninformed buyers [3]. Jackson (1991) tries to deal with the Grossman and Stiglitz paradox: when buyers are price takers and can acquire information about quality while prices are fully revealing, information acquisition does not occur in equilibrium. He relaxes the price-taking assumption to solve the paradox and get fully revealing prices and costly info acquisition. Bester and Ritzberger (2001) study a trade model with multiple buyers who have the option of perfect information acquisition upon observing the price of the good. They prove the Grossman–Stiglitz Paradox in the equilibrium with pure strategies under the Intuitive Criterion of Cho and Kreps, but show that mixed strategies allow to approach the full information prices. Gertz (2016) analyses a similar model with a single buyer who can choose the amount (precision) of acquired information. He shows that information acquisition can lead to more efficient outcomes, but that there are limits to how much information can be acquired. Martinez-Gorricho (2020) studies the same problem where information acquisition is costless but always imperfect. Voorneveld and Weibull (2011) also study a case with incomplete information on both sides of the market where the buyer does not know the exact quality of the good but receives a private costless signal. Similarly, Figueroa and Guadalupi (2015) study a bilateral trade model where the quality of the good is not known to either but the seller holds a more informative prior about the quality while the buy can acquire a costly but imprecise signal about the quality. Hauswald and Marquez (2006) analyze a screening problem where competing decision-makers can invest in acquiring private information about the agents. Argenziano et al. (2016) study an environment where uninformed decision-makers can acquire private information from a biased expert showing that it is not always more efficient than direct information acquisition by the decision-maker. Persico (2000) studies the acquisition of information about the value of the object for sale in an auction by the bidders and sees in which auction format such acquisition occurs. Martin (2017) studies the pricing strategies when sellers face two types of irrational buyers who can be rationally inattentive and naive and finds that in both cases there are equilibria where information revelation fails.

---

[3] Dranove and Jin (2010) provide a comprehensive review of quality disclosure and certification literature up to 2010.

Similarly, Gabaix et al. (2006) study how the information acquisition process differs in limited rationality models compared to standardly assumed agents with full rationality and test the predictions in a lab experiment. Roesler and Szentes (2017) analyze a bilateral trade environment with take-it-or-leave-it offers where the buyer is uncertain about her valuation of the object but can be signaled about it. Stahl and Strauz (2017) add a third-party certification device available both for the seller and the buyer but not at the same time and find that seller-induced certification is more efficient in terms of total welfare when transparency of the market is beneficial for the welfare. Metthews and Persico (2005) study a bilateral trade of refundable goods, that is, a dissatisfied buyer can be reimbursed for a good if she is not satisfied with the quality. Bergemann and Välimäki (2002) see if information can be acquired efficiently in a mechanism before agents decide to participate in it finding that it is generally achievable in private-value environments but not with common values. Moscarini and Ottaviani (2001) characterize prices and profits in separating, pooling, and mixed-strategy equilibria in a Bertrand competition with the buyer's private information about her valuation when there is a common prior shared by the sellers and the buyer and a private signal only for the buyer. Liu (2011) studies a dynamic game where uninformed buyers can costly learn about sellers' product quality from their past trades. Guan and Chen (2017) analyze a two-sided information asymmetry in the producer-retailer problem: the producer knows the exam quality of the product while the retailer holds a prior about it; the buyer privately knows the preference of consumers for quality. The producer can acquire information (publicly) about the preferences at one cost and disclose information truthfully about the quality at a different cost. For a more detailed review see Capozza et al. (2021) who provide a survey on information acquisition literature across all the subfields of economics. Finally, one of the two closest papers to our model is Bester et al. (2019) who study a competitive labor market application where both the workers can provide a costly (unlike our model) signal of their ability a la Spence while the firms can learn at a cost (audit) the exact ability of the workers. Although their model also endows both sides with costly information technologies, its competitive signaling framework and the nature of the signals differ fundamentally from our bilateral-trade setting with partially verifiable seller disclosure and noisy buyer authentication.

Contribution and Relevance This chapter contributes to the literature on both evidence provision and costly information acquisiton by embedding both seller-provided disclosure and buyer-purchased authentication in a single adverse-selection framework. Unlike earlier analyses that treat these channels in isolation or assume perfectly precise tests, our model lets each technology's accuracy vary from uninformative to perfect and traces out exactly when seller evidence alone can achieve full separation and when—even flawless—authentication cannot substitute for credible disclosure. In particular, while previous work has shown that a costly, perfectly precise authentication can be acquired only with probability less than one (Bester and Ritzberger 2001; Stahl and Strausz 2017), we demonstrate that a fallible authentication technology may in fact be purchased with certainty in pure-strategy equilibria. Moreover, by fully characterizing pooling, separating, and mixed-strategy outcomes under varying combinations of evidence and authentication precision, we pinpoint the precise parameter regions in which each regime prevails and clarify how they discipline fraudulent sellers.

The relevance of this work lies in its normative implications. Our findings make clear that improving the verifiability of seller disclosures is the only reliable route to eliminating fraud, since authentication—even if perfect—cannot on its own guarantee full market efficiency. At the same time, we uncover the surprising fact that when seller evidence is very noisy, perfect testing can actually lower total welfare compared to a no-authentication regime, because suppressing some surplus-generating trades outweighs the gains from reduced fraud. By quantifying how evidence precision and authentication accuracy jointly shape welfare and fraud incidence, our analysis offers concrete guidance for platforms and regulators seeking the optimal balance between verification rigor and trade volume.

The rest of the paper is structured as follows. In Section 2, we present the model and introduce the necessary notation. In Section 3, we describe the optimal behavior of the seller and the buyer. Section 4 provides the equilibrium analysis and the main results. In Section 5, we perform the welfare-fraud analysis under different authentication regimes. Section 6 includes comments and conclusions. All proofs are relegated to the Appendix.

1.2 Model

We consider an asymmetric information model with a buyer and a seller. The latter holds private information over the actual quality of a good. While setting the price, the seller can also provide some evidence about the quality of the product which generates a quality signal prior to the trade. This evidence technology is costless and monotone in quality so that, for a given level of evidence, a high-quality signal is more likely when the actual quality is high than when it is low. We will make this assumption explicit in the next paragraphs.

Upon observing the price and the evidence, the buyer has a few decisions to make. Firstly, she can order, for a fee, a quality authentication from an intermediary to check, at least to some extent, the value of the good. We consider two different types of authentication technology: perfect authentication, which always signals the actual quality, and imperfect authentication, which entails some degree of error. Then, whether she has acquired the authentication or not, she decides if she wants to buy the good.

In an extension of the game, we will consider an augmented game in which a third player, a profit-maximizer authenticator, can commit to an authentication structure, defining both the precision and the price of the authentication, before the seller-buyer subgame takes place.[4]

We now provide a formal outline of the setting and the assumptions of the model.

Preliminaries. We consider a good whose quality can be either low or high. Let $\theta \in \Theta = \{\theta_l, \theta_h\}$ denote the low and high quality, respectively, with $0 < \theta_l < \theta_h < \infty$, and $\Delta\theta := \theta_h - \theta_l$, representing the quality difference between the two types of good. Both the seller and the buyer are risk-neutral. The seller does not value the good per se but is only interested in selling the good at the highest price p possible, while the buyer gets utility from the value of the product minus the price she pays, net of any other additional costs. The reservation utility in the case of non-trade is zero for both players. There is no cost of production for the good of either quality. Furthermore, we normalize both the costs of evidence disclosure and authentication provision to zero. Clearly, efficiency would require that the trade always takes place.

---

[4] Results and proofs will be provided in the definitive version of the paper.

The realization of the quality θ is private information of the seller, while the buyer holds a prior over its distribution, with $\mu_0 := \mathbb{P}(\theta = \theta_h)$. Moreover, the buyer uses Bayesian updating to form her posterior belief about the seller's type.

Seller's strategy and partially verifiable evidence. After observing the quality of the good, the seller needs to choose both a price $p \in \mathbb{R}_+$ and an evidence level $e \in \{e_l, e_h\}$ where e_i represents evidence of the good being of quality θ_i. Allowing for mixed strategies, we define a strategy for the seller as a probability distribution over price-evidence vectors, and we denote by $\sigma_i(p, e) := \sigma(p, e|\theta_i)$ the probability assigned by the seller strategy to the vector (p, e) by the seller of type θ_i. Providing any type of evidence is costless regardless of the actual quality. However, the evidence can, in certain conditions, reveal the type of the seller, or equivalently, be verified.

We model this idea by assuming that the buyer not only observes the (p, e) vector prescribed by the seller strategy but also the realization of a binary (evidence) signal $S^e \in \{S^e_l, S^e_h\}$ whose probability distribution depends jointly on the actual seller's quality θ and on the evidence e provided by the seller. In particular, we denote the conditional probability distributions of S^e by

$$\mathbb{P}(S^e_h|e_h, \theta_h) = \phi \qquad \mathbb{P}(S^e_h|e_h, \theta_l) = \chi$$
$$\mathbb{P}(S^e_h|e_l, \theta_h) = \psi \qquad \mathbb{P}(S^e_h|e_l, \theta_l) = \omega$$

and, accordingly, $\mathbb{P}(S^e_l|e, \theta) = 1 - \mathbb{P}(S^e_h|e, \theta)$ for every e and θ. Without loss of generality, we assume $\phi \in \left[\frac{1}{2}, 1\right]$ so that when the high-type seller provides high evidence, S^e_h is more likely to occur than S^e_l.[5] As stated above, we also assume the evidence technology to be monotone in the actual quality, i.e. $\phi \geq \chi$ and $\psi \geq \omega$. For the same level of evidence, the high-type seller sends the high signal (weakly) more often than the low-type seller.

To simplify the analysis, we assume that when high evidence is provided, the high-type seller always conveys the high signal, while the low-type seller does not. Moreover, we assume the low evidence to be completely uninformative.[6]

---

[5] Otherwise, we could simply relabel the signal realizations.

[6] This last assumption is not substantial, as the main results will hold without it, but it drastically increases readability.

Assumption 1 (Evidence signal)**.** *The conditional distribution of the evidence signal S^e is characterized by:*

(i) $\chi < \phi$ *or, equivalently,* $\chi = \alpha\phi$ *with* $\alpha \in (0, 1)$*,*

(ii) $\phi = 1$*,*

(iii) $\psi = \omega$*.*

We can interpret providing evidence e_l as if the seller does not disclose evidence, meaning that the buyer cannot get any additional information about the good's quality by observing e_l. On the contrary, e_h generates a different probability of a high signal depending on the actual type. In particular, the buyer is able to detect a low seller providing inconsistent high evidence whenever S^e_l realizes, which happens with probability $\mathbb{P}(S^e_l|e_h, \theta_l) = 1 - \alpha$ since the high-quality seller never induces a low signal with high evidence. We will see that the parameter α plays a key role in the analysis.

Beliefs. Given a strategy of the seller (σ_h, σ_l) and an evidence signal realization S^e, the buyer updates his belief about the good being of high quality following Bayes' rule whenever possible. Hence, the posterior beliefs upon observing a price p, given (σ_h, σ_l) are

$$
\begin{aligned}
\mu(p, e_h, S^e_h) &= \frac{\mu_0\sigma_h(p, e_h)\phi}{\mu_0\sigma_h(p, e_h)\phi + (1-\mu_0)\sigma_l(p, e_h)\chi} \\
\mu(p, e_h, S^e_l) &= \frac{\mu_0\sigma_h(p, e_h)(1-\phi)}{\mu_0\sigma_h(p, e_h)(1-\phi) + (1-\mu_0)\sigma_l(p, e_h)(1-\chi)} \\
\mu(p, e_l, S^e_h) &= \frac{\mu_0\sigma_h(p, e_l)\psi}{\mu_0\sigma_h(p, e_l)\psi + (1-\mu_0)\sigma_l(p, e_l)\omega} \\
\mu(p, e_l, S^e_l) &= \frac{\mu_0\sigma_h(p, e_l)(1-\psi)}{\mu_0\sigma_h(p, e_l)(1-\psi) + (1-\mu_0)\sigma_l(p, e_l)(1-\omega)}
\end{aligned}
$$

depending on the evidence provided and its signal realization whenever $\sigma_h(p,e)+\sigma_l(p,e)>0$ for the corresponding (p,e). Given assumption 1, the posterior beliefs simplify to

$$\mu(p,e_h,S^e_h)=\frac{\mu_0\sigma_h(p,e_h)}{\mu_0\sigma_h(p,e_h)+(1-\mu_0)\sigma_l(p,e_h)\alpha} \tag{1.1}$$

$$\mu(p,e_h,S^e_l)=0 \tag{1.2}$$

$$\mu(p,e_l,S^e_h)=\mu(p,e_l,S^e_l)=\frac{\mu_0\sigma_h(p,e_l)}{\mu_0\sigma_h(p,e_l)+(1-\mu_0)\sigma_l(p,e_l)} \tag{1.3}$$

whenever they are defined. Clearly, low evidence e_l is uninformative on the type of the seller since the posterior belief is independent of the signal realization. On the other hand, high evidence can be fully informative whenever the signal realization is low, indication a low-quality seller.

Quality authentication and buyer's strategies. In addition to the evidence provided by the seller, the buyer has available a product authentication technology for a fee $c>0$, which reveals additional information on the quality of the good. In particular, if the buyer purchases authentication, she receives a binary authentication signal $S^a\in\{S^a_l,S^a_h\}$ whose distribution depends only on θ and is independent of the realization of S^e. The conditional distribution of S^a is defined by:

$$\mathbb{P}(S^a_h|\theta_h)=\eta \qquad\qquad \mathbb{P}(S^a_h|\theta_l)=\varepsilon$$

Similarly to the evidence-signal case, we assume, without loss of generality $\eta\in\left[\frac{1}{2},1\right]$, and the authentication technology to be monotone in the good's quality, i.e. $\eta\geq\varepsilon$. Moreover, we limit our analysis to the case in which the authentication of the high-quality good always reveals the actual high type, while we consider two different outcomes when the authentication is ordered on a low-quality good.

Assumption 2 (Authentication signal)**.** *The conditional distribution of the authentication signal S^a is characterized by $\eta=1$ and, alternatively,*

(i) $\varepsilon=0$, in which case we define the authentication technology as perfect.

(ii) $\varepsilon\in(0,1)$, in which case we define the authentication technology as imperfect.

Clearly, when $\varepsilon = 0$, acquiring the authentication fully reveals the type of seller to the buyer. When ε is positive, instead, there is a positive probability that a low-quality good passes the test, sending an authentication signal S^a_h. In this regard, we can consider ε as the authentication imprecision level, or error term.

Given this additional technology, the buyer has three alternatives upon observing (p, e, S^e): buying the good without authentication, action b, not buying the good, action n, acquiring the authentication technology and then buying the good if and only if $S^a = S^h$, action ba.[7] Therefore, given any realization (p, e, S^e) and the corresponding posterior $\mu(p, e, S^e)$, we define a strategy for the buyer as a probability distribution over the set of actions $\{b, n, ba\}$, and we denote by $\beta(x|p, \mu)$ the probability assigned by the buyer's strategy to action $x \in \{b, n, ba\}$.

Timing, payoffs, and solution concept. The game unfolds as follows:

1. Nature chooses the quality of the good θ according to the prior distribution μ_0.
2. The seller observes θ and chooses a probability distribution over vectors (p, e).
3. The evidence signal S^e is realized according to its conditional probability distribution.
4. The buyer observes (p, e, S^e), (possibly) updates her belief according to (1.1), (1.2), and (1.3), and chooses a probability distribution over her available action set $\{b, n, ba\}$.
5. If $x = ba$, the authentication signal S^a is realized.
6. Finally, buyer's payoff u and seller's profits π are realized as follows:
    - If $x = b$, the trade takes place and payoffs are: $u = \theta - p$ and $\pi = p$.
    - If $x = n$, there is no trade and payoffs are: $u = \pi = 0$.
    - If $x = ba$ and $S^a = S^a_h$, the buyer bears the price of the authentication, the trade takes place and payoffs are: $u = \theta - p - c$ and $\pi = p$.

---

[7] As already pointed out in Stahl and Strausz (2017) for a perfect authentication technology environment, any other action, namely, acquiring the authentication and always buying the good, acquiring the authentication and never buying the good, acquiring the authentication and buying if and only if $S^a = S^l$ are dominated by the first three alternatives and hence are disregarded in the analysis.

- If $x = ba$ and $S^a = S^a_l$, the buyer bears the price of the authentication but there is no trade and payoffs are $u = -c$ and $\pi = 0$.

The game can be solved by backward induction and our analysis considers (weak) Perfect Bayesian Equilibria (PBE) which specifies the optimal strategies of the agents and the system of equilibrium beliefs consistent with the strategies and Bayes' rule.

We focus first on pure-strategy equilibria, and then move our focus to a specific class of mixed-strategy equilibria. The procedure adopted to solve the model involves a series of steps. We specify candidate price-evidence strategies for low- and high-type sellers. Then, we determine the equilibrium posterior upon observing the equilibrium price-evidence vector of the buyer and her optimal strategy. This allows us to determine the optimal pricing posted by the two types of sellers in equilibrium. Finally, we check for possible deviations and, potentially, find necessary conditions determining the subset of the parameter space allowing for such equilibria.

1.3 Buyer's Behaviour and Seller's Profits

Let us first analyze what happens in the last stage of the game, when the buyer is called to play. Given any vector realization (p, e, S^e) and seller strategy (σ_h, σ_l), the buyer holds beliefs $\mu(p, e, S^e)$ on the good being of high quality according to (1.1), (1.2), and (1.3). Similarly to Gertz (2014), we define the expected quality of the good given some belief $\hat{\mu}$ as

$$\bar{\theta}_{\hat{\mu}} := \hat{\mu}\theta_h + (1-\hat{\mu})\theta_l$$

and, accordingly, we denote the unconditional average quality given prior μ_0 as $\bar{\theta}_{\mu_0}$.

Assume that the buyer holds posterior $\hat{\mu}$ after observing (p, e, S^e). Then, the expected utility attainable by each action of the buyer is:

$$u(b|p,\hat{\mu}) = \bar{\theta}_{\hat{\mu}} - p \tag{1.4}$$

$$u(n) = 0 \tag{1.5}$$

$$u(ba|p,\hat{\mu}) = \hat{\mu}(\theta_h - p) + (1-\hat{\mu})\varepsilon(\theta_l - p) - c \tag{1.6}$$

While (1.4) and (1.5) are straightforward, let us interpret (1.6). When the buyer buys the good requiring the authentication, the probability of receiving a high authentication signal S_h^a and the trade taking place is 1 when $\theta = \theta_h$, while it is ε when $\theta = \theta_l$. When the authentication technology detects a low quality good, which happens with probability $1-\varepsilon$ when the quality is actually low, there is no trade. Note that we can rewrite $u(ba|p,\hat{\mu})$ as

$$\bar{\theta}_{\hat{\mu}} - p + (1-\hat{\mu})(1-\varepsilon)(p-\theta_l) - c.$$

Comparing the latter to $u(b|p,\hat{\mu})$ highlights the benefit of authentication with respect to an unauthenticated purchase: it permits to avoid a deceiving acquisition of a low-quality good thus saving the amount $p-\theta_l$ whenever the good is perceived to be be of low quality and the authentication recognize it as such, event with probability $(1-\hat{\mu})(1-\varepsilon)$, by paying the authentication price c. Clearly, the authentication value decreases in the precision of the S^a signal and so decreases in ε.

The following lemma characterizes the optimal strategy of the buyer as a function of p and belief $\hat{\mu}$ held after observing a vector (p, e, S^e).

Lemma 1. *Suppose the strategy β is optimal for the buyer given price p and belief $\hat{\mu}$. Then,*

$$\beta(b|p,\hat{\mu}) \geq 0 \iff p \leq \min\left\{\bar{\theta}_{\hat{\mu}},\ \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)}\right\} \tag{1.7}$$

$$\beta(ba|p,\hat{\mu}) \geq 0 \iff p \in \left[\theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)},\ \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right] \tag{1.8}$$

$$\beta(n|p,\hat{\mu}) \geq 0 \iff p \geq \max\left\{\bar{\theta}_{\hat{\mu}},\ \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right\} \tag{1.9}$$

While the proof of the lemma is straightforward, its implications are twofold. First, there exist three price thresholds that determine the optimal strategy of the buyer. Second, there exists an upper bound for the authentication price above which acquiring the authentication is not optimal independently of the good's price.

Corollary 1. *Acquiring the authentication can be optimal for a buyer holding posterior $\hat{\mu}$ only if*

$$c \leq \hat{\mu}(1-\hat{\mu})(1-\varepsilon)\Delta\theta, \tag{1.10}$$

or, equivalently, $\varepsilon \leq 1 - \frac{c}{\hat{\mu}(1-\hat{\mu})\Delta\theta}$.

When the price of the authentication is too high, in a sense specified by the corollary, acquiring the authentication is not optimal independently of the price, and the best of the other two alternatives, simply buying or not buying, yields a higher payoff to the buyer. The threshold for c depends positively on three different elements: i) the quality differential, which measures the opportunity cost of receiving θ_l instead of θ_h; ii) the posterior diffusion, which measures uncertainty on the true quality; iii) the precision of the signal, $1-\varepsilon$, as the authentication technology becomes more valuable to the buyer. Clearly, whenever the uncertainty about the type is resolved, i.e. $\hat{\mu} \in \{0,1\}$, buying the authentication is dominated since the buyer is already fully informed on the quality of the good.

The expression in Corollary 1 determines two thresholds for the buyer's posterior belief. This ensures that purchasing authentication is not strictly dominated, provided the posterior belief falls within that specific interval. Specifically, buying the authentication *can* be optimal for a buyer holding posterior $\hat{\mu}$ only if $\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, where

$$\underline{\mu} = \frac{1}{2} - \sqrt{\frac{1}{4} - \frac{c}{(1-\varepsilon)\Delta\theta}} \qquad \bar{\mu} = \frac{1}{2} + \sqrt{\frac{1}{4} - \frac{c}{(1-\varepsilon)\Delta\theta}} \tag{1.11}$$

The set $\left[\underline{\mu}, \bar{\mu}\right]$ is nonempty if and only if

$$c \leq \frac{(1-\varepsilon)\Delta\theta}{4} =: \bar{c} \tag{1.12}$$

or $\varepsilon \leq \frac{4c}{\Delta\theta} =: \bar{\varepsilon}$. The quantity $\bar{c}$ represents the maximum feasible authentication price, regardless of the good's price posted by the seller. Note that $\bar{c}$ corresponds to the threshold specified in Corollary 1 when $\hat{\mu} = \frac{1}{2}$: the value of authentication for the buyer increases with the uncertainty about the seller's type, reaching its maximum when the posterior probability is one half. Moreover, the expression for $\bar{c}$ reveals the negative relationship between ε and c: there is a trade-off between the imprecision of the authentication signal and its price to ensure that authentication can be potentially optimal. Finally, let us highlight the difference between conditions (1.10) and (1.12): condition (1.10) implies that a generic $\hat{\mu}$ lies in $\left[\underline{\mu}, \bar{\mu}\right]$, while condition (1.12) implies that $\left[\underline{\mu}, \bar{\mu}\right]$ is non-empty. Clearly, the first one implies the second one.

Unless stated otherwise, in what follows, we will consider condition (1.12) to be satisfied, as the opposite would be equivalent to a setting without the authentication option.

With this in mind, the buyer's optimal behavior can be conveniently expressed as a function of her posterior belief as follows.

Lemma 2. *Suppose the strategy β is optimal for the buyer given price p and belief $\hat{\mu}$. Then, (i) if $\hat{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$,*

$$\begin{aligned} \beta(b|p,\hat{\mu}) \geq 0 &\iff p \leq \bar{\theta}_{\hat{\mu}}, \\ \beta(n|p,\hat{\mu}) \geq 0 &\iff p \geq \bar{\theta}_{\hat{\mu}}; \end{aligned}$$

(ii) if $\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$,

$$\begin{aligned} \beta(b|p,\hat{\mu}) \geq 0 &\iff p \leq \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)} \\ \beta(ba|p,\hat{\mu}) \geq 0 &\iff p \in \left[\theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)}, \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right] \\ \beta(n|p,\hat{\mu}) \geq 0 &\iff p \geq \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon} \end{aligned}$$

The lemma underlines the importance of the posterior beliefs in determining the optimal action of the buyer and proves to be extremely convenient in the analysis. In particular, it specifies the buyer's optimal behavior based on price p given a fixed belief $\hat{\mu}$. *Extreme* posterior beliefs, $\hat{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, do not allow for authentication to be optimal, so that the maximum price attainable by the seller is the conditional average quality $\bar{\theta}_{\hat{\mu}}$. When beliefs are *intermediate*, $\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, authentication is possibly optimal and two fundamental prices arises: a price for which the buyer is indifferent between b and ba, $\theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)}$, and another price for which she is indifferent between ba and n, $\frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}$. Let us define those quantities as

$$p_b(\hat{\mu}) := \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)}$$

and

$$p_{ba}(\hat{\mu}) := \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}$$

as the maximum price for which the buyer holding posterior $\hat{\mu}$ buys without authentication and with authentication, respectively, where we make explicit the dependency of these two threshold prices on a generic posterior $\hat{\mu} \in [\underline{\mu}, \bar{\mu}]$. Clearly, both are increasing in θ_h,

θ_l, and $\hat{\mu}$ since all these parameters determine a higher expected quality of the good for sale. However, $p_{ba}(\hat{\mu})$ decreases, while $p_b(\hat{\mu})$ increases, in ε and c since a lower precision and a higher fee decrease the value of the authentication technology for the buyer, making the other two alternatives, b and n more profitable.[8]

Finally, note that $\underline{\mu}$ and $\bar{\mu}$ determine two quantities: $\bar{\theta}_{\underline{\mu}}$ and $\bar{\theta}_{\bar{\mu}}$, i.e., the expected quality of the good given posterior belief $\underline{\mu}$ and $\bar{\mu}$, respectively. In the appendix, we provide the counterpart of Lemma 2, which describes the optimal behavior of the buyer based on her posterior belief given a fixed price. Symmetrically to Lemma 2, when the price is either *sufficently* high or low, $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, acquiring the authentication is never optimal. When, instead, the price is *intermediate*, $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, all the available actions of the buyer can be optimal depending on her posterior belief.[9] In particular, the buyer is indifferent among her available actions only for the specific combinations of price-posterior $(\bar{\theta}_{\underline{\mu}}, \underline{\mu})$ and $(\bar{\theta}_{\bar{\mu}}, \bar{\mu})$. In fact, $p_{ba}(\hat{\mu}) = p_b(\hat{\mu}) = \bar{\theta}_{\bar{\mu}}$ when $\hat{\mu} = \bar{\mu}$, and $p_{ba}(\hat{\mu}) = p_b(\hat{\mu}) = \bar{\theta}_{\underline{\mu}}$ when $\hat{\mu} = \underline{\mu}$.

Consider now a seller with strategy (σ_l, σ_h) and a buyer with belief system μ and strategy β. Choosing a price-evidence vector (p, e) when the object quality is θ, yields the following expected profits:

$$\pi(p, e, \mu|\beta, \theta) = \sum_{S^e \in \{S^e_l, S^e_h\}} \mathbb{P}(S^e|e, \theta)\left[\beta(b|p, \mu(p, e, S^e)) + \mathbb{P}(S^a_h|\theta)\beta(ba|p, \mu(p, e, S^e))\right] p \tag{1.13}$$

which depends on the conditional probabilities of the evidence signal, the probability assigned to b by the buyer strategy, the conditional probability of the high authentication signal times the probability assigned to ba by the buyer strategy, and the price p. Given our assumptions 1 and 2 and the expressions for the posteriors derived in (1.1), (1.2), and

---

[8] Note that $\frac{\partial}{\partial \varepsilon}\left(\frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right) = \frac{(1-\hat{\mu})(c - \hat{\mu}\Delta\theta)}{[\hat{\mu} + (1-\hat{\mu})\varepsilon]^2} \lesseqqgtr 0$, but $\left.\frac{(1-\hat{\mu})(c - \hat{\mu}\Delta\theta)}{[\hat{\mu} + (1-\hat{\mu})\varepsilon]^2}\right|_{\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]} < 0$.

[9] We rely on this to establish the restrictions on out-of-equilibrium beliefs.

(1.3) we have

$$\begin{aligned}
\pi(p, e_h, \mu|\beta, \theta_h) &= \left[\beta(b|p, \mu(p, e_h, S_h^e)) + \beta(ba|p, \mu(p, e_h, S_h^e))\right] p \\
\pi(p, e_l, \mu|\beta, \theta_h) &= \left[\beta(b|p, \mu(p, e_l, S_h^e)) + \beta(ba|p, \mu(p, e_l, S_h^e))\right] p \\
\pi(p, e_h, \mu|\beta, \theta_l) &= \alpha \left[\beta(b|p, \mu(p, e_h, S_h^e)) + \varepsilon\beta(ba|p, \mu(p, e_h, S_h^e))\right] p \\
&\quad + (1-\alpha)\left[\beta(b|p, 0) + \varepsilon\beta(ba|p, 0)\right] p \\
\pi(p, e_l, \mu|\beta, \theta_l) &= \left[\beta(b|p, \mu(p, e_l, S_h^e)) + \varepsilon\beta(ba|p, \mu(p, e_l, S_h^e))\right] p
\end{aligned}$$

The high-quality seller is able to trade the object whenever the buyer decides to acquire it, with or without authentication. Hence her profits correspond to the product of the probabilities of b and ba prescribed by the buyer's strategy β and the selling price p. The profit function of the low-quality seller is more complicated and depends on the kind of evidence e provided. If a low seller provides evidence of high quality, trade takes place in three cases: i) S_h^e realizes, which happens with probability α, and β prescribes to buy the good without authentication with positive probability; ii) S_h^e realizes, which happens with probability α, β prescribes to buy the good along with authentication with positive probability and S_h^a realizes, event with probability ε; iii) S_l^e realizes with probability $1-\alpha$, the low quality of the good is detected and β gives positive probability on buying when μ is equal to 0. If, instead, the low seller provides low-quality evidence, she is able to sell to the buyer when β prescribes to buy without authentication or to buy with authentication and the authentication test fails to detect the low quality good, which happens with probability ε.

By adopting a strategy (σ_h, σ_l) the high and the low types expected profits are then equal to:

$$\pi(\sigma_h|\beta, \theta_h) = \sum_{(p,e)} \sigma_h(p, e)\pi(p, e, \mu|\beta, \theta_h) \qquad \pi(\sigma_l|\beta, \theta_l) = \sum_{(p,e)} \sigma_l(p, e)\pi(p, e, \mu|\beta, \theta_l)$$

Finally, we say that $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ represent a PBE of the game if σ_h^*, and σ_l^* maximizes the expected payoffs of the high and the low-quality seller respectively, β^* maximizes the expected payoffs of the buyer, and the system of belief μ^* is consistent with σ_h^*, σ_l^*, β^* and Bayes' rule whenever possible.

1.4 Equilibrium Analysis

1.4.1 Benchmark: no evidence or authentication

Before we proceed with the fully-fledged analysis with evidence and authentication, let us describe the equilibrium in pure strategies where no such technology is available. Firstly, under complete information, the two types clearly set their prices equal to their respective qualities: with an abuse of notation, we say $\sigma_h(\theta_h) = \sigma_l(\theta_l) = 1$ and the buyer buys whenever the price offered is weakly lower than the observed quality. The trade always takes place and the outcome is efficient.

Under incomplete information, separation is no longer possible and any price weakly below the unconditional expected quality can be sustained in a pooling equilibrium.

Proposition 1. *In equilibrium, under no evidence and no authentication, which is equivalent to* $a = \varepsilon = 1$*, it holds:* $\sigma_h(p^*) = \sigma_l(p^*) = 1$ *with* $p^* \in [\theta_l, \bar{\theta}_{\mu_0}]$ *and posterior belief* $\mu(p) \leq \frac{p-\theta_l}{\Delta\theta}$*, for every* $p > p^*$*.*

Indeed, separation cannot be achieved since the low type will always mimic the high type by setting a higher price. Pooling on a price lower than θ_l or higher than the unconditional average quality would result in forgone sure profits in the first case, and rejection of the offer and zero profits in the second case. On the other hand, any price between these two can be sustained in equilibrium by a system of posterior beliefs that are sufficiently pessimistic.

We now turn to the equilibrium analysis in the original model. First, we analyze equilibria in pure strategy, with both perfect and imperfect authentication technology. Then we considered the same cases when extending the possibility of mixing between strategies.

1.4.2 Pure-strategy equilibria with perfect authentication technology

In the first part of this section, we study the model where the authentication technology acquired by the buyer is perfect, that is $\varepsilon = 0$. When the buyer requires the authentication, a low-quality good is detected with probability one since $\mathbb{P}(S_l^a|\theta_l) = 1$. We begin by characterizing the pure-strategy equilibria and providing the necessary conditions for their existence. We consider first separating and then pooling PBEs. The main result is that authentication is never bought in pure-strategy equilibria. Moreover, the sets of possible equilibrium prices charged by the high-quality seller in separating and pooling-on-high-evidence equilibria are not overlapping.

A note on out-of-equilibrium beliefs: To fully characterize the equilibria of the model, we must specify which restrictions apply to the buyer's belief for price-evidence vectors that are never chosen in equilibrium. First, note that for any vector of price and high evidence (p, e_h), it holds that $\mu(p, e_h, S_l^e) = 0$, regardless of whether (p, e_h) is played with positive probability by either type of seller in equilibrium. Second, by (1.3) we have that $\mu(p, e_l, S_h^e) = \mu(p, e_l, S_l^e)$ for any vector of price and high evidence (p, e_l) since the evidence technology is ineffective when low evidence is provided. Therefore, when discussing out-of-equilibrium beliefs for (p, e) such that $\sigma_h(p, e) + \sigma_l(p, e) = 0$, we will always be referring to $\mu(p, e, S_h^e)$ with $e \in \{e_l, e_h\}$. To ease the notation, we define

$$\mu_p^{e_h} := \mu(p, e_h, S_h^e) \qquad\qquad \mu_p^{e_l} := \mu(p, e_l, S_h^e)$$

as the out-of-equilibrium belief over price-evidence vectors when the evidence provided is high and low, respectively.[10] This means that the buyer evaluates the likelihood of facing a high-quality good off-path given both the price and the evidence provided. This is consistent with the fact that not only the evidence disclosed by the seller but also the price can, in principle, reveal information about the quality of the good, as is standard in the literature.

Finally, let us define equilibrium specific $\bar{\mu}_p^{e_h}$ and $\bar{\mu}_p^{e_l}$ as the maximum out-of-equilibrium belief that does not make (p, e_h) and (p, e_l), respectively, a profitable deviation from the equilibrium strategies for any type of seller. This means that in any equilibrium, it must hold $\mu_p^{e_h} < \bar{\mu}_p^{e_h}$ and $\mu_p^{e_l} < \bar{\mu}_p^{e_l}$ for every out-of-equilibrium price p. In other words, any equilibrium requires out-of-equilibrium beliefs over (p, e_h) and (p, e_l) to be at least as

---

[10] Note that analyzing out-of-equilibrium beliefs is not strictly necessary in a signaling setting such as the one we consider, as it is always possible to sustain a large number of equilibria by defining *sufficiently* pessimistic out-of-equilibrium beliefs. However, specific conditions on out-of-equilibrium beliefs are required for a full characterization of the equilibria and will prove useful when adapting the refinement criterion developed in Bester and Ritzberger (2001) to this setting.

pessimistic as $\bar{\mu}_p^{e_h}$ and $\bar{\mu}_p^{e_l}$, respectively, to be sustainable.

Separating equilibrium

In a separating equilibrium, the two types of sellers provide different price-evidence combinations, which reveal the quality of the good to the buyer. Let (p^h, e^h) and (p^l, e^l) represent the equilibrium vectors for the high-type and low-type sellers, respectively, i.e., $\sigma_h(p^h, e^h) = \sigma_l(p^l, e^l) = 1$.[11] Then, in a separating equilibrium, it must hold that $(p^h, e^h) \neq (p^l, e^l)$ so that $\mu(p^h, e^h) = 1$ and $\mu(p^l, e^l) = 0$. For this and all the following types of equilibrium we present, we relegate the out-of-equilibrium belief requirements to the proofs in the appendix.

A natural candidate for a separating equilibrium is the one in which the high-type seller provides high-quality evidence, and the low-type seller provides low-quality evidence. We will now characterize such separating equilibria.

Proposition 2. *The following strategies of the agents and beliefs of the buyer characterize separating PBEs when* $\varepsilon = 0$*:* $\sigma_h(p^*, e_h) = \sigma_l(\theta_l, e_l) = 1$*,* $\mu(p^*, e_h, S^e) = 1$*,* $\mu(\theta_l, e_l, S^e) = 0$*,* $\forall S^e \in \{S_h^e, S_l^e\}$*,* $\beta(b|p^*, 1) = \beta(b|\theta_l, 0) = 1$*, with* $p^* \in \left(\theta_l, \min\{\theta_l/\alpha, \theta_h\}\right]$*.*

First, a separating equilibrium in which the two types provide truthful evidence always exists whenever the evidence technology is able to detect a lying low-type seller with positive probability, i.e., $\alpha < 1$. The buyer can identify the quality of the good so that the low type is forced to charge a price equal to θ_l. In contrast, any price p^* charged by the high type is sustainable in equilibrium as long as it is incentive-compatible for the low type, given sufficiently pessimistic out-of-equilibrium beliefs on prices above p^*. The incentive compatibility constraint of the low type requires $\pi(\sigma_l|\theta_l) \geq \pi(\sigma_h|\theta_l)$ or

$$\theta_l \geq \alpha p^*, \tag{1.14}$$

where the right-hand side represents the expected utility of the low type deviating to (p^*, e_h): he passes the evidence test and receives p^* with probability α and gets zero otherwise. The more precise the evidence technology, i.e., the smaller α, the larger the

---

[11] Note that e^i refers to the evidence provided by the seller of the i-type in equilibrium, while e_i refers to the i-type of evidence, $i \in \{h, l\}$.

set of prices for the high type that can be sustained in equilibrium. Finally, equilibrium prices can never exceed θ_h since they make not buying strictly dominant for the buyer.

Secondly, authentication is not acquired in equilibrium. This is not surprising since perfectly updated posteriors make buying the authentication strictly dominated, as specified by Lemma 2. The authentication parameters only enter the out-of-equilibrium belief restrictions and do not shape the behavior of either the buyer or the seller. In a sense, the product authentication fails to discipline the seller in a separating equilibrium, and the same results would hold in a setting without the authentication option.

Finally, we left the full interpretation of the out-of-equilibrium belief restrictions for the appendix, but let us highlight two facts. First, since deviating to evidence of the high-type, e_h, is *costly* only for the low type, the general out-of-equilibrium restrictions on $\mu_p^{e_h}$ are solely determined by the no-deviation conditions of the high type and no restrain is necessary for prices below p^*, as they are not profitable for the high type. Second, the restrictions on out-of-equilibrium beliefs when low evidence is provided, are determined by the low type as long as the deviating price is below p^* and by the high type in the remaining case. In general, the low-type finds deviating less appealing than the high type for two reasons: i) it is more costly since she can be detected when providing high-quality evidence, thus getting zero; ii) it is (weakly) less profitable since he cannot propose prices triggering the authentication purchase, which are higher than those for which the buyer simply buys the good, as she would be detected with probability one and the trade would not take place.

The main result is that partially verifiable evidence can (partially) overcome the Grossman-Stiglitz paradox: prices can be perfectly informative of the quality of the good for sale, even if no information is acquired in equilibrium, provided that sellers have the option to voluntarily disclose verifiable information. In fact, a fully informative equilibrium exists as long as the buyer can detect false information about the quality of the good with positive probability.

The next proposition clarifies that the set of separating equilibria characterized in Proposition 2 is not the unique set of such equilibria. Indeed, equivalent separating equilibria in which both types provide e_h are possible given the same restriction on the price posted by the high-quality seller in equilibrium.

Proposition 3. *In any separating equilibrium, the high-type seller provides high evidence, e_h. In particular, the set of separating equilibria is fully characterized by Proposition 2, up to the quality of evidence provided by the low-type seller in equilibrium.*

In summary, the low type is constrained to post a price equal to θ_l while the high type charges a higher price in any separating equilibrium. Therefore, the high type cannot offer low-quality evidence and set a higher price than the low type since the latter would always find it profitable and safe to mimic the high type's strategy. On the other hand, the low type's choice of evidence is irrelevant for the buyer in a separating equilibrium since she can detect the type of the seller independently of the evidence signal realization. The equilibrium conditions and the out-of-equilibrium analysis remain the same when both types of sellers provide e_h.

Pooling Equilibria

We now turn our attention to equilibria in which the two types of sellers provide the same price-evidence vector. Unlike the separating case, the buyer might have the incentive to buy the authentication when purchasing the good in order to avoid disguised low-quality goods. However, the perfectly precise nature of the authentication technology prevents this from happening in equilibrium. We show that this result holds for pooling on high-quality evidence equilibria first, and then we do the same for the pooling on low-quality evidence counterpart.

Pooling on high-quality evidence. First, we consider pooling equilibria in which both types of sellers present evidence of high quality. In this case, the evidence technology sends signal S_h^e with probability one if the seller is of the high type, and with probability α if the seller is of the low type. Bayes' rule then implies that the posterior after observing (p, e_h, S_h^e) is

$$\mu(p, e_h, S_h^e) = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\alpha}$$

which we denote by $\tilde{\mu}$.

Proposition 4. *The following strategies of the agents and beliefs of the buyer represent the set of pooling PBEs on e_h when $\varepsilon = 0$: $\sigma_h(p^*, e_h) = \sigma_l(p^*, e_h) = 1$, $\mu(p^*, e_h, S_h^e) = \tilde{\mu}$, $\mu(p^*, e_h, S_l^e) = 0$, $\beta(b|p^*, \tilde{\mu}) = \beta(n|p^*, 0) = 1$, with $p^* \in \left[\frac{\theta_l}{\alpha}, \bar{p}\right]$, and $\bar{p} = \bar{\theta}_{\tilde{\mu}}$ if $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$*

and $\bar{p} = \theta_l + \frac{c}{1-\tilde{\mu}}$ *if* $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$. *The necessary condition for the existence of such equilibria is* $\alpha \geq \theta_l/\bar{p}$.

Pooling equilibria in which both types send evidence of high quality exist as long as the evidence technology fails to detect a lying low-type seller with a relatively high probability. When providing high-quality evidence, the low type passes the test with probability α so the posterior belief when observing S^e_h equals $\tilde{\mu}$ as defined in the proposition. It is determined by μ_0 and α, and univocally pins down the maximum price $\bar{p}$ that can be set in equilibrium by both types. When $\tilde{\mu}$ is *extreme*, i.e., $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, the buyer either buys the good without authentication or not so the maximum price attainable is the expected quality given posterior μ, i.e. $\bar{\theta}_{\tilde{\mu}}$. Instead, when $\tilde{\mu}$ is *intermediate*, i.e., $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, the threat of the authentication is credible, and the sellers cannot set a price above $\theta_l + \frac{c}{1-\tilde{\mu}}$ in equilibrium, as any price above the threshold would induce an acquisition with authentication or no trade at all, leaving the low type with zero profits in any case.

Differently from the separating case, the equilibrium price p^* is bounded below by the participation constraint of the low type. In fact, the profit of both sellers cannot be lower than θ_l, since the buyer would always find optimal to buy at that price independently of her beliefs. Therefore, $\pi(\sigma_l|\theta_l) \geq \theta_l$ or

$$\alpha p^* \geq \theta_l, \tag{1.15}$$

since the low type gets p^* when $S^e = S^e_h$, which happens with probability α, and 0 otherwise. The necessary condition is that $\alpha \geq \frac{\theta_l}{\bar{p}}$, which means that the probability of the low type passing the evidence test when lying about the quality of the good needs to be sufficiently high, otherwise it would be beneficial to ask for and surely get the lowest price θ_l. Note, however, that $\frac{\partial \bar{p}}{\partial \alpha}$ is negative: clearly, if the chance of facing a low-quality good when observing high evidence is higher, the quality the buyer expects from the good decreases, and so does the maximum price she is willing to pay. There is then a trade-off for the low-type seller between the probability of selling when providing false evidence and getting a higher price in equilibrium. However, since both $\bar{\theta}_{\tilde{\mu}}$ and $\theta_l + \frac{c}{1-\tilde{\mu}}$ are weakly larger than θ_l, there always exists a non-empty set of values for α that guarantees condition (1.15) holds.

The authentication is not acquired in this case either, even when $\tilde{\mu}$ is not extreme. Since the authentication is perfect, a price inducing the acquisition of authentication is not sustainable by the low type in equilibrium. We will see that this need not be the case when the authentication technology is not perfect, meaning it might fail to detect a lying low seller with positive probability. Finally, note that, in this type of equilibria, the buyer always extracts some rent whenever her posterior is intermediate. In fact, the price charged in equilibrium is necessarily lower than the good's expected quality, $\bar{\theta}_{\tilde{\mu}}$. Therefore, the possibility of acquiring authentication is beneficial to the buyer, although indirectly. In this sense, the authentication technology represents a threat and serves to discipline the low-type seller's pricing strategy.

Finally, the restrictions on the out-of-equilibrium beliefs are determined by the no-deviation conditions of the high type when deviating to evidence of the high type, as in the separating case, since deviation is always costlier for the low type. The same applies to deviation to the low-type evidence, except for the prices above αp^* and below the equilibrium price p^*: those prices are appealing only to the low type, so the conditions are less stringent, i.e., beliefs need not be as pessimistic, or low, compared to those prices representing a profitable deviation also for the high-quality seller.

However, pooling on high evidence is not the unique pooling strategy that can arise in equilibrium.

Pooling on low-quality evidence. We characterize the pooling equilibria in which high and low-type sellers provide evidence of low quality.

Proposition 5. *The following strategies of the agents and beliefs of the buyer represent the set of pooling PBEs on* e_l *when* $\varepsilon = 0$*:* $\sigma_h(p^*, e_l) = \sigma_l(p^*, e_l) = 1$, $\mu(p^*, e_l, S^e) = \mu_0$, $\forall S^e \in \{S^e_h, S^e_l\}$, $\beta(b|p^*, \mu_0) = 1$, *with* $p^* \in [\theta_l, \bar{p}]$, *and* $\bar{p} = \bar{\theta}_{\mu_0}$ *if* $\mu_0 \notin \left[\underline{\mu}, \bar{\mu}\right]$ *and* $\bar{p} = \theta_l + \frac{c}{1-\mu_0}$ *if* $\mu_0 \in \left[\underline{\mu}, \bar{\mu}\right]$.

The equilibrium outcome when pooling on the low-quality evidence is similar to the previous pooling case. The main difference is that the evidence technology is ineffective since the low type provides low-quality evidence and cannot be detected. Therefore, the evidence signal does not carry any information, and the posterior upon observing price p^* and low evidence is equal to the prior μ_0. Consequently, the upper bound of the equilibrium price is lower than the pooling on high-evidence case since $\mu_0 < \tilde{\mu}$. In turn, the low-quality seller is never detected, so he enjoys the same expected utility as the high-quality seller. This implies that the out-of-equilibrium conditions are fully determined by the high type, for whom deviations are always less costly than for the low type. Again, perfect precision of the authentication technology prevents its acquisition by the buyer independently of the prior, as in the previous cases, but guarantees a positive rent to the buyer whenever she holds intermediate beliefs.

Notice that, by combining the results of all pure-strategy equilibria, we derive a series of important implications. First, authentication is never acquired in equilibrium since it is always sub-game dominated for the buyer. The intuition is simple: if any equilibrium price induces the buyer to acquire authentication, it is necessarily a price set by the high type, as the low type is always detected by the authentication technology and would always find it profitable to deviate to the lowest price $\theta_l > 0$. This implies that upon observing a price triggering the acquisition of the authentication, the buyer attributes probability one to the good being of high quality and thus finds it optimal to buy the good without authentication. Nonetheless, the authentication serves the role of disciplining the low type from choosing a higher price in both in and off-equilibrium paths, and benefits the buyer whenever the equilibrium posterior is not extreme.

Second, for any given set of model parameters, the intersection of the set of prices that can be posted by the high type in a separating equilibrium and the set of prices that can be posted by the high type in a pooling equilibrium on the high evidence is always empty. In fact conditions (1.14) and (1.15) are mutually exclusive so that *low* prices can be sustained only in separating equilibria, and *high* prices only in separating equilibria, regardless of the restrictions on out-of-equilibrium beliefs.[12] The reason is that the incentive to deviate from a separating equilibrium for a low-quality seller increases together with the price proposed by the high type, while the profitability of providing high evidence in a pooling equilibrium decreases with it. This potentially explains why providing verifiable evidence is more effective in markets with moderately expensive objects with a relatively little quality variation (e.g. clothes in online re-sale platforms), compared to markets of more valuable goods (e.g. used cars, etc.) where other separating mechanisms are necessary.

---

[12] Here, low (high) prices are defined as those prices that are below (above) the threshold $\frac{\theta_l}{\alpha}$.

Third, two preliminary welfare measures to rank equilibria naturally arise.[13]. Since trade is always optimal in our setting, we can first rank the different equilibrium types by the probability that trade takes place. However, trades are not all equivalent, as the distribution of the surplus between seller and buyer can vary significantly. In particular, we focus on those equilibria generating a negative utility for the buyer, i.e., *fraudulent* equilibria. Following Martinez-Gorricho (2020), we define a *fraud* as the act of selling a good at a price which exceeds the buyer's willingness to pay for it under full information. In our setting, a fraud occurs whenever the low-quality good is sold for a price above θ_l.[14] When authentication is perfectly precise, the ex-ante probability of trade is equal to 1 for separating and pooling on low-evidence equilibria, while it is equal to $\mu_0 + (1 - \mu_0)\alpha$ for pooling on high-evidence equilibria. Moreover, the ex-ante probability of fraud is equal to 0 for separating equilibria, $(1 - \mu_0)\alpha$ for pooling on high-evidence equilibria, and $1 - \mu_0$ for pooling on low-evidence equilibria.

Separating and pooling on the low-type equilibria always lead to trade since the low type does not provide evidence of high quality, and the buyer does not acquire the authentication. Instead, when a low-quality seller provides high evidence in a pooling equilibrium, the evidence signal indicates a low-quality seller with probability $1 - \alpha$, so that the transaction only occurs if the seller is of the high type or the evidence technology fails to detect a lying low type, which happens with probability $\mu_0 + (1 - \mu_0)\alpha$. However, pooling on the low-evidence equilibria induce a fraudulent outcome whenever the seller is of the low type, so with probability $1 - \mu_0$, while pooling on the low-evidence equilibria are frauds with probability $(1 - \mu_0)\alpha$. Clearly, separating equilibria are optimal from a social point of view when taking into account the fraudulent nature of an equilibrium outcome.

The first important implication of the results in this section is that achieving separation is generally possible if the evidence technology is relatively precise, or if the quality dispersion of the products for sale is not too high. Separation is optimal, and any enhancement in the verifiability of the information provided by the seller can potentially increase welfare. This is because it increases the set of parameter that allows for separating equilibria and reduces the set allowing for pooling equilibria. Nonetheless, the model

---

[13] We will study the welfare implications of the model in depth in Section 5

[14] This is the case since there cannot be an equilibrium price above θ_h; thus, the high-quality good can never be the object of a fraud.

suffers from the usual multiplicity of equilibria common in the signaling-games literature. In particular, the equilibria in which both types pool on low-quality evidence seem unrealistic since it is costless for the high type to provide high-quality evidence and signal the high quality of the good to the buyer. Therefore, drawing punctual welfare implications would benefit from a reduction of the possible equilibria by means of an appropriate refinement criterion. We will provide insights on this in the next sections.

The second implication relates to the impossibility result regarding the acquisition of authentication in equilibrium, which seems rather counterintuitive. Why should the authentication option be offered if it is never purchased? Who offers the authentication and why? In the next section, we show that authentication can possibly be bought in equilibrium if the authentication technology is imperfect.

1.4.3 Pure-strategy equilibria with imperfect authentication technology

The main result when considering a setting in which the authentication technology is not perfect, i.e. $\varepsilon > 0$, is that the authentication can be acquired in equilibrium with positive probability. Purchasing the authentication remains a dominated action in separating equilibria even in this context, since the equilibrium beliefs are degenerate. However, pooling equilibria on high (and low) evidence are characterized by two different outcomes. The first one is the imperfect-technology counterpart of the pooling equilibria defined in Proposition 4 (and 5) in which the buyer buys the good without authentication. The second outcome, instead, allows for authentication acquisition in equilibrium.

Proposition 6. *Any pure-strategy PBE under imperfect authentication, i.e. $\varepsilon > 0$, belongs to one of the following categories:*

(i) ***Separating****, where $\sigma_h(p^*, e_h) = \sigma_l(\theta_l, e^l) = 1$, $e^l \in \{e_l, e_h\}$, $\beta(b|p^*, 1) = \beta(b|\theta_l, 0) = 1$, and $p^* \in (\theta_l, \min\{\theta_l/\alpha, \theta_h\}]$.*

(ii) ***Pooling(e_h)-b****, where $\sigma_h(p^*, e_h) = \sigma_l(p^*, e_h) = 1$, $\beta(b|p^*, \tilde{\mu}) = 1$, and $p^* \in [\theta_l/\alpha, \bar{\theta}_{\tilde{\mu}}]$ if $\tilde{\mu} \notin [\underline{\mu}, \bar{\mu}]$, and $p^* \in [\theta_l/\alpha, p_b(\tilde{\mu})]$ if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$; The necessary condition for existence is $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ if $\tilde{\mu} \notin [\underline{\mu}, \bar{\mu}]$, and $\alpha \geq \theta_l/p_b(\tilde{\mu})$ if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$.*

(iii) ***Pooling(e_h)-ba****, where* $\sigma_h(p^*, e_h) = \sigma_l(p^*, e_h) = 1$, $\beta(ba|p^*, \tilde{\mu}) = 1$, $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$, *and* $p^* \in [\max\{p_b(\tilde{\mu}), \theta_l/(\alpha\varepsilon)\}, p_{ba}(\tilde{\mu})]$*; The necessary conditions for existence are* $\alpha \geq \theta_l/(\varepsilon p_{ba}(\tilde{\mu}))$*, and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$.

(iv) ***Pooling(e_l)-b****, where* $\sigma_h(p^*, e_l) = \sigma_l(p^*, e_l) = 1$, $\beta(b|p^*, \mu_0) = 1$*, and* $p^* \in [\theta_l, \bar{\theta}_{\mu_0}]$ *if* $\mu_0 \notin [\underline{\mu}, \bar{\mu}]$*, and* $p^* \in [\theta_l, p_b(\mu_0)]$ *if* $\mu_0 \in [\underline{\mu}, \bar{\mu}]$.

(v) ***Pooling(e_l)-ba****, where* $\sigma_h(p^*, e_l) = \sigma_l(p^*, e_l) = 1$, $\beta(ba|p^*, \mu_0) = 1$, $\mu_0 \in [\underline{\mu}, \bar{\mu}]$*, and* $p^* \in [\max\{p_b(\mu_0), \theta_l/\varepsilon\}, p_{ba}(\mu_0)]$*. The necessary condition for existence are* $\varepsilon \geq \theta_l/p_{ba}(\mu_0)$*, and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$.

Separating equilibria, as characterized in point (*i*) are always possible under imperfect authentication because neither the authentication's precision nor its price plays any role when posterior beliefs perfectly determine the seller's type. Hence, the outcome remains identical to the perfect authentication technology case (i.e., $\varepsilon = 0$): the low type charges the minimum price θ_l regardless of the evidence provided, while the high type charges a larger price p^* which is weakly lower than θ_l/α, preventing any profitable deviation of the low type. The buyer recognizes the quality of the good and buys it without authentication.

Points (*ii*) and (*iv*) generalize the results of Proposition 4 and 5. Pooling equilibria in which the buyer buys without authentication are always feasible when pooling on e_l. In this case, the posterior equals the prior μ_0 and the equilibrium price is bounded above by $p_b(\mu_0)$ or $\bar{\theta}_{\mu_0}$. An equivalent outcome is attainable when sellers pool on e_h. Pooling on e_h triggers the evidence technology so that the low type is not detected with probability α. The posterior is then $\tilde{\mu} > \mu_0$, and the common price is constrained by an upper bound that depends on $\tilde{\mu}$. Specifically, if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$, the price cannot exceed $p_b(\tilde{\mu})$, the price for which the buyer is indifferent between buying with and without authentication when holding posterior $\tilde{\mu}$; otherwise, it cannot exceed $\bar{\theta}_{\tilde{\mu}}$, the expected quality given posterior $\tilde{\mu}$ which corresponds to the price for which the buyer is indifferent between buying and not buying when holding posterior $\tilde{\mu}$ for which authentication is domiated. However, to prevent the low type from deviating to a safe price of θ_l, the equilibrium price must be weakly above θ_l/α.

The main implication following the proposition is that authentication can be bought in equilibrium when the authentication precision is not perfect and when the posterior beliefs is intermediate, i.e., it belongs to $[\underline{\mu}, \bar{\mu}]$. On the one hand, the value of the authentication for the buyer decreases since the informativeness of the authentication is lower. On the other hand, the low type may be willing to face the authentication process since the equilibrium price when the buyer acquires the authentication is higher than the price without authentication for the same posterior, since $p_{ba}(\mu) > p_b(\mu)$ for any $\mu \in (\underline{\mu}, \bar{\mu})$. In particular, an equilibrium price is sustainable provided it is high enough to prevent the low-quality seller from deviating, yet low enough so as not to deter the buyer from acquisition. Points (*iii*) and (*v*) specify the novel outcome under imperfect technology highlighting the distinct pricing requirements and the different posterior based on the evidence provided by the two types. If sellers pool on e_h, the equilibrium price cannot be lower than $\theta_l/(\alpha\varepsilon)$ since the low type faces the scrutiny of both the evidence and the authentication technology and receives the common price with probability $\alpha\varepsilon$. When pooling on e_l instead, only the authentication test binds and the equilibrium price needs to weakly exceeds θ_l/ε.

Proposition 6 implies that equilibria in which no test takes place are feasible for any parameter configuration; these include separating and pooling(e_l)-b equilibria. In contrast, when either the authentication or the evidence technology is triggered—as in pooling(e_h)-ba, pooling(e_h)-b, and pooling(e_l)-ba equilibria—certain necessary conditions must be met for existence. These conditions, specified in Proposition 6, guarantee that the set of equilibrium prices for these equilibria is nonempty. In particular, the conditions can be specified as requirements that the parameters α and ε be sufficiently large to prevent the low-type seller from deviating to a price—and, consequently, obtaining a payoff—of θ_l, which is always attainable by both types. In fact, both technologies must be sufficiently imprecise for the low-quality seller to be willing to face them, potentially enabling her to secure the common price $p^* > \theta_l$.

Let us focus on the pooling(e_h) equilibria.[15] As will become clear in the next sections, it is convenient to express the necessary conditions for each equilibrium in terms of the parameter α. The proof of Proposition 6 includes the explicit expression for the lower bound for α necessary for pooling(e_h)-b equilibria. For $\tilde{\mu} \notin [\underline{\mu}, \bar{\mu}]$ the condition is the same as for pooling on e_h when $\varepsilon = 0$ (Proposition 4), while, if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$, it is equivalent to the perfect technology case once we account for the general definition of $p_{ba}(\mu)$. The

---

[15] We will see in the next section that pooling(e_l) equilibria do not survive the out-of-equilibrium belief restriction we impose.

result is that for any combination of the other parameters, is it always possible to find an α for which the necessary condition for pooling(e_h)-b is satisfied. However, the same result does not hold true for pooling(e_h)-ba equilibria.

Corollary 2. *Under imperfect authentication, i.e., $\varepsilon > 0$, pooling(e_h)-ba equilibria exists only if $c \leq \bar{c}_{ba}$ where*

$$\bar{c}_{ba} = \min\left\{\tilde{\mu}(1-\tilde{\mu})(1-\varepsilon)\Delta\theta, \tilde{\mu}\theta_h - \frac{\tilde{\mu}+(1-\tilde{\mu})(1-\alpha\varepsilon)\varepsilon}{\alpha\varepsilon}\theta_l\right\}. \tag{1.16}$$

The corollary provides the necessary condition for pooling(e_h)-ba equilibria in terms of c. On the one hand, the authentication price negatively affects the maximum price that can be posted so that the good is bought together with the authentication in equilibrium, $p_{ba}(\tilde{\mu})$. On the other hand, it negatively affects the range of posterior beliefs for which authentication is not dominated, $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$, since $\partial\underline{\mu}/\partial c > 0$ and $\partial\bar{\mu}/\partial c < 0$. The first term of the min function defining $\bar{c}_{ba}$ ensures $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$ so that authentication is potentially optimal, while the second term guarantees that α is greater than $\theta_l/\varepsilon p_{ba}(\tilde{\mu})$, so that the low type has no incentive to deviate. The non-negativity condition on the authentication price c requires that both thresholds be weakly positive. Although the first threshold always satisfies this condition, the second does not, so it is not possible to guarantee the existence of pooling(e_h)-ba equilibria in general. However, we prove that, given any set of parameters $(\mu_0, \theta_l, \theta_h)$, it is always possible to find some pair (α, ε) such that $\bar{c}_{ba}$ is positive, which implies that both the necessary conditions for pooling(e_h)-ba hold.

Proposition 7. *For any given primitives $(\mu_0, \theta_l, \theta_h)$, there exists a pair $(\alpha, \varepsilon) \in (0,1)^2$ such that $\bar{c}_{ba} > 0$.*

A pooling(e_h)-ba equilibrium is not sustainable for all parameter combinations. However, we proved that given any $(\mu_0, \theta_l, \theta_h)$ there always exists $\bar{c}_{ba} > 0$, such that for any $c \in [0, \bar{c}_{ba}]$ we can find some combination of valid α and ε for which that type of equilibrium is feasible. Note also that the first term of the min function in the definition of $\bar{c}_{ba}$, ensures $\bar{c}_{ba} < \bar{c}$.

The authentication price plays a pivotal role in making authentication acquisition optimal. All else being equal, lower authentication prices increase its value for the buyer, thereby increasing the instances in which it is chosen in equilibrium, and enabling the seller to charge higher prices that induce authentication acquisition. However, a low authentication price is not sufficient to ensure authentication acquisition, but it relaxes the requirements on the imprecision level of the two tests, α and ε.

The main implication of this section is that when the authentication's signal is imperfect, authentication may not only serve to discipline the low type from choosing higher prices when posterior beliefs are not extreme—pooling(e_h)-b and pooling(e_l)-b equilibria—but can also be acquired in equilibrium—as in the pooling(e_h)-ba and pooling(e_l)-ba equilibria—, thereby revealing low types misrepresenting the quality of their good with a positive probability equal to $1-\varepsilon$. Nonetheless, the latter type of outcome can be sustained only if the signals of the two tests are sufficiently imprecise so as not to incentivize the low type to deviate toward safe option, θ_l. Therefore, the welfare implications are not obvious. Moreover, without stronger requirements on the out-of-equilibria beliefs and the equilibrium prices, it is not simple to draw welfare conclusions given the multiplicity of equilibria that arises in this setting. However, we can extend the conclusions from the previous section, observing that the three equilibrium types under $e=0$, i.e., separating, pooling(e_h) and pooling(e_l), fall within the definition of their direct counterpart under $e>0$, i.e, separating, pooling(e_h)-b and pooling(e_l)-b, once we account for the general definitions of $p_{ba}(\mu)$ and $p_b(\mu)$.

Under imperfect authentication technology, the ex-ante probability of trade is equal to one for separating and pooling(e_l)-b equilibria, $\mu_0+(1-\mu_0)\alpha$ for pooling(e_h)-b equilibria, $\mu_0+(1-\mu_0)\varepsilon$ for pooling(e_l)-ba equilibria, and $\mu_0+(1-\mu_0)\alpha\varepsilon$ for pooling(e_h)-ba equilibria. Moreover, the ex-ante probability of fraud is equal to 0 for separating equilibria, $(1-\mu_0)\alpha\varepsilon$ for pooling(e_h)-ba equilibria, $(1-\mu_0)\varepsilon$ for pooling(e_l)-ba equilibria, $(1-\mu_0)\alpha$ for pooling(e_h)-b equilibria, and $1-\mu_0$ for pooling(e_l)-b equilibria.

Regardless of the imprecision of the authentication signal, separating equilibria maximize the probability of trade and preclude any fraud, thanks to their fully revealing nature. Consequently, they represent the preferred outcome on both fronts. In contrast, pooling equilibria exhibit an obvious trade-off between the probability of trade and the probability of legitimate exchange: any instance in which nature selects the low type constitutes a fraud attempt, since the common equilibrium price is always strictly higher than θ_l. The higher the probability that this attempt is successful, the higher both the probability of trade and the incidence of fraud. Therefore, depending on the welfare criterion of interest, we could be interested in ranking each type of pooling equilibrium based on the precision (or imprecision) of the evidence and/or authentication signals. Note, however, that each equilibrium type may have necessary conditions that are not always satisfied, which makes direct comparisons difficult. We will provide a full welfare analysis in Section 5.

Overall, extending the analysis to the imperfect authentication case provides new insights into the possible outcomes that can arise in equilibrium. When the authentication signal is not perfectly precise, buying authentication can be an equilibrium outcome when the sellers pool to a common equilibrium. All the results from the special case $\varepsilon = 0$ still hold, so that separating equilibria are always possible and remain the most desirable both in terms of probability of trade and probability of fraud. In addition, when imperfect, authentication acquisition can actually improve the outcome in terms of fraud incidence whenever represents a feasible equilibrium otcome.

Despite some sharp results, the multiplicity of equilibria and the indeterminacy of equilibrium prices drastically reduce the model's predictive power. This issue, typical of signaling games, is further exacerbated by the large parameter space under consideration. However, some equilibrium outcomes are sustained by rather implausible out-of-equilibrium beliefs. To enhance the robustness of our predictions, we rely on two different refinement criteria, which allow us to limit equilibrium indeterminacy and perform a more thorough welfare analysis.

1.4.4 Refinement Criteria

In order to narrow down the set of possible equilibria, we rely on two distinct refinement criteria. The first restricts the buyer's out-of-equilibrium beliefs to prevent implausible outcomes from being sustained by extreme and unreasonable beliefs. The second selects only those equilibrium outcomes that are Pareto optimal for the sellers both within each equilibrium type and across different equilibrium types (among those that are feasible for a given set of model parameters). Combined, these two criteria ensure a unique equilibrium outcome for most parameter configurations.

First, we adapt Bester and Ritzberger (2001)'s extension of the Cho and Kreps (1987) intuitive criterion to intermediate posterior beliefs that are not in $\{0, 1\}$ (as also used in Stahl and Strauz (2017)) by accounting for the richer strategy space available to the seller in our model, which includes both price and quality evidence. Since the definition of PBE is silent regarding out-of-equilibrium beliefs, many equilibrium outcomes—some involving rather implausible pricing strategies—can be sustained in signaling games by assuming sufficiently pessimistic beliefs about unplayed strategies. The following refinement limits this arbitrariness by disallowing somewhat *unjustified* pessimism, i.e., overestimating the likelihood of the seller being of the low type.

Assumption 3 (Belief restriction)**.** *A PBE* $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ *satisfies the belief restriction if, for any posterior* $\hat{\mu} \in [0, 1]$ *and any out-of-equilibrium price-evidence pair (p,e),* $\pi(p, e, \hat{\mu}|\beta^*, \theta_h) > \pi(\sigma_h^*|\beta^*, \theta_h)$ *and* $\pi(p, e, \hat{\mu}|\beta^*, \theta_l) < \pi(\sigma_l^*|\beta^*, \theta_l)$ *implies* $\mu^*(p, e, S_h^e) \geq \hat{\mu}$.

The intuition behind the belief restriction is as follows: if, under belief $\hat{\mu}$, deviating to a price-evidence pair (p, e) yields a profitable deviation relative to the equilibrium payoff only for the high type, then the actual belief should not fall below $\hat{\mu}$, meaning it should not be even more pessimistic. The next lemma provides a simple criterion to determine which equilibria do not survive the belief restrictions in terms of equilibrium payoffs of the two types of sellers.

Lemma 3. *A PBE exhibiting equilibrium payoff* π_h^* *for the high type and* π_l^* *for the low type does not survive the refinement criterion of Assumption 3 if and only if at least one of the following holds:*

(i) $\pi_h^* < \min\left\{\dfrac{\pi_l^*}{\alpha}, \theta_h\right\}$*;*

(ii) $\pi_h^* \in [\theta_l, \bar{\theta}_{\underline{\mu}}]$ *and* $\dfrac{\pi_l^*}{\alpha\varepsilon} > \bar{\theta}_{\underline{\mu}}$*;*

(iii) $\pi_h^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ *and* $\dfrac{\pi_l^*}{\alpha\varepsilon} > \pi_h^*$*.*

In addition to providing a sharp criterion for determining which equilibrium does not survive the belief restrictions, the lemma offers important insights regarding the equilibrium payoffs of the two seller types. First, the high type can exploit the evidence technology to secure a higher payoff than the low type, and the gap between their payoffs widens as α decreases. The belief restrictions underscore the role of evidence as an efficient signaling device, since any price below π_l^*/α is unprofitable for the low type (when required to provide high evidence). It follows that, in equilibrium, the high type must always choose high evidence (since it is costless for her) to secure a superior profit. Moreover, all equilibria in which the high-quality seller's payoff is at least $\bar{\theta}_{\bar{\mu}}$, survive, because $\bar{\theta}_{\bar{\mu}}$ represents the maximum price at which authentication acquisition can be optimal, regardless of the buyer's beliefs. Consequently, for prices above this threshold, the high type cannot further leverage the authentication technology to obtain additional profits by demonstrating that they are the only type willing to undergo authentication scrutiny. Instead, if the high type's profit falls below $\bar{\theta}_{\bar{\mu}}$, then either both types obtain relatively low payoffs—making deviations to prices that trigger authentication also profitable for the low type—or the high type's profit exceeds that of the low type by such a margin that any price in $[\pi^*, \bar{\theta}_{\bar{\mu}})$ becomes a profitable deviation despite the authentication test.

The next proposition characterizes the equilibrium categories specified in Proposition 6 that survives the belief restriction assumption.

Proposition 8. *Any pure-strategy PBE satisfying the belief restriction (Assumption 3) belongs to one of the following categories:*

- *(i)* ***Separating***, $p^l = \theta_l$, *either* $p^* = \theta_l/\alpha$, *with necessary condition* $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, *or* $p^* = \min\{\theta_l/\alpha, \theta_h\}$, *with necessary condition* $\alpha \leq \theta_l/\bar{\theta}_{\bar{\mu}}$.
- *(ii)* ***Pooling(e_h)-b***:
  - *(a) if* $\tilde{\mu} \in (0, \underline{\mu})$, $p^* \in [\theta_l/\alpha, \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\}]$, *with necessary condition* $\alpha \geq \max\{\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$.
  - *(b) if* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$, $p^* \in [\theta_l/\alpha, \varepsilon\bar{\theta}_{\underline{\mu}}]$, *with necessary condition* $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$.
  - *(c) if* $\tilde{\mu} \in [\bar{\mu}, 1)$, $p^* \in [\theta_l/\alpha, \varepsilon\bar{\theta}_{\underline{\mu}}]$, *with necessary condition* $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, *or* $p^* \in [\max\{\theta_l/\alpha, \bar{\theta}_{\bar{\mu}}\}, \bar{\theta}_{\tilde{\mu}}]$ *with necessary condition* $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$.
- *(iii)* ***Pooling(e_h)-ba***, $p^* \in [\max\{p_b(\tilde{\mu}), \theta_l/(\alpha\varepsilon)\}, p_{ba}(\tilde{\mu})]$, *with necessary conditions for existence* $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$, *and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$.

Equilibrium strategies are those provided in Proposition 6.

The proposition highlights three main implications of the out-of-equilibrium belief restriction. First, pooling(e_l) equilibria do not survive the refinement criterion. These equilibria require counterintuitively pessimistic beliefs for out-of-equilibrium combinations of high evidence and prices slightly above the common equilibrium price, p^*. In pooling(e_l) equilibria, both types obtain a payoff equal to the common equilibrium price p. When providing high evidence, only prices above p^*/α represent a profitable deviation for the low type, while any price above p would be profitable for the high type, as she is not deterred by the evidence technology. Therefore, pooling(e_l) equilibria require *unjustified* pessimistic beliefs over pairs (p, e_h) for $p \in (p^*, p^*/\alpha)$, and thus they do not survive the belief restriction. In equilibrium, the high type always provides costless high-quality evidence.

Moreover, the out-of-equilibrium belief requirement tends to reduce price indeterminacy and, consequently, imposes much stricter necessary conditions for the existence of each equilibrium. Under the belief refinement, the equilibrium price of separating equilibria is uniquely determined and equals the minimum between θ_l/α and θ_h. In fact, any price below θ_l/α would require unjustified pessimistic beliefs, since such prices represent a profitable deviation only for the high type when high evidence is provided. Furthermore, the restriction on price in separating equilibria is even stronger: not only is the price uniquely determined, but it must also lie either above $\bar{\theta}_{\bar{\mu}}$ or below $\varepsilon\bar{\theta}_{\underline{\mu}}$. This is because any price $p \in (\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ can be sustained in equilibrium only by pessimistic out-of-equilibrium beliefs over prices in $(p, \theta_l/(\alpha\varepsilon))$—prices that represent a profitable deviation for the high type and not for the low type when they trigger authentication acquisition. A similar argument implies that the minimum threshold for p^* to be $\varepsilon\bar{\theta}_{\underline{\mu}}$. Therefore, under the belief restriction, separating equilibria are possible only for extreme values of α, that is, for $\alpha \notin [\theta_l/\bar{\theta}_{\bar{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}]$. On the one hand, small α is required to allow the high type to receive a payoff larger than the maximum price for which authentication can be optimally purchases; on the other large α limit the high type payoff so that any price within the authentication range are attractive for the low type as well.

Finally, while all pooling(e_h)-ba equilibria survive the refinement criterion, Pooling(e_h)-b equilibria are strongly limited by it: common equilibrium prices must be above θ_l/α, to prevent deviations of the low type to θ_l, and below $\varepsilon\bar{\theta}_{\underline{\mu}}$ whenever $\tilde{\mu} \in (0, \bar{\mu})$, following the same reasoning as for separating equilibria. When $\tilde{\mu} \in [\bar{\mu}, 1)$, higher prices can be sustained since, for posterior beliefs above $\bar{\mu}$, authentication acquisition is never optimal and out-of-equilibrium restrictions for both the high and the low type coincide.

Overall, the out-of-equilibrium belief refinement narrows the set of possible equilibria to those that highlight many interesting features of the model, eliminating many counterintuitive outcomes and significantly reducing multiplicity. Note that all surviving equilibria characterized in Proposition 8 depend on certain necessary conditions involving α and, possibly, $\tilde{\mu}$—the posterior belief held by the buyer after observing (p^*, e_h, S_h^e). [16] For example, when $\alpha \leq \theta_l/\theta_h$ the unique outcome is separating, regardless of the posterior; the same is true for pooling(e_h)-ba when $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/Pba(\tilde{\mu})]$ and $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$. However, for other combinations of parameters, the belief restriction is not sufficient to uniquely determine equilibrium outcome. For example, if $\tilde{\mu} \in (0, \underline{\mu})$ and $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\}$ then both pooling(e_h)-b and separating are feasible outcomes. Similarly, if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$ and $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ all surviving outcomes are possible. Moreover, price indeterminacy still affects pooling equilibria since the belief restriction is not able to pin down exact pricing strategies when prices lie (far) outside of the authentication range, as in b-equilibria, or all prices above the common equilibrium price inducing authentication represent a profitable deviation for both types, as in ba-equilibria.

Given these considerations, we adopt an additional refinement criterion to strengthen the model's predictive power and reduce price indeterminacy. The rationale is that some equilibrium outcomes, whether within or across categories, appear counterintuitive. As a motivating example, consider the case in which all equilibrium outcomes are technically possible—that is, when $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$ and $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. However, the profits for each type differ across these outcomes. In a separating equilibrium, the low type is forced to accept θ_l while the high type earns at most $\varepsilon\bar{\theta}_{\underline{\mu}}$; in a pooling($e_h$)-b equilibrium, the low type does slightly better, achieving a profit that is weakly higher than θ_l but still no more than $\alpha\varepsilon\bar{\theta}_{\underline{\mu}}$, while the high type remains capped at $\varepsilon\bar{\theta}_{\underline{\mu}}$. By contrast, pooling($e_h$)-ba equilibria guarantee strictly larger payoffs for both seller types, with $\pi_l^* \geq \alpha\varepsilon p_b(\tilde{\mu}) > \alpha\varepsilon\bar{\theta}_{\underline{\mu}}$ and $\pi_h^* \geq p_b(\tilde{\mu}) > \bar{\theta}_{\underline{\mu}}$. Moreover, suppose $\tilde{\mu} \in [\bar{\mu}, 1)$ and $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. In this case, two outcomes are possible within the pooling(e_h)-b category, which constitutes the unique feasible equilibrium in that specific parameter region. In both outcomes, the high type

---

[16] The second observation rests on the fact that for any equilibrium outcome surviving the belief restriction, either the buyer holds a perfect posterior—as in separating equilibria or when she observes (p, e_h, S_l^e) in a pooling(e_h) equilibrium (thereby recognizing a cheating low type), so that the prior plays no role—or the buyer holds a posterior equal to $\tilde{\mu}$, as when she observes (p, e_h, S_h^e) in a pooling(e_h) equilibrium (since pooling(e_l) equilibria do not survive the belief restriction).

receives p^* and the low type receives αp^*; however, in one outcome the common price p^* lies in $[\theta_l/\alpha, \varepsilon\bar{\theta}_{\underline{\mu}}]$, while in the other it is at least as high as $\bar{\theta}_{\bar{\mu}}$, yielding a necessarily larger payoff. The following refinement restricts the analysis to equilibria that are not strictly payoff-dominated from the perspective of the sellers.

Assumption 4 (Seller-optimality restriction)**.** *A PBE satisfies the seller-optimality criterion if there does not exist another feasible PBE, either within or across equilibrium categories, that yields a strictly higher payoff for at least one seller type while providing at least as high a payoff for the other seller type.*

This second criterion enables sharper analytical conclusions by explicitly identifying outcomes that sellers would collectively prefer. Although not innocuous, focusing on the equilibria that are most beneficial from the sellers' perspective allows us to discard some unnecessarily negative and implausible outcomes. In particular, it eliminates the *inefficient* separating equilibrium in which the high type is forced to post a price much closer to θ_l than the θ_h she could attain, even though the buyer perfectly recognizes her as the high-quality seller. Note also that such separating equilibria are not guaranteed to exist, as they require ε to be sufficiently large so that the minimum threshold for α, namely $\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, is less than 1. Moreover, selecting the equilibrium with the highest payoff within a pooling category is both intuitive and inconsequential for the analysis, as all outcomes within that category share the same fundamental characteristics.[17] Finally, focusing on the seller-preferred equilibria appears especially reasonable given the timing structure of the game: first, sellers choose their price-evidence combinations, and subsequently the buyer responds. This approach could also be interpreted as capturing the sellers' first-mover advantage.

Given the above refinement criteria, it is possible to provide a complete and sharp characterization of the pure-strategy equilibria that survive the restrictions.

Proposition 9. *Consider any pure-strategy PBE that satisfies both the belief restriction and the seller-optimality restriction (Assumptions 3 and 4). Then:*

(i) If $\alpha \le \min\{\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\}$[18]*, it is a separating equilibrium with* $p^* = \min\{\theta_l/\alpha, \theta_h\}$*.*

---

[17] Clearly, this criterion is distinct from a Pareto efficiency argument, as maximizing the payoff for both seller types entails minimizing the payoff for the buyer.

[18] With strict inequality for $\theta_l/\bar{\theta}_{\tilde{\mu}}$.

(ii) *If* $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ *and* $\tilde{\mu} \in [\bar{\mu}, 1)$*, it is a pooling(*$e_h$*)-b equilibrium with* $p^* = \bar{\theta}_{\tilde{\mu}}$*.*

(iii) *If* $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$ *and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$*, it is a pooling(*$e_h$*)-ba equilibrium with* $p^* = p_{ba}(\tilde{\mu})$*.*

(iv) *If* $\alpha \geq \max\{\theta_l/\varepsilon\bar{\theta}_{\underline{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\}$ *and* $\tilde{\mu} \in (0, \underline{\mu})$*, it is a pooling(*$e_h$*)-b equilibrium with* $p^* = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}\}$*.*

(v) *If* $\alpha \in [\theta_l/\varepsilon\bar{\theta}_{\underline{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}]$ *it is a separating equilibrium with* $p^* = \theta_l/\alpha$*.*

Moreover, these conditions are also necessary.

Overall, the combination of the two refinement criteria allows for sharp predictions based on the parameters α and $\tilde{\mu}$. A separating equilibrium is attainable—regardless of the buyer's prior— when the evidence technology is sufficiently accurate, that is, when α is small. In this case, the low type is forced to post the minimum price θ_l while the high type earns a payoff that is weakly larger than $\bar{\theta}_{\tilde{\mu}}$. The *inefficient* separating—the counterintuitive outcome that requires a large α and yields a small payoff for the high type—residually survives the criterion as the only feasible equilibrium when $\tilde{\mu}$ is too low to sustain a pooling equilibrium. When $\tilde{\mu}$ is sufficiently high (namely, larger than $\bar{\mu}$) but a separating equilibrium is not feasible, the sellers pool at a price equal to the expected quality given the posterior $\tilde{\mu}$ (i.e., $\bar{\theta}_{\tilde{\mu}}$), and the buyer opts not to purchase authentication. This equilibrium is sustained by large prior probabilities of the high type, particularly when the evidence technology is imprecise. Conversely, when both the precision of the evidence signal and the posterior are sufficiently low, the only outcome is a pooling equilibrium with a common price bounded above by $\varepsilon\bar{\theta}_{\underline{\mu}}$; here, the buyer's posterior is too low to justify purchasing authentication, yet she still secures a rent whenever the expected quality exceeds the price. Finally, whenever the posterior lies in the interval where authentication can be optimal (i.e., $(\underline{\mu}, \bar{\mu})$) and the imprecision of the evidence technology precludes separation, the unique outcome is a pooling equilibrium in which the seller charges the maximum price for authentication (given the posterior) and the buyer opts to purchase it.

In general, the higher the initial probability that quality is high—and hence the higher the posterior $\tilde{\mu}$—the higher the equilibrium payoffs for both seller types, if separating is not feasible. Moreover, the high type always benefits from a reduction in α, since this not only increases the likelihood of reaching the separating region but also raises the equilibrium common price in pooling, which the high type fully enjoys, thereby shifting the posterior $\tilde{\mu}$ toward more profitable outcomes. Conversely, an increase in the precision of the evidence technology is not always harmful to the low type, as the reduced likelihood of passing the evidence test can be offset by both higher common prices and a larger posterior, leading to more profitable equilibria. Overall, the effect of α on the low type's payoff is non-monotonic, at least as long as α remains above the threshold that results in a separating equilibrium.

Finally, the authentication technology strongly affects the equilibrium outcome. When its precision level is high, in the limit $\varepsilon = 0$, the only feasible outcomes are separating equilibria and pooling equilibria without authentication acquisition (although the set of posterior values allowing for pooling decreases, since $\partial\bar{\mu}/\partial\varepsilon > 0$). This result is consistent with our findings in the perfect authentication technology section and generalizes those results to cases where ε is sufficiently small. In contrast, when the authentication precision is sufficiently low (i.e., for larger ε), pooling equilibria with authentication acquisition become feasible because they make the low type willing to face authentication in order to secure a larger payoff. The price of authentication has effects similar to those of imprecision: lower authentication prices expand the range of $\tilde{\mu}$ values that lead to authentication acquisition, since authentication is cheaper; however, lower fees also require a higher level of α to sustain a pooling-ba equilibrium, because a reduced authentication fee negatively affects $p_{ba}(\tilde{\mu})$.

Note that Proposition 9 is silent on what happens in some parameter regions. In particular, pure-strategy equilibria are not sustainable for some combinations of medium to low posteriors $\tilde{\mu}$ and medium to high evidence imprecision levels α. In these cases, α is too high to achieve separation but too low to prevent deviations of the low type from pooling. The next section fills this gap by providing a characterization of a specific class of mixed-strategy equilibria—consistent with those found in the literature—in which the high type chooses a specific price along with high-quality evidence with probability one, the low type randomizes between the high-type combination and the lowest price θ_l, and the buyer randomizes among her optimal actions, which are determined by the seller's strategies.

1.4.5 Mixed-strategy Equilibria

Following Bester and Ritzberger (2001), Stahl and Strausz (2017), and Martinez-Gorricho (2020), we examine a specific class of mixed-strategy equilibria and illustrate how they integrate into our overall analysis of pure-strategy equilibria. In particular, Stahl and Strausz (2017) show that in a comparable setting with perfect authentication (and without evidence) the equilibrium outcome is unique. In that equilibrium, the high-quality seller deterministically sets the highest price that makes the buyer indifferent among all her actions while the low-quality seller randomizes between the lowest price and the price set by the high type, and the buyer mixes between purchasing authentication and not buying the good at all. A similar result arises in Bester and Ritzberger (2001), while Martinez-Gorricho (2020) shows that this type of mixed-strategy equilibrium is one of the possible outcomes in a binary-price setting when an imprecise signal about the quality of the good is received by the buyer before taking action.

The contemporary presence of potentially imprecise evidence and authentication technologies in our model yields a richer set of mixed-strategy equilibria in addition to the pure-strategy equilibria analyzed in previous sections. Thus, the class we present now—consistent with those previously found in the signaling literature—is not exhaustive of all possible mixed-strategy equilibria. Nevertheless, any alternative type of mixed-strategy equilibrium either fails to simultaneously satisfy both the belief and seller-payoff restrictions or is equivalent to a corresponding, simpler mixed equilibrium. We will present the formal results and provide some supportive examples in the next subsection. For now, we begin by defining and examining mixed-strategy equilibria in which the high type does not mix, the low type randomizes, and the buyer (possibly) randomizes.[19] We refer to these as *standard* mixed equilibria.

Definition 1. *A standard mixed equilibrium is a PBE* $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ *in which:* $\sigma_h(p^*, e_h) = 1$*,* $\sigma_l(p^*, e_h) \in (0,1)$*,* $\sigma_l(\theta_l, e) + \sigma_l(p^*, e_h) = 1$*, with* $e \in \{e_l, e_h\}$*, and* $\sum_x \beta(x|p^*, \mu^*(p^*, e_h, S_h^e)) = 1$*, with* $x \in \{b, ba, n\}$*.*

The high type sets a price p^* and provides high evidence with certainty, while the low

---

[19] Differently from the literature, we observe that under certain parameter configurations the buyer might not mix between her actions. This occurs for specific combinations of α and ε, such that choosing ba with certainty leaves the low type seller indifferent between cheating and pricing correctly.

type randomizes between (p^*, e_h) and (θ_l, e), where e is irrelevant since $\mu^*(\theta_l, e, S^e) = 0$ because only the low type charges θ_l with positive probability. Upon observing (p^*, e_h, S^e_h), the buyer (possibly) randomizes among a set of available actions that is determined by p^*. Mixing, to be optimal, requires an agent to be indifferent among the actions she randomizes between. Therefore, the expected payoff of the low type in this equilibrium must equal θ_l since charging θ_l leads the buyer to buy the good with certainty. To make the low type indifferent between cheating and pricing honestly, the buyer must use a strategy that limits the probability that the low type receives the high price p^*. Conversely, the low type's pricing strategy must ensure that the buyer's posterior—upon observing (p^*, e_h, S^e_h)—leaves her indifferent between the actions she randomizes between.

The set of actions on which a buyer randomizes depends on p^*. If $p^* \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then the buyer can only be indifferent between n and b. This occurs when $p^* = \bar{\theta}_{\mu'} := \mu'\theta_h + (1 - \mu')\theta_l$ and the buyer holds posterior μ'. In contrast, if $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, the buyer can be indifferent either between n and ba (if $p^* = p_{ba}(\mu')$) or between b and ba (if $p^* = p_{ba}(\mu')$), provided that in both cases she holds posterior μ'. Finally, if $p^* = \bar{\theta}_{\underline{\mu}}$ or $p^* = \bar{\theta}_{\bar{\mu}}$, the buyer is indifferent among all her three actions when she holds posterior $\underline{\mu}$ or $\bar{\mu}$, respectively. The next lemma characterizes the equilibrium strategies in standard mixed equilibria depending on the price p^*.

Lemma 4. *Suppose* $\{\sigma^*_h, \sigma^*_l, \beta^*, \mu^*\}$ *represents a standard mixed equilibrium with* $\sigma^*_h(p^*, e_h) = 1$*, where* $p^* \in (\theta_l, \theta_h)$ *and let* μ' *be the posterior belief for which the buyer is indifferent between at least two of her actions given* p^**. Then,* $\sigma_l(p^*, e_h) = \frac{\tilde{\mu}(1-\mu')}{\mu'(1-\tilde{\mu})}$ *and:*

(i) If $p^* \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, $p^* = \bar{\theta}_{\mu'}$*, and* $\beta(b|p^*, \mu') = \theta_l/(\alpha p^*)$, $\beta(ba|p^*, \mu') = 0$.

(ii) If $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$*, either*

(a) $p^* = p_{ba}(\mu')$*, and* $\beta(b|p^*, \mu') = 0$, $\beta(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$*; or*

(b) $p^* = p_b(\mu')$*, and* $\beta(b|p^*, \mu') = \frac{\theta_l - \alpha\varepsilon p^*}{\alpha(1-\varepsilon)p^*}$, $\beta(ba|p^*, \mu') = \frac{\alpha p^* - \theta_l}{\alpha(1-\varepsilon)p^*}$.

(iii) If $p^* = \bar{\theta}_{\underline{\mu}}$, $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$, $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) \in \left[0, \min\left\{\frac{\alpha\bar{\theta}_{\underline{\mu}} - \theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}\right]$.

(iv) If $p^* = \bar{\theta}_{\bar{\mu}}$, $\beta(b|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \theta_l/(\alpha\bar{\theta}_{\bar{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu})$, $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) \in \left[0, \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}} - \theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}\right]$.

The proof of the lemma is straightforward once we note that the low-quality seller's indifference condition between (θ_l, e^l) and (p^*, e_h) requires

$$\theta_l = \alpha[\beta(b|p^*, \mu') + \varepsilon\beta(ba|p^*, \mu')]p^* \tag{1.17}$$

while the buyer's indifference condition requires her posterior to equal the belief for which randomization given p^* is optimal, i.e.,

$$\mu' = \frac{\mu_0}{\mu_0 + (1-\mu_0)\alpha\sigma_l(p^*, e_h)} = \frac{\tilde{\mu}}{\tilde{\mu} + (1-\tilde{\mu})\sigma_l(p^*, e_h)}.$$

Note that the last condition requires $\tilde{\mu} < \mu'$ for $\sigma_l(p^*, e_h)$ to be in $(0, 1)$, so that any standard mixed-strategy equilibrium can be sustained only if $\tilde{\mu}$ is lower than the equilibrium posterior that makes the buyer indifferent between some of her actions given p^*.

Similarly to the pure-strategy case, without further restrictions on out-of-equilibrium beliefs and agents' payoffs, any price $p^* \in (\theta_l, \theta_h)$ can be sustained in equilibrium in a standard mixed equilibrium, yielding an infinite number of possible outcomes. We therefore restore to Assumptions 3 and 4 to limit our attention to the more plausible standard mixed equilibria yielding the highest payoff to both sellers. The following proposition characterizes the standard mixed-strategy equilibria that survive both restrictions.

Proposition 10. *The standard mixed-strategy PBEs that satisfy both the belief restriction and the seller-optimality restriction (Assumptions 3 and 4) are the following:*

(i) ***Mixed($\bar{\theta}_{\underline{\mu}}$)****:* $\sigma_h(\bar{\theta}_{\underline{\mu}}, e_h) = 1$, $\sigma_l(\bar{\theta}_{\underline{\mu}}, e_h) = \frac{\tilde{\mu}(1-\underline{\mu})}{\underline{\mu}(1-\tilde{\mu})}$, $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = 0$, $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\underline{\mu}}}$, *with necessary conditions:* $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ *and* $\tilde{\mu} \leq \underline{\mu}$*. Moreover,* $\pi_h^* = \theta/(\alpha\varepsilon) \in [\theta_l/\varepsilon, \bar{\theta}_{\underline{\mu}}]$ *and* $\pi_l^* = \theta_l$*.*

(ii) ***Mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_{ba}$****:* $p^* = p_{ba}(\mu') \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ *for some* $\mu' \in (\underline{\mu}, \bar{\mu})$, $\sigma_h(p^*, e_h) = 1$, $\sigma_l(p^*, e_h) = \frac{\tilde{\mu}(1-\mu')}{\mu'(1-\tilde{\mu})}$, $\beta(b|p^*, \mu') = 0$, $\beta(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$, *with necessary conditions:* $\alpha > \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$ *and* $\tilde{\mu} \leq \mu'$*. Moreover,* $\pi_h^* = \theta/(\alpha\varepsilon) \in [\theta_l/\varepsilon, \bar{\theta}_{\bar{\mu}})$ *and* $\pi_l^* = \theta_l$*.*

(iii) ***Mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_b$****:* $p^* = \theta_l/(\alpha\varepsilon) = p_b(\mu') \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ *for some* $\mu' \in (\underline{\mu}, \bar{\mu})$, $\sigma_h(p^*, e_h) = 1$, $\sigma_l(p^*, e_h) = \frac{\tilde{\mu}(1-\mu')}{\mu'(1-\tilde{\mu})}$, $\beta(b|p^*, \mu') = 0$, $\beta(ba|p^*, \mu') = 1$, *with necessary conditions:* $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$ *and* $\tilde{\mu} \leq \mu'$*. Moreover,* $\pi_h^* = \theta/(\alpha\varepsilon) \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ *and* $\pi_l^* = \theta_l$*.*

(iv) ***Mixed($\bar{\theta}_{\bar{\mu}}$):*** $\sigma_h(\bar{\theta}_{\bar{\mu}}, e_h) = 1$, $\sigma_l(\bar{\theta}_{\bar{\mu}}, e_h) = \frac{\tilde{\mu}(1-\bar{\mu})}{\bar{\mu}(1-\tilde{\mu})}$, $\beta(b|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}} - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu})$, $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}$, *with necessary conditions* $\alpha \geq \theta_l/\bar{\theta}_{\bar{\mu}}$ *and* $\tilde{\mu} \leq \bar{\mu}$. *Moreover,* $\pi_h^* = \theta/(\alpha\varepsilon) \in [\theta_l/\varepsilon, \bar{\theta}_{\bar{\mu}}]$ *and* $\pi_l^* = \theta_l$

(v) ***Mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$):*** $p^* = \bar{\theta}_{\mu'} \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ *for some* $\mu' \in (\bar{\mu}, 1)$, $\sigma_h(p^*, e_h) = 1$, $\sigma_l(p^*, e_h) = \frac{\tilde{\mu}(1-\mu')}{\mu'(1-\tilde{\mu})}$, $\beta(b|p^*, \mu') = \frac{\theta_l}{\alpha\bar{\theta}_{\mu'}}$, $\beta(ba|p^*, \mu') = 0$, *with necessary conditions:* $\alpha \in (\theta_l/\bar{\theta}_{\mu'}, \theta_l/\bar{\theta}_{\bar{\mu}})$ *and* $\tilde{\mu} \leq \mu'$. *Moreover,* $\pi_h^* = \theta/\alpha \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ *and* $\pi_l^* = \theta_l$.

Each standard mixed equilibrium category is defined by its semipooling price p^*, as the name of the category indicates either the unique price p^* or the set to which p^* belongs. Moreover, when $p^* \in (\underline{\mu}, \bar{\mu})$, we further specify whether the price renders the buyer indifferent between ba and n, p_{ba}, or between b and ba, p_b, similar to the pooling equilibria discussed in the previous section. In general, the beliefs and seller-optimality requirements restrict the set of sustainable prices p^* to the interval $[\bar{\theta}_{\underline{\mu}}, \theta_h)$, with the exact range depending on α. When α is sufficiently high, namely $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$ all prices in $[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ are feasible, as even low prices can be attractive enough for the low type given the higher probability of passing the evidence test. As α decreases, however, the minimum feasible semipooling price increases because a low p^* might not be sustainable even under the most optimistic buyer behavior from the seller's perspective—that is, when the buyer purchases the good at p^* with the highest probability. In fact, in every standard mixed equilibrium, the buyer's behavior forces the low type to be indifferent between setting a high price p^* and facing the evidence test, and posting the minimum price θ_l. If the probability of passing the test decreases, then the reward must be sufficiently high; otherwise, θ_l becomes the better option even under the most favorable buyer's behavior. Then for α above $\theta_l/\bar{\theta}_{\bar{\mu}}$, but below $\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$, the only feasible price is $\bar{\theta}_{\bar{\mu}}$, since the unique feasible standard mixed equilibrium is the mixed($\bar{\theta}_{\bar{\mu}}$).[20] Finally, for sufficiently small values of α, namely when $\alpha \in (\theta_l/\bar{\theta}_{\mu'}, \theta_l/\bar{\theta}_{\bar{\mu}})$, the feasible set of prices p^* lies within $(\bar{\theta}_{\bar{\mu}}, \theta_h)$, with the lower bound decreasing in α, following the same reasoning as above.

When considering standard mixed equilibria, the restrictions usually fail to yield a uniquely determined equilibrium price p^*, in particular when α is sufficiently large, $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$), or sufficiently small, $\alpha \in (\theta_l/\bar{\theta}_{\mu'}, \theta_l/\bar{\theta}_{\bar{\mu}})$. This is because, unlike in the pure-

---

[20] This result is consistent with Stahl and Strausz (2017), as in a setting with perfectly precise authentication, i.e., $\varepsilon = 0$, and without evidence technology, i.e., $\alpha = 1$, the unique outcome of the model, when the prior is below $\bar{\mu}$, is the equivalent of a mixed($\bar{\theta}_{\bar{\mu}}$).

strategy case, the payoff of the high type does not correspond to the equilibrium price p^*. In fact, the indifference condition for the low-quality seller shapes the buyer's purchasing behavior such that the high type becomes indifferent between different equilibrium prices. In fact, as the price increases, the probability that the buyer purchases the good (whether with or without authentication) decreases, thereby affecting both seller types. Consequently, the equilibrium payoff of the high type is given by $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$ whenever the semipooling price p^* lies in the interval $[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, and by $\pi_h^* = \theta_l/\alpha$ if $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$. Therefore, the payoff refinement criterion, which is based on equilibrium payoffs and not on equilibrium prices, lacks bite in most cases, and price indeterminacy arises across different categories. Note also that while π_l^* is fixed at θ_l, π_h^* is weakly increasing in α. In particular, when p^* lies within the *authentication interval* (i.e., values at which authentication is potentially optimal), the high type can exploit both technologies to achieve a payoff that increases with the precision of the signal (i.e., as α and ε decrease), resulting in even larger gains.

Overall, the two refinement criteria seem to have limited power in reducing the set of feasible standard mixed equilibria, with some noteworthy exceptions. First, standard mixed equilibria featuring $p^* \in (\theta_l, \bar{\theta}_{\underline{\mu}})$ do not survive the payoff restrictions because they yield a strictly lower payoff for the high type than other equilibria with higher semipooling price p^*. Second, the seller-optimality requirement forces the buyer to assign the highest possible probability to ba (i.e., the probability of acquiring authentication) whenever the buyer's purchasing behavior is not uniquely determined—as in mixed($\bar{\theta}_{\underline{\mu}}$) and mixed($\bar{\theta}_{\bar{\mu}}$) equilibria. The reason lies in the indifference condition (1.17) for the low type: although the condition holds for various combinations of $\beta(b|p^*, \mu')$ and $\beta(ba|p^*, \mu')$, a decrease in $\beta(b|p^*, \mu')$ must be offset by a larger increase in $\beta(ba|p^*, \mu')$, since when a cheating low type acquires authentication, a trade occurs only with probability ε. Moreover, since b and ba are equivalent from the high type's perspective, the combination of probabilities that maximizes the weight on ba yields the highest payoff for the high type.[21] Third, the belief restriction imposes stricter conditions on the values of α required to sustain each equilibrium. In particular, it determines an upper bound for α in the case of mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$)

---

[21] In any mixed equilibrium with semipooling price p^*, the high-type payoff is given by $\pi_h^* = [\beta(b|p^*, \mu') + \beta(ba|p^*, \mu')]p^*$. Condition (1.17) implies $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$. Substituting this into the payoff expression yields $\pi_h^* = \theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) + (1-\varepsilon)\beta(ba|p^*, \mu')]p^*$, which is maximized when $\beta(ba|p^*, \mu')$ is maximized.

equilibria. The rationale is that higher values of α imply a lower probability assigned to b and a higher probability assigned to n in order to satisfy the low type's indifference condition. This, in turn, results in a lower payoff for the high type. If π_h^* were below $\bar{\theta}_{\bar{\mu}}$, then the high type would have a greater incentive than the low type to deviate to a price slightly higher than π_h^*, which triggers authentication acquisition. Therefore, in this case, α cannot exceed $\theta_l/\bar{\theta}_{\bar{\mu}}$.

A final consideration is that, for any equilibrium to be sustainable, the posterior belief when both types provide high evidence and charge the same price, $\tilde{\mu}$, must be lower than the *target* posterior belief μ', which is the posterior at which the buyer is indifferent between some of her actions when the price is p^* (i.e. $p^* \in \{\bar{\theta}_{\mu'}, p_{ba}(\mu'), p_b(\mu')\}$). In particular, the buyer's indifference condition requires that the probability that the low type assigns to pricing dishonestly, $\sigma_l(p^*, e_h)$, decreases as the gap between the initial posterior $\tilde{\mu}$ and the target posterior μ' widens; that is, the buyer will only be willing to purchase at a price p^* much higher than the maximum she would pay—whether with or without authentication—when holding posterior $\tilde{\mu}$, if the probability that the low type cheats is sufficiently small.

Despite the limited effectiveness on the price side, the refinement criteria yield an important result: we can identify three distinct regions in the $(\tilde{\mu}, \alpha)$ space, each characterized by a unique standard mixed equilibrium type. In fact, when $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$, all feasible standard mixed equilibria that survive the belief and payoff restrictions—namely, mixed($\bar{\theta}_{\underline{\mu}}$), mixed($\bar{\theta}_{\bar{\mu}}$), mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_b$ and mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_{ba}$—are equivalent in terms of both the buyer's purchasing strategy and the sellers' payoffs, once the different equilibrium price p^* is taken into account. We therefore define an encompassing category of standard mixed equilibria, mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, with $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, characterized by purchasing strategy β such that $\beta(b|p^*, \mu') = 0$ and $\beta(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$.

Corollary 3. *Given Proposition 10:*

(i) if $\tilde{\mu} \in (0,1)$ *and* $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\bar{\mu}})$*, the unique standard mixed equilibrium category is mixed(*$\bar{\theta}_{\bar{\mu}}, \theta_h$*) with* $\pi_h^* = \theta_l/\alpha \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$*.*

(ii) if $\tilde{\mu} \in (0, \bar{\mu})$ *and* $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}))$*, the unique standard mixed equilibrium is mixed(*$\bar{\theta}_{\bar{\mu}}$*) with* $\pi_h^* = \bar{\theta}_{\bar{\mu}}$*.*

(iii) *if* $\tilde{\mu} \in (0, \bar{\mu})$ *and* $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), 1]$*, the unique standard mixed equilibrium category is mixed*$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ *with* $\pi_h^* = \theta_l/(\alpha\varepsilon) \in [\theta_l/\varepsilon, \bar{\theta}_{\bar{\mu}}]$.

The main result of the corollary—that the surviving standard mixed equilibria, for sufficiently large α, are equivalent—is only apparently straightforward, since the four equilibrium types within the mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ category have very different characteristics. When the semipooling price p^* is either $\bar{\theta}_{\underline{\mu}}$ or $\bar{\theta}_{\bar{\mu}}$, an indifferent buyer randomizes among all three available actions; when $p^* = p_{ba}(\mu')$, she randomizes between *ba* and *b*; and when $p^* = p_b(\mu')$, she randomizes between *ba* and *n*. Nevertheless, the two refinement criteria force the buyer to assign maximum probability to *ba* in order to maximize the high type's payoff and/or ensure the plausibility of the out-of-equilibrium belief sustaining an equilibrium. As a result, the equilibrium strategies and payoffs are equivalent across all equilibria in this category, despite their different natures.

Finally, by combining the results for the pure-strategy equilibria (Proposition 9) and the standard mixed-strategy equilibria (Corollary 3) that satisfy both the belief and payoff restrictions, we can fully characterize the equilibrium outcome for each region of the $(\tilde{\mu}, \alpha)$ space.

Proposition 11. *Consider all pure-strategy and standard mixed-strategy PBEs that satisfy both the belief restriction and the seller-optimality restriction (Assumptions 3 and 4). Then:*

(i) *If* $\alpha \leq \theta_l/\theta_h$ *it is a separating equilibrium with* $p^* = \theta_h$.

(ii) *If* $\alpha \in (\theta_l/\theta_h, \min\{\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\})$ *it is either a separating equilibrium with* $p^* = \theta_l/\alpha$ *or a mixed*$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ *equilibrium with* $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$.

(iii) *If* $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ *and* $\tilde{\mu} \in [\bar{\mu}, 1)$*, it is a pooling(*$e_h$*)-b equilibrium with* $p^* = \bar{\theta}_{\tilde{\mu}}$.

(iv) *If* $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}))$ *and* $\tilde{\mu} \in (0, \bar{\mu})$ *it is a mixed*$(\bar{\theta}_{\bar{\mu}})$ *with* $p^* = \bar{\theta}_{\bar{\mu}}$.

(v) *If* $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$ *and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$*, it is a pooling(*$e_h$*)-ba equilibrium with* $p^* = p_{ba}(\tilde{\mu})$.

(vi) *If* $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\})$ *and* $\tilde{\mu} \in (0, \underline{\mu})$*, or* $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/\varepsilon p_{ba}(\tilde{\mu}))$ *and* $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$*, it is a mixed*$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ *with* $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$.

(vii) If $\alpha \in [\max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\}, \max\{\theta_l/(\varepsilon^2\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}})\})$ *and* $\tilde{\mu} \in (0, \underline{\mu})$ *it is either a pooling(*e_h*)-b equilibrium with* $p^* = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}\}$ *or a mixed*$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]$ *equilibrium with* $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]$*.*

(viii) If $\alpha \geq \max\{\theta_l/(\varepsilon^2\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}})\}$ *and* $\tilde{\mu} \in (0, \underline{\mu})$*, it is a pooling(*$e_h$*)-b equilibrium with* $p^* = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}\}$*.*

Standard mixed equilibria complete the characterization of equilibrium outcomes in the $(\tilde{\mu}, \alpha)$ space, ensuring that at least one equilibrium exists in every region. They arise when α is intermediate to high—that is, when the evidence technology is sufficiently imprecise—and when the posterior $\tilde{\mu}$ is sufficiently low, with lower $\tilde{\mu}$ becoming increasingly necessary as α increases. Under these conditions, the no-deviation condition for the low-quality seller cannot be satisfied in pooling equilibria: either α is too low to ensure a high enough probability of passing the evidence test, or the pooling price is unattractive because a low $\tilde{\mu}$ limits the buyer's maximum willingness to pay when confidence in quality is low. Conversely, the evidence technology is not precise enough to sustain a separating equilibrium, because a high α fails to prevent the low type from deviating to the high type's price p^*. The insights from Proposition 9 remain applicable: for intermediate $\tilde{\mu}$ and moderate values of α, pooling with authentication acquisition is the only feasible outcome; when $\tilde{\mu}$ is large, pooling without authentication emerges; and when α is small, separating equilibria prevail, regardless of the posterior belief.[22]

Therefore, we can interpret the standard mixed equilibria as semi-separating equilibria: the buyer's purchasing strategy—by assigning maximum weight to authentication acquisition—disciplines the low type by confining him to the minimum payoff, θ_l, while allowing the high type to secure a higher payoff. However, this payoff is lower than that of a fully separating equilibrium, since the buyer does not purchase the good with certainty as in pure-strategy equilibria, thereby preventing the high type from always receiving the full price p^*.

The refinement criteria provide an almost fully sharp prediction about the equilibrium outcome, with two exceptions—regions where a pure-strategy equilibrium and a standard mixed equilibrium overlap for different reasons. First, as highlighted in point

---

[22] Note that the inefficient separating equilibrium for high values of α and low values of $\tilde{\mu}$ does not satisfy the payoff requirement when compared to standard mixed equilibria.

(ii) of the proposition, when $\alpha \in (\theta_l/\theta_h, \min\{\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\})$ the equilibrium outcome can be either separating or a mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ equilibrium. Remarkably, the mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ equilibrium replicates the payoffs of the efficient separating equilibrium (in the region where it yields strictly less than θ_h) by requiring a semipooling price p^* that is higher than the price charged by the high type in a separating equilibrium, and by inducing the buyer to purchase the good with a probability that decreases as the gap between θ_l and p^* widens. Second, when $\alpha \in \big[\max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\}, \max\{\theta_l/(\varepsilon^2\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}})\}\big)$ and $\tilde{\mu} \in (0, \underline{\mu})$, the equilibrium outcome can be either a pooling(e_h)-b equilibrium or a mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium. Note that standard mixed equilibria always yield a payoff of θ_l to the low type, whereas pooling equilibria yield a payoff that is weakly larger. Consequently, according to Assumption 4, the pooling equilibrium always satisfies the payoff condition when compared to the mixed equilibrium. However, in the specified interval, the mixed equilibrium yields a higher payoff for the high type, so it also meets the payoff condition—and both outcomes are actually feasible.

One final remark before we move to the characterization of non-standard mixed equilibria. In our overview, we have described equilibrium feasibility in the $(\tilde{\mu}, \alpha)$-space while holding the other primitives $(\mu_0, \theta_h, \theta_l, \varepsilon, c)$ fixed.[23] Although one could use other parameterizations, this choice aligns directly with the existence conditions we derived throughout this section and will prove especially convenient for the welfare comparisons in Section 5.

Notice also that not every equilibrium type is guaranteed to be feasible for any $(\tilde{\mu}, \alpha)$. Our overview has assumed an implicit "balance" among the primitives.[24] The separating equilibrium is clearly one of the exceptions. For example, if $\varepsilon < \theta_l/\bar{\theta}_{\bar{\mu}}$, then the lower bound on α required for a mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium exceeds one, making both that equilibrium and pooling(e_h)-ba infeasible across the entire $(\tilde{\mu}, \alpha)$ domain. In the next subsection, we will show how these equilibrium regions expand, contract, or disappear as we vary key parameters, before proceeding with the welfare analysis.

Non-standard mixed equilibria

As mentioned above, *standard* mixed equilibria—where the high type posts price p^* and provides e_h with certainty while the low type randomizes between pricing θ_l and mimicking

---

[23] Note that $\tilde{\mu}$ is itself a function of α.

[24] The separating equilibrium is clearly one of the exceptions.

the high type—are not the only mixed-strategy equilibria that can arise in the model. Alternative mixed equilibria may vary along different dimensions; first, the high type might not provide high evidence along with p^*; second, the high type might (possibly as well) randomize among two or more different price-evidence combinations. We will show in this section that non-standard mixed equilibria cannot satisfy both the belief and payoff restrictions unless they replicate a corresponding standard mixed equilibrium—that is, unless they are payoff equivalent and share the same fundamental nature.

Let us first consider a couple of examples of PBEs that follow the same structure as a standard mixed equilibrium, except that the high type provides low evidence.

Example 1. *Let* $p^* = \bar{\theta}_{\mu'} \in (\theta_l, \bar{\theta}_{\underline{\mu}})$ *for some* $\mu' \in (0, \underline{\mu})$, *with* $\sigma_h(p^*, e_l) = 1$, $\sigma_l(p^*, e_l) = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)}$, *and* $\sigma_l(\theta_l, e) = \frac{\mu'-\mu_0}{\mu'(1-\mu_0)}$ *for* $e \in \{e_l, e_h\}$. *Moreover, let* $\beta(b|p^*, \mu') = \frac{\theta_l}{p^*}$, $\beta(n|p^*, \mu') = \frac{p^*-\theta_l}{p^*}$ *and* $\beta(b|\theta_l, 0) = 1$. *Then, for* $\mu_0 < \mu'$, *these strategies* σ *and* β *can be sustained in a PBE yielding payoffs* $\pi_h^* = \pi_l^* = \theta_l$. *However, by Lemma 3, any such PBE does not survive the belief refinement since* $\pi_h^* = \theta_l < \theta_l/\alpha = \pi_l^*/\alpha$.

In this case, by providing low evidence, the high type cannot exploit the evidence technology to secure a higher payoff and ends up receiving θ_l. Moreover, any price in $(\theta_l, \theta_l/\alpha)$ is a profitable deviation exclusively for the high type—even under perfectly optimistic out-of-equilibrium beliefs—so the PBE fails to satisfy the belief restriction. Nonetheless, providing e_l does not automatically imply a failure to meet the belief restriction as it does in the case of pure-strategy equilibria, as the next example shows.

Example 2. *Let* $p^* = p_{ba}(\mu') \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ *for some* $\mu' \in (\underline{\mu}, \bar{\mu})$, $\sigma_h(p^*, e_l) = 1$, $\sigma_l(p^*, e_l) = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)}$, $\sigma_l(\theta_l, e) = 1 - \sigma_l(p^*, e_l)$ *for* $e \in \{e_l, e_h\}$. *Moreover, let* $\beta(ba|p^*, \mu') = \frac{\theta_l}{\varepsilon p^*}$, $\beta(n|p^*, \mu') = \frac{\varepsilon p^*-\theta_l}{\varepsilon p^*}$ *and* $\beta(b|\theta_l, 0) = 1$. *Then, for* $\mu_0 < \mu'$, *these strategies* σ *and* β *can be sustained in a PBE yielding payoffs* $\pi_h^* = \theta_l/\varepsilon$ *and* $\pi_l^* = \theta_l$. *If* $\alpha \geq \{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \varepsilon\}$, *any such PBE survives the belief refinement, but it does not survive the payoff restriction when compared to any feasible mixed*$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, *pooling*$(e_h)$*-b or pooling*$(e_h)$*-ba equilibrium, which yield at least* $\pi_h^* = \theta_l/(\alpha\varepsilon)$.

In this second example, the fact that the semipooling price falls within the authentication range $(\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ enables a PBE with semipooling vector (p^*, e_l) to satisfy the belief refinement criterion, yielding $\pi_h^* = \theta_l/\varepsilon$, for sufficiently low ε so that the high type's payoff exceeds θ_l/α. However, this outcome fails to satisfy the payoff restriction because, in the same parameter region, the high type can secure a profit of at least $\theta_l/(\alpha\varepsilon)$ in any alternative feasible equilibrium (which depends on the prior μ_0) by fully leveraging the evidence technology through the provision of high evidence. The following lemma formalizes this intuitive result.

Lemma 5. *Any mixed-strategy PBE* $\{\sigma_h, \sigma_l, \beta, \mu\}$ *in which* $\sigma_h(p^*, e_l) = 1$ *fails to simultaneously satisfy both the belief and the payoff restrictions.*

While ruling out equilibria in which the high type does not randomize and provides e_l is relatively straightforward, there is no clear, intuitive reason why instances in which the high type randomizes among different prices should fail to satisfy any of the refinement criteria.

Suppose there exists an equilibrium in which the high type randomizes among $k \geq 2$ different prices $p_1 < p_2 < ... < p_k$, and define $P_i = \{p \mid \exists e \in \{e_l, e_h\} : (p, e) \in \text{supp}(\sigma_i)\}$ as the set of prices on which type i assigns positive probability. Then, the high type's indifference condition, necessary for having her mixing among different prices, is given by

$$[\beta(b|p_1,\mu_1)+\beta(ba|p_1,\mu_1)]p_1 = [\beta(b|p_2,\mu_2)+\beta(ba|p_2,\mu_2)]p_2 = ... = [\beta(b|p_k,\mu_k)+\beta(ba|p_k,\mu_k)]p_k \tag{1.18}$$

where μ_j is the buyer's belief that renders her indifferent among some of her actions at price $p_j \in P_h$.

The next lemma derives some properties of the equilibrium outcome that hold in any PBE in which the high type randomizes.

Lemma 6. *In any mixed-strategy PBE* $\{\sigma_h, \sigma_l, \beta, \mu\}$ *with* $P_h = \{p_1, p_2, ..., p_k\}$*, where* $p_1 < p_2 < ... < p_k$ *and* $k \geq 2$*, we have:*

(i) The low type mixes at least among the $k-1$ *highest prices:* $\{p_2, ..., p_k\} \subseteq P_l$*.*

(ii) In general (unless $\alpha = \varepsilon$*), either* $\sigma_h(p_j, e_h) > 0$ *for all* $p_j \in P_h \cap P_l$*, or* $\sigma_h(p_j, e_l) > 0$ *for all* $p_j \in P_h \cap P_l$*.*

(iii) $\beta(b|p_2, \mu_2)p_2 = ... = \beta(b|p_k, \mu_k)p_k$ *and* $\beta(ba|p_2, \mu_2)p_2 = ... = \beta(ba|p_k, \mu_k)p_k$*.*

(iv) Either $\beta(b|p_2, \mu_2) = ... = \beta(b|p_k, \mu_k) = 0$*, or* $\beta(ba|p_2, \mu_2) = ... = \beta(ba|p_k, \mu_k) = 0$*.*

(v) If $\sigma_h(p_j, e_h) > 0$ for some $p_j \in P_h \cap P_l$, then $\tilde{\mu} < \mu_j$ if and only if $\sigma_h(p_j, e_h) > \sigma_l(p_j, e_h)$, and $\tilde{\mu} > \mu_j$ if and only if $\sigma_h(p_j, e_h) < \sigma_l(p_j, e_h)$. Likewise, if $\sigma_h(p_j, e_l) > 0$ for some $p_j \in P_h \cap P_l$, then the same relationship holds with μ_0 and e_l in place of $\tilde{\mu}$ and e_h, respectively.

The lemma highlights some important characteristics of any PBE in which the high type randomizes among $k \geq 2$ different prices. First, the low type must necessarily randomize over at least the $k-1$ highest prices $p_2, \ldots, p_k$. Suppose, to the contrary, that there exists a price $p_j > \min P_h$ such that $p_j \notin P_l$. Then, when the buyer observes p_j, she recognizes the good as being of high quality and buys it without authentication (i.e., $\beta(b|p_j, 1) = 1$). However, condition (1.18) would then be violated, since any price lower than p_j in P_h can yield a profit equal to p_j only if the purchasing probability exceeds one, which is impossible.

Second, the indifference conditions of the two types can be simultaneously satisfied only if, for every common price $p_j \in P_h \cap P_l$, both types provide one and only one evidence type—either e_l or e_h. The reason is that the high type secures the charged price regardless of the evidence provided, whereas the low type's payoff is diluted by the probability of successfully passing the evidence and authentication tests. The only exception is when $\alpha = \varepsilon$: only in that case, is it possible to construct strategies in which both evidence levels are provided by requiring purchasing with authentication for prices posted along low evidence, and aquisition with authentication when the price is accompanied by high evidence.[25] Therefore, the indifference conditions for the low type can be written as

$$[\beta(b|p_2, \mu_2) + \varepsilon\beta(ba|p_2, \mu_2)]p_2 = \ldots = [\beta(b|p_k, \mu_k) + \varepsilon\beta(ba|p_k, \mu_k)]p_k, \tag{1.19}$$

since for all prices only one between e_l and e_h is provided: if e_l is provided, the low type faces only the authentication technology if the buyer selects ba with positive probability; whereas if e_h is provided, she also encounters the evidence technology, and the payoff at each price is multiplied by α, which can therefore be factored out giving the same result

---

[25] In general, α is not equal to ε. Moreover, any mixed-strategy equilibrium in which sellers provide a mix of evidence levels has a payoff-equivalent counterpart in which only e_h is provided. Therefore, we assume from now on that $\alpha \neq \varepsilon$. This assumption does not affect the subsequent results, except for the uniqueness of the evidence level. Finally, note that Proposition (??) remains valid even without this restrictive assumption.

as in the e_l-case.

Moreover, this and the fact that the considered price and posterior belief can make the buyer indifferent between different sets of actions imply two additional features of this type of equilibrium. First, we have $\beta(b|p_2, \mu_2)p_2 = \cdots = \beta(b|p_k, \mu_k)p_k$ and $\beta(ba|p_2, \mu_2)p_2 = \cdots = \beta(ba|p_k, \mu_k)p_k$, meaning that the buyer's probability of purchasing the good—whether with or without authentication—must decrease as the price increases. Second, for every price $p_j \in P_h \cap P_l$, either $\beta(b|p_j, \mu_j) = 0$ or $\beta(ba|p_j, \mu_j) = 0$. Consequently, equilibria in which the buyer randomizes between ba and b are unsustainable.

Finally, the buyer's indifference condition requires that the strategies σ_h and σ_l satisfy, for any $p_j \in P_h \cup P_l$, either

$$\mu_j = \frac{\tilde{\mu}\sigma_h(p_j, e_h)}{\tilde{\mu}\sigma_h(p_j, e_h) + (1 - \tilde{\mu})\sigma_l(p_j, e_h)} \tag{1.20}$$

if $\sigma_h(p_j, e_h) > 0$, or

$$\mu_j = \frac{\mu_0\sigma_h(p_j, e_l)}{\mu_0\sigma_h(p_j, e_l) + (1 - \mu_0)\sigma_l(p_j, e_l)} \tag{1.21}$$

if $\sigma_h(p_j, e_l) > 0$, where μ_j is the belief that satisfies $p_j = p_{ba}(\mu_j)$ or $p_j = p_b(\mu_j)$ if $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, or $p_j = \bar{\theta}_{\mu_j}$ if $p_j \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$. Clearly, if $\sigma_h(p_j, e_h) > \sigma_l(p_j, e_h)$, then condition (1.20) can only be satisfied if $\tilde{\mu} < \mu_j$; conversely, if $\sigma_h(p_j, e_h) < \sigma_l(p_j, e_h)$, then it must be that $\tilde{\mu} > \mu_j$. The same relationship holds symmetrically for μ_0 when e_l is provided instead of e_h.

Building on these observations, consider the implications for the low type's equilibrium strategy when the high type randomizes among different prices. Lemma 6 establishes that in any PBE the low type must randomize among at least the $k - 1$ highest prices $p_2 < \cdots < p^k$, but it does not pinpoint the exact set P_l. The additional prices that the low type can potentially randomize over are limited to two: $p_1 = \min P_h$ and θ_l. In fact, any other price would result in the buyer assigning a posterior belief of zero, thereby ensuring no acquisition and zero profit. Depending on whether the low type assigns positive probability solely to these $k - 1$ prices or also to p_1 and θ_l, three distinct cases arise in equilibrium, as specified by the following lemma.

Lemma 7. *In any mixed-strategy PBE* $\{\sigma_h, \sigma_l, \beta, \mu\}$ *with* $P_h = \{p_1, p_2, ..., p_k\}$*, where* $p_1 < p_2 < ... < p_k$ *and* $k \geq 2$*, we have that* $P_l \subseteq P_h \cup \{\theta_l\}$ *and either:*

(i) $P_l = P_h \setminus \{p_1\}$ *with necessary condition* $\tilde{\mu} > \mu_2$ *(*$\mu_0 > \mu_2$*) if* e_h *(*e_l*) is provided;*

(ii) $P_l = P_h$ *with necessary condition* $\tilde{\mu} \in (\mu_1, \mu_k)$ *(*$\mu_0 \in (\mu_1, \mu_k)$*) if* e_h *(*e_l*) is provided;*

(iii) $P_l = P_h \cup \{\theta_l\}$ *with necessary condition* $\tilde{\mu} < \mu_k$ *(*$\mu_0 < \mu_k$*) if* e_h *(*e_l*) is provided.*

In addition to characterizing the set of prices over which the low type can mix when the high type also randomizes, the lemma explicitly states conditions on the corresponding buyer belief, namely, $\tilde{\mu}$ when the common evidence provided is e_h, and μ_0 when it is e_l, that are necessary to sustain each type of equilibrium. In particular, for equilibria in which the low type assigns greater probability weight than the high type to common prices (case (i)), the initial belief must exceed the induced belief that renders the buyer indifferent at the lowest common price. That is, a sufficiently high initial belief is required to sustain mixed-strategy equilibria in which the low type mixes more than the high type on common prices. Conversely, when the high type allocates more probability weight to common prices than the low type (case (iii)), the equilibrium can be sustained provided that the initial belief is lower than the belief that makes the buyer indifferent at the highest common price,—even if the initial belief is arbitrarily low.

We are now ready to provide the counterpart to Lemma 5 for the case where the high type assigns positive probability weight to more than one price. Specifically, this result shows that under the belief and the payoff restrictions, the high type always provides high evidence in equilibrium.

Proposition 12. *Any mixed-strategy PBE* $\{\sigma_h, \sigma_l, \beta, \mu\}$ *in which* $\sigma_h(p, e_l) > 0$*, for any* $p \in P_h$*, fails to simultaneously satisfy both the belief and the payoff restrictions.*

The intuition behind the proposition is that providing e_l prevents the high type from taking full advantage of the costless evidence technology thus maximizing her payoff. Even if some equilibrium of this type satisfies the belief restriction, when the authentication technology is sufficiently precise—thus offering a high payoff for the high type—it will fail the payoff restriction compared to the case where the high type chooses e_h.

The following examples examine cases where the high type provides high evidence and randomizes among multiple prices, while the low type—whether she randomizes or not—always satisfies $P_l \subset P_h$. In other words, the low type assigns zero probability to θ_l.

Example 3. *Let* $p' = \bar{\theta}_{\mu'}$ *and* $p'' = \bar{\theta}_{\mu''}$ *for some* $\mu' < \mu''$ *with* $\mu', \mu'' \in (\bar{\mu}, 1)$*, so that* $\bar{\theta}_{\bar{\mu}} < p' < p'' < \theta_h$*. Assume that* $\sigma_h(p'', e_h) = \frac{\mu''(1-\tilde{\mu})}{\tilde{\mu}(1-\mu'')}$*,* $\sigma_h(p', e_h) = \frac{\tilde{\mu}-\mu''}{\tilde{\mu}(1-\mu'')}$*, and* $\sigma_l(p'', e_h) = 1$*. Moreover, let* $\beta(b|p'', \mu'') = \frac{p'}{p''}$*,* $\beta(n|p'', \mu'') = \frac{p''-p'}{p''}$ *and* $\beta(b|p', \mu') = 1$*. Then, for* $\tilde{\mu} > \mu''$ *and* $\alpha \geq \theta_l/\bar{\theta}_{\mu'}$*, these strategies* σ *and* β *can be sustained in a PBE yielding payoffs* $\pi_h^* = p' = \bar{\theta}_{\mu'}$ *and* $\pi_l^* = \alpha p' = \alpha\bar{\theta}_{\mu'}$*. Any such PBE survives the belief refinement, since* $\pi_h^* = \pi_l^*/\alpha > \bar{\theta}_{\bar{\mu}}$*, but it does not survive the payoff restriction when compared to the feasible pooling(*e_h*)-b equilibrium which yields* $\pi_h^* = \bar{\theta}_{\tilde{\mu}} > \bar{\theta}_{\mu'}$ *and* $\pi_l^* = \alpha\bar{\theta}_{\tilde{\mu}} > \alpha\bar{\theta}_{\mu'}$ *since* $\tilde{\mu} > \mu'' > \mu'$*.*

In this example, the high type randomizes between two prices, p' and p'', while the low type always chooses the higher price, p''. When the buyer observes the lower price, she recognizes the good's high quality and buys with certainty. To make the high type indifferent between posting p' and p'', the buyer must purchase at p'' with a reduced probability, namely p'/p'', which also deters the low type from deviating to p''. However, the high type's payoff equals the lower price p' while the low type receives $\alpha p'$. Since the buyer's indifference condition requires $\tilde{\mu} > \mu'' > \mu'$, this outcome fails the payoff condition: without mixing, the high type could secure a payoff of $\bar{\theta}_{\tilde{\mu}}$ (which is greater than $\bar{\theta}_{\mu'} = p'$) by pooling with the low type on $\bar{\theta}_{\tilde{\mu}}$. A similar result arises in the final example where both sellers randomize.

Example 4. *Let* $p' = p_{ba}(\mu')$ *and* $p'' = p_{ba}(\mu'')$ *for some* $\mu' < \mu''$ *with* $\mu', \mu'' \in [\underline{\mu}, \bar{\mu}]$*, so that* $\bar{\theta}_{\underline{\mu}} \leq p' < p'' \leq \bar{\theta}_{\bar{\mu}}$*. Assume that* $\sigma_h(p'', e_h) = \frac{\mu''(\tilde{\mu}-\mu')}{\tilde{\mu}(\mu''-\mu')}$*,* $\sigma_h(p', e_h) = \frac{\mu'(\mu''-\tilde{\mu})}{\tilde{\mu}(\mu''-\mu')}$*,* $\sigma_l(p'', e_h) = \frac{(1-\mu'')(\tilde{\mu}-\mu')}{(1-\tilde{\mu})(\mu''-\mu')}$*, and* $\sigma_l(p', e_h) = \frac{(1-\mu')(\mu''-\tilde{\mu})}{(1-\tilde{\mu})(\mu''-\mu')}$*. Moreover, let* $\beta(ba|p'', \mu'') = \beta(ba|p', \mu')\frac{p'}{p''}$*, with* $\beta(ba|p', \mu') \in (0, 1]$*. Then, for* $\tilde{\mu} \in (\mu', \mu'')$ *and* $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu'))$*, these strategies* σ *and* β *can be sustained in a PBE yielding payoffs* $\pi_h^* = \beta(b|p', \mu')p' = \beta(b|p', \mu')p_{ba}(\mu')$ *and* $\pi_l^* = \alpha\varepsilon\beta(b|p', \mu')p' = \alpha\varepsilon\beta(b|p', \mu')p_{ba}(\mu')$*. Any such PBE survives the belief refinement, since* $\pi_h^* = \pi_l^*/(\alpha\varepsilon)$*, and meets the intra-category payoff restriction when* $\beta(b|p', \mu') = 1$*, which implies that* $\pi_h^* = p' = p_{ba}(\mu')$ *and* $\pi_l^* = \alpha\varepsilon p' = \alpha\varepsilon p_{ba}(\mu')$*. However, it does not satisfy the inter-category payoff restriction when compared to the feasible pooling(*e_h*)-ba equilibrium, which yields* $\pi_h^* = p_{ba}(\tilde{\mu}) > p_{ba}(\mu')$ *and* $\pi_l^* = \alpha\varepsilon p_{ba}(\tilde{\mu}) > \alpha\varepsilon p_{ba}(\mu')$ *because* $\tilde{\mu} > \mu'$*.*

When both seller types mix over two prices p', p'' above the authentication range, the indifference conditions for the low and high types coincide, leaving a single degree of freedom in determining the buyer's optimal strategy. Accordingly, we can define the equilibrium purchasing strategy $\beta(ba|p'', \mu'')$ as the product of $\beta(b|p', \mu')$ and $\frac{p'}{p''}$, ensuring that the purchasing probability decreases as the price increases. To satisfy the payoff restrictions, we then require $\beta(ba|p', \mu') = 1$, which maximizes the payoff for both seller types. Nevertheless, even the highest possible payoffs in this outcome, $p_{ba}(\mu')$ for the high type and $\alpha\varepsilon p_{ba}(\mu')$ for the low type, are lower than what the sellers would obtain by simply pooling on $p_{ba}(\tilde{\mu})$. This is analogous to the third example, and it is feasible since the proposed equilibrium requires that $\tilde{\mu} \in (\mu', \mu'')$.

Examples 3 and 4 illustrate that mixed-strategy equilibria where the low type does not mix at θ_l—and hence obtains a weakly larger payoff—are suboptimal for the seller whenever they are feasible. In fact, for any parameter configuration that permits such equilibria, there exists an alternative pure-strategy pooling equilibrium based on the posterior belief $\tilde{\mu}$, either with or without authentication acquisition, that yields a higher payoff for both types. The intuition is that whenever the high type randomizes and the low type excludes θ_l from her support, the sellers' equilibrium payoffs are capped at the smallest common price, say p_1. However, because the initial belief requirement for such equilibria is that $\tilde{\mu} > \mu_1$, these non-standard mixed strategies can be interpreted as a sort of "backward" randomization: instead of pooling exclusively on the highest possible price given $\tilde{\mu}$—i.e., $p_{ba}(\tilde{\mu})$ or $\bar{\theta}_{\tilde{\mu}}$—the sellers randomize over at least one lower price, namely p_1, which nonetheless limits their maximum attainable payoff. Consequently, these equilibria are not rationalizable under the proposed seller-optimality restriction. The following proposition formalizes this result.

Proposition 13. *Any mixed-strategy PBE* $\{\sigma_h, \sigma_l, \beta, \mu\}$ *with* $P_h = \{p_1, p_2, ..., p_k\}$, *where* $p_1 < p_2 < ... < p_k$ *and* $k \geq 2$, *that satisfies both the belief and the payoff restrictions exhibits:* $P_l = P_h \cup \{\theta_l\}$, $\sigma_h(p_j, e_h) > 0$ *for every* $p_j \in P_h$, $\pi_l^* = \theta_l$, *and necessary condition* $\tilde{\mu} < \mu_1$. *Moreover:*

(i) either $p_j \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ *for every* $p_j \in P_h$, $\beta(b|p_j, \mu_j) = \frac{\theta_l}{\alpha p_j}$, $\beta(ba|p_j, \mu_j) = 0$, $\pi_h^* = \theta_l/\alpha$, *with necessary condition* $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\bar{\mu}})$;

(ii) or $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ *for every* $p_j \in P_h$, $\beta(ba|p_j, \mu_j) = \frac{\theta_l}{\alpha\varepsilon p_j}$, $\beta(b|p_j, \mu_j) = 0$, $\pi_h^* = \theta_l/(\alpha\varepsilon)$, *with necessary condition* $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), 1]$.

The only non-standard mixed-strategy PBEs that satisfy both restrictions are those in which the low type assigns positive probability to θ_l. Their essence is very different from that of the non-standard mixed equilibria illustrated in Examples 3 and 4, where the low type never reveals her type but always mimics the high-quality seller. Similarly to the standard mixed equilibria, we can interpret the surviving non-standard mixed equilibria of Proposition 13 as "semi-separating," emphasizing their role in differentiating the two types when pure separating equilibria are not feasible.

Moreover, the similarity between these selected non-standard mixed equilibria and the previously considered standard mixed is not merely superficial but essential. In fact, by examining Corollary 3 (points (i) and (iii)), it is clear that any non-standard equilibrium with $|P_h| \geq 2$ described in Proposition 13, whenever feasible, is equivalent to a corresponding standard mixed equilibrium in which the high type assigns probability one to the minimum common price in P_h, i.e. p_1; in particular, the buyer's strategy, relative to (p_1, μ_1), and the payoffs are identical in the two equilibria.[26]

Corollary 4. *For any non-standard mixed-strategy PBE* $\{\hat{\sigma}_h, \hat{\sigma}_l, \hat{\beta}, \hat{\mu}\}$ *with* $|P_h| \geq 2$ *that satisfies both the belief and the payoff restrictions, there exists a corresponding standard mixed-strategy PBE* $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ *with* $\sigma_h^*(p_1, e_h) = 1$, $\beta^*(b|p_1, \mu_1) = \hat{\beta}(b|p_1, \mu_1)$, $\beta^*(ba|p_1, \mu_1) = \hat{\beta}(ba|p_1, \mu_1)$, *where* $p_1 = \min P_h$, *that is payoff-equivalent, i.e.,* $\pi_h^* = \hat{\pi}_h$ *and* $\pi_l^* = \hat{\pi}_l$.

In particular, the non-standard mixed equilibria identified in Proposition 13 point (i) are equivalent to the standard mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ equilibria, while those defined in point (ii) correspond to the mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibria. Note that standard mixed$(\bar{\theta}_{\bar{\mu}})$ equilibria do not have a non-standard counterpart. This is because the simultaneous presence of indifference conditions for both seller types restricts the buyer's purchasing strategy from assigning positive probability to both *ba* and *b* at the same price—a feature that is, instead, characteristic of those standard mixed equilibria.

In what follows, we will take advantage of Corollary 4 and, for simplicity, refer only to pure-strategy and standard mixed-strategy equilibria, since the latter encompass all surviving non-standard mixed-strategy equilibria.

---

[26] Note that p^*, the price uniquely posted by the high type in standard mixed equilibria, need not be strictly equal to p_1, the minimum price posted with positive probability by the high type in non-standard mixed equilibria, for the two equilibrium types to be equivalent: in fact, any price within the same interval to which p_1 belongs, either $[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ or $(\bar{\theta}_{\bar{\mu}}, \theta_h)$, would suffice.

1.5 Welfare Analysis

Having identified the equilibria that survive our refinements and mapped out their feasibility regions, we now turn to their economic consequences. In particular, we are interested in understanding how the precision of the evidence and authentication technologies influence welfare and fraud incidence, and, ultimately, which equilibria maximize total welfare and minimize fraud and under which conditions. We define welfare as the sum of buyer and seller surplus, and fraud probability as the chance that a low-quality good sells at a price above its full-information value (Martinez-Gorricho).[27]

Definition 2. *For any PBE $E = \{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ characterized in Proposition 11, total welfare, i.e., the ex-ante equilibrium gains from trade, is given by*

$$\begin{aligned} W_E =& \mu_0 \left[\beta^*(b|p^*, \mu^*) + \beta^*(ba|p^*, \mu^*)\right] \theta_h + (1 - \mu_0)\sigma_l^*(\theta_l, e)\theta_l \\ &+ (1 - \mu_0)\sigma_l^*(p^*, e_h)\alpha \left[\beta^*(b|p^*, \mu^*) + \varepsilon\beta^*(ba|p^*, \mu^*)\right] \theta_l, \end{aligned} \tag{1.22}$$

while the probability of fraud is given by

$$F_E = (1 - \mu_0)\sigma_l^*(p^*, e_h)\alpha \left(\beta^*(b|p^*, \mu^*) + \varepsilon\beta^*(ba|p^*, \mu^*)\right). \tag{1.23}$$

Full efficiency corresponds to welfare equal to $\mu_0\theta_h + (1 - \mu_0)\theta_l$ and fraud probability zero. This occurs only under a separating equilibrium, i.e. $\beta^*(b \mid p^*, \mu^*) = 1$ and $\sigma_l^*(\theta_l, e) = 1$, which forces $\sigma_l^*(p^*, e_h = 0$. When the outcome differs from full separation, meaning $\sigma_l^*(p^*, e_h) > 0$, fraud probability is strictly positive, and welfare falls below its maximum. However, the relationship between welfare and fraud varies with the type of equilibrium under consideration. Using the equilibrium characterizations from the previous section, we now derive expressions for both welfare and fraud probability under each equilibrium type. From now on, we will label pooling equilibria solely by the buyer's equilibrium action—"pooling–b" when she declines authentication and "pooling–ba" when she purchases it—omitting the explicit "e_h" denomination, given that no pooling equilibria where sellers provide low evidence satisfy the belief refinement.

---

[27] Here, welfare excludes any authentication price paid by the buyer or provision costs; the latter are normalized to zero throughout the paper, the former are omitted since they represent pure transfers, as it will become clear once we introduce the authenticator as a strategic player.

Lemma 8. *For any PBEs characterized in Proposition 11, welfare and probability of fraud are given by:*

(i) separating: $W_s = \mu_0\theta_h + (1-\mu_0)\theta_l$, $F_s = 0$;

(ii) pooling-ba: $W_{ba} = \mu_0\theta_h + (1-\mu_0)\alpha\varepsilon\theta_l$, $F_{ba} = (1-\mu_0)\alpha\varepsilon$;

(iii) pooling-b: $W_b = \mu_0\theta_h + (1-\mu_0)\alpha\theta_l$, $F_b = (1-\mu_0)\alpha$;

(iv) mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$: $W_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]} = \frac{\mu_0\theta_l\theta_h}{\alpha\varepsilon p_{ba}(\mu')} + (1-\mu_0)\theta_l - \frac{\mu_0(1-\mu')(p_{ba}(\mu')-\theta_l)\theta_l}{\alpha\mu' p_{ba}(\mu')}$, $F_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]} = \frac{\mu_0(1-\mu')\theta_l}{\alpha\mu' p_{ba}(\mu')}$;

(v) mixed$(\bar{\theta}_{\bar{\mu}})$: $W_{(\bar{\theta}_{\bar{\mu}})} = \mu_0\theta_h + (1-\mu_0)\theta_l - \frac{\mu_0(1-\bar{\mu})\Delta\theta}{\alpha\bar{\theta}_{\bar{\mu}}}\theta_l$, $F_{(\bar{\theta}_{\bar{\mu}})} = \frac{\mu_0(1-\bar{\mu})\theta_l}{\alpha\bar{\mu}\bar{\theta}_{\bar{\mu}}}$;

(vi) mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$: $W_{(\bar{\theta}_{\bar{\mu}},\theta_h)} = \frac{\mu_0\theta_l}{\alpha} + (1-\mu_0)\theta_l$, $F_{(\bar{\theta}_{\bar{\mu}},\theta_h)} = \frac{\mu_0(1-\mu')\theta_l}{\alpha\mu'\bar{\theta}_{\mu'}}$,

where $p' = p_{ba}(\mu')$ *is the semi-pooling price in standard mixed strategy equilibria, and* μ' *is the corresponding target posterior belief.*

Although Proposition 11 pinpoints a unique equilibrium for most $(\tilde{\mu}, \alpha)$ pairs, the fact that multiple equilibrium types occupy different regions makes comparative-statics nontrivial: the same parameters we wish to vary—in particular the precision levels α and ε—also appear inside the very existence conditions for each equilibrium. For example, lowering authentication precision (higher ε) shrinks the range of $\tilde{\mu}$ in which a pooling–ba equilibrium exists, while at the same time loosening the minimum α needed for that equilibrium. Thus, any statement about how welfare or fraud responds to a change in one parameter must be read as "all else equal—and within the same equilibrium regime."

Let us consider now how welfare and fraud probability respond as evidence precision (α) and authentication precision (ε) vary. Separating equilibrium is full efficient independently of either precision parameter. Instead, pooling outcomes exhibit a similar comparative statics. Raising the precision of the relevant signal (evidence in pooling–b; evidence and authentication in pooling–ba) reduces the chance that a cheating low type succeeds, so fraud falls—but because every high-quality good trades with probability one, welfare also falls: in a pooling equilibrium more low-type sales always correspond to more fraud, and hence to higher welfare. Therefore, maximizing welfare comes at the expense of increased fraud.[28] Crucially, note that, since authentication is never purchased in equilibrium, in pooling-b equilibria neither welfare nor fraud ever depends on its precision.

---

[28] This insight is consistent with the results in Martinez-Gorricho (2020).

Mixed-strategy equilibria display the reverse pattern on fraud, confirming their "semi-separating" nature. As evidence becomes more precise (lower α) fraud increases. The reason for this counterintuitive result lies in the fact that a more precise evidence signal reduces the probability that a cheating low type passes the test. However, it also increases both the probability of purchasing the good (with or without authentication depending on the specific mixed equilibrium) by the buyer, and the probability of cheating: a more precise evidence increases buyer confidence, which in turn induces the seller to pool more aggressively so that she reveals his type less often.

Improved authentication precision, by contrast, only reduces fraud for mixed($\bar{\theta}_{\bar{\mu}}$), because there the buyer always buys the good, either with or without authentication (with a probability of buying it with authentication that increase in its precision). Therefore, when authentication signal becomes more accurate, the increased probability of acquiring the authentication does not compensate and the overall probability for a cheating low type of selling the good decreases, meaning lower frauds.

The effect of evidence precision on welfare is subtler; in particular higher evidence accuracy increases fraud and possibly the probability of buying without authentication, while it reduces the incentives for the low quality seller to behave fraudulently. The latter dominate, so that welfare decreases, for mixed($\bar{\theta}_{\bar{\mu}}$) where the high-quality good is always bought and a increased probability of a plain purchase carries no effects; however, it is dominated, so welfare increases, for mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$) since the high-quality is traded more often.

Notably, greater authentication precision (lower ε) never reduces welfare in any mixed equilibrium and in fact strictly raises welfare in both mixed($\bar{\theta}_{\bar{\mu}}$)) and mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibria while leaving mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$) equilibria unaffected since in that case authentication remains off the equilibrium path.

Lemma 9. *In both mixed($\bar{\theta}_{\bar{\mu}}$) and mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibria, welfare is strictly increasing in the authentication precision, i.e.,* $\frac{\partial W_{(\bar{\theta}_{\bar{\mu}})}}{\partial \varepsilon}, \frac{\partial W_{[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]}}{\partial \varepsilon} < 0$.

A naive welfare–fraud comparison, in that it abstracts from each equilibrium's feasibility constraints, delivers three immediate insights. First, the separating equilibrium strictly dominates on both dimensions: it maximizes welfare and drives fraud to zero. Second, between the two pooling outcomes, pooling–b yields both higher welfare and a higher fraud probability than pooling–ba. Third, matching pooling–b's welfare under pooling–ba would require $\varepsilon = 1$—but Corollary 1 shows that $\varepsilon = 1$ is incompatible with any nontrivial authentication technology unless $c = 0$, in which case pooling–ba collapses into pooling–b. Hence, for all admissible $\varepsilon < 1$, pooling–ba unambiguously implies a lower welfare level relative to pooling–b. Combining this observation with Lemma 9 suggests that total welfare is maximized only at the extremes of the authentication precision spectrum: either with perfectly precise authentication ($\varepsilon = 0$) or with authentication so noisy—namely $\varepsilon > 1 - \frac{4c}{\Delta\theta}$—that it is effectively never purchased, or equivalently not present.

We now pose the planner's problem. Take as given the prior μ_0 and the quality levels θ_l, θ_h. A social planner would like to choose the two imprecision parameters—evidence imprecision α and authentication imprecision ε— and the authentication price—c— so as to maximize total welfare and minimize fraud. Two trivial corner solutions guarantee full efficiency: either reduce α below the separation threshold θ_l/θ_h, in which case the separating equilibrium obtains regardless of ε; or set $\varepsilon = 0$ and $c = 0$, so that perfectly precise and free authentication restores separation. Both of these, however, sidestep the real trade-offs when authentication carries a strictly positive fee. Accordingly, we focus on the more interesting and realistic case in which evidence precision $\alpha > \theta_l/\theta_h$ and the authentication price $c > 0$ are exogenously given, and ask: how should the planner choose authentication precision ε to best improve welfare and cut fraud?[29] In the next subsection we will motivate this assumption—and endogenize both ε and c—by introducing an authenticator as a strategic player.

As noted, varying model parameters shifts feasibility thresholds and can alter the equilibrium outcome. This sensitivity is especially pronounced for authentication precision, ε. We therefore focus on two extreme cases: perfect authentication, where $\varepsilon = 0$; and effectively no authentication—i.e. $\varepsilon > 1 - \frac{4c}{\Delta\theta}$, which ensures that authentication is never acquired in equilibrium.

In the perfect-authentication case, $\bar{\mu}$ and $\bar{\theta}_{\bar{\mu}}$ attain their maximal values; denote these by $\bar{\mu}^*$ and $\bar{\theta}_{\bar{\mu}^*}$. Then, by Proposition 11, in addition to the separating equilibrium for

---

[29] Equivalently, one could allow the planner to subsidize authentication (driving c toward zero) or to invest in evidence, but holding both α and c fixed sharpens the trade-off over ε and highlights the limits of authentication as a stand-alone remedy.

$\alpha \leq \frac{\theta_l}{\theta_h}$, the equilibrium outcome is:

(i) pooling-b, whenever $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ and $\tilde{\mu} \in [\bar{\mu}^*, 1)$;

(ii) mixed($\bar{\theta}_{\bar{\mu}^*}$), whenever $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}^*}, 1]$ and $\tilde{\mu} \in (0, \bar{\mu}^*]$;

(iii) mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$) or a separating, whenever $\alpha \in (\theta_l/\theta_h, \min\{\theta_l/\bar{\theta}_{\bar{\mu}^*}, \theta_l/\bar{\theta}_{\tilde{\mu}}\})$.

In the no-authentication case, $\underline{\mu}$, $\bar{\mu}$, $\bar{\theta}_{\underline{\mu}}$ and $\bar{\theta}_{\bar{\mu}}$ are not defined, and clearly there is no authentication-region for the prior. Proposition 11 then implies, in addition to the separating region for $\alpha \leq \theta_l/\theta_h$, the equilibrium outcome to be:

(i) pooling-b, whenever $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$;

(ii) mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$) or a separating, whenever $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\tilde{\mu}})$.

Because equilibrium is overall indeterminate when $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\tilde{\mu}})$, we henceforth restrict our analysis to $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$.[30] Moreover, since the outcome is equivalently a pooling-b when $\tilde{\mu} \in [\bar{\mu}^*, 1)$, we will restrict our analysis to $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ and $\tilde{\mu} \in [0, \bar{\mu}^*]$.

One might hope that an authentication technology of intermediate precision could outperform both the perfectly precise and the no-authentication cases. However, when $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$ and $\tilde{\mu} \in [0, \bar{\mu}^*]$, Proposition 11 limits the feasible equilibria to exactly four types—pooling-b, pooling-ba, mixed($\bar{\theta}_{\bar{\mu}}$) and mixed[$\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$]. It can be shown that even the best possible mixed[$\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$] outcome delivers no more welfare than the mixed($\bar{\theta}_{\bar{\mu}^*}$) equilibrium or the pooling-ba equilibrium at their respective boundary values, and that whenever the authentication-fee c is strictly positive, pooling-ba is strictly welfare-inferior to mixed($\bar{\theta}_{\bar{\mu}^*}$).[31] Lemma 9 together with this observation implies that the mixed($\bar{\theta}\_\bar{\mu}^*$) equilibrium delivers strictly higher welfare than any equilibrium with intermediate authentication precision. Indeed, reducing ε both raises total welfare and raises the maximum α allowed for mixed($\bar{\theta}_{\bar{\mu}}$), thereby expanding its feasible region. Consequently, no authentication technology short of perfect accuracy can outperform flawless authentication in terms of welfare.

---

[30] Note that in this omitted region, perfect authentication may deliver higher welfare and lower fraud incidence than mixed($\bar{\theta}_{\bar{\mu}}, \theta_h$), or it may be dominated—both in welfare and in fraud minimization—by the separating equilibrium. Given this strong indeterminacy, we avoid further detail.

[31] Here "best" mixed[$\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$] refers to choosing $p_{\mu'} = \theta_l/(\alpha\varepsilon)$, and "best" pooling-ba corresponds to setting $\varepsilon = 1 - \frac{c}{\bar{\mu}(1-\bar{\mu})\Delta\theta}$ in each case maximizing welfare subject to the equilibrium's existence conditions.

Finally, we state our main welfare result: perfect authentication is optimal for minimizing fraud, while in certain parameter regions a no-authentication regime can achieve strictly higher total welfare.

Proposition 14. *Given* $(\mu_0, \theta_h, \theta_l, \alpha, c)$ *such that* $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$, $\tilde{\mu} \in [0, \bar{\mu}^*]$, *and* $c > 0$, *then:*

(i) the perfect authentication technology, $\varepsilon = 0$, *outcome, mixed(*$\bar{\theta}_{\bar{\mu}^*}$*), minimizes fraud probability;*

(ii) the no-authentication outcome, pooling-b, maximizes welfare if and only if

$$\alpha \geq 1 - \frac{\tilde{\mu}(1 - \bar{\mu}^*)\Delta\theta}{(1 - \tilde{\mu})\bar{\theta}_{\bar{\mu}^*}}. \tag{1.24}$$

Our welfare analysis delivers two complementary policy insights. First, perfect authentication unambiguously minimizes the probability of fraud—no other level of accuracy can reduce cheating from the low-quality buyer more. Second, however, when the evidence technology is too imprecise to sustain full separation, forgoing authentication entirely can yield strictly higher total welfare than any imperfect-authentication regime. In other words, while flawless authentication is best at deterring fraud, a no-authentication policy may better maximize surplus whenever seller-provided evidence remains uninformative. In particular, we have an even stronger welfare result whenever the evidence signal is completely uninformative.

Corollary 1. *If the evidence technology is completely uninformative,* $\alpha = 1$, *then perfect authentication yields strictly lower welfare than the no-authentication case.*

The key insight comes from comparing how welfare responds to rising fraud under the two equilibria: the semi-separating mixed($\bar{\theta}_{\bar{\mu}^*}$) outcome with perfect authentication, and the pooling-without-authentication outcome. In both cases, as the seller's evidence precision falls, cheating becomes easier, fraud rises, and total welfare increases. But the welfare gain is smaller under the mixed equilibrium than under pooling.

With perfect authentication, lower evidence precision not only makes it more likely for a low-quality seller to pass the test, but also dampens her incentive to cheat—because buyers acquire authentication more often and honest pricing becomes more attractive—so these opposing forces partially offset one another. By contrast, in a pooling equilibrium without authentication, a drop in evidence precision leaves buyer behavior and the seller's strategy unchanged; the only effect is more fraud, which mechanically boosts trade volume and welfare.

At the extreme, when evidence is entirely uninformative, pooling without authentication achieves the full-separation welfare benchmark—every trade occurs—but at the cost of maximal fraud. Perfect authentication can never reach that benchmark, since it always leads some goods to be rejected by the authentication technology. Therefore there exists a threshold level of evidence imprecision above which perfect authentication delivers strictly lower welfare than the no-authentication regime.

Note that the key assumption driving this result is that authentication must carry a strictly positive price for the buyer. We did not rely on any cost-minimization motive—indeed all production, evidence and authentication prices were normalized to zero—so allowing for any positive provision cost would only strengthen the welfare advantage of a no-authentication regime, since it avoids that extra outlay altogether.

These findings carry straightforward policy lessons. Once evidence precision and buyer priors fall into the region where perfect authentication still leaves welfare strictly below the no-authentication benchmark, a social planner would be better off banning paid authentication (or forcing it to be provided at zero price). Conversely, when evidence is precise enough to sustain at least semi-separating outcomes, ensuring access to a high-precision authentication technology (even at positive cost) becomes the optimal tool to minimize fraud. In this way, our analysis directly inform when regulators should promote, limit or eliminate buyer-purchased authentication.

1.6 Discussion and Conclusions

Our analysis reveals a subtle complementarity—and at times conflict—between seller disclosure and buyer authentication. Although perfect authentication is highly effective at deterring fraud, it does not necessarily maximize social surplus when seller-provided evidence remains imprecise. These findings provide a clear framework for platforms and regulators to calibrate authentication requirements alongside transaction volume, so as to control fraud without unduly suppressing beneficial trade.

In this chapter, we examined a simple bilateral-trade setting in which a seller, privately informed about her good's quality, may disclose noisy evidence, and a buyer may purchase a noisy third-party authentication before deciding whether to trade. Our goal was to tease out how these two imperfect information channels interact to shape equilibrium prices, the likelihood of fraud, and overall welfare.

We begin by showing that precise seller disclosure is irreplaceable: no matter how accurate or cheap the authentication technology, it cannot on its own sustain full separation unless the seller's evidence signal is already sharp enough to distinguish high from low quality. Authentication can only deter low-quality sellers if buyers buy the test with positive probability, and that very act undermines the test's credibility as a threat.

When seller evidence falls short of this precision threshold, introducing an imperfect authentication service does help reduce the incidence of fraud, but it can never restore the first-best outcome. In every equilibrium where disclosure is too noisy, total welfare remains strictly below its maximum, and some low-quality goods continue to sell at inflated prices regardless of how buyers pay to verify.

Perhaps most surprisingly, we find that perfect authentication—while minimizing fraud to the greatest extent possible—can actually harm welfare whenever seller evidence is highly uninformative. In that regime, aggregate surplus rises with total trade volume, even if much of it comes from fraudulent transactions, and any positive testing fee simply suppresses mutually beneficial exchanges. Thus a market without authentication can deliver strictly higher welfare than one offering flawless but costly verification.

To sharpen our comparative-statics and pin down exactly when each of these regimes arises, we refine equilibrium selection by ruling out implausible beliefs and discarding equilibria that both types of seller would jointly prefer to avoid. This exercise yields a nearly complete map of six robust equilibrium patterns—ranging from full separation through pure pooling to three semi-separating mixed-strategy outcomes—across the two key parameters governing evidence and authentication precision.

Taken together, these results offer clear guidance for platform designers and regulators. Improving the verifiability of seller disclosures is the only surefire way to eliminate fraud, whereas buyer-paid testing can at best serve as a partial backstop. When disclosure remains weak, perfect authentication is the most effective tool against deception but may come at the cost of lost surplus. Capping authentication fees can expand the range in which testing is both attractive to buyers and welfare-enhancing, up to the point where its marginal benefit in fraud reduction outweighs its marginal cost in foregone trade.

Looking ahead, it would be natural to endogenize both the precision of seller evidence and the accuracy and price of authentication by introducing profit-maximizing intermediaries or platforms. Extending the model to many sellers or a dynamic setting would also allow reputation and competition to shape incentives for disclosure and testing. Finally, enriching the quality distribution beyond two types or allowing more complex signal structures could bring the framework even closer to the realities of large-scale trade platforms.

By weaving together voluntary disclosure and buyer authentication in one unified framework, we hope to clarify when these two imperfect tools reinforce one another and when they work at cross-purposes, and to provide concrete lessons for striking the right balance between authentication and trade volume in markets beset by low transparency.

1.7 References

Argenziano, Rossella, Severinov, Sergei, and Squintani, Francesco. 2016. “Strategic Information Acquisition and Transmission.” *American Economic Journal: Microeconomics* 8(3):119-55.

Bergemann, Dirk, and Valimaki, Juuso. 2000. “Information Acquisition and Efficient Mechanism Design.” Yale University, Department of Economics Working Paper No. 1248.

Bester, Helmut, Lang, Matthias, and Li, Jianpei. 2019. “Signaling versus Auditing.” CESifo Working Paper Series No. 7183.

Bester, Helmut, and Ritzberger, Klaus. 2001. “Strategic Pricing, Signalling, and Costly Information Acquisition.” *International Journal of Industrial Organization* 19(9):1347-1361.

Capozza, Francesco, Haaland, Ingar, Roth, Christopher, and Wohlfart, Johannes. 2021. “Studying Information Acquisition in the Field: A Practical Guide and Review.” ECONtribute Discussion Papers Series 124, University of Bonn and University of Cologne, Germany.

Dranove, David, and Zhe Jin, Ginger. 2010. “Quality Disclosure and Certification: Theory and Practice.” *Journal of Economic Literature* 48(4):935-963.

Ellingsen, Tore. 1997. “Price Signals Quality: The Case of Perfectly Inelastic Demand.” *International Journal of Industrial Organization* 16(1):43-61.

Figueroa, Nicolás, and Guadalupi, Carla. 2022. “Price Signaling with Information Acquisition.” Working paper.

Gabaix, Xavier, Laibson, David, Moloche, Guillermo, and Weinberg, Stephen. 2006. “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model.” *American Economic Review* 96(4):1043-1068.

Gertz, Christopher. 2014. “Quality Uncertainty with Imperfect Information Acquisition.” Center for Mathematical Economics Working Papers 487, Center for Mathematical Economics, Bielefeld University.

Guan, Xiaolan, and Chen, Yehning. 2017. “The Interplay between Information Acquisition and Quality Disclosure.” *Production and Operations Management* 26:389-408.

Hong, Xianpei, Cao, Xinlu, Gong, Yeming, and Chen, Wanying. 2021. “Quality Information Acquisition and Disclosure with Green Manufacturing in a Closed-loop Supply Chain.” *International Journal of Production Economics* 232:107997.

Jackson, Matthew O. 1991. “Equilibrium, Price Formation, and the Value of Private Information.” *Review of Financial Studies* 4(1):1-16.

Marquez, Robert S., and Hauswald, Robert B.H. 2002. “Competition and Strategic Information Acquisition in Credit Markets.”

Martin, Daniel. 2017. “Strategic Pricing with Rational Inattention to Quality.” *Games and Economic Behavior* 104:131-145.

Martinez-Gorricho, S. 2020. “Signalling, Information and Consumer Fraud.” *Games* 11(3):29.

Matthews, Steven A., and Persico, Nicola G. 2005. “Information Acquisition and the Excess Refund Puzzle.”

Moscarini, Giuseppe, and Ottaviani, Marco. 2001. “Price Competition for an Informed Buyer.” *Journal of Economic Theory* 101:457-493.

Persico, N. 2000. “Information Acquisition in Auctions.” *Econometrica* 68(1):135-148.

Qingmin, Liu. 2011. “Information Acquisition and Reputation Dynamics.” *The Review of Economic Studies* 78(4):1400-1425.

Roesler, Anne-Katrin, and Szentes, Balázs. 2017. “Buyer-Optimal Learning and Monopoly Pricing.” *American Economic Review* 107(7):2072-2080.

Stahl, Konrad, and Strausz, Roland. 2017. “Certification and Market Transparency.” *Review of Economic Studies* 84(4):1842-1868.

Voorneveld, M., and Weibull, J.W. 2011. “A Scent of Lemon—Seller Meets Buyer with a Noisy Quality Observation.” *Games* 2(1):163-186.

1.8 Appendix

1.8.1 Buyer's optimal behavior given p

Lemma 2 specifies the buyer's optimal behavior based on price p given a fixed belief $\hat{\mu}$. The next lemma reverses the approach and describes the buyer's best response to any posterior belief given a specific price. Let us define

$$\underline{\mu}_{ba}(p) = \frac{c + \varepsilon(p - \theta_l)}{\theta_h - p + \varepsilon(p - \theta_l)} \text{ if } p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right], \tag{1.25}$$

and

$$\underline{\mu}_{b}(p) = \begin{cases} 1 - \dfrac{c}{(1-\varepsilon)(p - \theta_l)} & \text{if } p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \dfrac{p - \theta_l}{\Delta\theta} & \text{if } p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases} \tag{1.26}$$

The lemma below proves that $\underline{\mu}_{ba}(p)$ is the posterior, given a price p, for which the buyer is indifferent between ba and n, while $\underline{\mu}_{b}(p)$ for which she is indifferent (i) between b and n if $p \notin \left[\bar{\theta}_{\underline{\mu}}\bar{\theta}_{\bar{\mu}}\right]$ and (ii) between b and ba if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$.

Lemma 10. *Suppose the strategy β is optimal for the buyer given price p and belief $\hat{\mu}$. Then, (i) if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$,*

$$\beta(b|p,\hat{\mu}) \geq 0 \iff \hat{\mu} \geq \underline{\mu}_{b}(p),$$
$$\beta(n|p,\hat{\mu}) \geq 0 \iff \hat{\mu} \leq \underline{\mu}_{b}(p);$$

(ii) if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$,

$$\beta(b|p,\hat{\mu}) \geq 0 \iff \hat{\mu} \geq \underline{\mu}_{b}(p)$$
$$\beta(ba|p,\hat{\mu}) \geq 0 \iff \hat{\mu} \in \left[\underline{\mu}_{ba}(p), \underline{\mu}_{b}(p)\right]$$
$$\beta(n|p,\hat{\mu}) \geq 0 \iff \hat{\mu} \leq \underline{\mu}_{ba}(p)$$

Proof. Consider expressions (1.4), (1.5) and (1.6). First, $u(b|p,\hat{\mu}) \geq u(n)$ if and only if $p \leq \hat{\mu}\theta_h + (1-\hat{\mu})\theta_l$ or $\hat{\mu} \geq \frac{p-\theta_l}{\Delta\theta}$. Moreover, $u(b|p,\hat{\mu}) \geq u(ba|p,\hat{\mu})$ if and only if

$c \geq (1-\hat{\mu})(1-\varepsilon)(p-\theta_l)$ or $\hat{\mu} \geq 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}$. Hence, b is optimal whenever both conditions are satisfied:

$$\hat{\mu} \geq \max\left\{\frac{p-\theta_l}{\Delta\theta},\ 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}\right\}.$$

Second, $u(ba|p,\hat{\mu}) \geq u(n)$ if and only if $[\theta_h-\theta_l-(1-\varepsilon)(p-\theta_l)]\hat{\mu} \geq c+p-\theta_l-(1-\varepsilon)(p-\theta_l)$ or $\hat{\mu} \geq \frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}$. Moreover, $u(ba|p,\hat{\mu}) \geq u(b|p,\hat{\mu})$ if and only if $\hat{\mu} \leq 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}$. Hence, ba is optimal whenever both conditions are satisfied:

$$\hat{\mu} \in \left[\frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}, 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}\right]$$

Finally, $u(n) \geq u(b|p,\hat{\mu})$ if and only if $\hat{\mu} \leq \frac{p-\theta_l}{\Delta\theta}$, while $u(n) \geq u(ba|p,\hat{\mu})$ if and only if $\hat{\mu} \leq \frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}$. Hence, n is optimal whenever both conditions are satisfied:

$$\hat{\mu} \leq \min\left\{\frac{p-\theta_l}{\Delta\theta},\ \frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}\right\}.$$

Clearly, ba is optimal only if the interval $\left[\frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}, 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}\right]$ is not empty or equivalently, $\frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)} \leq 1-\frac{c}{(1-\varepsilon)(p-\theta_l)}$. The inequality of the second order in p is satisfied when $p \in [p_1, p_2]$ with $p_1 = \frac{\theta_h+\theta_l}{2}-\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\Delta\theta$ and $p_2 = \frac{\theta_h+\theta_l}{2}+\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\Delta\theta$. Condition 1.12 is the necessary and sufficient condition for the existence of p_1 and p_2. Note that the thresholds p_1 and p_2 are equivalent to the expected quality given posterior $\underline{\mu}$ and $\bar{\mu}$, respectively:

$$\begin{aligned}
\bar{\theta}_{\underline{\mu}} :&= \underline{\mu}\theta_h + (1-\underline{\mu})\theta_l \\
&= \left(\frac{1}{2}-\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\right)\theta_h + \left(\frac{1}{2}+\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\right)\theta_l \\
&= \frac{\theta_h+\theta_l}{2}-\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\Delta\theta \\
&= p_1,
\end{aligned}$$

and

$$\begin{aligned}\bar{\theta}_{\bar{\mu}} &:= \bar{\mu}\theta_h + (1-\bar{\mu})\theta_l\\ &= \left(\frac{1}{2}+\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\right)\theta_h + \left(\frac{1}{2}-\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\right)\theta_l\\ &= \frac{\theta_h+\theta_l}{2}+\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)\Delta\theta}}\Delta\theta\\ &= p_2.\end{aligned}$$

Therefore, when $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, ba is dominated, and b (n) is optimal if and only if $\hat{\mu} \geq (\ \leq\ )\frac{p-\theta_l}{\Delta\theta}$, thus proving point (ii). If, instead, $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ then

$$\begin{aligned}1-\frac{c}{(1-\varepsilon)(p-\theta_l)} &\geq \frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}\\ (1-\varepsilon)(p-\theta_l)[\theta_h-p+\varepsilon(p-\theta_l)]-c[\theta_h-p+\varepsilon(p-\theta_l)] &\geq (1-\varepsilon)(p-\theta_l)[c+\varepsilon(p-\theta_l)]\\ (1-\varepsilon)(p-\theta_l)(\theta_h-p) &\geq c\Delta\theta\end{aligned}$$

which implies both

$$\begin{aligned}1-\frac{c}{(1-\varepsilon)(p-\theta_l)} &\geq \frac{p-\theta_l}{\Delta\theta}\\ \theta_h-p &\geq \frac{c\Delta\theta}{(1-\varepsilon)(p-\theta_l)}\\ (1-\varepsilon)(p-\theta_l)(\theta_h-p) &\geq c\Delta\theta\end{aligned}$$

and

$$\begin{aligned}\frac{p-\theta_l}{\Delta\theta} &\geq \frac{c+\varepsilon(p-\theta_l)}{\theta_h-p+\varepsilon(p-\theta_l)}\\ (p-\theta_l)[\theta_h-p+\varepsilon(p-\theta_l)] &\geq c\Delta\theta+\varepsilon(p-\theta_l)\Delta\theta\\ (1-\varepsilon)(p-\theta_l)(\theta_h-p) &\geq c\Delta\theta\end{aligned}$$

which proves point (ii). □

1.8.2 Out-of-Equilibrium Beliefs Requirements

The next lemma specifies the general requirements for out-of-equilibrium beliefs necessary to sustain any equilibrium. We will resort to it extensively when characterizing the different equilibrium types and applying the refinement criterion over out-of-equilibrium beliefs.

Lemma 11. *Given equilibrium payoffs π_h^* and π_l^* for the high and the low type of seller, and any out-of-equilibrium price $p > \theta_l$, any PBE requires the following type-based restrictions on out-of-equilibrium beliefs $\mu_p^{e_l} = \mu(p, e_l, S_h^e)$ and $\mu_p^{e_h} = \mu(p, e_h, S_h^e)$:*

(i) (θ_h-e_l) If $p \leq \pi_h^*$, then $\mu_p^{e_l} \leq 1$; if $p > \pi_h^*$ and $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p > \pi_h^*$ and $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$.

(ii) (θ_h-e_h) If $p \leq \pi_h^*$, then $\mu_p^{e_h} \leq 1$; if $p > \pi_h^*$ and $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p > \pi_h^*$ and $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$.

(iii) (θ_l-e_l) If $p \leq \pi_l^*$, then $\mu_p^{e_l} \leq 1$; if $p > \pi_l^*$ and $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, or $p \in (\pi_l^*, \pi_l^*/\varepsilon] \cap [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p > \pi_h^*/\varepsilon$ and $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$.

(iv) (θ_l-e_h) If $p \leq \pi_l^*/\alpha$, then $\mu_p^{e_h} \leq 1$; if $p > \pi_l^*/\alpha$ and $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, or $p \in (\pi_l^*/\alpha, \pi_l^*/(\alpha\varepsilon)] \cap [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p > \pi_h^*/(\alpha\varepsilon)$ and $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$.

Proof. We apply Lemma 10 case by case.

(i) and (ii) Regardless of the precision of the two technologies and the type of evidence provided, the high-type seller gets p whenever the buyer chooses b or ba, and 0 otherwise. Therefore, the requirements for $\mu_p^{e_l}$ and $\mu_p^{e_h}$ must coincide, and any price representing a profitable deviation should induce n to the buyer. Any price $p \leq \pi_h^*$ is not a profitable deviation, hence $\mu_p^{e_l}, \mu_p^{e_h} \leq 1$, while any price $p > \pi_h^*$ requires $\mu_p^{e_l}, \mu_p^{e_h} < \underline{\mu}_b(p)$ if $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, and $\mu_p^{e_l}, \mu_p^{e_h} < \underline{\mu}_{ba}(p)$ if $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$.

(iii) When providing e_l, the low-type seller may face only the authentication test, which she passes with probability ε, i.e., the probability of receiving signal S_h^e when $\theta = \theta_l$. Therefore, deviations to (p, e_l) guarantees a payoff of p if the buyer chooses b, εp if the buyer chooses ba, and 0 otherwise. Clearly, any price $p \leq \pi_l^*$ is not a profitable deviation, hence $\mu_p^{e_l} \leq 1$. Prices above π_l^* for which authentication is never optimal, i.e. $p \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ requires n in order not to be profitable, so $\mu_p^{e_h} < \underline{\mu}_b(p)$. Any $p \in (\pi_l^*, \pi_l^*/\varepsilon]$ is a profitable deviation only if the buyer chooses b so they require $\mu_p^{e_h} < \underline{\mu}_b(p)$ whenever authentication can be optimal i.e. when $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$. Instead prices above π_l^*/ε represents a profitable deviation if the buyer chooses b or ba that they require $\mu_p^{e_h} < \underline{\mu}_b(p)$ if $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$.

(iv) Same as the previous case, but, since the low-type seller passes the evidence test with probability α, i.e, the probability of S_h^e when $e^l = e_h$, only prices above π_l^*/α might be profitable. In fact, deviations to (p, e_h) guarantees a payoff of α if the buyer chooses b, $\alpha\varepsilon p$ if the buyer chooses ba, and 0 otherwise. In particular, prices $p \in (\pi_l^*/\alpha, \pi_l^*/(\alpha\varepsilon)]$ are profitable only if authentication is not acquired, while prices above $\pi_l^*/(\alpha\varepsilon)$ provides a larger payoff regardless of whether the authentication is bought or not. The result follows immediately.

$\square$

Proofs

Proof of Lemma 1

Proof. Consider expressions (1.4), (1.5) and (1.6). First, $u(b|p,\hat{\mu}) \geq u(n)$ if and only if $p \leq \bar{\theta}_{\hat{\mu}}$. Moreover, $u(b|p,\hat{\mu}) \geq u(ba|p,\hat{\mu})$ if and only if $\bar{\theta}_{\hat{\mu}} - p \geq \bar{\theta}_{\hat{\mu}} - p - (1-\hat{\mu})(1-\varepsilon)(\theta_l - p) - c$ or $p \leq \theta_l + c(1-\hat{\mu})^{-1}(1-\varepsilon)^{-1}$. Hence, b is optimal whenever both conditions are satisfied:

$$p \leq \min\left\{\bar{\theta}_{\hat{\mu}},\ \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)}\right\}.$$

Second, $u(ba|p,\hat{\mu}) \geq u(n)$ if and only if $[\hat{\mu} + (1-\hat{\mu})\varepsilon]p \leq \hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c$ or $p \leq [\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c][\hat{\mu} + (1-\hat{\mu})\varepsilon]^{-1}$. Moreover, $u(ba|p,\hat{\mu}) \geq u(b|p,\hat{\mu})$ if and only if $p \geq \theta_l + c(1-\hat{\mu})^{-1}(1-\varepsilon)^{-1}$. Hence, ba is optimal whenever both conditions are satisfied:

$$p \in \left[\theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)},\ \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right]$$

Finally, $u(n) \geq u(b|p,\hat{\mu})$ if and only if $p \geq \bar{\theta}_{\hat{\mu}}$, while$u(n) \geq u(ba|p,\hat{\mu})$ if and only if $p \geq [\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c][\hat{\mu} + (1-\hat{\mu})\varepsilon]^{-1}$ Hence, n is optimal whenever both conditions are satisfied:

$$p \geq \max\left\{\bar{\theta}_{\hat{\mu}},\ \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon}\right\}.$$

□

Proof of Corollary 1

Proof. Given Lemma 1, the set of values of p for which ba is optimal is non-empty if and only if:

$$\begin{aligned}
\theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)} &\leq \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon} \\
(1-\hat{\mu})(1-\varepsilon)[\hat{\mu} + (1-\hat{\mu})\varepsilon]\theta_l + [\hat{\mu} + (1-\hat{\mu})\varepsilon]c &\leq (1-\hat{\mu})(1-\varepsilon)[\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c] \\
\hat{\mu}(1-\hat{\mu})(1-\varepsilon)\theta_l + [\hat{\mu} + (1-\hat{\mu})\varepsilon]c &\leq (1-\hat{\mu})(1-\varepsilon)[\hat{\mu}\theta_h - c] \\
c &\leq \hat{\mu}(1-\hat{\mu})(1-\varepsilon)\Delta\theta
\end{aligned}$$

which proves the claim. □

Proof of Lemma 2

Proof. From corollary 1, we have that ba can be optimal, when holding posterior $\hat{\mu}$, only if $\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$. Hence, when $\hat{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, the only options available are b and n. Expressions (1.4) and (1.5) immediately prove point (i). Instead, if $\hat{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, then $c \leq \hat{\mu}(1-\hat{\mu})(1-\varepsilon)\Delta\theta$ which implies both

$$\begin{aligned} \bar{\theta}_{\hat{\mu}} &\geq \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)} \\ \theta_l + \hat{\mu}\Delta\theta &\geq \theta_l + \frac{c}{(1-\hat{\mu})(1-\varepsilon)} \\ \hat{\mu}(1-\hat{\mu})(1-\varepsilon)\Delta\theta &\geq c \end{aligned}$$

and

$$\begin{aligned} \bar{\theta}_{\hat{\mu}} &\leq \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon} \\ \theta_l + \hat{\mu}\Delta\theta &\leq \frac{\hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c}{\hat{\mu} + (1-\hat{\mu})\varepsilon} \\ \hat{\mu}\theta_l + (1-\hat{\mu})\varepsilon\theta_l + \hat{\mu}^2\Delta\theta + \hat{\mu}(1-\hat{\mu})\Delta\theta &\leq \hat{\mu}\theta_h + (1-\hat{\mu})\varepsilon\theta_l - c \\ c &\leq \hat{\mu}\Delta\theta - \hat{\mu}^2\Delta\theta - \hat{\mu}(1-\hat{\mu})\Delta\theta \\ c &\leq \hat{\mu}\Delta\theta(1-\hat{\mu}-\varepsilon-\mu\varepsilon) \\ c &\leq \hat{\mu}(1-\hat{\mu})(1-\varepsilon)\Delta\theta. \end{aligned}$$

This, together with Lemma 1, proves point (ii). □

Proof of Proposition 1

Proof. Let us denote by p^h and p^l the price proposed by the high-type and the low-type seller respectively, so that $\sigma_h(p^h) = \sigma_l(p^l) = 1$. Suppose $p^h \neq p^l$ so that $\mu(p^h) = 1$ and $\mu(p^l) = 0$. The buyer buys if $p \leq \bar{\theta}_\mu$ which means $p^h \leq \theta_h$ and $p^l \leq \theta_l$. Clearly, any price strictly below θ_l and above θ_h is dominated for either type, so we limit our analysis to the price interval $[\theta_l, \theta_h]$. This implies $p^l = \theta_l$ and $p^h \in (\theta_l, \theta_h]$ which is not sustainable since the low type has an incentive to mimic the high type and deviate to p^h. Hence, any PBE must be a pooling equilibrium in which $p^h = p^l = p^*$. Then, $\mu(p^*) = \mu_0$ and the buyers only buys for prices $p \leq \bar{\theta}_{\mu_0}$ so it must hold $p^* \in \left[\theta_l, \bar{\theta}_{\mu_0}\right]$. Any deviation to a price $p \leq p^*$ would not be profitable for either seller. Deviating to $p > p^*$ is not profitable whenever the buyer decides not to buy the product at price p. Therefore, it must hold $\bar{\theta}_{\mu(p)} - p < 0$ for every $p > p^*$, where $\mu(p)$ is the posterior belief on the high-type assigned after observing price p. The condition can be written as:

$$\begin{aligned}
\bar{\theta}_{\mu(p)} - p &< 0 \\
\mu(p)\theta_h + (1 - \mu(p))\theta_l &< p \\
\theta_l + \mu(p)\Delta\theta &< p \\
\mu(p) &< \frac{p - \theta_l}{\Delta\theta}
\end{aligned}$$

□

Note: The following proofs characterize the equilibrium outcomes when $\varepsilon = 0$*. In such a case, the following hold:*

$$\begin{aligned}
p_b(\hat{\mu}) &= \theta_l + \frac{c}{1 - \hat{\mu}} \\
p_{ba}(\hat{\mu}) &= \theta_h - \frac{c}{\hat{\mu}} \\
\underline{\mu}_{ba}(p) &= \frac{c}{\theta_h - p} \\
\underline{\mu}_b(p) &= \begin{cases} 1 - \dfrac{c}{p - \theta_l} & \text{if } p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \dfrac{p - \theta_l}{\Delta\theta} & \text{if } p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}
\end{aligned}$$

Proof of Proposition 2

Proof. Consider a separating PBE in which the high-quality seller chooses (p^h, e_h), and the low-quality seller chooses (p^l, e_l). Given the assumptions on the evidence technology, (p^h, e_h) induces signal S_h^e with probability one, while (p^l, e_l) induces either S_h^e or S_l^e. The posterior beliefs of the buyer upon observing these two vectors and the evidence signal are $\mu(p^h, e_h, S_h^e) = 1$ and $\mu(p^l, e_l, S_h^e) = \mu(p^l, e_l, S_l^e) = 0$, respectively. The evidence signal is irrelevant since the two types choose different price-evidence combinations. From Lemma 2 we have that for any price p and posterior $\hat{\mu}$, in a pure-strategy equilibrium with $\varepsilon = 0$

it must hold $\beta(x|p,\hat{\mu}) = 1$ with

$$x = \begin{cases} b & \text{if } p \leq \bar{\theta}_{\hat{\mu}}, \\ n & \text{if } p > \bar{\theta}_{\hat{\mu}}; \end{cases} \qquad \text{if } \hat{\mu} \notin \Big[\underline{\mu}, \bar{\mu}\Big] \tag{1.27}$$

and

$$x = \begin{cases} b & \text{if } p \leq \theta_l + \frac{c}{1-\hat{\mu}} \\ ba & \text{if } \theta_l + \frac{c}{1-\hat{\mu}} < p \leq \theta_h - \frac{c}{\hat{\mu}} \\ n & \text{if } p > \theta_h - \frac{c}{\hat{\mu}} \end{cases} \qquad \text{if } \hat{\mu} \in \Big[\underline{\mu}, \bar{\mu}\Big] \tag{1.28}$$

Since both posterior beliefs are *extreme*, i.e. they do not belong to the interval $\Big[\underline{\mu}, \bar{\mu}\Big]$, action *ba* is dominated. Hence, the best response of the buyer to (p^h, e_h) and $\mu(p^h, e_h, S^e_h) = 1$ requires

$$x = \begin{cases} b & \text{if } p^h \leq \theta_h \\ n & \text{if } p^h > \theta_h \end{cases}$$

while the best response to (p^l, e_l) and $\mu(p^l, e_l, S^e) = 0$ with $S^e \in \{S^e_h, S^e_l\}$ requires

$$x = \begin{cases} b & \text{if } p^l \leq \theta_l \\ n & \text{if } p^l > \theta_l \end{cases}$$

Since we assume no cost of production for either good, any price above θ_h or below θ_l is strictly dominated. Any price above θ_h leads to no trade and zero profits, while all prices below θ_l lead to trade without authentication. Thus, both types could always deviate to (θ_l, e_l) and get $\theta_l > 0$ for sure, so that all equilibrium prices must lie in the interval $[\theta_l, \theta_h]$. Therefore, the optimal strategy for the seller requires $p^h = p^* \in (\theta_l, \theta_h]$ and $p^l = \theta_l$, which imply $\pi(\sigma_h|\beta, \theta_h) = \pi^*_h = p^*$ and $\pi(\sigma_l|\beta, \theta_l) = \pi^*_l = \theta_l$. Note that, in principle, any price of the high-quality seller above θ_l can be sustained in equilibrium given sufficiently pessimistic out-of-equilibrium beliefs. However, neither type of seller should have an incentive to deviate.

On-path deviations for the high type are clearly not profitable since mimicking the low type would lead to a profit of $\pi(\sigma_l|\beta, \theta_h) = \theta_l < p^* = \pi(\sigma_h|\beta, \theta_h)$. However the expected profit of the low type from deviating to (p^*, e_h) is $\pi(\sigma_h|\beta, \theta_l) = \alpha p^* + (1-\alpha)0$ since with probability $1-\alpha$ the evidence technology detects the low quality of the good leading the buyer not to buy the good. The incentive compatibility constraint for the low type, then, is $\theta_l \geq \alpha p^*$ which implies an upper bound to the equilibrium prices charged by the high type in equilibrium. Specifically, this upper bound for p^* is the minimum between θ_l/α and θ_h.

Let us now consider the off-equilibrium belief restrictions. The equilibrium payoffs for each seller type are given by $\pi_h^* = p^*$ and $\pi_l^* = \theta_l$. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are $\mu_p^{e_l} < \underline{\mu}_b(p)$ if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$ if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ for all $p > p^*$; there are no restrictions for prices $p \leq p^*$, i.e., $\mu_p^{e_l} \leq 1$. The exact same conditions apply to $\mu_p^{e_h}$ on the high type side. Instead, since $\varepsilon = 0$, the conditions for the low-type seller are simply $\mu_p^{e_l} < \underline{\mu}_b(p)$ for all $p > \theta_l$, and $\mu_p^{e_h} < \underline{\mu}_b(p)$ for all $p > \theta_l/\alpha$.

Since $p^* < \theta_l/\alpha$ and $\underline{\mu}_{ba}(p) < \underline{\mu}_b(p)$ on $[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, $\bar{\mu}_p^{e_h}$ is determined by the high-type restrictions for $p > p^*$ and is unrestricted, i.e. equals 1, otherwise. Conversely, $\bar{\mu}_p^{e_l}$ is determined by the restrictions of the low type for any $p \leq p^*$, and by the high type restrictions for any $p > p^*$, so that:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

$$\bar{\mu}_p^{e_l} = \begin{cases} \underline{\mu}_b(p) & \text{if } p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \text{ or } p \in [\theta_l, p^*] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

□

Proof of Proposition 3

Proof. A separating equilibrium in pure-strategy requires $(p^h, e^h) \neq (p^l, e^l)$ with $\sigma_h(p^h, e^h) = \sigma_l(p^h, e_h)$. Hence, there are four possible separating alternatives when considering the type of evidence provided by the two types of sellers in equilibrium:

1. (p^h, e_h), (p^l, e_l),
2. (p^h, e_l), (p^l, e_h),
3. (p^h, e_h), (p^l, e_h), with $p^h \neq p^l$,
4. (p^h, e_l), (p^l, e_l), with $p^h \neq p^l$,

where the first vector corresponds to (p^h, e^h) and the second one to (p^l, e^l). As shown in the proof of Proposition 2, the optimal strategy of the buyer in any separating equilibrium is characterized as follows: regardless of the actual realization of the evidence signal S^e, the best response to (p^h, e^h) and $\mu(p^h, e^h, S^e) = 1$ is the optimal action

$$x = \begin{cases} b & \text{if } p^h \leq \theta_h \\ n & \text{if } p^h > \theta_h \end{cases}$$

while the best response to (p^l, e^l) and $\mu(p^l, e^l, S^e) = 0$ is

$$x = \begin{cases} b & \text{if } p^l \leq \theta_l \\ n & \text{if } p^l > \theta_l \end{cases}$$

This holds independently of the price posted or the evidence provided by the two types, since the difference in the two price-evidence combinations allows for full disclosure of the quality type to the buyer. Therefore, the equilibrium prices must satisfy $p^h \in (\theta_l, \theta_h]$ and $p^l = \theta_l$. Let us consider the four cases above separately.

1. The first case corresponds to the equilibria characterized by Proposition 2.
2. Suppose the equilibrium price-evidence combinations of the high-quality and low-quality seller are (p^h, e_l), (p^l, e_h), respectively. The equilibrium payoffs are $\pi(\sigma_h|\beta, \theta_h) = p^*$ and $\pi(\sigma_l|\beta, \theta_l) = \theta_l$. However, the incentive compatibility constraint for the low type cannot be satisfied since $\pi(\sigma_h|\beta, \theta_l) = p^* > \theta_l = \pi(\sigma_l|\beta, \theta_l)$. If the low type deviates to (p^h, e_l) the evidence technology cannot detect him and therefore he gets p^* with probability one. Therefore, there cannot be an equilibrium where both high- and low-quality sellers provide untruthful evidence.

3. Suppose the equilibrium price-evidence combinations of the high-quality and low-quality seller are (p^h, e_l), (p^l, e_l), respectively, with $p^h \neq p^l$. The equilibrium payoffs are $\pi(\sigma_h|\beta, \theta_h) = p^*$ and $\pi(\sigma_l|\beta, \theta_l) = \theta_l$. Similarly to the previous case, the incentive compatibility constrained for the low type cannot be satisfied since $\pi(\sigma_h|\beta, \theta_l) = p^* > \theta_l = \pi(\sigma_l|\beta, \theta_l)$. Even in this case, if the low type deviates to (p^h, e_l) the evidence technology cannot detect him and he thus gets p^* with probability one. Therefore, there cannot be an equilibrium where both high- and low-quality sellers provide low-quality evidence. Together with the conclusion in point 2, we conclude that there cannot be a separating equilibrium in which the high type provides e_l.

4. Suppose the equilibrium price-evidence combinations of the high-quality and low-quality seller are (p^h, e_h), (p^l, e_h), respectively, with $p^h \neq p^l$. The equilibrium payoffs are $\pi(\sigma_h|\beta, \theta_h) = p^*$ and $\pi(\sigma_l|\beta, \theta_l) = \theta_l$. If the low type deviates to (p^h, e_h), he gets p^* if $S^e = S^e_h$, which happens with probability α, and 0 otherwise. Therefore, the incentive compatibility constraint of the low type is

$$\pi(\sigma_h|\beta, \theta_l) = \alpha p^* \leq \theta_l = \pi(\sigma_l|\beta, \theta_l)$$

   which holds as long as $p^* \leq \theta_l/\alpha$. This is the same condition for the separating equilibrium described in Proposition 2. Moreover, the rest of the equilibrium analysis is identical to that presented in the proof of Proposition2, thus proving the claim that separating equilibria in which both types of seller provide e_h can exist.

Since we exhausted the set of possible separating equilibria in pure strategies, we proved the proposition.[32] □

Proof of Proposition 4

Proof. Consider a pooling PBE in which both high- and low-quality sellers choose to provide high-quality evidence in equilibrium so that $(p^h, e^h) = (p^l, e^l) = (p^*, e_h)$ where p^* is the common equilibrium price. Given the assumptions on the evidence technology,

---

[32] Note: In the previous version of the model, we assumed that untruthful reporting by the low-quality seller would lead to no trade whenever detected by the evidence technology, even when the equilibrium price-evidence vector differentiated the two types of sellers. With this stricter assumption, the only surviving equilibria would be those characterized by Proposition 2.

(p^*, e_h) induces signal S_h^e with probability one when $\theta = \theta_h$; instead, when $\theta = \theta_l$, (p^*, e_h) induces signal S_h^e with probability α and S_l^e with probability $1-\alpha$. Therefore, the posterior belief of the buyer upon observing (p^*, e_h, S_h^e) is

$$\mu(p^*, e_h, S_h^e) = \frac{\mu_0}{\mu_0 + (1-\mu_0)\alpha} = \tilde{\mu},$$

while after observing (p^*, e_h, S_l^e), it is $\mu(p^*, e_h, S_l^e) = 0$. In this case, the evidence technology is able to detect a low-type providing untruthful evidence with positive probability. From the proof of Proposition 2, we have that for any price p and posterior $\hat{\mu}$, in a pure-strategy equilibrium with $\varepsilon = 0$, the optimal β satisfies (1.27) and (1.28). Clearly, the best response of the buyer to $\mu(p^*, e_h, S_l^e) = 0$ requires

$$x = \begin{cases} b & \text{if } p^* \leq \theta_l \\ n & \text{if } p^* > \theta_l \end{cases}$$

We will show that θ_l is not a possible value for the equilibrium price so the optimal action of the buyer after observing S_l^e is always not to buy the good. The best response to $\mu(p^*, e_h, S_h^e)$, instead, is

$$x = \begin{cases} b & \text{if } p^* \leq \bar{\theta}_{\tilde{\mu}}, \\ n & \text{if } p^* > \bar{\theta}_{\tilde{\mu}}; \end{cases} \qquad \text{if } \tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$$

and

$$x = \begin{cases} b & \text{if } p^* \leq \theta_l + \frac{c}{1-\tilde{\mu}} \\ ba & \text{if } \theta_l + \frac{c}{1-\tilde{\mu}} < p^* \leq \theta_h - \frac{c}{\tilde{\mu}} \\ n & \text{if } p^* > \theta_h - \frac{c}{\tilde{\mu}} \end{cases} \qquad \text{if } \tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$$

If $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, inducing b over n is optimal for the seller so any price $p^* \in \left[\theta_l, \bar{\theta}_{\tilde{\mu}}\right]$ can potentially be an equilibrium price. If, instead, $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, prices inducing the acquisition of the authentication cannot be sustained in equilibrium: since the authentication technology is perfect, the low type would always be detected thus receiving zero profit. Hence, in this case, potential common equilibrium prices must satisfy $p^* \in [\theta_l, p_b(\tilde{\mu})] = \left[\theta_l, \theta_l + \frac{c}{1-\tilde{\mu}}\right]$. In summary, p^* needs to be lower than $\bar{p}$ with

$$\bar{p} = \begin{cases} \bar{\theta}_{\tilde{\mu}} & \text{if } \tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right] \\ \theta_l + \frac{c}{1-\tilde{\mu}} & \text{if } \tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right] \end{cases}$$

The equilibrium payoffs for the high- and the low-quality seller are, respectively, $\pi(\sigma_h|\beta, \theta_h) = \pi_h^* = p^*$ and $\pi(\sigma_l|\beta, \theta_l) = \pi_l^* = \alpha p^*$. Neither type should have an incentive to deviate. In particular, the profit of the low type should be at least θ_l, which is the seller's lower bound for profits, always achievable by simply setting a price equal to θ_l. Therefore, it must hold $\pi(\sigma_l|\beta, \theta_l) \geq \theta_l$ or $p^* \geq \theta_l/\alpha$. The necessary condition for this type of equilibrium is $\alpha \geq \frac{\theta_l}{\bar{p}}$, ensuring that the set of equilibrium prices is nonempty. If $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, this requires

$$\alpha \geq \frac{\theta_l}{\bar{\theta}_{\tilde{\mu}}}, \tag{1.29}$$

which, becomes

$$\frac{\theta_l}{\alpha} \leq \frac{\mu_0\theta_h + (1-\mu_0)\alpha\theta_l}{\mu_0 + (1-\mu_0)\alpha}$$

or

$$(1-\mu_0)\theta_l\alpha^2 + [\mu_0\theta_h - (1-\mu_0)\theta_l]\alpha - \mu_0\theta_l \geq 0$$

The second-degree inequality in α has solutions $\alpha \leq \alpha_1$ and $\alpha \geq \alpha_2$ with

$$\alpha_1 = \frac{(1-\mu_0)\theta_l - \mu_0\theta_h - \sqrt{[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2}}{2(1-\mu_0)\theta_l}$$
$$\alpha_2 = \frac{(1-\mu_0)\theta_l - \mu_0\theta_h + \sqrt{[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2}}{2(1-\mu_0)\theta_l}$$

Since $\sqrt{[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2} > |\mu_0\theta_h - (1-\mu_0)\theta_l|$, the sign of the numerator corresponds to the sign of the square root term. Thus, $\alpha_1 < 0$ and $\alpha \leq \alpha_1$ has no valid solutions. However, α_2 is lower than 1 since:

$$\begin{aligned}
\alpha_2 &< 1 \\
\sqrt{[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2} &< \mu_0\theta_h + (1-\mu_0)\theta_l \\
[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2 &< [\mu_0\theta_h + (1-\mu_0)\theta_l]^2 \\
-2\mu_0(1-\mu_0)\theta_h\theta_l + 4\mu_0(1-\mu_0)\theta_l^2 &< 2\mu_0(1-\mu_0)\theta_h\theta_l \\
4\mu_0(1-\mu_0)\theta_l^2 &< 4\mu_0(1-\mu_0)\theta_h\theta_l \\
\theta_l &< \theta_h.
\end{aligned}$$

Therefore, when $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, pooling equilibria on e_h exists only if $\alpha \geq \alpha_2$.

If, instead, $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, the necessary condition becomes

$$\alpha \geq \frac{\theta_l}{p_b(\tilde{\mu})} \tag{1.30}$$

Under $\varepsilon = 0$, this corresponds to

$$\begin{aligned}
\frac{\theta_l}{\alpha} &\leq \theta_l + \frac{c}{1-\tilde{\mu}} \\
\frac{1-\alpha}{\alpha}\theta_l &\leq \frac{c}{1-\tilde{\mu}} \\
\frac{1-\alpha}{\alpha}\theta_l &\leq \frac{\mu_0 + (1-\mu_0)\alpha}{(1-\mu_0)\alpha}c \\
(1-\alpha)(1-\mu_0)\theta_l &\leq [\mu_0 + (1-\mu_0)\alpha]c \\
(1-\mu_0)(\theta_l + c)\alpha &\geq (1-\mu_0)\theta_l - \mu_0 c \\
\alpha &\geq \frac{(1-\mu_0)\theta_l - \mu_0 c}{(1-\mu_0)(\theta_l + c)} =: \alpha_3
\end{aligned}$$

which is lower than 1 if $c > 0$, and equal to 1 otherwise. Therefore, when $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, pooling equilibria on e_h exist only if $\alpha \geq \alpha_3$. These results imply $\alpha_2 = \theta_l/\bar{\theta}_{\tilde{\mu}(\alpha_2)}$ and $\alpha_3 = \theta_l/p_b(\tilde{\mu}(\alpha_3))$. When $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, we have $\bar{\theta}_{\tilde{\mu}} < p_b(\tilde{\mu})$ implying $\alpha_2 > \alpha_3$. Since α_2 is the relevant threshold, in this case, we can write the condition as $\alpha \geq \alpha_2 > \alpha_3$. Symmetrically, when $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, we have $p_b(\tilde{\mu}) < \bar{\theta}_{\tilde{\mu}}$ implying $\alpha_3 > \alpha_2$. Since α_3 is the relevant threshold, in this case, we can write the condition as $\alpha \geq \alpha_3 > \alpha_2$. Putting everything together we get that the overall necessary condition for the existence of pooling equilibria on e_h when $\varepsilon = 0$ is $\alpha \geq \max\{\alpha_2, \alpha_3\}$.

Let us now consider the off-equilibrium belief restrictions. The equilibrium's payoffs for the types of sellers are $\pi_h^* = p^*$ and $\pi_l^* = \alpha p^*$. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high type side.

Instead, since $\varepsilon = 0$, the conditions for the low-type seller are: if $p > \alpha p^*$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p > p^*$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$. There are no restrictions otherwise.

Since $\underline{\mu}_{ba}(p) < \underline{\mu}_b(p)$, $\bar{\mu}_p^{e_h}$ is determined by the high-type restrictions for $p > p^*$ and is unrestricted, otherwise. Conversely, $\bar{\mu}_p^{e_l}$ is determined by the restrictions of the low type for any $p \in (\alpha p^*, p^*]$, by restrictions of the high type for any $p > p^*$, and is unrestricted otherwise. Overall, this implies:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

and

$$\bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, \alpha p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (\alpha p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in (\alpha p^*, p^*] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in [p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

□

Proof of Proposition 5

Proof. Consider a pooling PBE in which both high- and low-quality sellers choose to provide low-quality evidence in equilibrium so that $(p^h, e^h) = (p^l, e^l) = (p^*, e_l)$ where p^* is the common equilibrium price. Given the assumptions on the evidence technology, (p^*, e_l) induces an uninformative evidence signal since $\mathbb{P}(S_h^e | e_l, \theta_h) = \mathbb{P}(S_h^e | e_l, \theta_l)$. Therefore, the posterior belief of the buyer upon observing (p^*, e_l, S_h^e) or (p^*, e_l, S_l^e) is simply the prior μ_0, since the evidence technology is not able to differentiate high- and low-type of sellers. From the proof of Proposition 2, we have that for any price p and posterior $\hat{\mu}$, in a pure-strategy equilibrium with $\varepsilon = 0$, the optimal β satisfies (1.27) and (1.28). Clearly, the best response of the buyer to posteriors $\mu(p^*, e_l, S_h^e) = \mu(p^*, e_l, S_h^e) = \mu_0$ is

$$x = \begin{cases} b & \text{if } p^* \leq \bar{\theta}_{\mu_0}, \\ n & \text{if } p^* > \bar{\theta}_{\mu_0}; \end{cases} \quad \text{if } \mu_0 \notin \left[\underline{\mu}, \bar{\mu}\right]$$

and

$$x = \begin{cases} b & \text{if } p^* \leq \theta_l + \frac{c}{1-\mu_0} \\ ba & \text{if } \theta_l + \frac{c}{1-\mu_0} < p^* \leq \theta_h - \frac{c}{\mu_0} \\ n & \text{if } p^* > \theta_h - \frac{c}{\mu_0} \end{cases} \quad \text{if } \mu_0 \in \left[\underline{\mu}, \bar{\mu}\right]$$

If $\mu_0 \notin \left[\underline{\mu}, \bar{\mu}\right]$, inducing b over n is optimal for the seller so any price $p^* \in \left[\theta_l, \bar{\theta}_{\mu_0}\right]$ can potentially be an equilibrium price. If, instead, $\mu_0 \in \left[\underline{\mu}, \bar{\mu}\right]$, prices inducing the acquisition of the authentication cannot be sustained in equilibrium: since the authentication technology is perfect, the low type would always be detected thus receiving zero profit. Hence, in this case, potential common equilibrium prices must satisfy $p^* \in [\theta_l, p_b(\mu_0)] = \left[\theta_l, \theta_l + \frac{c}{1-\mu_0}\right]$. In summary, p^* needs to be lower than $\bar{p}$ with

$$\bar{p} = \begin{cases} \bar{\theta}_{\mu_0} & \text{if } \mu_0 \notin \left[\underline{\mu}, \bar{\mu}\right] \\ \theta_l + \frac{c}{1-\mu_0} & \text{if } \mu_0 \in \left[\underline{\mu}, \bar{\mu}\right] \end{cases}$$

The equilibrium payoff is the same for the high- and the low-quality seller and equals $\pi(\sigma_h|\beta, \theta_h) = \pi(\sigma_l|\beta, \theta_l) = p^*$ since any equilibrium price induce the buyer to buy without authentication and the evidence technology is not able to detect the low type. Differently from the previous equilibrium types, there are no in-equilibrium profitable deviations for either type, since the common equilibrium profit is (weakly) larger than θ_l. This implies that there are no necessary conditions for the existence of pooling equilibria on e_l, so that this kind of equilibrium exists for any combination of the parameters of the model.

We are left to determine the restrictions on the out-of-equilibrium beliefs necessary to sustain the equilibrium. The equilibrium's payoffs are $\pi_h^* = \pi_l^* = p^*$. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high type side.

Instead, since $\varepsilon = 0$, the conditions for the low-type seller are: if $p > \alpha p^*$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p > p^*\alpha$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$. There are no restrictions otherwise.

Since $\underline{\mu}_{ba}(p) < \underline{\mu}_b(p)$ and $p^* < p^*/\alpha$, both $\bar{\mu}_p^{e_h}$ and $\bar{\mu}_p^{e_l}$ are fully determined by the high-type restrictions for $p > p^*$ and are unrestricted, i.e. equals 1, otherwise. Therefore,

$$\bar{\mu}_p^{e_h} = \bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

□

Proof of Proposition 6

Proof. (i) The proof of the optimality of the proposed strategies follows the same argument as in Proposition 2, once we account for the general definitions of $p_{ba}(\mu)$, $p_b(\mu)$ with $\varepsilon > 0$, and the fact the he evidence provided by the low type in equilibrium, e^l could be either e_l or e_h. Equilibrium payoffs are $\pi_h^* = p^*$ and $\pi_l^* = \theta_l$, with equilibrium price $p^* \in (\theta_l, \min\{\theta_l/\alpha, \theta_h\}]$. No conditions are required to ensure existence.

From Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are as follows: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high-type side. Instead, since $\varepsilon > 0$, the conditions for the low-type seller are as follows. First, consider deviations to out-of-equilibrium (p, e_l): for all $p > \theta_l$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\theta_l, \theta_l/\varepsilon]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > \theta_l/\varepsilon$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$. Second, consider deviations to out-of-equilibrium (p, e_h): for all $p > \theta_l/\alpha$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\theta_l/\alpha, \theta_l/(\alpha\varepsilon)]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > \theta_l/(\alpha\varepsilon)$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \theta_l/\alpha$ (i.e., $\mu_p^{e_h} \leq 1$). Overall, it implies:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

$$\bar{\mu}_p^{e_l} = \begin{cases} \underline{\mu}_b(p) & \text{if } p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \text{ or } p \in [\theta_l, \min\{p^*, \theta_l/\varepsilon\}] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (\min\{p^*, \theta_l/\varepsilon\}, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

Finally, note that the same reasoning that led to Proposition 3 applies when $\varepsilon > 0$, so that in every separating equilibrium, the high type provides high evidence e_h, while the low type is indifferent between e_l and e_h, since the different prices posted by the two types lead to a degenerate posterior regardless of the evidence provided by the low type.

(ii) The proof of the optimality of the proposed strategies follows the same argument as in Proposition 4, once we account for the general definitions of $p_{ba}(\mu)$ and $p_b(\mu)$ with $\varepsilon > 0$. The only difference is that when $\varepsilon = 0$, authentication acquisition cannot be sustained in equilibrium, whereas in a pooling(e_h) equilibrium with $\varepsilon > 0$, it is one of the two possible equilibrium outcomes. In this case we are considering $\beta(b|p^*, \tilde{\mu}) = 1$. Equilibrium payoffs are $\pi_h^* = p^*$ and $\pi_l^* = \alpha p^*$, with equilibrium price $p^* \in [\theta_l/\alpha, \bar{p}]$ with $\bar{p} = \bar{\theta}_{\tilde{\mu}}$ if $\tilde{\mu} \notin [\underline{\mu}, \bar{\mu}]$, and $\bar{p} = p_b(\tilde{\mu})$ if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$. The necessary and sufficient condition for equilibrium, ensuring the equilibrium price set is non-empty, is $\frac{\theta_l}{\alpha} \leq \bar{p}$. If $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$, the condition becomes $\alpha \geq \frac{\theta_l}{\bar{\theta}_{\tilde{\mu}}}$, which corresponds to condition (1.29) in the proof of Proposition 4. From there, we have that the condition is satisfied if and only if $\alpha \geq \alpha_2$, with

$$\alpha_2 = \frac{(1-\mu_0)\theta_l - \mu_0\theta_h + \sqrt{[\mu_0\theta_h - (1-\mu_0)\theta_l]^2 + 4\mu_0(1-\mu_0)\theta_l^2}}{2(1-\mu_0)\theta_l},$$

which we proved to be strictly lower than one. If $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, the necessary condition corresponds to

$$\alpha \geq \frac{\theta_l}{p_b(\tilde{\mu})} \tag{1.31}$$

or

$$\begin{aligned}
\frac{\theta_l}{\alpha} &\leq \theta_l + \frac{c}{(1-\tilde{\mu})(1-\varepsilon)} \\
\frac{1-\alpha}{\alpha}\theta_l &\leq \frac{c}{(1-\tilde{\mu})(1-\varepsilon)} \\
\frac{1-\alpha}{\alpha}\theta_l &\leq \frac{\mu_0 + (1-\mu_0)\alpha}{(1-\mu_0)(1-\varepsilon)\alpha}c \\
(1-\alpha)(1-\mu_0)(1-\varepsilon)\theta_l &\leq [\mu_0 + (1-\mu_0)\alpha]c \\
(1-\mu_0)(1-\varepsilon)\theta_l - \mu_0 c &\leq [(1-\mu_0)c + (1-\mu_0)(1-\varepsilon)\theta_l]\alpha \\
\alpha &\geq \frac{(1-\mu_0)(1-\varepsilon)\theta_l - \mu_0 c}{(1-\mu_0)\left[(1-\varepsilon)\theta_l + c\right]} := \alpha_4
\end{aligned}$$

which is lower than one whenever $c > 0$, and one otherwise. Therefore, when $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$, pooling($e_h$)-b equilibria exist only if $\alpha \geq \alpha_4$. Unsurprisingly, α_3, from the proof of Proposition 4, and α_4 are equivalent when $\varepsilon = 0$. Moreover, no conditions on ε are required to guarantee the existence of pooling(e_h)-b equilibria. Similarly to the proof of Proposition 4, it can be shown that the overall necessary condition for the existence of pooling equilibria on e_h when $\varepsilon > 0$ is $\alpha \geq \max\{\alpha_2, \alpha_4\}$, which encompasses both $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$ and $\tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$.

From Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are as follows: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high-type side. Instead, since $\varepsilon > 0$, the conditions for the low-type seller are as follows. First, consider deviations to out-of-equilibrium (p, e_l): for all $p > \alpha p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\alpha p^*, \alpha p^*/\varepsilon]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > \alpha p^*/\varepsilon$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \alpha p^*$ (i.e., $\mu_p^{e_l} \leq 1$). Second, consider deviations to out-of-equilibrium (p, e_h): for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (p^*, p^*/\varepsilon]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*/\varepsilon$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_h} \leq 1$). Overall, it implies:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

and

$$\bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, \alpha p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (\alpha p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in [\alpha p^*, \min\{p^*, \alpha p^*/\varepsilon\}] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (\min\{p^*, \alpha p^*/\varepsilon\}, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

(iii) Consider a pooling PBE with $\varepsilon > 0$, in which both high- and low-quality sellers provides e_h in equilibrium so that $(p^h, e^h) = (p^l, e^l) = (p^*, e_h)$ where p^* is the common equilibrium price. The posterior belief of the buyer upon observing (p^*, e_h, S_h^e) is $\mu(p^*, e_h, S_h^e) = \frac{\mu_0}{\mu_0 + (1-\mu_0)\alpha} = \tilde{\mu}$, while after observing (p^*, e_h, S_l^e), it is $\mu(p^*, e_h, S_l^e) = 0$. From Lemma 2 we have that the best response of the buyer to $\mu(p^*, e_h, S_l^e) = 0$ requires

$$x = \begin{cases} b & \text{if } p^* \leq \theta_l \\ n & \text{if } p^* > \theta_l \end{cases}$$

We will see that equilibrium must be strictly above θ_l, so the optimal action following S_l^e is n. The best response to $\mu(p^*, e_h, S_h^e) = \tilde{\mu}$, instead, is

$$x = \begin{cases} b & \text{if } p^* \leq \bar{\theta}_{\tilde{\mu}}, \\ n & \text{if } p^* \geq \bar{\theta}_{\tilde{\mu}}; \end{cases} \qquad \text{if } \tilde{\mu} \notin \left[\underline{\mu}, \bar{\mu}\right]$$

and

$$x = \begin{cases} b & \text{if } p \leq p_b(\tilde{\mu}) \\ ba & \text{if } p_b(\tilde{\mu}) \leq p \leq p_{ba}(\tilde{\mu}) \\ n & \text{if } p \geq p_{ba}(\tilde{\mu}) \end{cases} \quad \text{if } \tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right] \tag{1.32}$$

Let us consider the case in which $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$ and ba is the optimal strategy of the buyer. In fact, when $\varepsilon > 0$, prices inducing authentication acquisition can possibly be sustained in equilibrium; an imperfect authentication technology allows the low type not to be detected with positive probability when authentication is acquired. Potential common equilibrium prices belong to the interval $[p_b(\tilde{\mu}), p_{ba}(\tilde{\mu})] = \left[\theta_l + \frac{c}{(1-\tilde{\mu})(1-\varepsilon)}, \frac{\tilde{\mu}\theta_h + (1-\tilde{\mu})\varepsilon\theta_l - c}{\tilde{\mu} + (1-\tilde{\mu})\varepsilon}\right]$. The equilibrium payoff for the high-quality seller is $\pi(\sigma_h|\beta, \theta_h) = \pi_h^* = p^*$, while the equilibrium payoff of the low-quality seller is $\pi(\sigma_l|\beta, \theta_l) = \pi_l^* = \alpha\varepsilon p^*$. Indeed, a low type providing high evidence induces S_h^e with probability α, and authentication acquisition generates signal S_h^a with probability ε. The no-deviation condition for the low-quality seller requires a payoff of at least θ_l. Therefore, it must hold $\alpha\varepsilon p^* \geq \theta_l$. A necessary condition for the existence of pooling(e_h)-ba equilibria is then

$$\alpha \geq \theta_l / \varepsilon p_{ba}(\tilde{\mu}), \tag{1.33}$$

which corresponds to

$$\begin{aligned}
\frac{\theta_l}{\alpha\varepsilon} &\leq \frac{\tilde{\mu}\theta_h + (1-\tilde{\mu})\varepsilon\theta_l - c}{\tilde{\mu} + (1-\tilde{\mu})\varepsilon} \\
\frac{\theta_l}{\alpha\varepsilon} &\leq \frac{\dfrac{\mu_0\theta_h + (1-\mu_0)\alpha\varepsilon\theta_l}{\mu_0 + (1-\mu_0)\alpha} - c}{\dfrac{\mu_0 + (1-\mu_0)\alpha\varepsilon}{\mu_0 + (1-\mu_0)\alpha}} \\
\frac{\theta_l}{\alpha\varepsilon} &\leq \frac{\mu_0\theta_h + (1-\mu_0)\alpha\varepsilon\theta_l - [\mu_0 + (1-\mu_0)\alpha]c}{\mu_0 + (1-\mu_0)\alpha\varepsilon} \\
[\mu_0 + (1-\mu_0)\alpha\varepsilon]\theta_l &\leq \alpha\varepsilon\mu_0\theta_h + (1-\mu_0)\alpha^2\varepsilon^2\theta_l - \alpha\varepsilon[\mu_0 + (1-\mu_0)\alpha]c.
\end{aligned}$$

This, in turn, corresponds to the following second-degree inequality in α:

$$(1-\mu_0)\varepsilon(\varepsilon\theta_l - c)\alpha^2 + \varepsilon[\mu_0(\theta_h - c) - (1-\mu_0)\theta_l]\alpha - \mu_0\theta_l \geq 0$$

The corresponding equality has roots:

$$\alpha_5 = \frac{-\varepsilon[\mu_0(\theta_h - c) - (1-\mu_0)\theta_l] - \sqrt{[\varepsilon(\mu_0(\theta_h - c) - (1-\mu_0)\theta_l)]^2 + 4\varepsilon\mu_0(1-\mu_0)(\varepsilon\theta_l - c)\theta_l}}{2\varepsilon(1-\mu_0)(\varepsilon\theta_l - c)}$$
$$\alpha_6 = \frac{-\varepsilon[\mu_0(\theta_h - c) - (1-\mu_0)\theta_l] + \sqrt{[\varepsilon(\mu_0(\theta_h - c) - (1-\mu_0)\theta_l)]^2 + 4\varepsilon\mu_0(1-\mu_0)(\varepsilon\theta_l - c)\theta_l}}{2\varepsilon(1-\mu_0)(\varepsilon\theta_l - c)}$$

Since the sign of the coefficient of the squared term is undetermined, it is not clear a priori which solutions of the quadratic inequality are valid. We will now show that the only relevant threshold is α_6. To simplify notation, define $A = (1-\mu_0)\varepsilon(\varepsilon\theta_l - c)$, $B = \varepsilon[\mu_0(\theta_h - c) - (1-\mu_0)\theta_l]$ and $C = -\mu_0\theta_l$. There are three cases:

(a) Suppose $\varepsilon\theta_l - c > 0$, so that $A > 0$. Then, $\alpha_5 \leq \alpha_6$ and the solutions are $\alpha \leq \alpha_5$ and $\alpha \geq \alpha_6$. $A > 0$ implies $-4AC > 0$, so that $\sqrt{B^2 - 4AC} > |B|$, and ensures the discriminant of the quadratic equation is strictly positive. Moreover, $A > 0$ and $\sqrt{B^2 - 4AC} > |B|$ imply that the sign of the thresholds is determined by the sign of the square root. So $\alpha_5 < 0$, so that no valid solutions come from $\alpha \leq \alpha_5$, while $\alpha_6 > 0$, which requires $\alpha_6 \leq 1$ to have valid solutions.

(b) Suppose $\varepsilon\theta_l - c = 0$, so that the inequality simplifies to

$$\alpha \geq \frac{\mu_0\theta_l}{\varepsilon[\mu_0\theta_h - \mu_0 c - (1-\mu_0)\theta_l]}$$

The initial assumption implies $\varepsilon = \frac{c}{\theta_l}$. Substituting into the inequality yields

$$\alpha \geq \frac{\mu_0\theta_l}{\varepsilon[\mu_0\theta_h - \mu_0 c - (1-\mu_0)\theta_l]} = \frac{\mu_0\theta_l^2}{c[\mu_0\theta_h - \mu_0 c - (1-\mu_0)\theta_l]} = \lim_{\varepsilon \to c/\theta_l} \alpha_6.$$

In fact, by defining $K = \mu_0(\theta_h - c) - (1-\mu_0)\theta_l$ we can rewrite α_6 as

$$\alpha_6 = \frac{-\varepsilon K + \sqrt{\varepsilon^2 K^2 + 4\varepsilon\mu_0(1-\mu_0)(\varepsilon\theta_l - c)\theta_l}}{2\varepsilon(1-\mu_0)(\varepsilon\theta_l - c)}$$

As $\varepsilon \to c/\theta_l$, both the denominator and the numerator tend to 0 since $\varepsilon\theta_l - c = 0$ and $-c/\theta_l K + \sqrt{(c/\theta_l)^2 K^2} = 0$, since $K > 0$. Define the numerator and the denominator of α_6 as

$$\begin{aligned} N &= -\varepsilon K + \sqrt{\varepsilon^2 K^2 + 4\varepsilon\mu_0(1-\mu_0)(\varepsilon\theta_l - c)\theta_l}, \\ D &= 2\varepsilon(1-\mu_0)(\varepsilon\theta_l - c). \end{aligned}$$

Then,

$$\frac{\partial N}{\partial\varepsilon} = -K + \frac{2\varepsilon K^2 + 8\mu_0(1-\mu_0)\theta_l^2\varepsilon - 4c\mu_0(1-\mu_0)\theta_l}{2\sqrt{\varepsilon^2 K^2 + 4\varepsilon\mu_0(1-\mu_0)(\varepsilon\theta_l - c)\theta_l}},$$

and

$$\frac{\partial D}{\partial\varepsilon} = 4\varepsilon\theta_l(1-\mu_0) - 2c(1-\mu_0).$$

At $\varepsilon = \frac{c}{\theta_l}$, we have $\varepsilon\theta_l - c = 0$, so that substituting in $\frac{\partial N}{\partial\varepsilon}$ and $\frac{\partial D}{\partial\varepsilon}$ we get

$$\frac{\partial N}{\partial\varepsilon} = -K + \frac{2(c/\theta_l)K^2 + 4c\mu_0(1-\mu_0)\theta_l}{2(c/\theta_l)K} = -K + \frac{K^2 + 2\mu_0(1-\mu_0)\theta_l^2}{K} = \frac{2\mu_0(1-\mu_0)\theta_l^2}{K},$$

and

$$\frac{\partial D}{\partial\varepsilon} = 4c(1-\mu_0) - 2c(1-\mu_0) = 2c(1-\mu_0).$$

Therefore, by applying L'Hospital's rule, we have:

$$\lim_{\varepsilon\to c/\theta_l} \alpha_6 = \frac{\frac{\partial N}{\partial\varepsilon}}{\frac{\partial D}{\partial\varepsilon}} = \frac{\frac{2\mu_0(1-\mu_0)\theta_l^2}{K}}{2c(1-\mu_0)} = \frac{\mu_0\theta_l^2}{cK} = \frac{\mu_0\theta_l^2}{c\left[\mu_0(\theta_h - c) - (1-\mu_0)\theta_l\right]}.$$

Thus, case (b) is equivalent to case (a), so that the only valid solutions are $\alpha \geq \alpha_6$, requiring $\alpha_6 \leq 1$.

(c) Suppose $\varepsilon\theta_l - c < 0$, so that $A < 0$. Then $\alpha_5 = \frac{B+\sqrt{B^2-4AC}}{-2A}$ and $\alpha_6 = \frac{B-\sqrt{B^2-4AC}}{-2A}$ so that $\alpha_5 > \alpha_6$ whenever the discriminant $B^2 - 4AC \geq 0$. Therefore, the condition is satisfied if and only if $\alpha_6 \leq \alpha \leq \alpha_5$ and $B^2 - 4AC \geq 0$. To have valid solutions, it must hold that $\alpha_6 \leq 1$ and $\alpha_5 \geq 0$. Now, since $A < 0$, we have $-4AC < 0$, so that $\sqrt{B^2 - 4AC} < |B|$: whenever the discriminant $B^2 - 4AC \geq 0$, the sign of the thresholds is determined by the sign of B. Note, however, that if a solution exists, then it must be that $B^2 - 4AC \geq 0$, which implies $B > 0$, and, in turn, $\alpha_5 > 0$.

Therefore, the condition on α_5 is redundant. Thus, valid solutions satisfy $\alpha \geq \alpha_6$, which again requires that $\alpha_6 \leq 1$.

We proved there is only one relevant threshold for α: condition (1.35) is satisfied if and only if $\alpha \geq \alpha_6$. Note, however, that α_6 is not necessarily lower than 1 for all other parameter combinations, which is a necessary condition for the set of possible solutions to be nonempty.

We are left to determine the restrictions on the out-of-equilibrium beliefs necessary to sustain the equilibrium. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high type side.

Instead, the conditions for the low-type seller are as follows. First, consider deviations to out-of-equilibrium (p, e_l): for all $p > \alpha\varepsilon p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\alpha\varepsilon p^*, \alpha p^*]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > \alpha p^*$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \varepsilon\alpha\varepsilon p^*/$ (i.e., $\mu_p^{e_l} \leq 1$). Second, consider deviations to out-of-equilibrium (p, e_h): for all $p > \varepsilon p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\varepsilon p^*, p^*]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \varepsilon p^*$ (i.e., $\mu_p^{e_h} \leq 1$). Overall, it implies:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, \varepsilon p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (\varepsilon p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in (\varepsilon p^*, p^*] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

and

$$\bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, \alpha\varepsilon p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (\alpha\varepsilon p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in (\alpha\varepsilon p^*, \alpha p^*] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (\alpha p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

(iv) The proof of the optimality of the proposed strategies follows the same argument as in Proposition 5, once we account for the general definitions of $p_{ba}(\mu)$, $p_b(\mu)$ with $\varepsilon > 0$. The only difference is that when $\varepsilon = 0$, authentication acquisition cannot be sustained in equilibrium, whereas in a pooling(e_l) equilibrium with $\varepsilon > 0$, it is one of the two possible equilibrium outcomes. In this case we are considering $\beta(b|p^*, \mu_0) = 1$. Equilibrium payoffs are $\pi_h^* = \pi_l^* = p^*$, with equilibrium price $p^* \in [\theta_l, \bar{\theta}_{\mu_0}]$ if $\mu_0 \notin [\underline{\mu}, \bar{\mu}]$, and $p^* \in [\theta_l, p_b(\mu_0)]$ if $\mu_0 \in [\underline{\mu}, \bar{\mu}]$. No conditions are required to ensure existence.

We are left to determine the restrictions on the out-of-equilibrium beliefs necessary to sustain the equilibrium. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high type side.

Instead, since $\varepsilon > 0$, the conditions for the low-type seller are as follows. First, consider deviations to out-of-equilibrium (p, e_l): for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (p^*, p^*/\varepsilon]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*/\varepsilon$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*/$ (i.e., $\mu_p^{e_l} \leq 1$). Second, consider deviations to out-of-equilibrium (p, e_h): for all $p > p^*/\alpha$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (p^*/\alpha, p^*/(\alpha\varepsilon)]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*/(\alpha\varepsilon)$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*/\alpha$ (i.e., $\mu_p^{e_h} \leq 1$). Overall, it implies:

$$\bar{\mu}_p^{e_h} = \bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

(v) Consider a pooling PBE with $\varepsilon > 0$, in which both high- and low-quality sellers provides e_l in equilibrium so that $(p^h, e^h) = (p^l, e^l) = (p^*, e_l)$ where p^* is the common equilibrium price. Given the assumptions on the evidence technology, (p^*, e_l) induces an uninformative evidence signal since $\mathbb{P}(S_h^e|e_l, \theta_h) = \mathbb{P}(S_h^e|e_l, \theta_l)$. Therefore, the posterior belief of the buyer upon observing (p^*, e_l, S_h^e) or (p^*, e_l, S_l^e) is simply the prior μ_0. From Lemma 2 we have that the best response of the buyer to $\mu(p^*, e_l, S_h^e) = \mu(p^*, e_l, S_l^e) = \mu_0$, is

$$x = \begin{cases} b & \text{if } p^* \leq \bar{\theta}_{\mu_0}, \\ n & \text{if } p^* \geq \bar{\theta}_{\mu_0}; \end{cases} \quad \text{if } \mu_0 \notin \left[\underline{\mu}, \bar{\mu}\right]$$

and

$$x = \begin{cases} b & \text{if } p \leq p_b(\mu_0) \\ ba & \text{if } p_b(\mu_0) \leq p \leq p_{ba}(\mu_0) \\ n & \text{if } p \geq p_{ba}(\mu_0) \end{cases} \quad \text{if } \mu_0 \in \left[\underline{\mu}, \bar{\mu}\right] \tag{1.34}$$

Let us consider the case in which $\mu_0 \in \left[\underline{\mu}, \bar{\mu}\right]$ and ba is the optimal strategy of the buyer. In fact, when $\varepsilon > 0$, prices inducing authentication acquisition can possibly be sustained in equilibrium; an imperfect authentication technology allows the low type not to be detected with positive probability when authentication is acquired. Potential common equilibrium prices belong to the interval $[p_b(\mu_0), p_{ba}(\mu_0)] = \left[\theta_l + \frac{c}{(1-\mu_0)(1-\varepsilon)}, \frac{\mu_0\theta_h + (1-\mu_0)\varepsilon\theta_l - c}{\mu_0 + (1-\mu_0)\varepsilon}\right]$. The equilibrium payoff for the high-quality seller is $\pi(\sigma_h|\beta, \theta_h) = \pi_h^* = p^*$, while the equilibrium payoff of the low-quality seller is $\pi(\sigma_l|\beta, \theta_l) = \pi_l^* = \varepsilon p^*$. Indeed, when sellers pool on e_l, S^e is uninformative, while authentication acquisition generates signal S_h^a with probability ε. The no-deviation condition for the low-quality seller requires a payoff of at least θ_l. Therefore, it must hold $\varepsilon p^* \geq \theta_l$. A necessary condition for the existence of pooling(e_h)-ba equilibria is then

$$\varepsilon \geq \theta_l / p_{ba}(\mu_0), \tag{1.35}$$

which we do not study further since pooling(e_l) equilibria do not survive the first refinement criterion.[33]

We are left to determine the restrictions on the out-of-equilibrium beliefs necessary to sustain the equilibrium. Applying Lemma 11, we have that the necessary out-of-equilibrium restrictions for the high type are: for all $p > p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l}, < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq p^*$ (i.e., $\mu_p^{e_l} \leq 1$). The exact same conditions apply to $\mu_p^{e_h}$ on the high type side.

Instead, since $\varepsilon > 0$, the conditions for the low-type seller are as follows. First, consider deviations to out-of-equilibrium (p, e_l): for all $p > \varepsilon p^*$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\varepsilon p^*, p^*]$, then $\mu_p^{e_l} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*$, then $\mu_p^{e_l} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \varepsilon p^*$ (i.e., $\mu_p^{e_l} \leq 1$). Second, consider deviations to out-of-equilibrium (p, e_h): for all $p > \varepsilon p^*/\alpha$, if $p \notin \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$, then

---

[33] It can be shown that there exists values of ε such that the condition is satisfied only if $c \leq \mu_0 \Delta\theta$.

$\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p \in (\varepsilon p^*/\alpha, p^*/\alpha]$, then $\mu_p^{e_h} < \underline{\mu}_b(p)$; if $p \in \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]$ and $p > p^*/\alpha$, then $\mu_p^{e_h} < \underline{\mu}_{ba}(p)$. There are no restrictions for prices $p \leq \varepsilon p^*/\alpha$ (i.e., $\mu_p^{e_h} \leq 1$). Overall, it implies:

$$\bar{\mu}_p^{e_h} = \begin{cases} 1 & \text{if } p \in [\theta_l, \min\{\varepsilon p^*/\alpha, p^*\}] \\ \underline{\mu}_b(p) & \text{if } p \in (\min\{\varepsilon p^*/\alpha, p^*\}, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in (\min\{\varepsilon p^*/\alpha, p^*\}, p^*/\alpha] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*/\alpha, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

and

$$\bar{\mu}_p^{e_l} = \begin{cases} 1 & \text{if } p \in [\theta_l, \varepsilon p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (\varepsilon p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or } p \in (\varepsilon p^*, p^*] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases}$$

Since we exhausted the set of pure-strategies combinations that can be sustained in equilibrium, we proved the proposition. $\square$

Proof of Corollary 2

Proof. The necessary conditions for pooling(e_h)-ba equilibria are $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$ and $\tilde{\mu} \in \left[\underline{\mu}, \bar{\mu}\right]$. From the results following Corollary 1, the latter condition is equivalent to $c \leq \tilde{\mu}(1-\tilde{\mu})\Delta\theta = \frac{\mu_0(1-\mu_0)\alpha}{[\mu_0+(1-\mu_0)\alpha]^2}\Delta\theta$. The first condition, instead, corresponds to

$$\begin{aligned} \frac{\theta_l}{\alpha\varepsilon} &\leq \frac{\tilde{\mu}\theta_h + (1-\tilde{\mu})\varepsilon\theta_l - c}{\tilde{\mu} + (1-\tilde{\mu})\varepsilon} \\ c &\leq \tilde{\mu}\theta_h + (1-\tilde{\mu})\varepsilon\theta_l - \frac{\tilde{\mu} + (1-\tilde{\mu})\varepsilon}{\alpha\varepsilon}\theta_l \\ c &\leq \tilde{\mu}\theta_h - \frac{\tilde{\mu} + (1-\tilde{\mu})(1-\alpha\varepsilon)\varepsilon}{\alpha\varepsilon}\theta_l \\ c &\leq \frac{\mu_0\theta_h}{\mu_0 + (1-\mu_0)\alpha} - \frac{\mu_0 + (1-\mu_0)(1-\alpha\varepsilon)\alpha\varepsilon}{\mu_0\alpha\varepsilon + (1-\mu_0)\alpha^2\varepsilon}\theta_l. \end{aligned}$$

Since both necessary conditions must hold, c must be weakly lower than $\min\{\tilde{\mu}(1-\tilde{\mu})\Delta\theta, \tilde{\mu}\theta_h - \frac{\tilde{\mu}+(1-\tilde{\mu})(1-\alpha\varepsilon)\varepsilon}{\alpha\varepsilon}\theta_l\}$. $\square$

Proof of Proposition 7

Proof. Suppose $\alpha = 1-\delta_\alpha$ and $\varepsilon = 1-\delta_\varepsilon$, with $\delta_\alpha, \delta_\varepsilon > 0$ such that $(\alpha, \varepsilon) \in (0,1)^2$. Then,

$$\tilde{\mu} = \frac{\mu_0}{\mu_0 + (1-\mu_0)(1-\delta_\alpha)} \in (0,1)$$

which implies

$$\tilde{\mu}(1-\tilde{\mu})(1-\varepsilon)\Delta\theta = \tilde{\mu}(1-\tilde{\mu})\delta_\varepsilon\Delta\theta > 0$$

so that the first term of the min function in the definition of $\bar{c}_{ba}$ is strictly positive.

Now, we know from the proof of Corollary 2 that

$$\tilde{\mu}\theta_h - \frac{\tilde{\mu} + (1-\tilde{\mu})(1-\alpha\varepsilon)\varepsilon}{\alpha\varepsilon}\theta_l = \frac{\mu_0\theta_h}{\mu_0 + (1-\mu_0)\alpha} - \frac{\mu_0 + (1-\mu_0)(1-\alpha\varepsilon)\alpha\varepsilon}{\mu_0\alpha\varepsilon + (1-\mu_0)\alpha^2\varepsilon}\theta_l.$$

When $\delta_\alpha, \delta_\varepsilon \to 0$, we have $\alpha, \varepsilon \to 1$. Consequently, the above expression tends to

$$\frac{\mu_0\theta_h}{\mu_0 + (1-\mu_0)} - \frac{\mu_0 + (1-\mu_0)(1-1)}{\mu_0 + (1-\mu_0)}\theta_l = \mu_0\theta_h - \mu_0\theta_l = \mu_0\Delta\theta > 0.$$

By continuity of this expression with respect to α and ε, there always exist $\delta_\alpha, \delta_\varepsilon > 0$ such that the second term of the min function is strictly positive, which implies that $\bar{c}_{ba} > 0$.

□

Proof of Corollary ??

Proof. The proof is straightforward once we account for the definition of each equilibrium:

(i) In separating equilibria, the high type charges $p^* < \theta_h$ and the low type charges a price of θ_l. Therefore, the probability of fraud is 0. Moreover, the buyer has degenerate posterior and buys regardless of whether they are zero or one, so that the probability of trade is 1.

(ii) In pooling equilibria, the low type charges a price $p^* > \theta_l$, so any time a trade involving a low type takes place it is fraudulent:

  (a) In pooling(e_h)-b equilibria, the trade takes place if the type is high or the type is low and the evidence technology fails to recognize the low type (α).

Thus, the probability of trade is $\mu_0 + (1-\mu_0)\alpha$ and the probability of fraud is $(1-\mu_0)\alpha$.

(b) In pooling(e_h)-ba equilibria, the trade takes place if the type is high or the type is low and both the evidence and the authentication technologies fail to recognize the low type ($\alpha\varepsilon$). Thus, the probability of trade is $\mu_0 + (1-\mu_0)\alpha\varepsilon$ and the probability of fraud is $(1-\mu_0)\alpha\varepsilon$.

(c) In pooling(e_l)-b equilibria, the trade always takes place since no test is triggered. Thus, the probability of trade is 1, and the probability of fraud corresponds to the probability of the type being low, $(1-\mu_0)$.

(d) In pooling(e_l)-ba equilibria, the trade takes place if the type is high or the type is low and the authentication technology fails to recognize the low type (ε). Thus, the probability of trade is $\mu_0 + (1-\mu_0)\varepsilon$ and the probability of fraud is $(1-\mu_0)\varepsilon$.

□

Proof of Proposition ??

Proof. Given $(\mu_0, \theta_h, \theta_l, \alpha, \varepsilon)$, if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$ and $\alpha \geq \max\{\theta_l/\varepsilon p_{ba}(\tilde{\mu}), \theta_l/p_b(\tilde{\mu})\}$, then both pooling($e_h$)-ba and pooling($e_h$)-b are feasible. Given corollary **??** the probability of fraud for pooling(e_h)-ba equilibria is $(1-\mu_0)\alpha\varepsilon$ while the probability of fraud for pooling(e_h)-b equilibria is $(1-\mu_0)\alpha$, which prices the claim since $\varepsilon \in (0,1]$.

□

Proof of Lemma 3

Proof. Consider a PBE with maximum out-of-equilibrium requirements $\bar{\mu}^{e_h}$ and $\bar{\mu}^{e_l}$. From Assumption 3, an equilibrium does not survive the belief restriction if there exists a belief $\hat{\mu} \in [0,1]$ and an OOE pair (p, e) such that $\pi_h(p, e, \hat{\mu}) > \pi_h^*$, $\pi_l(p, e, \hat{\mu}) < \pi_l^*$ and $\mu^*(p, e, S_h^e) < \hat{\mu}$. Recall the definitions of $\bar{\mu}_p^{e_h}$ and $\bar{\mu}_p^{e_l}$, as the maximum, i.e., the least pessimistic, $\mu_p^{e_h} = \mu(p, e_h, S_h^e)$ and $\mu_p^{e_l} = \mu(p, e_l, S_h^e)$ necessary to sustain an equilibrium. From Lemma 11, we have that the restrictions on OOE beliefs for the high type are always of the form:

$$\bar{\mu}_p^{e_h}(\theta_h) = \bar{\mu}_p^{e_l}(\theta_h) = \begin{cases} 1 & \text{if } p \in [\theta_l, p^*] \\ \underline{\mu}_b(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \\ \underline{\mu}_{ba}(p) & \text{if } p \in (p^*, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases} \tag{1.36}$$

where we explicitly denote the dependency on the seller's type, and we use the fact that $\pi_h^* = p^*$ in every pure-strategy equilibrium. Note that necessary restrictions are independent of the evidence type. For the low type, however, the restrictions on OOE beliefs are less stringent for deviations to OOE pairs (p, e_h) than for those to (p, e_l); that is, for any OOE price p, $\mu_p^{e_h} \geq \mu_p^{e_l}$. This means it is sufficient to check for deviations to OOE (p, e_h) pairs since they require less pessimistic (i.e., larger) beliefs to sustain an equilibrium. Therefore, we focus on

$$\bar{\mu}_p^{e_h}(\theta_l) = \begin{cases} 1 & \text{if } p \in [\theta_l, \pi_l^*/\alpha] \\ \underline{\mu}_b(p) & \text{if } p \in (\pi_l^*/\alpha, \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right]^{\mathsf{c}} \text{ or p} \in (\pi_l^*/\alpha, \pi_l^*/(\alpha\varepsilon)] \cap [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}] \\ \underline{\mu}_{ba}(p) & \text{if } p \in (\pi_l^*/(\alpha\varepsilon), \theta_h] \cap \left[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}\right] \end{cases} \tag{1.37}$$

Importantly, the definitions of $\bar{\mu}_p^{e_h}(\theta_h)$ and $\bar{\mu}_p^{e_h}(\theta_l)$ imply that for any OOE p, if $\hat{\mu} \geq \bar{\mu}_p^{e_h}(\theta_h)$ then $\pi_h(p, e_h, \hat{\mu}) > \pi_h^*$, while conversely, if $\hat{\mu} < \bar{\mu}_p^{e_h}(\theta_h)$ then $\pi_l(p, e_h, \hat{\mu}) < \pi_l^*$ Thus, for any $\hat{\mu} \in (\bar{\mu}_p^{e_h}(\theta_h), \bar{\mu}_p^{e_h}(\theta_l))$ (if the set is nonempty), we have $\pi_h(p, e_h, \hat{\mu}) > \pi_h^*$ and $\pi_l(p, e_h, \hat{\mu}) < \pi_l^*$. Therefore, we can express the refinement condition as follows: an equilibrium does not survive the belief restriction if there exists an OOE p such that

$$\bar{\mu}_p^{e_h}(\theta_h) < \bar{\mu}_p^{e_h}(\theta_l).$$

A direct comparison of $\bar{\mu}_p^{e_h}(\theta_h)$ and $\bar{\mu}_p^{e_h}(\theta_l)$, together with the fact that $\underline{\mu}_{ba}(p) < \underline{\mu}_b(p)$ for $p \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ shows that this inequality holds if and only if at least one of the following conditions is met: (i) $p^* < \pi_l^*/\alpha$; (ii) $p^* \geq \pi_l^*/\alpha$ and the set $(p^*, \theta_h] \cap (\pi_l^*/(\alpha\varepsilon), \pi_l^*/(\alpha\varepsilon)] \cap [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ is nonempty. This nonemptiness requires $p^* < \min\{\bar{\theta}_{\bar{\mu}}, \pi_l^*/(\alpha\varepsilon)\}$ and $\pi_l^*/(\alpha\varepsilon) > \bar{\theta}_{\underline{\mu}}$, or, equivalently, $\pi_l^*/(\alpha\varepsilon) > \max\{p^*, \bar{\theta}_{\underline{\mu}}\}$ and $p^* < \bar{\theta}_{\bar{\mu}}$. Since $\bar{\theta}_{\underline{\mu}} < \bar{\theta}_{\bar{\mu}}$ and $\pi_h^* = p^*$, the statement of the lemma is proved. □

Proof of Proposition 8

Proof. We prove the proposition by applying Lemma 3 to the equilibrium types characterized in Proposition 6 and recalling that for all of them $\pi_h^* = p^*$.

(i) In separating equilibria, we have $p^* \in (\theta_l, \min\{\theta_l/\alpha, \theta_h\}]$ and $\pi_l^* = \theta_l$. According to point (i) of Lemma 3, the unique equilibrium price is $p^* = \min\{\theta_l/\alpha, \theta_h\}$. Moreover, since $\pi_h^* < \pi_l^*/(\alpha\varepsilon)$, point (iii) implies that p^* must not lie in $(\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$. This, combined with point (ii), implies that either $p^* \geq \bar{\theta}_{\bar{\mu}}$ or $p^* < \bar{\theta}_{\bar{\mu}}$ and $\theta_l/(\alpha\varepsilon) < \bar{\theta}_{\underline{\mu}}$. In the first case, we have $p^* = \min\{\theta_l/\alpha, \theta_h\} \in [\bar{\theta}_{\bar{\mu}}, \theta_h]$ with the necessary condition $\alpha \leq \theta_l/\bar{\theta}_{\bar{\mu}}$; in the second case, $p^* = \theta_l/\alpha \leq \bar{\theta}_{\underline{\mu}}$ with the necessary condition $\theta_l/(\alpha\varepsilon) < \bar{\theta}_{\bar{\mu}}$. Note that if $\theta_l/\alpha \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, then no separating equilibria survive.

(ii) In pooling(e_h)-b equilibria:

   (a) If $\tilde{\mu} \in (0, \underline{\mu})$, we have $p^* \in [\theta_l/\alpha, \bar{\theta}_{\tilde{\mu}}]$ with $\bar{\theta}_{\tilde{\mu}} < \bar{\theta}_{\underline{\mu}}$, and $\pi_l^* = \alpha p^*$. Points (i) and (iii) do not have bite. Given point (ii), it must be $\pi_l^*/(\alpha\varepsilon) = p^*/\varepsilon \leq \bar{\theta}_{\underline{\mu}}$ or $p^* \leq \varepsilon\bar{\theta}_{\underline{\mu}}$. Therefore, $p^* \in [\theta_l/\alpha, \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\}]$ with necessary condition $\alpha \geq \max\{\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$.

   (b) If $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$, we have $p^* \in [\theta_l/\alpha, p_b(\tilde{\mu})]$ with $\bar{\theta}_{\tilde{\mu}} \geq \bar{\theta}_{\underline{\mu}}$ and $\pi_l^* = \alpha p^*$. Point (i) does not have bite. Given point (iii), if $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ then it must be $\pi_l^*/(\alpha\varepsilon) = p^*/\varepsilon \leq \bar{\theta}_{\underline{\mu}}$ or $p^* \leq \varepsilon\bar{\theta}_{\underline{\mu}}$ which is impossible. Given point (ii), if $p^* \in [\theta_l, \bar{\theta}_{\underline{\mu}}]$ then it must be $\pi_l^*/(\alpha\varepsilon) = p^*/\varepsilon \leq \bar{\theta}_{\underline{\mu}}$ or $p^* \leq \varepsilon\bar{\theta}_{\underline{\mu}}$. Therefore, $p^* \in [\theta_l/\alpha, \varepsilon\bar{\theta}_{\underline{\mu}}]$ (since $\bar{\theta}_{\tilde{\mu}} \geq \bar{\theta}_{\underline{\mu}} \geq \varepsilon\bar{\theta}_{\underline{\mu}}$) with necessary condition $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ is the only surviving solution in this case.

   (c) If $\tilde{\mu} \in [\bar{\mu}, 1)$, we have $p^* \in [\theta_l/\alpha, \bar{\theta}_{\tilde{\mu}}]$. Now, if $\theta_l/\alpha < \varepsilon\bar{\theta}_{\underline{\mu}}$ then, following the same argument for (b), we find one solution namely $p^* \in [\theta_l/\alpha, \varepsilon\bar{\theta}_{\underline{\mu}}]$ with necessary condition $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. Instead, regardless of θ_l/α, there is always an additional solution namely $p^* \in [\max\{\theta_l/\alpha, \bar{\theta}_{\bar{\mu}}\}, \bar{\theta}_{\tilde{\mu}}]$ with necessary condition $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$. In fact, if $p^* \geq \bar{\theta}_{\bar{\mu}}$, the lemma does not have bite. This, together with the fact that $p^* \in [\theta_l/\alpha, \bar{\theta}_{\tilde{\mu}}]$ provide the second solution.

(iii) In Pooling(e_h)-ba equilibria, we have $p^* \in [\max\{\theta_l/(\alpha\varepsilon), p_b(\tilde{\mu})\}, p_{ba}(\tilde{\mu})]$, $\pi_l^* = \alpha\varepsilon p^*$, with $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$. Since, $[\max\{\theta_l/(\alpha\varepsilon), p_b(\tilde{\mu})\}, p_{ba}(\tilde{\mu})] \subset [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, only point (iii) applies. Moreover, $\pi_l^*/(\alpha\varepsilon) = p^*$ so that any equilibrium of this type survives the belief restriction. So, it implies $p^* \in [\max\{\theta_l/(\alpha\varepsilon), p_b(\tilde{\mu})\}, p_{ba}(\tilde{\mu})]$ with necessary conditions for existence $\alpha \geq \theta_l/(\varepsilon p_{ba}(\tilde{\mu}))$, and $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$.

(iv) In Pooling(e_l)-b, $p^* = \pi_l^* \leq \pi_l^*/\alpha$, therefore, given point (i), no equilibrium of this type survives the belief restriction.

(v) In Pooling(e_l)-ba, we have $p^* \in [\max\{p_b(\mu_0), \theta_l/\varepsilon\}, p_{ba}(\mu_0)]$ an $\pi_l^* = \varepsilon p^*$. Note that point (i) only has bite if $\varepsilon > \alpha$ as it would imply $\pi_l^*/\alpha > \varepsilon p^*/\alpha > p^*$. Suppose, then, $\varepsilon \geq \alpha$. Since, $[\max\{p_b(\mu_0), \theta_l/\varepsilon\}, p_{ba}(\mu_0)] \subset [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, only point (iii) applies, and the fact that $\pi_l^*/(\alpha\varepsilon) = p^*/\alpha > p^*$ implies no equilibrium of this type survives the belief restriction.

□

Proof of Proposition 9

Proof. Proposition 8 characterizes the equilibria that survive the belief restrictions. First, we examine which equilibria within each category satisfy the seller-optimal restrictions. Then, if more than one category is feasible for a given parameter combination, we select the equilibria that survive across categories. Clearly, within each category, the equilibrium that meets the seller's payoff restrictions is the one exhibiting the highest price p^*. Therefore, the surviving equilibria within each category are as follows:

(i) Separating: regardless of $\tilde{\mu}$, the equilibrium price is unique and depends on α: if $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ then $p^* = \theta_l/\alpha$; if $\alpha \geq \theta_l/\bar{\theta}_{\bar{\mu}}$, then $p^* = \min\{\theta_l/\alpha, \theta_h\}$. Let us denote them as the *inefficient* and *efficient* separating, respectively. Note that $p^l = \theta_l$ in either case.

(ii) Pooling(e_h)-b:

  (a) if $\tilde{\mu} \in (0, \underline{\mu})$ and $\alpha \geq \max\{\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$, then $p^* = \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\}$, with $p^* \geq \theta_l/\alpha$.

  (b) if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$ and $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, then $p^* = \varepsilon\bar{\theta}_{\underline{\mu}}$, with $p^* \geq \theta_l/\alpha$.

  (c) if $\tilde{\mu} \in [\bar{\mu}, 1)$ and $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, then $p^* = \varepsilon\bar{\theta}_{\underline{\mu}}$; if $\tilde{\mu} \in [\bar{\mu}, 1)$ and $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$, then $p^* = \bar{\theta}_{\tilde{\mu}}$. Since the second price is larger and the necessary condition is less stringent, the only equilibrium surviving the criterion is $p^* = \bar{\theta}_{\tilde{\mu}} \geq \theta_l/\alpha$, with $p^* \geq \theta_l/\alpha$, whenever $\tilde{\mu} \in [\bar{\mu}, 1)$ and $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$.

(iii) Pooling(e_h)-ba, if $\tilde{\mu} \in [\underline{\mu}, \bar{\mu}]$ and $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$, then $p^* = p_{ba}(\tilde{\mu})$.

There are four regions with unique equilibrium outcome: if $\alpha \in [0, \min\{\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\})$ it is an efficient separating; if $\alpha \in [\theta_l/\varepsilon p_{ba}(\tilde{\mu}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$ and $\tilde{\mu} \in (\underline{\mu}, \bar{\mu})$ it is pooling(e_h)-ba; if $\alpha \in (\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$ and $\tilde{\mu} \in (\bar{\mu}, 1)$ it is pooling(e_h)-b; if $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}})$, and $\tilde{\mu} \in \left[0, \frac{\varepsilon\bar{\theta}_{\underline{\mu}} - \theta_l}{\Delta\theta}\right]$, it is inefficient separating.

Finally, there are four parameter regions, in terms of α and $\tilde{\mu}$, with multiple equilibrium outcomes:

(i) If $\alpha \in [\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/\bar{\theta}_{\bar{\mu}}]$ and $\tilde{\mu} \in [\bar{\mu}, 1)$, the only feasible equilibrium outcomes are efficient separating and pooling(e_h)-b: for separating, $\pi_h^* = \min\{\theta_l/\alpha, \theta_h\} = \theta_l/\alpha$ and $\pi_l^* = \theta_l$; for pooling(e_h)-b $\pi_h^* = \bar{\theta}_{\tilde{\mu}} \geq \theta_l/\alpha$ and $\pi_l^* = \alpha\bar{\theta}_{\tilde{\mu}} \geq \theta_l$, with strict inequalities except at the boundary, i.e., when $\theta_l/\alpha = \bar{\theta}_{\tilde{\mu}}$. Thus, only pooling($e_h$)-b survives the seller-optimal criterion.

(ii) If $\alpha \in [\max\{\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}, 1]$ and $\tilde{\mu} \in (0, \underline{\mu})$, If $\alpha \in [\max\{\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}, 1]$, then the two feasible equilibrium outcomes are inefficient separating and pooling(e_h)-b: for separating, $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$; for pooling(e_h)-b $\pi_h^* = \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\} \geq \theta_l/\alpha$ and $\pi_l^* = \alpha \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\} \geq \theta_l$, with strict inequalities except at the boundary, i.e., when $\theta_l/\alpha = \min\{\bar{\theta}_{\tilde{\mu}}, \varepsilon\bar{\theta}_{\underline{\mu}}\}$. Thus, only pooling($e_h$)-b survives the seller-optimal criterion when they are both feasible.

(iii) If $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$ and $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$, the feasible equilibrium outcomes are inefficient separating, pooling(e_h)-b, and pooling(e_h)-ba: for separating, $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$; for pooling(e_h)-b $\pi_h^* = \varepsilon\bar{\theta}_{\underline{\mu}} \geq \theta_l/\alpha$ and $\pi_l^* = \alpha\varepsilon\bar{\theta}_{\underline{\mu}} \geq \theta_l$, with strict inequalities except at the boundary, i.e., when $\theta_l/\alpha = \varepsilon\bar{\theta}_{\underline{\mu}}$; for pooling($e_h$)-ba $p^* = p_{ba}(\tilde{\mu}) > \varepsilon\bar{\theta}_{\underline{\mu}}$ and $\pi_l^* = \alpha\varepsilon p_{ba}(\tilde{\mu}) \geq \alpha\varepsilon\bar{\theta}_{\underline{\mu}}$. Thus, only pooling($e_h$)-ba survives the seller-optimal criterion.

(iv) If $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$ and $\tilde{\mu} \in [\bar{\mu}, 1)$, the only feasible equilibrium outcomes are inefficient separating and pooling(e_h)-b: for separating, $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$; since $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}) > \theta_l/\bar{\theta}_{\tilde{\mu}}$, $\pi_h^* = \bar{\theta}_{\tilde{\mu}} > \theta_l/\alpha$ and $\pi_l^* = \alpha\bar{\theta}_{\tilde{\mu}} > \theta_l$. Thus, only pooling($e_h$)-b survives the seller-optimal criterion.

Considering all cases analyzed above for the (α, μ)-regions with and without multiple feasible outcomes, the statement of the proposition follows. $\square$

Proof of Lemma 4

Proof. Since $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ is a standard mixed equilibrium with $\sigma_h(p^*, e_h) = 1$, with $\mu' = \mu^*(p^*, e_h, S_h^e)$, the following indifferent conditions for the low type and the buyer must hold:

(1) $\mu' = \frac{\tilde{\mu}}{\tilde{\mu}+(1-\tilde{\mu})\sigma_l(p^*,e_h)}$ or $\sigma_l(p^*, e_h) = \frac{\tilde{\mu}(1-\mu')}{\mu'(1-\tilde{\mu})} = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}$,

(2) $\theta_l = \alpha[\beta(b|p^*, \mu') + \varepsilon\beta(ba|p^*, \mu')]p^*$, or $\beta(b|p^*, \mu') = \theta_l/(\alpha p^*) - \varepsilon\beta(ba|p^*, \mu')$.

Condition (1) *forces* the buyer to hold the specific belief for which mixing is optimal given p^*. Condition (2) equates the expected utility for the low type of posting (θ_l, e_l), which guarantees θ_l, with that of posting (p^*, e_h), where the payoff is p^* multiplied by the overall probability of receiving p^*. This overall probability depends on the probability of passing both the evidence and the authentication tests, as well as the relative probabilities that the buyer purchases the good with and without authentication.

The first condition pins down $\sigma_l(p^*, e_h)$ for any price p^* and requires $\tilde{\mu} < \mu'$. From the second condition, we can derive the buyer's optimal strategy as a function of p^*.

(i) If $p^* \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, ba is dominated, i.e., $\beta(ba|p^*, \mu') = 0$, and the buyer can optimally randomize over b and n only if she holds posterior μ' and $p^* = \bar{\theta}_{\mu'}$. In that case, (2) and $\beta(ba|p^*, \mu') = 0$ imply $\beta(b|p^*, \mu') = \theta_l/(\alpha p^*)$ and $\beta(n|p^*, \mu') = 1 - \beta(b|p^*, \mu')$.

(ii) If $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, the buyer randomizes either between ba and n, only if she holds posterior μ' and $p^* = p_{ba}(\mu')$, or ba and b, only if she holds posterior μ' and $p^* = p_b(\mu')$. In the first case, (2) and $\beta(b|p^*, \mu') = 0$ imply $\beta(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$ and $\beta(n|p^*, \mu') = 1 - \beta(ba|p^*, \mu')$. In the second case, (2) and $\beta(b|p^*, \mu') + \beta(ba|p^*, \mu') = 1$ imply $\beta(b|p^*, \mu') = \frac{\theta_l - \alpha\varepsilon p^*}{\alpha(1-\varepsilon)p^*}$ and $\beta(ba|p^*, \mu') = \frac{\alpha p^* - \theta_l}{\alpha(1-\varepsilon)p^*}$.

(iii) If $p^* = \bar{\theta}_{\underline{\mu}}$, the buyer is indifferent among all her actions only if she holds posterior $\mu' = \underline{\mu}$. In that case, (2), $\beta(b|p^*, \mu') + \beta(ba|p^*, \mu') \leq 1$, and weak positivity constraints require that $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$, given $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) \in \left[0, \min\left\{\frac{\alpha\bar{\theta}_{\underline{\mu}} - \theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}\right]$, and $\beta(n|p^*, \underline{\mu}) = 1 - \beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) - \beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$.

(iv) If $p^* = \bar{\theta}_{\bar{\mu}}$, the statement is proved following the same reasoning as in case (iii) by substituting $\bar{\theta}_{\bar{\mu}}$ and $\bar{\mu}$ in place of $\bar{\theta}_{\underline{\mu}}$ and $\underline{\mu}$, respectively.

□

Proof of Proposition 10

Proof. First, recall that in any standard mixed equilibria, $\pi_l^* = \theta_l$ by construction, while $\pi_h^* = [\beta(b|p^*, \mu') + \beta(ba|p^*, \mu')]p^*$. Then, from Lemma 4, and Lemma 3, a standard mixed equilibrium surviving the belief restriction of Assumption 3 must belong to one of the following categories depending on p^*:

(i) Mixed(θ_l,$\bar{\theta}_{\underline{\mu}}$), if $p^* \in (\theta_l, \bar{\theta}_{\underline{\mu}})$. From Lemma 4 we have $p^* = \bar{\theta}_{\mu'}$, $\beta(b|p^*, \mu') = \theta_l/(\alpha p^*)$, and $\beta(ba|p^*, \mu') = 0$. A necessary condition is $p^* \geq \theta_l/\alpha$ or $\alpha \geq \theta_l/\bar{\theta}_{\mu'}$. Moreover, $\pi_h^* = \theta_l/(\alpha p^*)p^* = \theta_l/\alpha \leq \bar{\theta}_{\mu'} < \bar{\theta}_{\underline{\mu}}$ since $p^* \in (\theta_l, \bar{\theta}_{\underline{\mu}})$. Therefore, the belief restriction requires $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ or equivalently $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. Taking both conditions into account, we have the overall necessary condition: $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu'}\}$.

(ii) Mixed($\bar{\theta}_{\underline{\mu}}$), if $p^* = \bar{\theta}_{\underline{\mu}}$. From Lemma 4 we have that $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$, with $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) \in \left[0, \min\left\{\frac{\alpha\bar{\theta}_{\underline{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}\right]$. A necessary condition is $\bar{\theta}_{\underline{\mu}} \geq \theta_l/\alpha$ or $\alpha \geq \theta_l/\bar{\theta}_{\underline{\mu}}$. Moreover, $\pi_h^* = [\theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) + \beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})]\bar{\theta}_{\underline{\mu}} = [\theta_l/(\alpha\bar{\theta}_{\underline{\mu}}) + (1-\varepsilon)\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})]\bar{\theta}_{\underline{\mu}} \in [\theta_l/\alpha, \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\underline{\mu}}\}]$ and π_h^* is strictly increasing in $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$. Since $\pi_h^* \leq \bar{\theta}_{\underline{\mu}}$, the belief restriction requires $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ or equivalently $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. Taking both the latter and the initial condition into account, we have the overall necessary condition: $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\underline{\mu}}\} = \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$.

(iii) Mixed($\bar{\theta}_{\underline{\mu}}$,$\bar{\theta}_{\bar{\mu}}$), if $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$. From Lemma 4 we know there are two cases. If $p^* = p_{ba}(\mu')$, $\beta(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$ and $\beta(b|p^*, \mu') = 0$. A necessary condition is $p_{ba}(\mu') \geq \theta_l/(\alpha\varepsilon)$ or $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu'))$. Moreover, $\pi_h^* = \theta_l/(\alpha\varepsilon p_{ba}(\mu'))p_{ba}(\mu') = \theta_l/(\alpha\varepsilon) = \pi_l^*/(\alpha\varepsilon)$ so that the belief restriction does not have bite. If $p^* = p_b(\mu')$, $\beta(b|p^*, \mu') = \frac{\theta_l - \alpha\varepsilon p^*}{\alpha(1-\varepsilon)p^*}$ and $\beta(ba|p^*, \mu') = \frac{\alpha p^* - \theta_l}{\alpha(1-\varepsilon)p^*}$, which require $p_b(\mu') \in [\theta_l/\alpha, \theta_l/(\alpha\varepsilon)]$. Moreover, $\pi_h^* = p_b(\mu') \in [\theta_l/\alpha, \theta_l/(\alpha\varepsilon)]$. However, since $\pi_h^* \in (\bar{\theta}_{\underline{\mu}}$,$\bar{\theta}_{\bar{\mu}})$, the belief restriction requires $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \pi_h^*$ which, together with the restrictions on $p_b(\mu')$ means $\pi_h^* = p_b(\mu') = \theta_l/(\alpha\varepsilon)$. Therefore, the necessary condition in this case is $p_b(\mu') = \theta/(\alpha\varepsilon) \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$ or $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$.

(iv) Mixed($\bar{\theta}_{\bar{\mu}}$), if $p^* = \bar{\theta}_{\bar{\mu}}$. From Lemma 4 we have that $\beta(b|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \theta_l/(\alpha\bar{\theta}_{\bar{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu})$, with $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) \in \left[0, \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}\right]$. A necessary condition is $\bar{\theta}_{\bar{\mu}} \geq \theta_l/\alpha$ or $\alpha \geq \theta_l/\bar{\theta}_{\bar{\mu}}$. Moreover, $\pi_h^* = [\theta_l/(\alpha\bar{\theta}_{\bar{\mu}}) - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) + \beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu})]\bar{\theta}_{\bar{\mu}} = [\theta_l/(\alpha\bar{\theta}_{\bar{\mu}}) + (1-\varepsilon)\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu})]\bar{\theta}_{\bar{\mu}} \in [\theta_l/\alpha, \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}]$ and π_h^* is strictly increasing in $\beta(ba|\bar{\theta}_{\bar{\mu}}, \underline{\mu})$. Since $\pi_h^* \leq \bar{\theta}_{\bar{\mu}}$, there are two possibilities. If $\pi_h^* > \bar{\theta}_{\underline{\mu}}$ the belief restriction requires that either $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \pi_h^*$ or $\pi_h^* = \bar{\theta}_{\bar{\mu}}$ which implies $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$ since $\pi_h^* \leq \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$. This means that if $\theta_l/(\alpha\varepsilon) > \bar{\theta}_{\underline{\mu}}$ or, equivalently, $\alpha < \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, then $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}$ so that π_h^* is maximum. If, instead, $\pi_h^* \leq \bar{\theta}_{\underline{\mu}}$ the belief restriction requires $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ or $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. Overall, considering also the initial necessary condition, we have that: either $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\bar{\mu}}\} = \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ or $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$ with $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$ which requires $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}$.

(v) Mixed($\bar{\theta}_{\bar{\mu}}$,θ_h), if $p^* \in (\bar{\theta}_{\underline{\mu}}, \theta_h)$. From Lemma 4 we have $p^* = \bar{\theta}_{\mu'}$, and $\beta(b|p^*, \mu') = \theta_l/(\alpha p^*)$, $\beta(ba|p^*, \mu') = 0$. A necessary condition is then $p^* \geq \theta_l/\alpha$ or $\alpha \geq \theta_l/\bar{\theta}_{\mu'}$. Moreover, $\pi_h^* = \theta_l/(\alpha p^*)p^* = \theta_l/\alpha$. Now, there are two possibilities. If $\pi_h^* = \theta_l/\alpha \geq \bar{\theta}_{\bar{\mu}}$ or $\alpha \leq \theta_l/\bar{\theta}_{\bar{\mu}}$ the belief restriction does not have bite, so that overall requirement is $\alpha \in [\theta_l/\bar{\theta}_{\mu'}, \theta_l/\bar{\theta}_{\bar{\mu}}]$. If, instead, $\pi_h^* = \theta_l/\alpha < \bar{\theta}_{\bar{\mu}}$, the belief restriction requires $\pi_l^*/(\alpha\varepsilon) = \theta_l/(\alpha\varepsilon) \leq \max\{\pi_h^*, \bar{\theta}_{\underline{\mu}}\} = \max\{\theta_l/\alpha, \bar{\theta}_{\underline{\mu}}\}$ or equivalently $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, since $\theta_l/(\alpha\varepsilon) \geq \theta_l/\alpha$. Taking the latter and the initial necessary condition into account, we have that in the second case it must hold $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu'}\} = \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ since $p^* = \bar{\theta}_{\mu'} \in (\bar{\theta}_{\underline{\mu}}, \theta_h)$. Therefore, either $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ or $\alpha \in [\theta_l/\bar{\theta}_{\mu'}, \theta_l/\bar{\theta}_{\bar{\mu}}]$. Since $p^* \in (\bar{\theta}_{\underline{\mu}}, \theta_h)$, the second condition becomes $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\bar{\mu}})$.

Then we apply the seller's optimality restriction within each equilibrium type. Recall $\pi_l^* = \theta_l$ in every equilibrium, so we focus on π_h^*. Note that mixed(θ_l, $\bar{\theta}_{\underline{\mu}}$), mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_{ba}, and mixed($\bar{\theta}_{\bar{\mu}}$, θ_h) equilibria can be sustained by different prices—as specified by their category names—but yield fixed profits for both the high and the low types. Thus, any equilibrium satisfying the type-specific necessary conditions survives the profit restriction. The same holds for mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_b equilibria, which, however, are characterized by a unique high price $\pi_h^* = p_b(\mu') = \theta_l/(\alpha\varepsilon)$. Instead, in mixed($\bar{\theta}_{\underline{\mu}}$) and mixed($\bar{\theta}_{\bar{\mu}}$) equilibria, π_h^* depends positively on $\beta(ba|p^*, \mu')$, so that the only equilibrium in each category that survives the seller's optimality restriction is the one for which $\beta(ba|p^*, \mu')$, and hence π_h^*, is maximized. In particular, the surviving equilibria are mixed($\bar{\theta}_{\underline{\mu}}$) with $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\underline{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}$ and $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\underline{\mu}}\}$, and mixed($\bar{\theta}_{\bar{\mu}}$) with $\beta(ba|\bar{\theta}_{\bar{\mu}}, \bar{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}$ and $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$, respectively.

Finally, comparing the high-type payoff π_h^* of the surviving equilibria in each category at each level of α, we obtain the following:

- If $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$, then all types of standard mixed equilibria, except mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_b, are feasible, assuming they satisfy the posterior requirement. Mixed(θ_l, $\bar{\theta}_{\underline{\mu}}$) and mixed($\bar{\theta}_{\underline{\mu}}$,$\theta_h$) yield $\pi_h^* = \theta_l/\alpha$; mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_{ba} yields $\pi_h^* = \theta_l/(\alpha\varepsilon)$; mixed($\bar{\theta}_{\underline{\mu}}$) and mixed($\bar{\theta}_{\bar{\mu}}$) yield $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\underline{\mu}}\}$ and $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\}$, respectively, which both equal $\theta_l/(\alpha\varepsilon)$ when $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$. Therefore, the last three equilibrium types yield a higher π_h^* and survive the seller's optimality restriction.

- If $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$, three types of equilibria are feasible and all yield the same payoff for the high type, namely $\pi_h^* = \theta_l/(\alpha\varepsilon)$: mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_{ba}, mixed($\bar{\theta}_{\underline{\mu}}$, $\bar{\theta}_{\bar{\mu}}$)-p_b and mixed($\bar{\theta}_{\bar{\mu}}$) since $\min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\} = \theta_l/(\alpha\varepsilon)$ when $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})]$.

- If $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})]$, the only feasible mixed equilibrium is mixed($\bar{\theta}_{\underline{\mu}}$) yielding $\pi_h^* = \min\{\theta_l/(\alpha\varepsilon), \bar{\theta}_{\bar{\mu}}\} = \bar{\theta}_{\bar{\mu}}$.

- If $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\bar{\mu}})$, the only feasible mixed equilibrium is mixed($\bar{\theta}_{\bar{\mu}}$, θ_h) yielding $\pi_h^* = \theta_l/\alpha$.

□

Proof of Corollary 3

Proof. Points (i) and (ii) follow directly from the proof of Proposition 10, after recalling that $\tilde{\mu}$ must be lower than μ', i.e., the target posterior belief, which lies in $(\bar{\mu}, 1)$ for mixed($\bar{\theta}_{\bar{\mu}}$, θ_h) equilibria, and $\bar{\mu}$ for the mixed($\bar{\theta}_{\bar{\mu}}$) equilibrium. We now prove point (iii).

If $\alpha = \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$, we know from Proposition 10 that the unique feasible mixed equilibrium is the mixed($\bar{\theta}_{\bar{\mu}}$) which is characterized by $p^* = \bar{\theta}_{\bar{\mu}}$, $\mu' = \bar{\mu}$, $\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \min\left\{\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}, \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}\right\}$, and $\beta(b|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}} - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu})$. When $\alpha = \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$, we get $\frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}} = \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}} = 1$, so we can write $\beta(ba|p^*,\mu') = \frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}} = \frac{\theta_l}{\alpha\varepsilon p^*}$ and $\beta(b|p^*,\mu') = \frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}} - \varepsilon\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}} - \varepsilon\frac{\theta_l}{\alpha\varepsilon\bar{\theta}_{\bar{\mu}}} = 0$.

If $\alpha \in (\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}))$ we know from Proposition 10 that the only feasible mixed equilibria are mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_b$ and mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_{ba}$. The latters are characterized by $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, $\mu' \in (\underline{\mu}, \bar{\mu})$, $\beta(ba|p^*,\mu') = \frac{\theta_l}{\alpha\varepsilon p^*}$, and $\beta(b|p^*,\mu') = 0$. The formers are characterized by $p^* = \theta_l/(\alpha\varepsilon) \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, $\mu' \in (\underline{\mu}, \bar{\mu})$, $\beta(b|p^*,\mu') = 0$, $\beta(ba|p^*,\mu') = 1$. However, $p^* = \theta_l/(\alpha\varepsilon)$ implies $\theta_l/(\alpha\varepsilon p^*) = 1$ so we can write $\beta(ba|p^*,\mu') = \frac{\theta_l}{\alpha\varepsilon p^*}$.

If $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), 1]$ we know from Proposition 10 that the only feasible mixed equilibria are mixed($\bar{\theta}_{\underline{\mu}}$) and mixed($\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}$)-$p_{ba}$. We know from the previous paragraph that the latter satisfy $\beta(ba|p^*,\mu') = \frac{\theta_l}{\alpha\varepsilon p^*}$, and $\beta(b|p^*,\mu') = 0$. The former, is characterized by $p^* = \bar{\theta}_{\underline{\mu}}$, $\mu' = \underline{\mu}$, $\beta(b|\bar{\theta}_{\underline{\mu}},\underline{\mu}) = 0$, and $\beta(ba|\bar{\theta}_{\underline{\mu}},\underline{\mu}) = \frac{\theta_l}{\alpha\varepsilon\theta_{\underline{\mu}}}$. Thus, we can write $\beta(ba|p^*,\mu') = \frac{\theta_l}{\alpha\varepsilon p^*}$, and $\beta(b|p^*,\mu') = 0$.

Therefore, any standard mixed equilibrium feasible in the range $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), 1]$ exhibits the same strategy for both the buyer and the seller when the semipooling price is p^* and the target posterior is μ'. Moreover, Proposition 10 implies that $\pi_h^* = \theta_l/(\alpha\varepsilon)$ in every case. Hence, all the feasible standard mixed equilibria in this range of α are not only payoff equivalent but also strategically equivalent. ☐

Proof of Proposition 11

Proof. From Proposition 10 and Corollary 3, we can identify four regions in the $(\tilde{\mu}, \alpha)$ space with a unique feasible equilibrium outcome, either pure-strategy or standard mixed-strategy:

(a) if $\alpha \leq \theta_l/\theta_h$, it is a separating equilibrium with $p^* = \min\{\theta_l/\alpha, \theta_h\} = \theta_h$.

(b) if $\alpha \geq \theta_l/\bar{\theta}_{\bar{\mu}}$ and $\tilde{\mu} \in [\bar{\mu}, 1)$, it is a pooling($e_h$)-b equilibrium with $p^* = \bar{\theta}_{\tilde{\mu}}$.

(c) if $\alpha \in [\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}))$ and $\tilde{\mu} \in (0, \bar{\mu})$ it is a mixed($\bar{\theta}_{\bar{\mu}}$) with $p^* = \bar{\theta}_{\bar{\mu}}$.

(d) if $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \min\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon p_{ba}(\tilde{\mu}))\})$ and $\tilde{\mu} \in (0, \bar{\mu})$, it is a mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium with $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$.

Moreover, there are five regions in the $(\tilde{\mu}, \alpha)$ space with two feasible equilibrium outcomes, one pure-strategy and one standard mixed-strategy:

(e) if $\alpha \in (\theta_l/\theta_h, \min\{\theta_l/\bar{\theta}_{\bar{\mu}}, \theta_l/\bar{\theta}_{\tilde{\mu}}\})$ it is either a separating equilibrium with $p^* = \theta_l/\alpha$ or a mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ equilibrium with $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$. Since both yield payoffs $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$, both equilibria survive the payoff restriction.

(f) if $\alpha \in [\theta_l/\bar{\theta}_{\tilde{\mu}}, \theta_l/\bar{\theta}_{\bar{\mu}}]$ and $\tilde{\mu} \in [\bar{\mu}, 1)$, it is either a pooling$(e_h)$-b equilibrium with $p^* = \bar{\theta}_{\tilde{\mu}}$ or a mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$ equilibrium with $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$. The former yields payoffs $\pi_h^* = p^* = \bar{\theta}_{\tilde{\mu}}$ and $\pi_l^* = \alpha p^* = \alpha\bar{\theta}_{\tilde{\mu}}$, while the latter yields payoffs $\pi_h^* = \theta_l/\alpha \leq \bar{\theta}_{\tilde{\mu}}$ and $\pi_l^* = \theta_l \leq \alpha\bar{\theta}_{\tilde{\mu}}$ since $\alpha \geq \theta_l/\bar{\theta}_{\tilde{\mu}}$. Therefore, only the pooling(e_h)-b equilibrium survives the payoff restriction.

(g) if $\alpha \geq \theta_l/\varepsilon p_{ba}(\tilde{\mu})$ and $\tilde{\mu} \in [\underline{\mu}, \bar{\mu})$, it is either a pooling$(e_h)$-ba equilibrium with $p^* = p_{ba}(\tilde{\mu})$ or mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium with $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$. The former yields payoffs $\pi_h^* = p^* = p_{ba}(\tilde{\mu})$ and $\pi_l^* = \alpha\varepsilon p^* = \alpha\varepsilon p_{ba}(\tilde{\mu})$, while the latter yields payoffs $\pi_h^* = \theta_l/(\alpha\varepsilon) \leq p_{ba}(\tilde{\mu})$ and $\pi_l^* = \theta_l \leq \alpha\varepsilon p_{ba}(\tilde{\mu})$ since $\alpha \geq \theta_l/(\varepsilon p_{ba}(\tilde{\mu}))$. Therefore, only the pooling(e_h)-ba equilibrium survives the payoff restriction.

(h) if $\alpha \in [\theta_l/\varepsilon\bar{\theta}_{\underline{\mu}}, \theta_l/\bar{\theta}_{\bar{\mu}}]$ it is either a mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium with $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ or separating equilibrium with $p^* = \theta_l/\alpha$. The former yields payoffs $\pi_h^* = \pi_h^* = \theta_l/(\alpha\varepsilon)$ and $\pi_l^* = \theta_l$, while the latter yields payoffs $\pi_h^* = \theta_l/\alpha \leq \theta_l/(\alpha\varepsilon)$ and $\pi_l^* = \theta_l$. Therefore, only the mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium survives the payoff restriction.

(j) if $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\}$, it is either a pooling$(e_h)$-b equilibrium with $p^* = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}\}$ or a mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium with $p^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$. The former yields the low type payoff $\pi_l^* = \alpha p^* = \alpha \min\{\varepsilon\bar{\theta}_{\underline{\mu}}\}$, while the former yields $\pi_l^* = \theta_l \leq \alpha \min\{\varepsilon\bar{\theta}_{\underline{\mu}}\}$ since $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\tilde{\mu}}\}$. Therefore, the pooling$(e_h)$-b equilibrium survives the payoff restrictions. However, when comparing the high type's payoff does not yield a unique result. The pooling(e_h)-b equilibrium yields $\pi_h^* = p^* = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}\}$, whereas the mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium yields $\pi_h^* = \theta_l/(\alpha\varepsilon)$. The pooling equilibrium yields a larger payoff for the high type than the mixed equilibrium only if $\varepsilon\bar{\theta}_{\underline{\mu}} \geq \theta_l/(\alpha\varepsilon)$ and $\bar{\theta}_{\tilde{\mu}} \geq \theta_l/(\alpha\varepsilon)$, or equivalently, if $\alpha \geq \max\{\theta_l/(\varepsilon^2\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}})\}$. Therefore, if $\alpha \geq \max\{\theta_l/(\varepsilon^2\bar{\theta}_{\underline{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}})\}$ only the pooling(e_h)-b equilibrium survives the payoff restrictions; otherwise, both equilibria survive.

Finally, (a) proves point (i); (e) proves point (ii); (b) and (f) prove point (iii); (c) proves point (iv); (g) proves point (v); (d) and (h) prove point (vi); (j) proves points (vii) and (viii). ☐

Proof of Lemma 5

Proof. First, since we are considering mixed-strategy equilibria, the low type must randomize. In this case, the only set of prices available for the low type to randomize over is $\{\theta_l, p^*\}$. This is because any price–evidence pair $(p, e) \neq (p^*, e_l)$ would imply that the buyer's belief $\mu(p, e, S^e) = 0$ (for any $e \in \{e_l, e_h\}$ and corresponding signal S^e)—thereby forcing the buyer to choose not to buy—unless $p = \theta_l$. Consequently, it must be that $\sigma_l(\theta_l, e) > 0$ and $\sigma_l(p^*, e_l) > 0$, with $\sigma_l(\theta_l, e) + \sigma_l(p^*, e_l) = 1$ for each $e \in \{e_l, e_h\}$.

Second, using the same reasoning as in the proof of Lemma 4, the following indifference conditions must hold in equilibrium for the low type and the buyer:

(1) $\mu' = \frac{\mu_0}{\mu_0 + (1-\mu_0)\sigma_l(p^*, e_l)}$ or $\sigma_l(p^*, e_l) = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)}$,

(2) $\theta_l = [\beta(b|p^*, \mu') + \varepsilon\beta(ba|p^*, \mu')]p^*$, or $\beta(b|p^*, \mu') = \theta_l/p^* - \varepsilon\beta(ba|p^*, \mu')$.

where μ' is the posterior belief for which the buyer is indifferent between at least two of her actions given p^*. The first condition requires $\mu_0 < \mu'$. The second condition determines buyer's optimal strategy as a function of p^*.

(i) If $p^* \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, ba is dominated, i.e., $\beta(ba|p^*, \mu') = 0$, and the buyer can optimally randomize over b and n only if she holds posterior μ' and $p^* = \bar{\theta}_{\mu'}$. In that case, $\beta(b|p^*, \mu') = \theta_l/p^*$ and $\pi_h^* = \pi_l^* = \beta(b|p^*, \mu')p^* = \theta_l$ which means the PBE does not survive the belief restriction by Lemma 3 point (i).

(ii) If $p^* \in (\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}})$, the buyer randomizes either between ba and n, only if she holds posterior μ' and $p^* = p_{ba}(\mu')$, or ba and b, only if she holds posterior μ' and $p^* = p_b(\mu')$.

In the first case, $\beta(ba|p^*, \mu') = \theta_l/(\varepsilon p^*)$, $\pi_h^* = \beta(ba|p^*, \mu')p^* = \theta_l/\varepsilon$ and $\pi_l^* = \theta_l$, with necessary condition $\theta_l/\varepsilon \leq p_{ba}(\mu') < \bar{\theta}_{\bar{\mu}}$. From Lemma 3, the PBE survives the belief restrictions only if $\theta_l/\varepsilon \geq \theta_l/\alpha$ and $\theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ which requires $\alpha \geq \max\{\varepsilon, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$. However, when $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ and $\mu_0 < \mu'$, any alternative feasible equilibrium (which depends on the value of $\tilde{\mu}$) yields at least $\theta_l/(\alpha\varepsilon) > \theta_l/\varepsilon$ for the high type and at least θ_l for the low type. Therefore, the PBE under consideration does not survive the payoff restriction.

In the second case, $\beta(b|p^*, \mu') = \frac{\theta_l - \varepsilon p^*}{(1-\varepsilon)p^*}$ and $\beta(ba|p^*, \mu') = \frac{p^* - \theta_l}{(1-\varepsilon)p^*}$ which requires $p^* \in [\theta_l, \theta_l/\varepsilon]$. Then, $\pi_h^* = [\beta(b|p^*, \mu') + \beta(ba|p^*, \mu')]p^* = p^* < \bar{\theta}_{\bar{\mu}}$ so that, by Lemma 3, this equilibrium survives the belief restrictions only if $\theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$. However, this condition is impossible since, by construction, $\pi_h^* \leq \theta_l/\varepsilon < \theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$, while simultaneously $\pi_h^* = p_b(\mu') > \bar{\theta}_{\underline{\mu}}$. Therefore, the PBE under consideration does not survive the belief restriction.

(iii) If $p^* = \bar{\theta}_{\underline{\mu}}$, the buyer is indifferent among all her actions only if she holds posterior $\mu' = \underline{\mu}$. In that case, $\beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \theta_l/\bar{\theta}_{\underline{\mu}} - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$, given $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) \in \left[0, \min\left\{\frac{\bar{\theta}_{\underline{\mu}} - \theta_l}{(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}\right]$. Then, $\pi_h^* = [\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) + \beta(b|\bar{\theta}_{\underline{\mu}}, \underline{\mu})]\bar{\theta}_{\underline{\mu}}$ which is equal to $[\theta_l/\bar{\theta}_{\underline{\mu}} - \varepsilon\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) + \beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})]\bar{\theta}_{\underline{\mu}} = [\theta_l/\bar{\theta}_{\underline{\mu}} + (1-\varepsilon)\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})]\bar{\theta}_{\underline{\mu}} \in [\theta_l, \min\{\theta_l/\varepsilon, \bar{\theta}_{\underline{\mu}}\}]$. As for the corresponding standard mixed equilibrium, π_h^* is strictly increasing in $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu})$; therefore, to satisfy the intra-category payoff restriction we require $\beta(ba|\bar{\theta}_{\underline{\mu}}, \underline{\mu}) = \min\left\{\frac{\bar{\theta}_{\underline{\mu}} - \theta_l}{(1-\varepsilon)\bar{\theta}_{\underline{\mu}}}, \frac{\theta_l}{\varepsilon\bar{\theta}_{\underline{\mu}}}\right\}$, which implies $\pi_h^* = \min\{\theta_l/\varepsilon, \bar{\theta}_{\underline{\mu}}\}$. Consequently, as in the first part of point (ii), the belief restriction requires $\alpha \geq \max\{\varepsilon, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$ since $\theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ implies $\theta_l/\varepsilon < \bar{\theta}_{\underline{\mu}}$ so that eventually $\pi_h^* = \theta_l/\varepsilon$. However, when $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$ and $\mu_0 < \mu'$, any alternative feasible equilibrium (which depends on the value of $\tilde{\mu}$) yields at least $\theta_l/(\alpha\varepsilon) > \theta_l/\varepsilon$ for the high type and at least θ_l for the low type. Therefore, the PBE under consideration does not survive the payoff restriction.

(iv) If $p^* = \bar{\theta}_{\bar{\mu}}$, the same reasoning as in case (iii) applies by substituting $\bar{\theta}_{\bar{\mu}}$ and $\bar{\mu}$ in place of $\bar{\theta}_{\underline{\mu}}$ and $\underline{\mu}$, respectively. After applying the intra-category payoff restriction, we have $\pi_h^* = \min\{\theta_l/\varepsilon, \bar{\theta}_{\bar{\mu}}\}$, and again the belief restriction requires $\alpha \geq \max\{\varepsilon, \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})\}$ since $\theta_l/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ implies $\theta_l/\varepsilon < \bar{\theta}_{\underline{\mu}}$, so that eventually $\pi_h^* = \theta_l/\varepsilon$. The same reasoning as the previous points leads to the conclusion that the PBE under consideration does not survive the payoff restriction.

$\square$

Proof of Lemma 6

Proof. (i) Suppose there exists $p_j \in P_h \backslash \{p_1\}$ (so that $p_j > p_1$) such that the low type does not randomize, i.e., $p_j \notin P_l$. Then, since $\sigma_h(p_j, e) > 0$ for some $e \in \{e_l, e_h\}$, the buyer would assign a belief of 1 at (p_j, e), i.e., $\mu(p_j, e, S^e) = 1$, regardless of S^e, which implies $\beta(b|p_j, 1) = 1$. But then, condition (1.18) requires $[\beta(b|p_1, \mu_1) + \beta(ba|p_1, \mu_1)] = p_j/p_1 > 1$ which is impossible. Therefore, the low type must assign positive probability to $p_2, ..., p_k$.

(ii) If $|P_h \cap P_l| = 1$ the claim is trivially proved. Assume $|P_h \cap P_l| \geq 2$ and suppose, contrary to the claim, there exist $p_i, p_j \in P_h \cap P_l$ such that $\sigma_h(p_i, e_h), \sigma_l(p_i, e_h) > 0$ and $\sigma_h(p_j, e_l), \sigma_l(p_j, e_l) > 0$. To ease the notation let us write $b_i = \beta(b|p_i, \mu_i)$, $ba_i = \beta(ba|p_i, \mu_i)$, $b_j = \beta(b|p_j, \mu_j)$, and $ba_j = \beta(ba|p_j, \mu_j)$. Depending on how p_i and μ_i relate, we have three scenarios: (a) if $p_i = \bar{\theta}_{\mu_i}$ then $b_i > 0$ and $ba_i = 0$; (b) if $p_i = p_{ba}(\mu_i)$ then $b_i = 0$ and $ba_i > 0$; (c) if $p_i = p_b(\mu_i)$ then $b_i \geq 0$ and $ba_i \geq 0$ and $b_i + ba_i = 1$, so that $ba_i = 1 - b_i$. Clearly, equivalent conditions hold for p_j and μ_j. In total there are 9 possible combinations given the nature of p_i and p_j. Since we are assuming that p_i is associated to e_h while p_j is associated to e_l, the indifferent conditions of the high and the low types are, respectively:

$$(b_i + ba_i)p_i = (b_j + ba_j)p_j, \tag{1.38}$$

$$\alpha(b_i + \varepsilon ba_i)p_i = (b_j + \varepsilon ba_j)p_j. \tag{1.39}$$

First, consider $p_i < p_j$ and note that if $b_j + ba_j = 1$ then condition (1.38) has no solutions; hence, we only need to check the following six remaining combinations.

1. $b_i > 0$, $ba_i = 0$, $b_j > 0$, $ba_j = 0$: from (1.38) we have $b_i = b_j p_j / p_i$, and substituting into (1.39) yields $\alpha b_j p_j = b_j p_j$, which is impossible for $\alpha \in (0, 1)$.

2. $b_i = 0$, $ba_i > 0$, $b_j > 0$, $ba_j = 0$: from (1.38) we have $ba_i = b_j p_j / p_i$, and substituting into (1.39) yields $\alpha \varepsilon b_j p_j = b_j p_j$, which is impossible for $\alpha \in (0, 1)$ and $\varepsilon \in (0, 1)$.

3. $b_i + ba_i = 1$, $b_j > 0$, $ba_j = 0$: from (1.38) we have $b_j = p_i / p_j$, and substituting into (1.39) yields $\alpha[b_i + \varepsilon(1 - b_i)]p_i = p_i$, which is impossible for $\alpha \in (0, 1)$ and $\varepsilon \in (0, 1)$.

4. $b_i > 0$, $ba_i = 0$, $b_j = 0$, $ba_j > 0$: we have $b_i = ba_j p_j / p_i$, and substituting into (1.39) yields $\alpha ba_j p_j = \varepsilon ba_j p_j$, which is not true in general (unless $\alpha = \varepsilon$).

5. $b_i = 0$, $ba_i > 0$, $b_j = 0$, $ba_j > 0$: from (1.38) we have $ba_i = ba_j p_j / p_i$, and substituting into (1.39) yields $\alpha\varepsilon ba_j p_j = \varepsilon ba_j p_j$, which is impossible for $\alpha \in (0,1)$.

6. $b_i + ba_i = 1$, $b_j = 0$, $ba_j > 0$: from (1.38) we have $ba_j = p_i/p_j$, and substituting into (1.39) yields $\alpha[b_i + \varepsilon(1 - b_i)]p_i = p_i$, which is impossible for $\alpha \in (0,1)$ and $\varepsilon \in (0,1)$.

Therefore, p_i cannot be lower than p_j.

Second, consider $p_i > p_j$ and note that if $b_i + ba_i = 1$ then condition (1.38) has no solutions; hence, we only need to check the following six remaining combinations.

1. $b_i > 0$, $ba_i = 0$, $b_j > 0$, $ba_j = 0$: from (1.38) we have $b_j = b_i p_i / p_j$, and substituting into (1.39) yields $\alpha b_i p_i = b_i p_i$, which is impossible for $\alpha \in (0,1)$.

2. $b_i > 0$, $ba_i = 0$, $b_j = 0$, $ba_j > 0$: we have $ba_j = b_i p_i / p_j$, and substituting into (1.39) yields $\alpha b_i p_i = \varepsilon b_j p_i$, which is not true in general (unless $\alpha = \varepsilon$).

3. $b_i > 0$, $ba_i = 0$, $b_j + ba_j = 1$: from (1.38) we have $b_i = p_j / p_i$, and substituting into (1.39) yields $p_j = [b_j + \varepsilon(1 - b_j)]p_j$, which is impossible for $\alpha \in (0,1)$ and $\varepsilon \in (0,1)$.

4. $b_i = 0$, $ba_i > 0$, $b_j > 0$, $ba_j = 0$: from (1.38) we have $b_j = ba_i p_i / p_j$, and substituting into (1.39) yields $\alpha\varepsilon ba_i p_i = ba_i p_i$, which is impossible for $\alpha \in (0,1)$ and $\varepsilon \in (0,1)$.

5. $b_i = 0$, $ba_i > 0$, $b_j = 0$, $ba_j > 0$: we have $ba_j = ba_i p_i / p_j$, and substituting into (1.39) yields $\alpha\varepsilon ba_i p_i = \varepsilon ba_i p_i$, which is impossible for $\alpha \in (0,1)$.

6. $b_i = 0$, $ba_i > 0$, $b_j + ba_j = 1$: from (1.38) we have $ba_i = p_j / p_i$, and substituting into (1.39) yields $\alpha\varepsilon p_j = [b_j + \varepsilon(1 - b_j)]p_j$, which would require $b_j = \frac{\alpha\varepsilon - \varepsilon}{1 - \varepsilon} < 0$ which is impossible.

Therefore, p_i cannot be higher than p_j. Since $p_i \neq p_j$, we reached a contradiction.

(iv) Given (ii), either all prices $\{p_2, ..., p_k\} = P_h \cap P_l$ are posted with high evidence, or all are posted with low evidence. In both cases the indifference condition of the low type is

$$[\beta(b|p_2, \mu_2) + \varepsilon\beta(ba|p_2, \mu_2)]p_2 = ... = [\beta(b|p_k, \mu_k) + \varepsilon\beta(ba|p_k, \mu_k)]p_k. \tag{1.40}$$

If $|P_h \cap P_l| = 1$ the claim is trivially proved. Assume $|P_h \cap P_l| \geq 2$ and consider any two different prices $p_i, p_j \in P_h \cap P_l$. Then, using the same notation as the previous point, the indifferent conditions for the high and the low type imply:

$$(b_i + ba_i)p_i = (b_j + ba_j)p_j, \tag{1.41}$$

$$(b_i + \varepsilon ba_i)p_i = (b_j + \varepsilon ba_j)p_j. \tag{1.42}$$

which, in turn, imply

$$\frac{p_i}{p_j} = \frac{b_j + ba_j}{b_i + ba_i} = \frac{b_j + \varepsilon ba_j}{b_i + \varepsilon ba_i}.$$

By cross-multiplying the last two terms we get

$$\begin{aligned} b_ib_j + b_iba_j + \varepsilon ba_ib_j + \varepsilon ba_iba_j &= b_ib_j + \varepsilon b_iba_j + ba_ib_j + \varepsilon ba_iba_j \\ b_iba_j + \varepsilon ba_ib_j &= \varepsilon b_iba_j + ba_ib_j \\ (1-\varepsilon)b_iba_j &= (1-\varepsilon)ba_ib_j \\ b_iba_j &= ba_ib_j \end{aligned} \tag{1.43}$$

Now, without loss of generality, let us assume $p_i < p_j$ so that $b_j + ba_j = 1$ is not possible given condition (1.41). Then, we have the following six remaining combinations:

1. $b_i > 0$, $ba_i = 0$, $b_j > 0$, $ba_j = 0$: condition (1.43) is satisfied as $0 = 0$.
2. $b_i = 0$, $ba_i > 0$, $b_j > 0$, $ba_j = 0$: condition (1.43) is not satisfied since $ba_ib_j > 0$.
3. $b_i + ba_i = 1$, $b_j > 0$, $ba_j = 0$: condition (1.43) is satisfied as $0 = 0$ only if $ba_i = 0$.
4. $b_i > 0$, $ba_i = 0$, $b_j = 0$, $ba_j > 0$: condition (1.43) is not satisfied since $b_iba_j > 0$.
5. $b_i = 0$, $ba_i > 0$, $b_j = 0$, $ba_j > 0$: condition (1.43) is satisfied as $0 = 0$.
6. $b_i + ba_i = 1$, $b_j = 0$, $ba_j > 0$: condition (1.43) is satisfied as $0 = 0$ only if $b_i = 0$.

Therefore, the indifferent conditions (1.41) and (1.42) are simultaneously satisfied if and only if either $b_i = b_j = 0$ and $ba_i, ba_j > 0$, or $ba_i = ba_j = 0$ and $b_i, b_j > 0$. Since the conditions must hold between any two generic prices in $P_h \cap P_l$, it must be that either (a) $b_2 = ... = b_k = 0$ and $ba_1, ..., ba_k > 0$; or (b) $ba_2 = ... = ba_k = 0$ and $b_2, ..., b_k > 0$.

(iii) Given (iv), condition (1.40) yields either: $\beta(b|p_2, \mu_2)p_2 = ... = \beta(b|p_k, \mu_k)p_k > 0$ and $\beta(ba|p_2, \mu_2)p_2 = ... = \beta(ba|p_k, \mu_k)p_k = 0$, or $\beta(b|p_2, \mu_2)p_2 = ... = \beta(b|p_k, \mu_k)p_k = 0$ and $\beta(ba|p_2, \mu_2)p_2 = ... = \beta(ba|p_k, \mu_k)p_k > 0$, thus proving the claim.

(v) Suppose $\sigma_h(p_j, e_h)>0$ with $p_j \in P_h \cup P_l$. To ensure the buyer's indifference, strategies σ_h and σ_l must satisfy

$$\mu_j = \frac{\mu_0 \sigma_h(p_j, e_h)}{\mu_0 \sigma_h(p_j, e_h) + (1-\mu_0)\sigma_l(p_j, e_h)\alpha} = \frac{\tilde{\mu}\sigma_h(p_j, e_h)}{\tilde{\mu}\sigma_h(p_j, e_h) + (1-\tilde{\mu})\sigma_l(p_j, e_h)} \tag{1.44}$$

Suppose $\mu_j > \tilde{\mu}$ then

$$\begin{aligned} \frac{\tilde{\mu}\sigma_h(p_j, e_h)}{\tilde{\mu}\sigma_h(p_j, e_h) + (1-\tilde{\mu})\sigma_l(p_j, e_h)} &> \tilde{\mu} \\ \sigma_h(p_j, e_h) &> \tilde{\mu}\sigma_h(p_j, e_h) + (1-\tilde{\mu})\sigma_l(p_j, e_h) \\ (1-\tilde{\mu})\sigma_h(p_j, e_h) &> (1-\tilde{\mu})\sigma_l(p_j, e_h) \\ \sigma_h(p_j, e_h) &> \sigma_l(p_j, e_h) \end{aligned}$$

Therefore, $\sigma_h(p_j, e_h) > \sigma_l(p_j, e_h)$, if and only if $\tilde{\mu} < \mu_j$; conversely, $\sigma_h(p_j, e_h) < \sigma_l(p_j, e_h)$, if and only if $\tilde{\mu} > \mu_j$.

Suppose, instead, $\sigma_h(p_j, e_l)>0$ with $p_j \in P_h \cup P_l$. In this case, the buyer's indifference, requires:

$$\mu_j = \frac{\mu_0 \sigma_h(p_j, e_l)}{\mu_0 \sigma_h(p_j, e_l) + (1-\mu_0)\sigma_l(p_j, e_l)} \tag{1.45}$$

Analogous reasoning leads to the conclusion that $\sigma_h(p_j, e_l) > \sigma_l(p_j, e_l)$ if and only if $\mu_0 < \mu_j$; conversely, $\sigma_h(p_j, e_l) < \sigma_l(p_j, e_l)$ if and only if $\mu_0 > \mu_j$. □

Proof of Lemma 7

Proof. From Lemma 6 we know that $P_l \supseteq \{p_2, \ldots, p_k\}$ and that only one type of evidence can be provided along with all prices in $P_h \cap P_l$ in equilibrium. Let $e^* \in \{e_l, e_h\}$ be this unique evidence type. Suppose there exists some $p_j \in P_l$ such that $p_j \notin P_h \cup \{\theta_l\}$. Then, there is some $e \in \{e_l, e_h\}$ with $\sigma_l(p_j, e) > 0$. However, since $p_j \notin P_h$, the buyer's belief must satisfy $\mu(p_j, e, S^e) = 0$ for any signal S^e, implying that $\beta(n|p_j, 0) = 1$ (because $p_j > \theta_l$). Since the low type can always guarantee a payoff of at least θ_l in equilibrium, p_j cannot belong to P_l. Therefore, $P_l \subseteq P_h \cup \{\theta_l\}$. Now, there are three possibilities depending on weather p_1 and θ_l are played with positive probability by the low type:

(i) $\sigma_l(p_1, e') = \sigma_l(\theta_l, e) = 0$, for any $e', e \in \{e_l, e_h\}$: in this case $P_l \subset P_h$, and $\sigma_l(p_1, e') = 0$ implies $\beta(b|p_1, 1) = 1$. Consequently, by condition (1.18), $\pi_h^* = p_1$. From Lemma 6, the buyer's equilibrium strategy exhibits either: (a) $\beta(b|p_j\mu_j) > 0$ and $\beta(ba|p_j\mu_j) = 0$ or (b) $\beta(b|p_j\mu_j) = 0$ and $\beta(ba|p_j\mu_j) > 0$ for all $p_j \in P_h \cap P_l$. Then, by condition (1.18), for every $p_j \in P_h \setminus \{p_1\}$: in case (a), $\beta(b|p_j, \mu^j) = p_1/p_j$; if $e^* = e_h$, then $\pi_l^* = \alpha p_1$ with the necessary condition $\alpha \geq \theta_l/p_1$, while if $e^* = e_l$, $\pi_l^* = p_1$; in case (b), $\beta(ba|p_j, \mu^j) = p_1/p_j$; if $e^* = e_h$, then $\pi_l^* = \alpha\varepsilon p_1$ with the necessary condition $\alpha \geq \theta_l/(\varepsilon p_1)$, while if $e^* = e_l$, then $\pi_l^* = \varepsilon p_1$ with the necessary condition $\varepsilon \geq \theta_l/p_1$.

Now, consider $e^* = e_h$. Since condition (1.20) must hold for all $p_j \in P_h \setminus \{p_1\}$, a necessary condition for such equilibria is $\tilde{\mu} > \mu_2$. In fact, suppose to the contrary that $\tilde{\mu} \leq \mu_2$. Then, by Lemma 6 point (v) we have $\sigma_h(p_2, e_h) \geq \sigma_l(p_2, e_h)$. However, because $\sigma_h(p_1, e_h) > 0$ and $\sigma_l(p_1, e_h) = 0$, it follows that

$$\sum_{p_j \in P_h \setminus \{p_1\}} \sigma_h(p_j, e_h) = \sum_{p_j \in P_l} \sigma_h(p_j, e_h) < 1 = \sum_{p_j \in P_h} \sigma_l(p_j, e_h).$$

Combined with $\sigma_h(p_2, e_h) \geq \sigma_l(p_2, e_h)$, this implies there exists at least one $p_j \in P_h \setminus \{p_1, p_2\}$ such that $\sigma_h(p_j, e_h) < \sigma_l(p_j, e_h)$. This, in turn, implies $\tilde{\mu} > \mu^j > \mu_2 \geq \tilde{\mu}$, a contradiction. If instead we consider $e^* = e_l$ and condition (1.21), the same reasoning applies with μ_0 in place of $\tilde{\mu}$, so that a necessary condition is $\mu_0 > \mu_2$.

(ii) $\sigma_l(p_1, e^*) > 0, \sigma_l(\theta_l, e) = 0$, for some $e'' \in \{e_l, e_h\}$: in this case $P_h = P_l$, and $\sigma_l(p_1, e^*) > 0$ implies that condition (1.18) does not uniquely determine $\beta(b|p_1, \mu_1)$ or $\beta(ba|p_1, \mu_1)$. In fact, by condition (1.18), for every $p_j \in P_h = P_l$: in case (a), $\beta(b|p_j, \mu_j) = \beta(b|p_1, \mu_1)p_1/p_j$, with $\beta(b|p_1, \mu_1) \in (0, 1]$, $\pi_h^* = \beta(b|p_1, \mu_1)p_1$; if $e^* = e_h$, then $\pi_l^* = \alpha\beta(b|p_1, \mu_1)p_1$ with the necessary condition $\alpha \geq \theta_l/(\beta(b|p_1, \mu_1)p_1)$, while if $e^* = e_l$, $\pi_l^* = \beta(b|p_1, \mu_1)p_1$ with the necessary condition $\beta(b|p_1, \mu_1) \geq \theta_l/p_1$; in case (b), $\beta(ba|p_j, \mu_j) = \beta(ba|p_1, \mu_1)p_1/p_j$, with $\beta(ba|p_1, \mu_1) \in (0, 1]$, $\pi_h^* = \beta(ba|p_1, \mu_1)p_1$; if $e^* = e_h$, then $\pi_l^* = \alpha\varepsilon\beta(b|p_1, \mu_1)p_1$ with the necessary condition $\alpha \geq \theta_l/(\varepsilon\beta(b|p_1, \mu_1)p_1)$, while if $e^* = e_l$, $\pi_l^* = \varepsilon\beta(b|p_1, \mu_1)p_1$ with the necessary condition $\beta(b|p_1, \mu_1) \geq \theta_l/\varepsilon p_1$.

Following the same reasoning from point (i), we can show that, if $e^* = e_h$, in this case the requirement on the posterior belief is $\tilde{\mu} \in (\mu_1, \mu_k)$. Specifically, suppose to the contrary that $\tilde{\mu} \leq \mu_1$. Then, by Lemma 6 point (v) we have $\sigma_h(p_1, e_h) \geq \sigma_l(p_1, e_h)$ and, for some $p_j \in P_h \setminus \{p_1\}$, $\sigma_h(p_j, e_h) \leq \sigma_l(p_j, e_h)$ leading to the contradiction $\tilde{\mu} \geq \mu_j > \mu_1 \geq \tilde{\mu}$. Conversely, if $\tilde{\mu} \geq \mu_k$, then $\sigma_h(p_k, e_h) \leq \sigma_l(p_k, e_h)$ and, for some $p_j \in P_h \setminus \{p_k\}$, $\sigma_h(p_j, e_h) \geq \sigma_l(p_j, e_h)$ leading to the contradiction $\tilde{\mu} \leq \mu_j < \mu_k \leq \tilde{\mu}$. If instead we consider $e^* = e_l$ and condition (1.21), the same reasoning applies with μ_0 in place of $\tilde{\mu}$, so that a necessary condition is $\mu_0 \in (\mu_1, \mu_k)$.

(iii) $\sigma_l(p_1, e^*) > 0, \sigma_l(\theta_l, e) > 0$, for some $e \in \{e_l, e_h\}$: in this case $P_l \supset P_h$, and $\sigma_l(\theta_l, e) > 0$ implies $\beta(b|\theta_l, 0) = 1$ so that $\pi_l^* = \theta_l$. From condition (1.40), we determine the equilibrium purchasing strategy of the buyer and the payoff of the high type. In fact, for every $p_j \in P_l \setminus \{\theta_l\} = P_h$: in case (a), if $e^* = e_h$, then $\beta(b|p_j, \mu_j) = \theta_l/(\alpha p_j)$, $\pi_h^* = \theta_l/\alpha$, with the necessary condition $\alpha \geq \theta_l/p_1$, while if $e^* = e_l$, $\beta(b|p_j, \mu_j) = \theta_l/p_j$, and $\pi_h^* = \theta_l$; in case (b), if $e^* = e_h$, then $\beta(ba|p_j, \mu_j) = \theta_l/(\alpha\varepsilon p_j)$, $\pi_h^* = \theta_l/(\alpha\varepsilon)$, with the necessary condition $\alpha \geq \theta_l/(\varepsilon p_1)$, while if $e^* = e_l$, $\beta(ba|p_j, \mu_j) = \theta_l/(\varepsilon p_j)$, $\pi_h^* = \theta_l/\varepsilon$ with the necessary condition $\varepsilon \geq \theta_l/p_1$.

Similarly to the previous points, we derive conditions on the posterior belief of the buyer. Consider first $e^* = e_h$. Since condition (1.20) must hold for all $p_j \in P_h$, a necessary condition for such equilibria is $\tilde{\mu} < \mu_k$. In fact, suppose to the contrary that $\tilde{\mu} \geq \mu_k$. Then, by Lemma 6 point (v) we have $\sigma_h(p_k, e_h) \leq \sigma_l(p_k, e_h)$. However, because $\sigma_h(\theta_l, e) = 0$ and $\sigma_l(\theta_l, e) > 0$, we have

$$\sum_{p_j \in P_l \setminus \{\theta_l\}} \sigma_l(p_j, e_h) = \sum_{p_j \in P_h} \sigma_l(p_j, e_h) < 1 = \sum_{p_j \in P_h} \sigma_h(p_j, e_h).$$

Combined with $\sigma_h(p_k, e_h) \leq \sigma_l(p_k, e_h)$, this implies there exists at least one $p_j \in P_h \setminus \{p_k\}$ such that $\sigma_h(p_j, e_h) > \sigma_l(p_j, e_h)$, which in turn implies $\tilde{\mu} < \mu_j < \mu_k \leq \tilde{\mu}$, a contradiction. If instead we consider $e^* = e_l$ and condition (1.21), the same reasoning applies with μ_0 in place of $\tilde{\mu}$, so that a necessary condition is $\mu_0 < \mu_k$.

□

Proof of Proposition 12

Proof. Lemma 5 proves the claim when $|P_h| = 1$. Suppose instead that $P_h = \{p_1, p_2, \ldots, p_k\}$ with $k \geq 2$ and $p_1 < p_2 < \cdots < p_k$. From Lemma 6 point (ii), we know that the common prices in $P_h \cap P_l$ must be provided with the same type of evidence; and from point (iv) we deduce that for every $p_j \in P_h \cap P_l$ either (a) $\beta(b|p_j, \mu_j) > 0$ and $\beta(ba|p_j, \mu_j) = 0$, or (b) $\beta(b|p_j, \mu_j) = 0$ and $\beta(ba|p_j, \mu_j) > 0$. Moreover, by Lemma 7 we know that either $P_l = P_h \setminus \{p_1\}$, or $P_l = P_h$, or $P_l = P_h \cup \{\theta_l\}$. Combining these prescriptions, we see that sixteen possible strategy combinations can arise in equilibrium (with respect to the type of evidence provided with each price by the two types and the purchasing action of the buyer), of which ten include e_l. In each case we will use the characterization of the equilibrium strategies derived in the proof of Lemma 7. Recall we define e^* as the unique evidence type provided along with common prices in $P_h \cap P_l$.

(i) $\underline{P_l = P_h \setminus \{p_1\}}$: $\pi_h^* = p_1$, and in case (a), $\beta(b|p_j, \mu^j) = p_1/p_j$, while in case (b), $\beta(ba|p_j, \mu^j) = p_1/p_j$.

- (i.i) $\sigma_h(p_1, e_h), \sigma_h(p_j, e_h) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_h$): no type provides e_l.
- (i.ii) $\sigma_h(p_1, e_l), \sigma_h(p_j, e_h) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_h$): (a) $\pi_l^* = \alpha p_1$, which is not sustainable since the low type could deviate to (p_1, e_l) and obtain p_1; (b) $\pi_l^* = \alpha\varepsilon p_1$, which is not sustainable since the low type could deviate to (p_1, e_l) and obtain p_1.
- (i.iii) $\sigma_h(p_1, e_h), \sigma_h(p_j, e_l) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_l$): (a) $\pi_l^* = p_1 = \pi_h^*$, which is not consistent with the belief restriction since $\pi_h^* < \pi_l^*/\alpha$; (b) $\pi_l^* = \varepsilon p_1$, but the belief restriction requires $\pi_h^* \geq \pi_l^*/\alpha \iff p_1 \geq (\varepsilon p_1)/\alpha \iff \alpha \geq \varepsilon$. However, this in turn implies the low type could deviate to (p_1, e_l) and obtain $\alpha p_1 > \varepsilon p_1$ (since we are assuming $\alpha \neq \varepsilon$).
- (i.iv) $\sigma_h(p_1, e_l), \sigma_h(p_j, e_l) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_l$): (a) $\pi_l^* = p_1 = \pi_h^*$, which is not consistent with the belief restriction since $\pi_h^* < \pi_l^*/\alpha$; (b) $\pi_l^* = \varepsilon p_1$, which is not sustainable since the low type could deviate to (p_1, e_l) and obtain p_1.

(ii) $\underline{P_l = P_h}$: in case (a), $\beta(b|p_j, \mu_j) = \beta(b|p_1, \mu_1)p_1/p_j$, with $\beta(b|p_1, \mu_1) \in (0, 1]$, $\pi_h^* = \beta(b|p_1, \mu_1)p_1$, while in case (b), $\beta(ba|p_j, \mu_j) = \beta(ba|p_1, \mu_1)p_1/p_j$, with $\beta(ba|p_1, \mu_1) \in (0, 1]$, $\pi_h^* = \beta(ba|p_1, \mu_1)p_1$.

- (ii.i) $\sigma_h(p_j, e_h) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_h$): no type provides e_l.
- (ii.ii) $\sigma_h(p_j, e_l) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_l$): (a) $\pi_l^* = \beta(b|p_1, \mu_1)p_1 = \pi^*$ which is not consistent with the belief restriction since $\pi_h^* < \pi_l^*/\alpha$; (b) $\pi_l^* = \varepsilon\beta(b|p_1, \mu_1)p_1$, but the belief restriction requires $\pi_h^* \geq \pi_l^*/\alpha \iff \beta(b|p_1, \mu_1)p_1 \geq (\varepsilon\beta(b|p_1, \mu_1)p_1)/\alpha \iff \alpha \geq \varepsilon$. However, this in turn implies the low type could deviate to (p_1, e_l) and obtain $\alpha p_1 > \varepsilon p_1$ (since we are assuming $\alpha \neq \varepsilon$).

(iii) $\underline{P_l = P_h \cup \{\theta_l\}}$: $\pi_l^* = \theta_l$.

(iii.i) $\sigma_h(p_j, e_h) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_h$): no type provides e_l.

(iii.ii) $\sigma_h(p_j, e_l) > 0$ for every $p_j \in P_h \cap P_l$ ($e^* = e_l$): (a), $\beta(b|p_j, \mu_j) = \theta_l/p_j$, and $\pi_h^* = \theta_l = \pi_l^*$, which is not consistent with the belief restriction since $\pi_h^* < \pi_l^*/\alpha$; (b) $\beta(ba|p_j, \mu_j) = \theta_l/(\varepsilon p_j)$, $\pi_h^* = \theta_l/\varepsilon$ but the belief restriction requires both $\pi_h^* \geq \pi_l^*/\alpha \iff \theta_l/\varepsilon \geq \theta_l/\alpha \iff \alpha \geq \varepsilon$ and $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$. However, since we know that this type of equilibrium requires $\mu_0 < \mu_k$, it fails to satisfy the payoff restriction—because the high type could secure at least $\theta_l/(\alpha\varepsilon)$ in any alternative feasible equilibrium (be it a standard mixed, pooling-b, or pooling-ba equilibrium).

Since we exhausted the set of equilibria in which $\sigma_h(p, e_l) > 0$, for any $p \in P_h$, we proved the claim. □

Proof of Proposition 13

Proof. To ease the notation, let us write $b_j = \beta(b|p_j, \mu_j)$, $ba_j = \beta(ba|p_j, \mu_j)$ for $j = 1, 2, ..., k$. Moreover, recall from (1.25) and (1.26) the definition of $\underline{\mu}_{ba}(p)$ and $\underline{\mu}_b(p)$ as the posterior beliefs that makes the buyer indifferent between ba and n and between b and ba, respectively, at price p. By Lemma 7 and Proposition 12, we know that any non-standard mixed equilibrium with $|P_h| \geq 2$ that satisfies both the belief and the payoff restrictions must exhibit common evidence $e^* = e_h$ and either: (a) $P_l = P_h$, (b) $P_l = P_h \setminus \{p_1\}$, (c) $P_l = P_h \cup \{\theta_l\}$. Now, we fully characterize the PBEs with (a) and (b) ad show they fail to satisfy the payoff restrictions whenever they satisfy the belief restriction.

(a) First, consider a PBE with $P_l = P_h$ with $e^* = e_h$. The indifference conditions of the high and the low type are

$$(b_1 + ba_1)p_1 = (b_2 + ba_2)p_2 = ... = (b_k + ba_k)p_k, \tag{1.46}$$

$$\alpha(b_1 + \varepsilon ba_1)p_1 = \alpha(b_2 + \varepsilon ba_2)p_2 = ... = \alpha(b_k + \varepsilon ba_k)p_k. \tag{1.47}$$

Depending on $p_1 = \min P_h$, we have three possible cases:

(a.1) $\underline{p_1 \in (\theta_l, \bar{\theta}_{\underline{\mu}})}$: the buyer can be indifferent only between b and n, meaning $p_1 = \bar{\theta}_{\mu_1}$ with $\mu_1 = \underline{\mu}_b(p_1)$; hence, it must be $b_j > 0$ and $ba_j = 0$, for every $j = 1, 2, ..., k$.

Therefore, (1.46) and (1.47) are equivalent and imply $b_j = b_1 p_1/p_j$ for every $p_j \in P_h$. Note that $b_j < b_1$ necessarily mean $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$, and $p_j \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ for every $p_j \in P_h$. Indeed, for any $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, the requirements $b_j > 0$ and $ba_j = 0$ imply $p_j = p_b(\mu_j)$ and $b_j = 1$, which would contradict (1.46) for any $p_j > p_1$.

Then, $\pi_h^* = b_1 p_1$ and $\pi_l^* = \alpha b_1 p_1$. By Lemma 3, the belief restriction requires $\pi_h^* = b_1 p_1 \geq b_1 p_1 = \pi_l^*/\alpha$, which is always satisfied; and, since $\pi^* \in (\theta_l, \bar{\theta}_{\underline{\mu}})$, we must also have $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$, implying $b_1 \leq \varepsilon\bar{\theta}_{\underline{\mu}}/p_1$. Moreover, the payoff restriction requires b_1 to be maximum since both π_h^* and π_l^* are strictly increasing in b_1. Thus, we get $b_1 = \min\{\varepsilon\bar{\theta}_{\underline{\mu}}/p_1, 1\}$, so that $\pi_h^* = \min\{p_1, \varepsilon\bar{\theta}_{\underline{\mu}}\}$ and $\pi_l^* = \min\{\alpha p_1, \alpha\varepsilon\bar{\theta}_{\underline{\mu}}\}$ with necessary condition $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu_1}\}$ since $p_1 = \bar{\theta}_{\mu_1}$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_1$. However, this implies that whenever the considered PBE is feasible, there exists a feasible pooling(e_h) equilibrium—depending on $\tilde{\mu}$—that yields a larger payoff for both types. For example, if $\tilde{\mu} \in (\mu_1, \underline{\mu})$, a pooling($e_h$)-b equilibrium would be feasible under the necessary conditions described above, and it would guarantee $\bar{\theta}_{\tilde{\mu}}$ to the high type and $\alpha\bar{\theta}_{\tilde{\mu}}$ to the low type, respectively. Consequently, the considered PBE does not satisfy the payoff requirement.

(a.2) $p_1 \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$: the buyer can be indifferent either between ba and n, if $p_1 = p_{ba}(\mu_1)$ (with $\mu_1 = \underline{\mu}_{ba}(p_1)$), or between ba and b if $p_1 = p_b(\mu_1)$ (with $\mu_1 = \underline{\mu}_b(p_1)$).

First, we consider $p_1 = p_{ba}(\mu_1)$. In this case, it must be that $ba_1 > 0$ and $b_1 = 0$, which implies $p_j = p_{ba}(\mu_j)$, with $\mu_j = \underline{\mu}_{ba}(p_j)$ and $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ for every $p_j \in P_h$, as $p_j = p_b(\mu_j)$ would require $ba_j = 1$, thus contradicting (1.46) for any $p_j > p_1$. Then, $\pi_h^* = ba_1 p_1$ and $\pi_l^* = \alpha\varepsilon ba_1 p_1$. By Lemma 3, the belief restriction requires $\pi_h^* = ba_1 p_1 \geq \varepsilon ba_1 p_1 = \pi_l^*/\alpha$, which is always satisfied; and, since $\pi^* \in (\theta_l, \bar{\theta}_{\bar{\mu}}]$, we must also have $\pi_l^*/(\alpha\varepsilon) \leq \max\{\pi_h^*, \bar{\theta}_{\underline{\mu}}\}$, which is always satisfied since $\pi_l^*/(\alpha\varepsilon) = ba_1 p_1 = \pi_h^*$. Moreover, the payoff restriction requires ba_1 to be maximum since both π_h^* and π_l^* are strictly increasing in ba_1. Thus, we get $ba_1 = 1$, so that $\pi_h^* = p_1$ and $\pi_l^* = \alpha\varepsilon p_1$ with necessary condition $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu_1))$ since we are considering $p_1 = p_{ba}(\mu_1)$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_1$. However, this implies that whenever the considered PBE is feasible, there exists a feasible pooling(e_h) equilibrium—depending on $\tilde{\mu}$—that yields a larger payoff for both types. For example, if $\tilde{\mu} \in (\mu_1, \bar{\mu}]$, a pooling($e_h$)-ba equilibrium would be feasible under the necessary conditions above, and would guarantee $p_{ba}(\tilde{\mu})$ to the high type and $\alpha\varepsilon p_{ba}(\tilde{\mu})$ to the low type, respectively. Consequently, the considered PBE does not satisfy the payoff requirement.

Second, we consider $p_1 = p_b(\mu_1)$. Now, either $b_1 = 1$ or $ba_1 = 1$. In the first case, from (1.46) we have $b_j = p_1/p_j < 1$, for every $p_j \in P_h \setminus \{p_1\}$ which implies $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$, and $p_j \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ for every $p_j \in P_h \setminus \{p_1\}$, as proved in point (a.1). Then, $\pi_h^* = p_1$ and $\pi_l^* = \alpha p_1$. By Lemma 3, the belief restriction requires $\pi_h^* = p_1 \geq p_1 = \pi_l^*/\alpha$, which is always satisfied; and, since $\pi^* \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, we must also have $\pi_l^*/(\alpha\varepsilon) \leq \pi_h^*$, or $p_1/\varepsilon \leq p_1$ which is impossible. In the second case, we have $ba_j = p_1/p_j < 1$ and, by symmetric reasoning as above, it must be that $p_j = p_{ba}(\mu_j)$, with $\mu_j = \underline{\mu}_{ba}(p_j) > \underline{\mu}_b(p_1)$ and $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ for every $p_j \in P_h \setminus \{p_1\}$. Then, $\pi_h^* = p_1$ and $\pi_l^* = \alpha\varepsilon p_1$, which clearly satisfy the belief restriction, and require $\alpha \geq \theta_l/(\varepsilon p_1)$. Now the same reasoning for the case $p_1 = p_{ba}(\mu_1)$ with $\mu_1 = \underline{\mu}_{ba}(p_1)$ applies here: in fact, $p_1 = p_b(\underline{\mu}_b(p_1)) = p_{ba}(\underline{\mu}_{ba}(p_1))$ so that in this case Lemma 7 requires an even larger $\tilde{\mu}$ since $\tilde{\mu} > \underline{\mu}_b(p_1) > \underline{\mu}_{ba}(p_1)$ meaning even larger payoffs from alternative and feasible pooling equilibria. Therefore, the considered PBE does not satisfy the payoff requirement.

(a.3) <u>$p_1 \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$</u>: the buyer can be indifferent only between b and n, meaning $p_1 = \bar{\theta}_{\mu_1}$ with $\mu_1 = \underline{\mu}_b(p_1)$; hence, $b_j > 0$ and $ba_j = 0$, for every $j = 1, 2, ..., k$. From (1.46) and (1.47) we have $b_j = b_1 p_1/p_j$ for every $p_j \in P_h$. Note that $p_j > p_1 > \bar{\theta}_{\bar{\mu}}$ necessarily implies $p_j \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ and $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$, for every $p_j \in P_h$.

Then, $\pi_h^* = b_1 p_1$ and $\pi_l^* = \alpha b_1 p_1$. By Lemma 3, the belief restriction requires $\pi_h^* = b_1 p_1 \geq b_1 p_1 = \pi_l^*/\alpha$, which is always satisfied; and, either $\pi_h^* \geq \bar{\theta}_{\bar{\mu}}$, which requires $b_1 \geq \bar{\theta}_{\bar{\mu}}/p_1$ or $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$, implying $b_1 \leq \varepsilon\bar{\theta}_{\underline{\mu}}/p_1$. Clearly, the payoff restriction requires b_1 to be maximum. Thus, we get $b_1 = 1$, so that $\pi_h^* = p_1$ and $\pi_l^* = \alpha p_1$ with necessary condition $\alpha \geq \theta_l/\bar{\theta}_{\mu_1}$ since $p_1 = \bar{\theta}_{\mu_1}$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_1$. However, this implies that, under the necessary conditions described above, there exists a pooling(e_h)-b equilibrium which guarantees $\bar{\theta}_{\tilde{\mu}} > \bar{\theta}_{\mu_1}$ to the high type and $\alpha\bar{\theta}_{\tilde{\mu}} > \alpha\bar{\theta}_{\mu_1}$ to the low type, respectively. Consequently, the considered PBE does not satisfy the payoff requirement.

(b) Second, consider a PBE with $P_l = P_h \setminus \{p_1\}$ with $e^* = e_h$. The indifference conditions of the high and the low type are

$$(b_1 + ba_1)p_1 = (b_2 + ba_2)p_2 = ... = (b_k + ba_k)p_k, \tag{1.48}$$

$$\alpha(b_2 + \varepsilon ba_2)p_2 = ... = \alpha(b_k + \varepsilon ba_k)p_k. \tag{1.49}$$

The only difference with case (a) is that now $\sigma_l(p_1, \mu_1) = 0$ implies $b_1 = 1$ since the buyer recognizes the high type with certainty when she observes (p_1, e_h). Therefore, $\pi_h^* = p_1$. The full characterization of this type of equilibria depends, on the smallest common price, i.e., $p_2 = \min P_h \cap P_l$.

(b.1) $\underline{p_2 \in (\theta_l, \bar{\theta}_{\underline{\mu}})}$: as in point (a.1), the buyer can be indifferent only between b and n, hence, it must be $b_j > 0$ and $ba_j = 0$, for every $j = 2, ..., k$, and also $p_j = \bar{\theta}_{\mu_j}$ with $\mu_j = \underline{\mu}_b(p_j)$, and $p_j \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]$ for every $p_j \in P_h \setminus \{p_1\}$. Therefore, (1.46) and (1.47) are equivalent and imply $b_j = p_1/p_j$ for every $p_j \in P_h$.

Then, $\pi_h^* = p_1$ and $\pi_l^* = \alpha p_1$. By Lemma 3, the belief restriction requires $\pi_h^* \geq \pi_l^*/\alpha$, which is always satisfied, and $p_1 \leq \varepsilon\bar{\theta}_{\underline{\mu}}$ since it must hold $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$. Since, $\pi_l^* = \alpha p_1 \leq \alpha\varepsilon\bar{\theta}_{\underline{\mu}}$ with have as necessary condition $\alpha \geq \theta_l/p_1 = \theta_l/\bar{\theta}_{\mu_1} = \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu_1}\}$ since $p_1 = \bar{\theta}_{\mu_1} \leq \varepsilon\bar{\theta}_{\underline{\mu}}$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_2 > \mu_1$, which implies $\bar{\theta}_{\tilde{\mu}} > \bar{\theta}_{\mu_2} > p_1$. The same reasoning as in point (a.1) applies, so that the considered PBE does not satisfy the payoff requirement.

(b.2) $\underline{p_2 \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]}$: the buyer can be indifferent either between ba and n, if $p_2 = p_{ba}(\mu_2)$ (with $\mu_2 = \underline{\mu}_{ba}(p_2)$), or between ba and b if $p_2 = p_b(\mu_2)$ (with $\mu_2 = \underline{\mu}_b(p_2)$).

First, we consider $p_2 = p_{ba}(\mu_2)$. As in point (a.2), it must hold $ba_j > 0$ and $b_j = 0$, $p_j = p_{ba}(\mu_j)$, with $\mu_j = \underline{\mu}_{ba}(p_j)$ and $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]$ for every $p_j \in P_h \setminus \{p_1\}$. From (1.48), we have $ba_j = p_1/p_j$ or every $p_j \in P_h \setminus \{p_1\}$, and $\pi_h^* = p_1$ and $\pi_l^* = \alpha\varepsilon p_1$, since $b_1 = 1$. The belief restriction is always satisfied since $\pi_h^* = \pi_l^*/(\alpha\varepsilon)$. The no-deviation condition for the low type requires $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu_1))$ since we are considering $p_1 = p_{ba}(\mu_1)$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_2 > \mu_1$, which implies $p_{ba}(\tilde{\mu}) > p_{ba}(\mu_2) > p_1$. The same reasoning as in the first part of point (b.1) applies, so that the considered PBE does not satisfy the payoff requirement.

Second, we rule out $p_2 = p_b(\mu_2)$. Indeed, this would imply that either $b_2 = 1$ or $ba_2 = 1$, but then condition (1.48) would force $p_1 = p_2$, which is impossible.

(b.3) $\underline{p_2 \in (\bar{\theta}_{\bar{\mu}}, \theta_h)}$: the buyer can be indifferent only between b and n, and it must hold $p_j \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ and $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$, for every $p_j \in P_h \setminus \{p_1\}$. From (1.46) and (1.47) we have $b_j = p_1/p_j$ for every $p_j \in P_h \setminus \{p_1\}$.

Then, $\pi_h^* = p_1$ and $\pi_l^* = \alpha p_1$, and, by Lemma 3, the belief restriction requires that either $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ or $\pi_h^* \geq \bar{\theta}_{\bar{\mu}}$, which imply, respectively, $p_1 \leq \varepsilon\bar{\theta}_{\underline{\mu}}$ or $p_1 \geq \bar{\theta}_{\bar{\mu}}$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} > \mu_2 > \mu_1$ so that the level of $\tilde{\mu}$ determines which of the two conditions above holds. The first case is equivalent to (b.1). In the second case, by applying the same reasoning as in point (a.3), we conclude that the considered PBE does not satisfy the payoff requirement when compared to the feasible pooling(e_h)-b equilibrium.

(c) Finally, we consider PBEs with $P_l = P_h \cup \{\theta_l\}$ and show that, under certain specific conditions, they survive both the belief and the payoff restrictions. The indifference conditions of the high and the low type are

$$(b_1 + ba_1)p_1 = (b_2 + ba_2)p_2 = ... = (b_k + ba_k)p_k, \tag{1.50}$$

$$\theta_l = \alpha(b_1 + \varepsilon ba_1)p_1 = \alpha(b_2 + \varepsilon ba_2)p_2 = ... = \alpha(b_k + \varepsilon ba_k)p_k. \tag{1.51}$$

Depending on $p_1 = \min P_h \cap P_l = P_h$, we have three possible cases:

(c.1) $\underline{p_1 \in (\theta_l, \bar{\theta}_{\underline{\mu}})}$: in this case we know it must hold $b_1 > 0$ and $ba_1 = 0$ with $p_1 = \bar{\theta}_{\mu_1}$ and $\mu_1 = \underline{\mu}_b(p_1)$; in turn, this implies $b_j > 0$ and $ba_j = 0$, for every $j = 1, 2, ..., k$, and, by following previous arguments, it must be $p_j \notin [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ and $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$ for every $p_j \in P_h$. From (1.51) we get $b_j = \theta_l/(\alpha p_j)$ for every $p_j \in P_h$. Since b_j must be lower than one for every $j = 1, 2, ..., k$, a necessary condition is that $\alpha \geq \theta_l/p_1$ or, equivalently, $\theta_l/\alpha \leq p_1$.

Then, $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$. Note that By Lemma 3, the belief restriction requires

$\pi_h^* \geq \pi_l^*/\alpha$, which is always satisfied; and, since $\pi_h^* = \theta_l/\alpha \leq \bar{\theta}_{\underline{\mu}}$, we must also have $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$, implying $\alpha \geq \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$. Hence, the overall condition on α is $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu_1}\}$ since $p_1 = \bar{\theta}_{\mu_1}$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} < \mu_k$.

However, whenever the considered PBE is feasible, we know from Corollary 3 and Proposition 11 that there exists either a feasible standard mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium—yielding $\pi_h^* = \theta_l/(\alpha\varepsilon)$ when $\tilde{\mu} \in (0, \bar{\mu})$—or a feasible pooling($e_h$)-b equilibrium—yielding $\pi_h^* = \bar{\theta}_{\tilde{\mu}}$ when $\tilde{\mu} \in (0, \bar{\mu})$. In either case, the considered PBE fails to satisfy the payoff restriction.

(c.2) $p_1 \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$: the buyer can be indifferent either between ba and n, if $p_1 = p_{ba}(\mu_1)$ (with $\mu_1 = \underline{\mu}_{ba}(p_1)$), or between ba and b if $p_1 = p_b(\mu_1)$ (with $\mu_1 = \underline{\mu}_b(p_1)$).

First, we consider $p_1 = p_{ba}(\mu_1)$. In this case, it must be that $ba_1 > 0$ and $b_1 = 0$. From (1.51) we get $ba_j = \theta_l/(\alpha\varepsilon p_j)$ for every $p_j \in P_h$, which implies $p_j = p_{ba}(\mu_j)$, with $\mu_j = \underline{\mu}_{ba}(p_j)$ and $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ for every $p_j \in P_h$, as $p_j = p_b(\mu_j)$ would require $ba_j = 1$, thus contradicting (1.50) for any $p_j > p_1$. Since ba_j must be lower than one for every $j = 1, 2, ..., k$, a necessary condition is that $\alpha \geq \theta_l/(\varepsilon p_1) = \theta_l/(\varepsilon p_{ba}(\mu_1))$. The payoffs are $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$. Since $\pi_h^* = \pi_l^*/(\alpha\varepsilon)$, the belief restriction is always satisfied.

From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} < \mu_k$, and we know that $\mu_k = \underline{\mu}_{ba}(p_k) \leq \bar{\mu}$. Therefore, whenever $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu_1))$, if $\tilde{\mu} \in [\mu_1, \mu_k)$ a pooling(e_h)-ba equilibrium—yielding a larger payoff to both types, namely $p_{ba}(\tilde{\mu})$ for the high type and $\alpha\varepsilon p_{ba}(\tilde{\mu})$ for the low type—would be feasible, since $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu_1)) \geq \theta_l/(\varepsilon p_{ba}(\tilde{\mu})$. Thus, the considered PBE would fail the payoff requirement. Consider, instead, the case $\tilde{\mu} \in (0, \mu_1)$. If $\alpha \geq \theta_l/(\varepsilon p_{ba}(\mu_1))$, then the pair of $(\tilde{\mu}, \alpha)$ belongs to one of the regions (v), (vi), (vii), and (viii) defined in Proposition 11. If $(\tilde{\mu}, \alpha)$ falls in regions (v) or (viii), the considered PBE does not satisfy the payoff restriction, however, if it falls in regions (vi) or (vii), the PBE yields the same payoff as a standard mixed $[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium, thus satisfying both restrictions. In such cases, the non-standard mixed equilibrium survives.

Second, we consider $p_1 = p_b(\mu_1)$. Now, either $b_1 = 1$ or $ba_1 = 1$. In the first case, from (1.51) we necessarily have $p_1 = \theta_l/\alpha$, and $b_j = p_1/p_j = \theta_l/(\alpha p_j)$ for every

$p_j \in P_h \setminus \{p_1\}$ which implies $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$, and $p_j \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ for every $p_j \in P_h \setminus \{p_1\}$, as proved in point (a.1). Then, $\pi_h^* = p_1 = \theta_l/\alpha$ and $\pi_l^* = \alpha p_1 = \theta_l$. By Lemma 3, the belief restriction requires $\pi_h^* \geq \pi_l^*/\alpha$, which is always satisfied; however, since $\pi^* = \theta_l/\alpha = p_1 \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, we must also have $\pi_l^*/(\alpha\varepsilon) \leq \pi_h^*$, or $\theta_l/\varepsilon \leq \theta_l$ which is impossible. In the second case, from (1.51) we necessarily have $p_1 = \theta_l/(\alpha\varepsilon)$, and $ba_j = p_1/p_j = \theta_l/(\alpha\varepsilon p_j)$ for every $p_j \in P_h \setminus \{p_1\}$ which, again, implies $p_j = p_{ba}(\mu_j)$, with $\mu_j = \underline{\mu}_{ba}(p_j)$ and $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ for every $p_j \in P_h \setminus \{p_1\}$. Then, $\pi_h^* = p_1 = \theta_l/(\alpha\varepsilon)$ and $\pi_l^* = \alpha\varepsilon p_1 = \theta_l$, which clearly satisfy the belief restriction. Since $p_1 = \theta_l/(\alpha\varepsilon) \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, a necessary condition is $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})]$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} < \mu_k$, and we know that $\mu_k = \underline{\mu}_{ba}(p_k) \leq \bar{\mu}$. Now, if $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})]$, and $\tilde{\mu} \in (0, \bar{\mu})$, the pair $(\tilde{\mu}, \alpha)$ falls either within region (v) or within region (vi) as defined in Proposition 11. In this case, the fact that $p_1 = \theta_l/(\alpha\varepsilon)$ implies that α is exactly equal to $\theta_l/(\varepsilon p_1) = \theta_l/(\varepsilon p_b(\underline{\mu}_b(p_1))) = \theta_l/(\varepsilon p_{ba}(\underline{\mu}_{ba}(p_1)))$ with $\underline{\mu}_{ba}(p_1) < \underline{\mu}_b(p_1)$. This, in turn, implies that if $\tilde{\mu} \in [\underline{\mu}_{ba}(p_1), \mu_k)$, the pair $(\tilde{\mu}, \alpha)$ lies in region (v) where a pooling(e_h)-ba equilibrium—yielding a larger payoff to both types, namely $p_{ba}(\tilde{\mu})$ for the high type and $\alpha\varepsilon p_{ba}(\tilde{\mu})$ for the low type—would be feasible. Thus, the considered PBE would fail the payoff requirement. If instead, $\tilde{\mu} \in (0, \underline{\mu}_{ba}(p_1))$, then the pair $(\tilde{\mu}, \alpha)$ belongs to region (vi), where the only feasible outcome is a standard mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibrium. Since the proposed equilibrium and a standard mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ yield the same payoffs to the sellers, the non-standard mixed equilibrium satisfies both the belief and the payoff restrictions.

(c.3) $\underline{p_1 \in (\bar{\theta}_{\underline{\mu}}, \theta_h)}$: in this case it must hold $b_1 > 0$ and $ba_1 = 0$ with $p_1 = \bar{\theta}_{\mu_1}$ and $\mu_1 = \underline{\mu}_b(p_1)$; in turn, this implies $b_j > 0$ and $ba_j = 0$, for every $j = 1, 2, ..., k$, and, by following previous arguments, it must be that $p_j \in (\bar{\theta}_{\underline{\mu}}, \theta_h)$ and $p_j = \bar{\theta}_{\mu_j}$, with $\mu_j = \underline{\mu}_b(p_j)$ for every $p_j \in P_h$. From (1.51) we get $b_j = \theta_l/(\alpha p_j)$ for every $p_j \in P_h$. Since b_j must be lower than one for every $j = 1, 2, ..., k$, a necessary condition is that $\alpha \geq \theta_l/p_1$.

Then, $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$, and, by Lemma 3, the belief restriction requires that either $\pi_l^*/(\alpha\varepsilon) \leq \bar{\theta}_{\underline{\mu}}$ or $\pi_h^* \geq \bar{\theta}_{\bar{\mu}}$, which imply, respectively, $\alpha \geq \varepsilon\bar{\theta}_{\underline{\mu}}$ or $\alpha \leq \theta_l/\bar{\theta}_{\bar{\mu}}$. Hence, there are two possibilities for the PBE to survive the belief restriction. In the first, the overall condition on α is $\alpha \geq \max\{\theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}}), \theta_l/\bar{\theta}_{\mu_1}\} = \theta_l/(\varepsilon\bar{\theta}_{\underline{\mu}})$, since $p_1 = \bar{\theta}_{\mu_1}$ and $p_1 > \bar{\theta}_{\tilde{\mu}} > \varepsilon\bar{\theta}_{\underline{\mu}}$. Clearly, the same reasoning used in point (c.1) applies here so that in this case the PBE does not satisfy the payoff restriction. In the second case, the overall condition on α is $\alpha \in (\theta_l/p_1, \theta_l/\bar{\theta}_{\tilde{\mu}}) = (\theta_l/\bar{\theta}_{\mu_1}, \theta_l/\bar{\theta}_{\tilde{\mu}})$. From Lemma 7, we know that the buyer's indifference condition requires $\tilde{\mu} < \mu_k$. In this case, the pair $(\tilde{\mu}, \alpha)$ falls either within region (ii) or within region (iii) as defined in Proposition 11. In particular, if $\tilde{\mu} \in [\mu_1, \mu_k)$ and $\alpha \in (\theta_l/\bar{\theta}_{\mu_1}, \theta_l/\bar{\theta}_{\tilde{\mu}})$, then $(\tilde{\mu}, \alpha)$ falls in region (iii), where a pooling(e_h)-b equilibrium—yielding a higher payoff to both types, namely $\bar{\theta}_{\tilde{\mu}}$ for the high type and $\alpha\bar{\theta}_{\tilde{\mu}}$ for the low type—would be feasible. Thus, the considered PBE would fail the payoff restriction. However, if $\tilde{\mu} \in (0, \mu_1)$, then the pair $(\tilde{\mu}, \alpha)$ belongs to region (ii), where the two feasible outcomes, namely separating and mixed($\bar{\theta}_{\tilde{\mu}}, \theta_h$) equilibria, yield the same payoffs to the sellers as the proposed PBE. Therefore, in this case, the non-standard mixed equilibrium satisfies both the belief and the payoff restrictions.

Therefore, from (a) and (b) we obtain that any non-standard mixed-strategy PBE with $|P_h| \geq 2$ and $P_l = P_h$ or $P_l = P_h \setminus \{p_1\}$ does not satisfy both the belief and the payoff restrictions. Moreover, point (c) shows that non-standard mixed-strategy PBEs with $P_l = P_h \cup \{\theta_l\}$ imply $\pi_l^* = \theta_l$ and must exhibit: either $p_j \in (\bar{\theta}_{\tilde{\mu}}, \theta_h)$ for every $p_j \in P_h$, $\beta(b|p_j, \mu_j) = \theta_l/(\alpha p_j)$, $\beta(ba|p_j, \mu_j) = 0$, $\pi_h^* = \theta_l/\alpha$, with necessary condition $\alpha \in (\theta_l/\theta_h, \theta_l/\bar{\theta}_{\tilde{\mu}})$ since $p_1 \in (\bar{\theta}_{\tilde{\mu}}, \theta_h)$ (thus, proving (i)); or $p_j \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\tilde{\mu}}]$ for every $p_j \in P_h$, $\beta(ba|p_j, \mu_j) = \theta_l/(\alpha\varepsilon p_j)$, $\beta(b|p_j, \mu_j) = 0$, $\pi_h^* = \theta_l/(\alpha\varepsilon)$, with necessary condition $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\tilde{\mu}}), 1]$ (thus, proving (ii)). Note, in fact, that sub-case (c.2) when $p_1 = p_b(\underline{\mu}_b(p_1)) = \theta_l/(\alpha\varepsilon)$ requires $\beta(ba|p_j, \mu_j) = p_1/p_j = \theta_l/(\alpha\varepsilon p_j)$. Therefore, we have proved the claim. $\square$

Proof of Corollary 4

For any non-standard mixed-strategy PBE $\{\hat{\sigma}_h, \hat{\sigma}_l, \hat{\beta}, \hat{\mu}\}$ with $|P_h| \geq 2$ that satisfies both the belief and the payoff restrictions, there exists a corresponding standard mixed-strategy PBE $\{\sigma_h^*, \sigma_l^*, \beta^*, \mu^*\}$ with $\sigma_h^*(p_1, e_h) = 1$, $\beta^*(b|p_1, \mu_1) = \hat{\beta}(b|p_1, \mu_1)$, $\beta^*(ba|p_1, \mu_1) = \hat{\beta}(ba|p_1, \mu_1)$, where $p_1 = \min P_h$, that is payoff-equivalent, i.e., $\pi_h^* = \hat{\pi}_h$ and $\pi_l^* = \hat{\pi}_l$.

Proof. Suppose $\{\hat{\sigma}_h, \hat{\sigma}_l, \hat{\beta}, \hat{\mu}\}$ is a feasible non-standard mixed-strategy PBE satisfying both the belief and the payoff restriction with $P_h = \{p_1, p_2, ..., p_k\}$ and $p_1 = \min P_h$. Then, from Proposition 13, it must be that:

(a) either $\alpha \in (\theta_l/p_1, \theta_l/\bar{\theta}_{\bar{\mu}})$ and $\tilde{\mu} \in (0, \mu_1)$ with $p_1 \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ and $\mu_1 \in (\bar{\mu}, 1)$; equilibrium strategies exhibits $\hat{\beta}(b|p_j, \mu_j) = \theta_l/(\alpha p_j)$, $\hat{\beta}(ba|p_j, \mu_j) = 0$, for every $p_j \in P_h$; and equilibrium payoffs equal $\hat{\pi}_h^* = \theta_l/\alpha$ and $\hat{\pi}_l^* = \theta_l$;

(b) or $\alpha \in [\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}}), 1]$ and $\tilde{\mu} \in (0, \mu_1)$ with $p_1 \in [\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ and $\mu_1 \in [\underline{\mu}, \bar{\mu}]$; equilibrium strategies exhibits $\hat{\beta}(ba|p_j, \mu_j) = \theta_l/(\alpha\varepsilon p_j)$, $\hat{\beta}(b|p_j, \mu_j) = 0$, for every $p_j \in P_h$; and equilibrium payoffs equal $\hat{\pi}_h^* = \theta_l/(\alpha\varepsilon)$ and $\hat{\pi}_l^* = \theta_l$.

From Corollary 3 point (i), we know that when the conditions on α and $\tilde{\mu}$ of point (a) apply, the unique feasible standard mixed equilibrium category is mixed$(\bar{\theta}_{\bar{\mu}}, \theta_h)$, with $\pi_h^* = \theta_l/\alpha$ and $\pi_l^* = \theta_l$, which is characterized by $\sigma_h^*(p^*, e_h) = 1$, $\beta^*(b|p^*, \mu') = \theta_l/(\alpha p^*)$ and $\beta^*(ba|p^*, \mu') = 0$ for some $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ and $\mu' \in (\bar{\mu}, 1)$. Then, by taking $p^* = p_1$, we prove point (i) of the corollary. Moreover, from Corollary 3 point (iii), we know that when the conditions on α and $\tilde{\mu}$ of point (b) apply, the unique feasible standard mixed equilibrium category is mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$, with $\pi_h^* = \theta_l/(\alpha\varepsilon)$ and $\pi_l^* = \theta_l$, which is characterized by $\sigma_h^*(p^*, e_h) = 1$, $\beta^*(ba|p^*, \mu') = \theta_l/(\alpha\varepsilon p^*)$ and $\beta^*(b|p^*, \mu') = 0$ for some $p^* \in (\bar{\theta}_{\bar{\mu}}, \theta_h)$ and $\mu' \in [\underline{\mu}, \bar{\mu}]$. Then, by taking $p^* = p_1$, we prove point (i) of the corollary. □

Proof of Lemma 8

Proof. Given Definition 2 and the equilibrium strategies of separating, pooling-ba, and pooling-b, proving (i),(ii), and (iii) is straightforward. Regarding the remaining expressions:

(iv) For mixed$[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]$ equilibria, we have $\beta(ba|p_{ba}(\mu'), \mu') = \frac{\theta_l}{\alpha\varepsilon p_{ba}(\mu')}$, $\beta(b|p_{ba}(\mu'), \mu') = 0$, $\sigma_l(p_{ba}(\mu'), e_h) = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}$ and $\sigma_l(\theta_l, e) = 1 - \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha} = \frac{[\mu_0+(1-\mu_0)\alpha]\mu'-\mu_0}{\mu'(1-\mu_0)\alpha}$. Substituting into the definition of welfare gives:

$$\begin{aligned}
W_{[\bar{\theta}_{\underline{\mu}}, \bar{\theta}_{\bar{\mu}}]} &= \frac{\mu_0\theta_l\theta_h}{\alpha\varepsilon p_{ba}(\mu')} + (1-\mu_0)\frac{[\mu_0 + (1-\mu_0)\alpha]\mu' - \mu_0}{\mu'(1-\mu_0)\alpha}\theta_l + (1-\mu_0)\alpha\frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}\varepsilon\frac{\theta_l}{\alpha\varepsilon p_{ba}(\mu')}\theta_l \\
&= \frac{\mu_0\theta_l\theta_h}{\alpha\varepsilon p_{ba}(\mu')} + \frac{[\mu_0 + (1-\mu_0)\alpha]\mu' - \mu_0}{\mu'\alpha}\theta_l + \frac{\mu_0(1-\mu')}{\mu'}\frac{\theta_l^2}{\alpha p_{ba}(\mu')} \\
&= \frac{\mu_0\theta_l\theta_h}{\alpha\varepsilon p_{ba}(\mu')} + (1-\mu_0)\theta_l - \frac{\mu_0(1-\mu')(p_{ba}(\mu') - \theta_l)\theta_l}{\alpha\mu' p_{ba}(\mu')},
\end{aligned}$$

while substituting into the definition of the fraud probability gives:

$$F_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]} = (1-\mu_0)\alpha\frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}\varepsilon\frac{\theta_l}{\alpha\varepsilon p_{ba}(\mu')} = \frac{\mu_0(1-\mu')\theta_l}{\alpha\mu' p_{ba}(\mu')}$$

(v) For mixed$(\bar{\theta}_{\bar{\mu}})$ equilibria, we have $\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\alpha\bar{\theta}_{\bar{\mu}}-\theta_l}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}$, $\beta(b|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\theta_l-\alpha\varepsilon\bar{\theta}_{\bar{\mu}}}{\alpha(1-\varepsilon)\bar{\theta}_{\bar{\mu}}}$, $\sigma_l(\bar{\theta}_{\bar{\mu}},e_h) = \frac{\mu_0(1-\bar{\mu})}{\bar{\mu}(1-\mu_0)\alpha}$ and $\sigma_l(\theta_l,e) = 1-\frac{\mu_0(1-\bar{\mu})}{\bar{\mu}(1-\mu_0)\alpha} = \frac{[\mu_0+(1-\mu_0)\alpha]\bar{\mu}-\mu_0}{\bar{\mu}(1-\mu_0)\alpha}$. Note that $\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu})+\beta(b|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = 1$, while $\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu})+\varepsilon\beta(b|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}}$.

Substituting into the definition of welfare gives:

$$\begin{aligned}
W_{(\bar{\theta}_{\bar{\mu}})} &= \mu_0\theta_h + (1-\mu_0)\frac{[\mu_0+(1-\mu_0)\alpha]\bar{\mu}-\mu_0}{\bar{\mu}(1-\mu_0)\alpha}\theta_l + (1-\mu_0)\alpha\frac{\mu_0(1-\bar{\mu})}{\bar{\mu}(1-\mu_0)\alpha}\frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}}\theta_l \\
&= \mu_0\theta_h + \frac{[\mu_0+(1-\mu_0)\alpha]\bar{\mu}-\mu_0}{\bar{\mu}\alpha}\theta_l + \frac{\mu_0(1-\bar{\mu})}{\bar{\mu}}\frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}}\theta_l \\
&= \mu_0\theta_h + (1-\mu_0)\theta_l - \frac{\mu_0(1-\bar{\mu})(\bar{\theta}_{\bar{\mu}}-\theta_l)}{\alpha\bar{\mu}\bar{\theta}_{\bar{\mu}}}\theta_l \\
&= \mu_0\theta_h + (1-\mu_0)\theta_l - \frac{\mu_0(1-\bar{\mu})\Delta\theta}{\alpha\bar{\theta}_{\bar{\mu}}}\theta_l
\end{aligned}$$

since $\bar{\theta}_{\bar{\mu}} = \theta_l + \bar{\mu}\Delta\theta$. Substituting into the definition of fraud probability, instead, gives:

$$F_{(\bar{\theta}_{\bar{\mu}})} = (1-\mu_0)\alpha\frac{\mu_0(1-\bar{\mu})}{\mu'(1-\mu_0)\alpha}\frac{\theta_l}{\alpha\bar{\theta}_{\bar{\mu}}} = \frac{\mu_0(1-\bar{\mu})\theta_l}{\alpha\bar{\mu}\bar{\theta}_{\bar{\mu}}}$$

(vi) For mixed$(\bar{\theta}_{\bar{\mu}},\theta_h)$ equilibria, we have $\beta(ba|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = 0$, $\beta(b|\bar{\theta}_{\bar{\mu}},\bar{\mu}) = \frac{\theta_l}{\alpha\bar{\theta}_{\mu'}}$, $\sigma_l(\bar{\theta}_{\mu'},e_h) = \frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}$ and $\sigma_l(\theta_l,e) = 1-\frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha} = \frac{[\mu_0+(1-\mu_0)\alpha]\mu'-\mu_0}{\mu'(1-\mu_0)\alpha}$.

Substituting into the definition of welfare gives:

$$\begin{aligned}
W_{(\bar{\theta}_{\bar{\mu}},\theta_h)} &= \frac{\mu_0\theta_l\theta_h}{\alpha\bar{\theta}_{\mu'}} + (1-\mu_0)\frac{[\mu_0+(1-\mu_0)\alpha]\mu'-\mu_0}{\mu'(1-\mu_0)\alpha}\theta_l + (1-\mu_0)\alpha\frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}\frac{\theta_l}{\alpha\bar{\theta}_{\mu'}}\theta_l \\
&= \frac{\mu_0\theta_l\theta_h}{\alpha\bar{\theta}_{\mu'}} + \frac{[\mu_0+(1-\mu_0)\alpha]\mu'-\mu_0}{\mu'\alpha}\theta_l + \frac{\mu_0(1-\mu')}{\mu'}\frac{\theta_l^2}{\alpha\bar{\theta}_{\mu'}} \\
&= \frac{\mu_0\theta_l\theta_h}{\alpha\bar{\theta}_{\mu'}} + (1-\mu_0)\theta_l - \frac{\mu_0(1-\mu')(\bar{\theta}_{\mu'}-\theta_l)\theta_l}{\alpha\mu'\bar{\theta}_{\mu'}} \\
&= \frac{\mu_0\theta_l\theta_h}{\alpha\bar{\theta}_{\mu'}} + (1-\mu_0)\theta_l - \frac{\mu_0(1-\mu')\Delta\theta\theta_l}{\alpha\bar{\theta}_{\mu'}} \\
&= (1-\mu_0)\theta_l - \frac{\mu_0(\theta_l+\mu'\theta_h-\mu'\theta_l)\theta_l}{\alpha\bar{\theta}_{\mu'}} \\
&= (1-\mu_0)\theta_l + \frac{\mu_0\theta_l}{\alpha},
\end{aligned}$$

while substituting into the definition of the fraud probability gives:

$$F_{(\bar{\theta}_{\bar{\mu}},\theta_h)} = (1-\mu_0)\alpha\frac{\mu_0(1-\mu')}{\mu'(1-\mu_0)\alpha}\frac{\theta_l}{\alpha\bar{\theta}_{\mu'}} = \frac{\mu_0(1-\mu')\theta_l}{\alpha\mu'\bar{\theta}_{\mu'}}$$

□

Proof of Lemma 9

Proof. Given the expressions in Lemma 8, $\frac{\partial W_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]}}{\partial\varepsilon} < 0$ is obvious. Instead, for mixed($\bar{\theta}_{\bar{\mu}}$) equilibria we have

$$W_{(\bar{\theta}_{\bar{\mu}})} = \mu_0\theta_h + (1-\mu_0)\theta_l - \frac{\mu_0\theta_l(1-\bar{\mu})(\theta_h-\theta_l)}{\alpha\bar{\theta}_{\bar{\mu}}},$$

so that

$$\frac{\partial W_M}{\partial\varepsilon} = -\frac{\mu_0\theta_l(\theta_h-\theta_l)}{\alpha}\frac{\partial}{\partial\varepsilon}\left(\frac{1-\bar{\mu}}{\bar{\theta}_{\bar{\mu}}}\right).$$

Since $\bar{\theta}_{\bar{\mu}} = \bar{\mu}\theta_h + (1-\bar{\mu})\theta_l$, we have that $\frac{\partial\bar{\theta}_{\bar{\mu}}}{\partial\varepsilon} = \bar{\mu}'(\theta_h-\theta_l)$, where

$$\bar{\mu}' = \frac{\partial\bar{\mu}}{\partial\varepsilon} = -\frac{c}{2(1-\varepsilon)^2(\theta_h-\theta_l)\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)(\theta_h-\theta_l)}}}.$$

Therefore,

$$\frac{\partial}{\partial\varepsilon}\left(\frac{1-\bar{\mu}}{\bar{\theta}_{\bar{\mu}}}\right)=\frac{-\bar{\mu}'\bar{\theta}_{\bar{\mu}}-(1-\bar{\mu})\bar{\mu}'(\theta_h-\theta_l)}{\bar{\theta}^2_{\bar{\mu}}}=-\bar{\mu}'\frac{\theta_h}{\bar{\theta}^2_{\bar{\mu}}}$$

and

$$\frac{\partial W_{(\bar{\theta}_{\bar{\mu}})}}{\partial\varepsilon}=-\frac{\mu_0\theta_l(\theta_h-\theta_l)}{\alpha}\left(-\bar{\mu}'\frac{\theta_h}{\bar{\theta}^2_{\bar{\mu}}}\right)=-\frac{\mu_0\theta_l\theta_h c}{2\alpha\bar{\theta}^2_{\bar{\mu}}(1-\varepsilon)^2\sqrt{\frac{1}{4}-\frac{c}{(1-\varepsilon)(\theta_h-\theta_l)}}},$$

which is strictly positive. □

Proof of Proposition 14

Proof. (i) Since $\alpha\geq\theta_l/\bar{\theta}_{\tilde{\mu}}$, and $\tilde{\mu}\in[0,\bar{\mu}^*]$, the only feasible equilbrium outcomes are pooling-ba with $F_{ba}=(1-\mu_0)\alpha\varepsilon$, pooling-b with $F_b=(1-\mu_0)\alpha$, mixed$[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]$ with $F_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]}=\frac{\mu_0(1-\mu')\theta_l}{\alpha\mu' p_{ba}(\mu')}$, and mixed$(\bar{\theta}_{\bar{\mu}})$ with $F_{(\bar{\theta}_{\bar{\mu}})}=\frac{\mu_0(1-\bar{\mu})\theta_l}{\alpha\bar{\mu}\bar{\theta}_{\bar{\mu}}}$. Poling-b yields a strictly larger fraud, regardless of ε. Since mixed$[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]$ feasibility condition requires $\alpha\geq\theta_l/(\varepsilon\bar{\theta}_{\bar{\mu}})$ or equivalently, $\varepsilon\geq\theta_l/(\alpha\bar{\theta}_{\bar{\mu}})$. Then $F_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]}\geq\frac{\mu_0(1-\hat{\mu})\theta_l}{\hat{\alpha}_{ba}(\hat{\mu})}$, where $\hat{\mu}=\bar{\mu}(\varepsilon=\theta_l/(\alpha\bar{\theta}_{\bar{\mu}}))$ since $F_{[\bar{\theta}_{\underline{\mu}},\bar{\theta}_{\bar{\mu}}]}$ is decreasing in $\mu'<\bar{\mu}$.The expression is strictly lower than mixed$(\bar{\theta}_{\bar{\mu}})$ with $e=0$, i.e. $\bar{\mu}=\bar{\mu^*}$. Finally, feasibility conditions for pooling-ba require $\alpha\geq\theta_l/(\varepsilon p_{ba}(\tilde{\mu}))$ or equivalently, $\varepsilon\geq\theta_l/(\alpha p_{ba}(\tilde{\mu}))$. Then, $F_{ba}=(1-\mu_0)\alpha\varepsilon\geq(1-\mu_0)\alpha\theta_l/(\alpha p_{ba}(\tilde{\mu}))$, which is strictly larger than $F_{(\bar{\theta}_{\bar{\mu^*}})}=\frac{\mu_0(1-\bar{\mu^*})\theta_l}{\alpha\bar{\mu^*}\bar{\theta}_{\bar{\mu^*}}}$ since

$$\begin{aligned}(1-\mu_0)\frac{\theta_l}{p_{ba}(\tilde{\mu})}&>\frac{\mu_0(1-\bar{\mu^*})\theta_l}{\alpha\bar{\mu^*}\bar{\theta}_{\bar{\mu^*}}}\\(1-\mu_0)\alpha\frac{\theta_l}{p_{ba}(\tilde{\mu})}&>\frac{\mu_0(1-\bar{\mu^*})\theta_l}{\bar{\mu^*}\bar{\theta}_{\bar{\mu^*}}}\end{aligned}$$

and, by dividing both sides by $\mu_0+(1-\mu_0)\alpha$, we obtain

$$\begin{aligned}(1-\tilde{\mu})\frac{1}{p_{ba}(\tilde{\mu})}&>\frac{\tilde{\mu}(1-\bar{\mu^*})}{\bar{\mu^*}\bar{\theta}_{\bar{\mu^*}}}\\(1-\tilde{\mu})\mu^*\bar{\theta}_{\bar{\mu^*}}&>\tilde{\mu}(1-\bar{\mu^*})p_{ba}(\tilde{\mu}),\end{aligned}$$

which is always satisfied since $\mu^*>\tilde{\mu}$ and $\bar{\theta}_{\bar{\mu^*}}>p_{ba}(\tilde{\mu})$.

(ii) Given the non optimality of intermediate authentication precisions, to maximize welfare in the specified parameter region, we need to compare the perfect authentication outcome, mixed($\bar{\theta}_{\bar{\mu}^*}$) and pooling-b. The latter yields a larger welfare compared to the former if and only if

$$\begin{aligned}
\mu_0\theta_h + (1-\mu_0)\alpha\theta_l &\geq \mu_0\theta_h + (1-\mu_0)\theta_l - \frac{\mu_0(1-\bar{\mu}^*)\Delta\theta}{\alpha\bar{\theta}_{\bar{\mu}^*}}\theta_l \\
(1-\mu_0)\alpha\theta_l &\geq (1-\mu_0)\theta_l - \frac{\mu_0(1-\bar{\mu}^*)\Delta\theta}{\alpha\bar{\theta}_{\bar{\mu}^*}}\theta_l \\
(1-\mu_0)(1-\alpha) &\leq \frac{\mu_0(1-\bar{\mu}^*)\Delta\theta}{\alpha\bar{\theta}_{\bar{\mu}^*}} \\
(1-\mu_0)\alpha(1-\alpha) &\leq \frac{\mu_0(1-\bar{\mu}^*)\Delta\theta}{\bar{\theta}_{\bar{\mu}^*}}
\end{aligned}$$

and, by dividing both sides by $\mu_0 + (1-\mu_0)\alpha$, we obtain

$$\begin{aligned}
(1-\tilde{\mu})(1-\alpha) &\leq \frac{\tilde{\mu}(1-\bar{\mu}^*)\Delta\theta}{\bar{\theta}_{\bar{\mu}^*}} \\
\alpha &\geq 1 - \frac{\tilde{\mu}(1-\bar{\mu}^*)\Delta\theta}{(1-\tilde{\mu})\bar{\theta}_{\bar{\mu}^*}},
\end{aligned}$$

where the RHS is strictly lower than 1. □

1.8.3 Proof of Corollary

Proof. See the last sentence of the proof of Proposition 14. □

Chapter 2

Soft News Strategies in Networks

2.1 Introduction

Consider the social media manager of a company that sends *personalized* ads to different subscribers of a social media platform about the products their company produces depending on the quality of the products they have in stock at the moment. Assume that receivers know that the content of the ads depends on the level of quality of the product and share the ads among their friends. These are divided into different categories: some of them love the products, while others are not regular customers. How should the social media manager design the communication strategy to maximize the inflow of new clients while being sure of retaining all fond customers? How does the structure of the receivers' relationships influence the optimal strategy?

We consider a binary Bayesian persuasion model with a sender and multiple heterogeneous receivers. As standard in the persuasion literature, the sender is able to commit to a state-dependent communication strategy. We assume the sender can send private messages to receivers. However, receivers are connected to each other and share information among neighbors. The goal of the sender is to induce receivers to take the sender's favorite action while receivers would like to match their action to the true underlying state. Receivers differ in their initial beliefs about the state of the world and can be divided into *believers* – those who assign a relatively high probability to the sender's favorable state – and sceptics – those that would a priori take sender's unfavored action. A feature of the model is that the sender is somehow constrained in the distribution of messages she can send to receivers: in fact, the strategy requires specifying probability for binary messages and the simultaneous correlation among them. We will present this characteristic formally in Section 2.

Innocenti (2021) considers the problem faced by a sender when heterogeneous receivers do not share information and the signal is public. The main finding is that there are only two optimal strategies for the sender depending on the parameters of the model. The first one is the so-called *hard news* strategy that sends a persuading message with probability one in the favorable state. They are designed to maximize the probability of persuading sceptics at the cost of dissuading believers under some signal realizations in the bad state. The second type of strategy, the *soft news* strategies, ensures that believers take the sender's favorite action regardless of the state, at the cost of losing the sceptics in the good state with positive probability.

Liporace (2021) studies the problem in a network with private messages and information spillovers. The focus, however, is on the network counterparts of the hard news strategies defined before. This paper borrows the main assumptions of the model and builds on it in order to extend the analysis to the soft news strategies in the connected world.

We formally define soft news strategy in a network as those strategies that send some persuading messages to believers and some to sceptics in the favorable state of the world with probability one. We then provide some preliminary results about different types of strategies. In particular, we show that not all the hard news strategies in Liporace (2021) have their soft news counterpart.

Related literature. We provide a brief literature review strictly connected to this work. The seminal paper of Kamenica and Gentzkow (2011) paved the way for the significant literature on Bayesian persuasion. They study a simple model with one sender and one receiver,, and their findings represent the groundwork of this paper.

Recently, different authors have focused on Bayesian persuasion models in networks.

Alonso and Camara (2016) consider an information voting model, where a sender (politician) can influence receivers' (voters') choices by designing a policy experiment. Kerman et al. (2020), and Kerman and Tenev (2021) study a multiple-receiver Bayesian persuasion model with homogeneous receivers. The sender sends private (possibly) correlated messages to the receivers, who are in an exogenous and commonly known network, in order to get their positive vote. As in our setting, receivers can observe their neighbors' messages. They find that the sender's value from persuasion is not monotonically decreasing with the density of the network.

Egorov and Sonin (2020) analyze a model, in which each receiver can buy direct access to the sender's signal or rely on network spillovers to get the same information. They find that the density of the network has a non-monotonic effect on the optimal level of propaganda as well.

Innocenti (2021), considers a model with heterogenous priors but where receivers are not connected. He finds that soft news strategies perform better than hard news strategies when (a) the group of believers is predominant, (b) polarization is low, and (c) the sender assigns a high probability to the favorable state. In addition, he considers the case with multiple senders and finds the *echo chambers* arise endogenously even though they are harmful to receivers.

Finally, Liporace (2021) studies a Bayesian-persuasion problem on a fixed communication network, focusing on the network counterparts of the "hard-news" strategies first identified by Innocenti (2021). She shows that, although these hard-news mechanisms underperform a public signal when receivers' priors are homogeneous, they can strictly outperform public communication whenever priors are heterogeneous and the average degree of the group that should not be dissuaded is sufficiently low. In particular, in segregated networks—where within-group ties dominate—hard-news strategies both (a) exploit connectivity to concentrate persuasion on targeted subsets of receivers and (b) mitigate the usual trade-off between persuading sceptics and preserving believers.

By contrast, our chapter extends the analysis to soft-news strategies in the same networked setting. We characterize the full class of soft-news equilibria, and provide preliminary results on how to derive welfare and fraud comparisons to show not only when network-exploiting soft-news strategies outperform public signaling, but also to build the ground to comparison with previously studied hard-news strategies in networks.

This work builds mostly on the framework analyzed in Liporace (2021) and is intended to bridge the gap between her work and the findings in Innocenti (2021). The model in the paper resembles that of Liporace (2021) and different relevant findings on the optimal strategy of the sender in networks are reported here as they are considered instrumental to the analysis of soft news strategy. Different from her research, we focus on different types of strategies that have not been studied yet in the literature on Bayesian persuasion. Moreover, we extend the definitions in Innocenti (2021) to the case of receivers belonging to a network. We provide some preliminary results on classes of soft news strategies. Furthermore, we state some educated conjectures regarding the characterization of the set of soft news strategies in networks.

2.1.1 Outline of the Paper

In Section 2 we present the model and introduce the necessary notation. In Section 3 we provide previous results on the optimal strategy of the sender. In particular, we report the main findings on hard news strategies in networks. In Section 4 we define soft news strategies in networks and provide some preliminary results. Section 5 concludes with a brief discussion and with the future steps to take in order to enhance the analysis of soft news strategies in networks.

2.2 Model

We consider a Bayesian persuasion model with a sender (she) and multiple heterogeneous receivers (they/he) connected via a network. The sender observes the state of the world and sends a private message to each receiver in order to persuade him to take her favorite action. Instead, each receiver wants to match his action to the realized state, of which they are not aware. In order to update his belief over the state, he uses not only his own message but also the messages received by his neighbors, wich all together represent the receiver's *signal.* We now provide a formal outline of the setting and the assumptions of the model.

Preliminaries. Consider a binary state of the world. Let ω denote the state and $\Omega = \{0, 1\}$ the state space. At the beginning of the game, the state realization is observed by the sender, but not by the receivers.

Let $N = \{1, 2, ..., n\}$ denote the (finite) set of receivers, and $J = \{sender\} \cup N$ the set of players. Each agent $j \in J$ holds a prior belief with full support over the possible state $\mu_j(\cdot) \in \Delta^o(\Omega)$, where $\Delta^o(\Omega)$ denotes the set of strictly positive probability distributions over Ω. Given the binary nature of Ω, we can fully specify agent j's prior by setting $\mu_j := \mathbb{P}(\omega = 1) \in (0, 1)$, i.e., j's prior belief that the state is 1. As in Innocenti (2021), we drop the common prior assumption: the sender has prior μ_S, while receivers display prior's heterogeneity.

Assumption 1 (Heterogeneous receivers)**.** *Receivers are heterogeneous in their prior. In particular, we assume that $N = A \cup B$, where all members of group A (respectively, B) share a common prior μ_A (μ_B), with $\mu_A > \mu_B$.*

Let n_A and n_B denote the number of receivers in each set and define $g_A = \frac{n_A}{n}$ and $g_B = \frac{n_B}{n}$, with $g_A + g_B = 1$.

Finally, let X indicate the common binary set of actions of the receivers, and M a binary set of individual messages the sender can communicate to receivers. In order to ease the notation, we assume throughout $X = M = \Omega$.

Network. We assume that receivers are organized in an undirected network $\Gamma = \{N, L\}$, where $L \subset N \times N$ denotes the set of links, i.e., $(i, j) \in L$ if and only if receivers i and j are connected.[1] Write $\mathcal{N}_i = \{j \in N : (i, j) \in L\} \cup \{i\}$ to indicate the neighborhood set of receiver i, including i himself, and $\mathcal{N}_i^A := \mathcal{N}_i \cap A$ ($\mathcal{N}_i^B := \mathcal{N}_i \cap B$) i's neighbors of the A (B) type. We define i's G-degree as $d_i^G := |\mathcal{N}_i^G|$ to be the number of neighbors of i, including himself, that belong to group $G \in \{A, B\}$.

The network is assumed to be exogenous and locally finite so that $\max_i |\mathcal{N}_i^G|$ is finite for every $i \in N$ and G. Given the definitions above, the network is characterized by group-specific degree distribution $\{\delta_{AA}, \delta_{AB}, \delta_{BA}, \delta_{BB}\}$ where $\delta_{IG}(k) = \Pr\Big(d_i^G = k \mid i \in I\Big)$ denotes the fraction of receivers of group I with G-degree $d_i^G = k$.

Instead of requiring full knowledge of the network structure, we assume that the sender knows N, A, B, $\{\delta_{AA}, \delta_{AB}, \delta_{BA}, \delta_{BB}\}$, as in a random network setting. On the other hand,

---

[1] All the results of the paper can be extended to the more general case of a directed network in which $(i, j) \in L \not\Rightarrow (j, i) \in L$, for $i, j \in N$, by considering the in-degree distribution instead of the degree distribution.

receivers know $\mathcal{N}_i^A$ and $\mathcal{N}_i^B$. These assumptions allow us to focus on the sender's strategies targeting specific sets of nodes given their group-specific degrees. Notice, nonetheless, that the sender might be aware of the network structure under some specific distributions – up to the identity of receivers that belong to the same group.[2] In addition, the assumption implies that each receiver is aware of his neighbors' group affiliation.

Finally, let $m_i \in M$ denote the individual *message* observed by receiver i. We define a *signal* as a vector of messages as in Kerman and Tenev (2021). We assume that information spreads deterministically given L, so that every receiver in the network observes his neighbors' messages in addition to his own. [3]

Assumption 2 (Information spillovers)**.** *Each* $i \in N$ *observes his own message realization and the realizations of all his neighbors, i.e., he observes signal* $\bar{s}_i := \{m_j\}_{j \in \mathcal{N}_i}$.

Therefore, non-empty networks allow receivers to gather (weakly) more information about the state compared to what they would get from the same communication strategy under the empty network.

Preferences. The sender has state-independent preferences over the actions of the receivers. He wants them to take his *favorite* action $x^* \in X$ regardless of the state of the world. Without loss of generality, we assume the sender's utility from a receiver taking action x to be $v : X \to \mathbb{R}$ with

$$v(x) = \mathbb{1}_{\{x=1\}},$$

so that $x^* = 1$. Accordingly, we say that $\omega = 1$ ($\omega = 0$) is sender's *favorable* (*unfavorable*) state of the world. Let $\bar{x} = (x_i)_{i \in N} \in \{0,1\}^n$ be a profile of actions taken by the receivers. We denote the (normalized) overall utility of the sender from $\bar{x}$ by $V : X^n \to \mathbb{R}$ so that

$$V(\bar{x}) = \frac{1}{n} \sum_{i \in N} v(x_i) = g_A \sum_{i \in A} v(x_i) + g_B \sum_{j \in B} v(x_j),$$

corresponds to the fraction of receivers taking her favorite action. The objective of the sender is to maximize $\mathbb{E}[V(\cdot)]$, i.e., the probability that a receiver plays $x = 1$.

---

[2] For example, consider $\delta_{AB}(2) = \delta_{BA}(2) = 1$ with $|A| = |B| \geq 2$: in this case, the sender knows that receivers are arranged on an $ABAB$... ring.

[3] Egorov and Sonin (2020), instead, assume that information passes through any link connecting two nodes with probability $p \in (0,1)$.

Each receiver wants to match his own action to the actual state. Formally, we represent the utility of agent $i \in G$, with $G \in \{A, B\}$ taking action x in state ω as $u_{i,G} : X \times \Omega \to \mathbb{R}$ with

$$u_{i,G}(x|\omega) = u_G(x|\omega) = t_G \mathbb{1}_{\{x=\omega=0\}} + (1 - t_G)\mathbb{1}_{\{x=\omega=1\}}, \quad 0 < t_G < 1.$$

Notice that t_G represents the posterior belief threshold for members of group G so that receiver $i \in G$ takes action 1 if and only if his posterior belief about $\omega = 1$ is greater than t_G. Throughout the paper, we always assume $\mu_A > t_A$ and $\mu_B < t_B$. This implies that, without information from the sender, the default action is $x = 1$ for A-members, and $x = 0$ for B-members. Therefore, by taking the perspective of the sender, we can refer to the members of group A (B) as *believers* (*sceptics*), as in Innocenti (2021). Finally, we assume that when indifferent, receivers choose the sender's favorite action, $x = 1$.

Communication Strategy. The sender can influence receivers' actions by designing a state-dependent signaling policy that sends informative messages to the receivers about the state of the world. In particular, we allow the sender to commit in advance to a communication strategy that sends private, potentially correlated messages to receivers, conditional on both the realized state and receivers' type.

Assumption 3 (Signaling Policy)**.** [4] *Sender's signaling policy consists of four mappings, two for each group of receivers,* (π_G, ρ_G), *with* $G = A, B$ *with:*

(i) $\pi_G : \Omega \to \Delta(M)$,

(ii) $\rho_G : \Omega \to [0, 1]$ *where* $\rho_G(\omega) = Corr(m_i, m_j|\omega) \geq 0$, *for every* $i, j \in G$.

In words, the sender not only commits to privately send messages $s \in \{0, 1\}$ in state $\omega \in \{0, 1\}$ with probability $\pi_A(s|\omega)$ ($\pi_B(s|\omega)$) to receivers of the A-type (B-type), but also (possibly) imposes positive within-group message correlation equal to $\rho_A(\omega)$ ($\rho_B(\omega)$), depending on the actual state. Since $\Omega = M$ are binary, a signaling policy needs to specify four parameters for each group G:

$$\begin{aligned} p_1^G &:= \pi_G(1|1) \qquad & \rho_1^G &:= \rho^G(1) \\ p_0^G &:= \pi_G(1|0) \qquad & \rho_0^G &:= \rho^G(0) \end{aligned}$$

---

[4] Note: We are currently reassessing this assumption in order to understand provide results for a wider and more general set of signaling policies.

where the subscript indicates the corresponding state.

Remark 1. *The information content of signal* $\bar{s}_i$ *for receiver* i *is equivalent to that of the pair* $s_i := (s_i^A, s_i^B)$ *where* $s_i^A = \sum_{j \in \mathcal{N}_i^A} m_j$ *and* $s_i^B = \sum_{j \in \mathcal{N}_i^B} m_j$.

We assumed the set of messages M to be binary, and the communication strategies to be group-specific. Therefore, individual messages received from neighbors in one group are ex-ante identical for i. Thus, the information i receives from his neighbours in a group (say G) corresponds to the number of 1's he observes (s_i^G) out of the total number of messages he receives from his neighbors in that group (d_i^G). Hence, with an abuse of language, we will refer to $s_i := (s_i^A, s_i^B)$ as a signal. Notice that s_i^G, does not follow a standard Bernoulli distribution since messages sent to group G are (possibly) not independent.

Belief update. Receivers use Bayesian updating to form their posterior beliefs about the state. Let $\bar{\pi}_i$ denote the conditional joint distribution of signals that are receivable by receiver i which is induced by the sender's signaling policy.For any signal s_i, let $\bar{\pi}(s_i|\omega)$ be the probability that receiver $i \in G$ observes s_i when the state is ω. Then the posterior belief about state $\omega = 1$, denoted by η, will be

$$\eta_{i,G}(s_i) = \frac{\mu_G \bar{\pi}_i(s_i|1)}{\mu_G \bar{\pi}_i(s_i|1) + (1-\mu_G)\bar{\pi}_i(s_i|0)}.$$

Finally, we define a key element of our analysis.

Definition 1 (Persuasion level)**.** [5] *The persuasion level of receiver* i *of type* G *is the value* α_G *defined by*

$$\alpha_G := \frac{\mu_G}{1-\mu_G} \frac{u_G(1|1) - u_G(0|1)}{u_G(0|0) - u_G(1|0)} = \frac{\mu_G(1-t_G)}{(1-\mu_G)t_G}.$$

It is straight forward to show that receiver $i \in G$ takes action $x = 1$ after observing s_i if and only $\eta_{i,G}(s_i) \geq t_G$, or equivalently,

$$\bar{\pi}_i(s_i|0) \leq \alpha_G \bar{\pi}_i(s_i|1). \tag{2.1}$$

Notice that, given the above assumption on t_A and t_B, we have

$$\alpha_B < 1 < \alpha_A. \tag{2.2}$$

---

[5] Arieli and Babichenko (2019) provide an alternative definition that bounds α upward by 1.

For group B, the persuasion level α_B can be interpreted as the maximum probability with which the sender may send signal s_i in the unfavorable state $\omega = 0$ so that a B-receiver optimally takes her preferred action upon observing s_i, assuming that the sender always sends s_i in the favorable state $\omega = 1$. Moreover, the closer the prior belief μ_B of the sceptics is to their decision threshold t_B, the higher the persuasion level α_B—that is, the more the sender can "bluff" a B-receiver (Arieli and Babichenko, 2019).

Timing and solution concept. The game unfolds as follows:

1. The sender commits to a signaling policy $(\pi_G, \rho_G)_{G \in \{A,B\}}$.
2. Uncertainty over the state is resolved and ω realizes. Individual messages (m_i) are sent to receivers.
3. Each receiver $(i \in G)$ observes sender's signaling policy $(\pi_G, \rho_G)_{G \in \{A,B\}}$, and his neighbors' signal realizations s_i. Finally, he updates his belief and chooses action $x_i = 1$ if and only if Equation 2.1 holds.

The game can be solved by backward induction and the analysis considers Subgame Perfect Equilibria. Notice that, condition 2.1 already characterized the optimal behaviour of the receivers. We need to specify the optimal signaling policy of the sender, given the optimal choice rule of the receiver.

2.3 Benchmarks and Instrumental Results

This section provides previous results from the literature that are instrumental for the analysis.

2.3.1 Distribution of correlated messages

In order to characterize the joint distribution of signals, it is necessary to specify the probability distribution of s_i^A and s_i^B, i.e. the number of messages equal to 1, or *sucesses*, in i's A-neighborhood and B-neighborhood, respectively. The signaling policy of the sender allows for positive correlation between individual messages, which implies that s_i^A s_i^B, does not follow a standard Bernoulli distribution.

Liporace (2021) specifies the probability for k successes to occur among n simultaneous correlated Bernoulli trials. Take a random variable x defined as the number of successes out of n dependent simultaneous trials. If trials have probability of success $p \in (0, 1)$ and are correlated with correlation ρ, then x has probability mass function:[6]

$$\mathbb{P}(x = k|p, n, \rho) = \binom{n}{k}(1-\rho)p^k(1-p)^{n-k} + \rho\left[(1-p)\mathbb{1}_{\{k=0\}} + p\mathbb{1}_{\{k=n\}}\right]. \tag{2.3}$$

If $\rho = 0$, the trials are independent and x follows a standard Binomial(n, p) distribution. If $\rho = 1$, the trials are perfectly correlated so that either all succeed or all fail; in this case $x = ny$, with $y \sim$ Bernoulli(p), i.e. $\mathbb{P}(x = n) = p$ and $\mathbb{P}(x = 0) = 1 - p$.

Consider now the signal policy of the sender. It specifies the probability of sending message 1 and the value of simultaneous correlation among the messages for each group given the realized state. Hence, the probability distribution of $S_G := \sum_{j \in G} m_j$, the total number of messages equal to 1 sent to the members of group G, in the favorable state and in the unfavorable state, respectively equal

$$\mathbb{P}(S_G = k|p_1^G, n_G, \rho_1^G) = \begin{cases} \rho_G^1(1-p_1^G) + (1-\rho_1^G)(1-p_1^G)^{n_G} & \text{if } k = 0 \\ \binom{n_G}{k}(1-\rho_1^G)(p_1^G)^k(1-\rho_1^G)^{n_g-k} & \text{if } k \in \{1, 2, ..., n_G - 1\} \\ \rho_1^G p_1^G + (1-\rho_1^G)(p_1^G)^{n_G} & \text{if } k = n_G \end{cases} \tag{2.4}$$

$$\mathbb{P}(S_G = k|p_0^G, n_G, \rho_0^G) = \begin{cases} \rho_0^G(1-p_0^G) + (1-\rho_0^G)(1-p_0^G)^{n_G} & \text{if } k = 0 \\ \binom{n_G}{k}(1-\rho_0^G)(p_0^G)^k(1-\rho_0^G)^{n_g-k} & \text{if } k \in \{1, 2, ..., n_G - 1\} \\ \rho_0^G p_0^G + (1-\rho_0^G)(p_0^G)^{n_G} & \text{if } k = n_G \end{cases} \tag{2.5}$$

We will see in the next subsection that it is always (weakly) optimal for the sender to commit to a signaling policy that sends informative messages only to B-type receivers, i.e., those that need to be persuaded. This implies that the information content of a signal $s_i = (s_i^A, s_i^B)$ is fully captured by s_i^B. Thus, the probability distributions specified in (2.4) and (2.5), substituting for B, fully characterize the joint distribution of messages sent by

---

[6] Clearly, if $p \in \{0, 1\}$, then the trials are deterministic and therefore perfectly correlated, implying $\rho = 1$.

the sender in each state.

Finally, the distribution of signals s_i^B observed by i with B-degree equal to d_i^B, follows (2.4) and (2.5) in the respective state, substituting n_G with d_i^B, i.e, the total number of B-messages that i observes.

2.3.2 Unconnected case: Hard and Soft News Strategies

Let us first consider the optimal strategy of the sender in a world with no network, or equivalently, when the network is empty ($L = \varnothing$). The sender can either commit to sending a single public message to all receivers or send a private message to each receiver individually.

In the public-information case, the sender's strategy is characterized by two parameters: the probability of sending the public signal $m = 1$ in the favorable state, p_1, and the probability of sending $m = 1$ in the unfavorable state, p_0. Innocenti (2021) studies this framework with heterogeneous receivers, and proposes two types of strategies for the sender, *hard news* and *soft news* strategies. A public signal implies a trade-off for the sender when designing her strategy. She wants to (i) persuade agents in B, the sceptics, as often as possible and (ii) avoid dissuading agents in A, the believers. Increasing informativeness of the signal in the favorable state of the world favors (i) at the cost of making (ii) less likely.

Innocenti (2021) studies this framework with heterogeneous receivers, and proposes two types of strategies for the sender, *hard news* and *soft news* strategies. The hard news strategy is designed to convince both groups with probability one in the favorable state of the world, at the cost of having a positive probability of dissuading believers in the unfavorable state. The sender sets

$$p_1 = 1 \qquad p_0 = \alpha_B \tag{2.6}$$

so that all receivers take action 1 upon seeing message 1 and take action 0 otherwise. Given the prior of the sender over state 1, μ_S, her expected probability that a receiver takes action 1 – equivalently, the *value* – of such strategy is

$$V_{hn} = \mu_S + (1 - \mu_S)\alpha_B.$$

Hence, the strategy corresponds to the optimal strategy in Kamenica and Gentzkow (2011), the standard Bayesian persuasion framework in which the sender wants to persuade an agent with prior μ_B.

On the other hand, the objective of a soft news strategy is to never dissuade believers, with the downside of reducing the probability that sceptics take action 1 in the favorable state of the world. The optimal strategy, in this case, prescribes to specify

$$p_1 = \frac{\alpha_A - \alpha_B}{\alpha_A - 1} \qquad p_0 = \frac{\alpha_B(\alpha_A - \alpha_B)}{\alpha_A - 1}. \tag{2.7}$$

so that the posterior of an A-receiver equals t_A in the unfavorable state, while the posterior of a B-receiver equals t_B in the favorable state. Now, believers are persuaded when they observe either message, but message 1 is sent less often in the favorable state compared to the hard news strategy, implying that sceptics are persuaded with a lower probability. The value of the soft news strategy is then

$$V_{sn} = g_A + g_B \left[\mu_S \frac{\alpha_A - \alpha_B}{\alpha_A - 1} + (1 - \mu_S) \frac{\alpha_B(\alpha_A - \alpha_B)}{\alpha_A - 1} \right].$$

Clearly, the sender's optimal strategy depends on the parameters g_A, g_B, α_A, α_B, and μ_S. Intuitively, a soft news strategy is preferable when (i) the group A is relatively larger than group B, (ii) the sender assigns a sufficiently high probability to the unfavorable state of the world, (iii) polarization between the two groups—defined as $\alpha_A - \alpha_B$—is high enough (Innocenti, 2021).

Consider now the case in which the sender is allowed to send a private message to each receiver, i.e., she needs to specify the vector $(p_1^A, p_0^A, p_1^B, p_0^B)$. Then, it is optimal for her to target each group differently. Believers do not need to be persuaded, hence she will send an uninformative signal to the members of A. Instead, she will send the optimal persuading signal to members of B:

$$p_1^A = p_0^A \qquad p_1^B = 1 \qquad p_0^B = \alpha_B.$$

Private individual messages allow the sender to optimally target each group. Therefore, a hard news strategy performs (weakly) better than a soft news strategy since believers cannot access to possibly dissuading information and always take the sender's favorite action. and, hence, the private information strategy determines the upward threshold for the value of the sender:

$$V_{pr} = g_A + g_B \left[\mu_S + (1 - \mu_S)\alpha_B \right].$$

The value for the sender of the different strategies provided above represents natural benchmarks to compare how well a sender can do when receivers share information through a network. In particular, notice that both hard and soft news strategies can be replicated in a network by setting perfect correlations among signals ($\rho_0^A = \rho_0^B = \rho_1^A = \rho_1^B$), and the same probability for all receivers, according to the strategy-specific values specified above.[7] Therefore, the optimal value for the sender in a network is bounded below by $V_{pu} := \max\{V_{hn}, V_{sn}\}$ and above by V_{pr}.

In the next subsection, we turn to the network setting. First, we recall Liporace's (2021) extension of the *hard news* strategy to communication networks. We then introduce and analyze the corresponding notion of a *soft news* strategy in the same environment.

2.3.3 Connected case: Hard News Strategies

Consider now the problem of the sender when she can commit to sending private messages to receivers, but, at the same time, she faces information spillovers between neighbors. This corresponds to an intermediate case between the private and the public information settings.

In this subsection, we present some of the results and the definitions in Liporace (2021) that help us in the analysis of potential soft news strategies in a network. One important result, is that any optimal signaling strategy π must send uninformative messages to believers, i.e., it prescribes $p_0^A = p_1^A$, and $\rho_0^G = \rho_1^G$. The reason is that believers do not need to be persuaded, so any information they receive is weakly detrimental, as they would already take the sender's favorite action in the absence of any information. Moreover, relying on believers to inform sceptics is unnecessary: the sender can target

---

[7] The public values of information are exactly achievable whenever $\delta_{AB}(0) = 0$, i.e., each believer A shares at least one link with a sceptics B. Instead, if $\delta_{AB}(0) > 0$ then the value for the hard news strategy in networks equals $[1 - \delta_{AB}(0)]V_hn + \delta_{AB}(0)$, since believers whose signal does not include at least one message from a B member can never be dissuaded under the optimal strategy. We will explain this formally in the next subsection.

sceptics directly. Using believers as intermediaries would only expose them to potential dissuasion, without improving the probability of persuading sceptics beyond what direct targeting achieves.

Assuming uninformative messages for the believers implies that the signal observed by receiver i, $s_i := (s_i^A, s_i^B)$ is fully defined, in the sense of the information that it carries, by s_i^B. Hence, from now on we will refer to the signal received by i with B-degree d_i^B simply as s_i^B. Moreover, we will re-denote some variables to save on notation, since the parameters of the signaling strategy for A will no longer be considered. Hence, we write $d_i := d_i^B$, $p_1 := p_1^B$, $p_0 := p_1^B$, $\rho_1 := \rho_1^B$ and $\rho_0 := \rho_0^B$.

Another important insight is that the network structure implies an additional level of heterogeneity among receivers compared to the unconnected case. Now, receivers differ not only in their prior but also in their degree (more relevantly, in their B-degree). Therefore, the strategies of the sender needs to *target* specific sets of nodes that are defined by their prior and their B-degree in order to make them just indifferent given some specific signal. In the absence of a network, all nodes are targeted when the sender adopts a soft news strategy, since the strategy implies that each public message is persuading, i.e. making indifferent about the two actions, one of the two groups: the believers in the unfavorable state, and the sceptics in the favorable one. In a connected world, the targets of a strategy need to be formally specified.

Following Liporace (2021), we define a *persuading signal* for receiver i to be a signal s_i such that $\eta_{i,G}(s_i) = t_G$. Similarly, a *favorable signal* for i, is a signal s_i such that $\eta_{i,G}(s_i) > t_G$, while a *dissuading signal* for $i \in G$ is a signal for which $\eta_{i,G}(s_i) < t_G$. The *targets* of strategy π are the subset of receivers whose set of persuading signals under π is nonempty—i.e., for whom there exists at least one signal that renders them indifferent between the two actions. Moreover, denote the set of dissuading signals for an $i \in N$, given a signaling policy π, as $DS_{i,\pi}$ and the set of persuading signals for any $i \in N$ as $PS_{i,\pi}$. Finally, let S_i be the set of possible signal realizations of $i \in N$.

Given the definitions above, Liporace (2021) defines a hard news strategies in network as a signaling policy π such that $\mathbb{P}(s_i \in PS_{i,\pi} | \omega = 1) = 1$ for any targeted receiver i, and proves that there are only three possible strategies that satisfy the definition: the *unconnected*, the *multiple-message*, and the *network-specific* hard news strategy. The unconnected strategy is the network replication of the hard news strategy of the unconnected case. Correlation is set to 1 in both states, so that all B nodes can be targeted regardless of their B-degree, $p_1 = 1$ and $p_0 = \alpha_B$. All the sceptics are targeted by this strategy because the informativeness of a single-message signal equals that of a multi-message signal, since full correlation implies that all messages directed to B nodes are identical. The $(\bar{d})$-multiple-message strategy targets nodes with a specific B-degree, $\bar{d}$, which are persuaded upon observing signal $s_i = d_i$, i.e. all 1 messages in their B-neighborhood. Finally, the $(\bar{d})$-network-specific strategy is based on differences in the correlation parameters such that $\rho_1 = 1$. It targets B nodes with $d_i = \bar{d}$, which are persuaded upon seeing $s_i \ \in \{0, \bar{d}\}$.

Importantly, under both the multiple-message and the network-specific strategies, any sceptic $j \in B$ with $d_j < \bar{d}$ never receives a persuading message, whereas any sceptic $l \in B$ with $d_l \geq \bar{d}$ receives a favorable signal with probability that decreases in d_l; differently, A receivers are always susceptible of receiving a favorable message, regardless of their degree in both strategies. The unconnected strategy, instead, *levels* the degree heterogeneity among receivers, since it implies an identical message to all receivers so that the number of messages received is not informative.

We know move to the definition and

2.4 Soft News Strategies in a Network

In this section, we generalize the definition of a *soft news* strategy proposed in Innocenti (2021) to the network case. The main idea of soft news strategy in the unconnected case is that they ensure that believers A are never dissuaded. However, an optimal strategy in a network, should always aim at maximizing the expected probability that a receiver takes the favorite action regardless of his type. Hence, defining a soft news strategy as a signaling mechanism that never dissuades receivers in A is not sufficient.

Remark 2. *It is always possible for the sender to design a strategy ensuring that every $i \in A$ takes her favorite action since she can always design an uninformative signaling policy for both groups.*

If no information is introduced in the network, then no sceptic can be persuaded, but nor will any believer be dissuaded. However, we are interested in analyzing strategies that not only avoid dissuading believers but also persuade sceptics.

Definition 2 (Soft News Strategy)**.** *A soft news strategy in a network is a signaling policy* π *such that for any targeted node* $i \in A$*, and for any targeted node* $j \in B$ *satisfies:*

(i) $\mathbb{P}(s_i \in PS_{i,\pi}|\omega = 1) > 0$*, and* $\mathbb{P}(s_j \in PS_{j,\pi}|\omega = 1) > 0$*,*

(ii) $\mathbb{P}(s_i \in PS_{i,\pi}|\omega = 1) + \mathbb{P}(s_j \in PS_{j,\pi}|\omega = 1) = 1$*.*

The definition of soft news strategy we proposed requires that both sceptics and believers are targeted so that in the favorable state, they all have a strictly positive probability of observing a persuading signal. However, this is not enough, as the definition requires that in the good state of the world, at least one targeted node observes a persuading signal. In fact, given the differences in priors, it is not possible for a signal s_i to be persuading for both a believer and a sceptic. Therefore, condition (ii) is effectively a mutually exclusive condition: at least one and only one targeted set observes a persuading message in the favorable state.

Notice that the definition does not explicitly require for that targeted $i \in A$, $\mathbb{P}(s_i \in DS_{i,\pi}) = 0$ as we would expect by the definition of a soft news strategy. However, it does so implicitly by requiring that every possible signal in the favorable state must be persuading for at least one target. Since $\alpha_A > \alpha_B$, any signal s_i that is persuading for a sceptic is strictly favorable for a believer. Therefore, targeted A members cannot be dissuaded when the sender implements a soft news strategy.

Finally, given the joint distribution of signals specified in (2.4) the definition implies that targets A and B have the same B-degree $\bar{d}$, otherwise $\mathbb{P}(s_i \in PS_{i,\pi}|\omega = 1) + \mathbb{P}(s_j \in PS_{j,\pi}|\omega = 1) \neq 1$.

Now let us consider the soft news (SN) strategy *counterparts* of the hard news (HN) strategies specified in Liporace (2021).[8]

Definition 3. *An unconnected-SN strategy* π*, is a signaling policy* π *that targets all* $i \in N$ *regardless of their* B*-degree. For any receiver* $j \in B$*,* $PS_j = \{d_i\}$ *and for any receiver* $i \in A$*,* $PS_i = \{0\}$*.*

Definition 4. *A network-specific-SN strategy* π*, is a signaling policy* π *in which* $PS_{j,\pi} =$

---

[8] We will not consider the symmetric counterparts of this strategies, but it should keep in mind that the content of the messages is irrelevant so that all the strategies proposed could be implemented with inverted realizations.

$\{0, d_j\}$ for any targeted node $j \in B$, and $PS_{i,\pi} = S_i \setminus \{0, d_i\}$ for any targeted node $i \in A$ with $d_i = d_j = \bar{d}$.

Definition 5. *A multi-message-SN strategy π, is a signaling policy π in which $PS_{j,\pi} = \{d_j\}$ for any targeted node $j \in B$, and $PS_{i,\pi} = S_i \setminus \{d_i\}$ for any targeted node $i \in a$ with $d_i = d_j = \bar{d}$.*

The SN counterparts of the network-specific multi-message strategies defined above specify the set of persuading signals for the A targets residually, out of the possible signals observed by a receiver with degree $\bar{d}$. In fact, in the unconnected-NS case, will we see that for any $i \in N$, $S_i = \{0, d_i\}$. Notice that this is consistent with the definition of a soft news strategy since the union of the two persuading sets is the set of all signals observable by i, i.e., S_i. However, with Proposition 1, we characterize the unconnected-NS and the network-specific-SN, while we show that a multi-message-SN strategy cannot be implemented.

Proposition 1. *The three SN strategies defined above are characterized as follows:*

(i) An unconnected-SN strategy π is a signaling policy for which $\rho_1 = \rho_0 = 1$, $p_1 = \frac{\alpha_A - 1}{\alpha_A - \alpha_B}$ and $p_0 = \frac{\alpha_B(\alpha_A - 1)}{\alpha_A - \alpha_B}$, with associated value for the sender

$$V_u^{SN} = g_A + g_B \left[\mu_S \frac{\alpha_A - 1}{\alpha_A - \alpha_B} + (1 - \mu_S) \frac{\alpha_B(\alpha_A - 1)}{\alpha_A - \alpha_B} \right]$$

(ii) A network-specific-SN strategy π is a signaling policy for which $\rho_0 = 1 - \alpha_A(1 - \rho_1)$, $p_1 = p_0 = 0.5$, with associated value

$$\begin{aligned} V_{ns}^{SN} = & g_A + g_B \mu_S \sum_{d_i = \bar{d}}^{n_B} \delta_{BB}(d_i) \left[\rho_1 + 2^{1-d_i}(1 - \rho_1) \right] + \\ & + g_B(1 - \mu_S) \sum_{d_i = \bar{d}}^{n_B} \delta_{BB}(d_i) \left[\rho_0 + 2^{1-d_i}(1 - \rho_0) \right] \end{aligned}$$

(iii) The multi-message-NS strategy is not implementable.

Proof. By definition, a persuading signal s_i for i is such that $\eta_{i,g}(s_i) = t_G$, or equivalently,

$$\mathbb{P}(s_i | 0) = \alpha_G \mathbb{P}(s_i | 1) \tag{2.8}$$

which corresponds to condition (2.1) in our current setting, where signals are defined as the sum of messages received by B-neighbors.

(i) A strategy targeting all nodes in the network requires that messages are perfectly correlated in each state of the world, hence, $\rho_1 = \rho_0 = 1$, so to level any degree differences. Then, from the joint distribution of signals, we know that for any i

$$\mathbb{P}(0|0) = 1 - p_0 \qquad \mathbb{P}(0|1) = 1 - p_1$$
$$\mathbb{P}(d_i|0) = p_0 \qquad \mathbb{P}(d_i|1) = p_1$$

For each node $i \in B$, $PS_{i,\pi} = \{d_i\}$, hence $\mathbb{P}(d_i|0) = \alpha_B \mathbb{P}(d_i|1)$, or $p_0 = \alpha_B p_1$. For each node $i \in A$, $PS_{i,\pi} = \{0\}$, hence $\mathbb{P}(0|0) \leq \alpha_A \mathbb{P}(0|1)$, or $(1 - p_0) = \alpha_A(1 - p_1)$. The two conditions imply

$$p_1 = \frac{\alpha_A - 1}{\alpha_A - \alpha_B} \qquad p_0 = \frac{\alpha_B(\alpha_A - 1)}{\alpha_A - \alpha_B}.$$

Therefore, believers are persuaded when observing $s_i = 0$ and are not dissuaded when observing $s_i = d_i$ since $\mathbb{P}(d_i|0) \leq \alpha_B \mathbb{P}(d_i|1)$ implies $\mathbb{P}(d_i|0) < \alpha_A \mathbb{P}(d_i|1)$ as $\alpha_A > \alpha_B$. Hence, A's are never dissuaded. Sceptics are just persuaded when observing $s_i = d_i$, which happens with probability $\mu_S p_1 + (1 - \mu_S)p_0$. The expected probability for the sender is then

$$V_u^{SN} = g_A + g_B \left[\mu_S \frac{\alpha_A - 1}{\alpha_A - \alpha_B} + (1 - \mu_S)\frac{\alpha_B(\alpha_A - 1)}{\alpha_A - \alpha_B}\right]$$

.

(ii) In a network-specific-NS strategy targeting nodes $i \in N$ with $d_i = \bar{d}$, the persuading signals of sceptics are $\{0, d_i\}$, while the persuading signals of the believers are $\{1, 2, ..., \bar{d}\text{-}1\}$. This implies that ρ_1, ρ_0 are lower than one and $\mathbb{P}(s_i|\omega) > 0$ for every s_i and every ω.

Let us assume, from now on, that both $i \in A$ and $j \in B$ are targets, i.e. $d_i = d_j = \bar{d}$. We know that $PS_{i,\pi} = \{1, 2, ..., \bar{d}\text{-}1\}$, and $PS_{j,\pi} = \{0, \bar{d}\}$ hence we have two *aggregated* conditions that need to be satisfied:

$$\mathbb{P}(0|0)+\mathbb{P}(\bar{d}|0)=\alpha_B[\mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)]$$
$$\sum_{d_i=1}^{\bar{d}-1}\mathbb{P}(d_i|0)=\alpha_A\sum_{d_i=1}^{\bar{d}-1}\mathbb{P}(d_i|0)\quad\Leftrightarrow\quad 1-\mathbb{P}(0|0)+\mathbb{P}(\bar{d}|0)=\alpha_A[1-\mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)]$$

By substituting the first in the second we have,

$$\begin{aligned}1-\alpha_B[\mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)]&=\alpha_A[1-\mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)]\\ \mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)&=\frac{\alpha_A-1}{\alpha_A-\alpha_B}\\ \rho_1+(1-\rho_1)(p_1^{\bar{d}}+(1-p_1)^{\bar{d}})&=\frac{\alpha_A-1}{\alpha_A-\alpha_B}\end{aligned}$$

since $\mathbb{P}(0|1)+\mathbb{P}(\bar{d}|1)=\rho_1+(1-\rho_1)(p_1^{\bar{d}}+(1-p_1)^{\bar{d}})$. Then, from the first aggregated condition, we get

$$\begin{aligned}\mathbb{P}(0|0)+\mathbb{P}(\bar{d}|0)&=\alpha_B\frac{\alpha_B(\alpha_A-1)}{\alpha_A-\alpha_B}\\ \rho_0+(1-\rho_0)(p_0^{\bar{d}}+(1-p_0)^{\bar{d}})&=\alpha_B\frac{\alpha_B(\alpha_A-1)}{\alpha_A-\alpha_B}\end{aligned}$$

since $\mathbb{P}(0|0)+\mathbb{P}(\bar{d}|0)=\rho_0+(1-\rho_0)(p_0^{\bar{d}}+(1-p_0)^{\bar{d}})$. These conditions are substantially identical to the one of a standard soft news strategy with the probability of a persuading signal for the B's equal to $\frac{\alpha_A-1}{\alpha_A-\alpha_B}$ in the favorable state and $\frac{\alpha_B(\alpha_A-1)}{\alpha_A-\alpha_B}$ in the unfavorable one.

Now, in addition to the aggregated conditions, we need to consider the specific conditions for each signal since the individual realization must be persuading the corresponding set of targets. In particular, since $PS_{i,\pi}=\{1,2,...,\bar{d}\text{-}1\}$ and $PS_{j,\pi}=\{0,\bar{d}\}$ we have the following $\bar{d}$ conditions

$$\begin{aligned}&\mathbb{P}(0|0)=\alpha_B\mathbb{P}(0|1)\\ &\mathbb{P}(x|0)=\alpha_A\mathbb{P}(x|1)\qquad\text{for } x\in\{1,2,...,\bar{d}-1\}\\ &\mathbb{P}(\bar{d}|0)=\alpha_B\mathbb{P}(\bar{d}|1)\end{aligned}$$

which correspond to

$$\rho_0(1-p_0)+(1-\rho_0)(1-p_0)^{\bar{d}}=\alpha_B\left[\rho_1(1-p_1)+(1-\rho_1)(1-p_1)^{\bar{d}}\right] \tag{2.9}$$

$$\binom{\bar{d}}{x}(1-\rho_0)p_0^x(1-p_0)^{\bar{d}-x}=\alpha_A\binom{\bar{d}}{x}(1-\rho_1)p_1^x(1-p_1)^{\bar{d}-x} \qquad \text{for } x\in\{1,2,...,\bar{d}-1\} \tag{2.10}$$

$$\rho_0p_0+(1-\rho_0)p_0^{\bar{d}}=\alpha_B\left[\rho_1p_1+(1-\rho_1)p_1^{\bar{d}}\right] \tag{2.11}$$

Now for any $x\in\{1,2,...,\bar{d}-1\}$. Then it can be written

$$\frac{\binom{\bar{d}}{x}(1-\rho_0)p_0^x(1-p_0)^{\bar{d}-x}}{\binom{\bar{d}}{x}(1-\rho_1)p_1^x(1-p_1)^{\bar{d}-x}}=\frac{(1-\rho_0)p_0^x(1-p_0)^{\bar{d}-x}}{(1-\rho_1)p_1^x(1-p_1)^{\bar{d}-x}}=\alpha_A$$

However, this condition must hold for any $x\in\{1,2,...,\bar{d}-1\}$. Take any $y>x$, such that $x,y\in\{1,2,...,\bar{d}-1\}$. Then it must hold

$$\begin{aligned}\frac{(1-\rho_0)p_0^x(1-p_0)^{\bar{d}-x}}{(1-\rho_1)p_1^x(1-p_1)^{\bar{d}-x}}&=\frac{(1-\rho_0)p_0^y(1-p_0)^{\bar{d}-y}}{(1-\rho_1)p_1^y(1-p_1)^{\bar{d}-y}}\\ \frac{p_0^x(1-p_0)^{\bar{d}-x}}{p_1^x(1-p_1)^{\bar{d}-x}}&=\frac{p_0^y(1-p_0)^{\bar{d}-y}}{p_1^y(1-p_1)^{\bar{d}-y}}\\ \frac{(1-p_0)^{y-x}}{(1-p_1)^{y-x}}&=\frac{p_0^{y-x}}{p_1^{y-x}}\\ \frac{1-p_0}{1-p_1}&=\frac{p_0}{p_1}\end{aligned}$$

which implies $p_0=p_1$. Therefore, from any of the conditions (2.10) we get

$$(1-\rho_0)=\alpha_A(1-\rho_1)$$

which implies $\rho_0=1-\alpha_A(1-\rho_1)$. Now, by substituting the expression for ρ and $p_1=p_0$

into (11) we get[9]

$$\begin{aligned}
[1-\alpha_A(1-\rho_1)]p_1+\alpha_A(1-\rho_1)p_1^{\bar{d}} &= \alpha_B\left[\rho_1 p_1+(1-\rho_1)p_1^{\bar{d}}\right]\\
[1-\alpha_A(1-\rho_1)]+\alpha_A(1-\rho_1)p_1^{\bar{d}-1} &= \alpha_B\left[\rho_1+(1-\rho_1)p_1^{\bar{d}-1}\right]\\
(\alpha_A-\alpha_B)(1-\rho_1)p_1^{\bar{d}-1}+(\alpha_A-\alpha_B)\rho_1 &= \alpha_A-1\\
(\alpha_A-\alpha_B)[\rho_1+(1-\rho_1)p_1^{\bar{d}-1}] &= \alpha_A-1\\
\rho_1+(1-\rho_1)p_1^{\bar{d}-1} &= \frac{\alpha_A-1}{\alpha_A-\alpha_B}
\end{aligned}$$

Since we know that

$$\rho_1+(1-\rho_1)(p_1^{\bar{d}}+(1-p_1)^{\bar{d}})=\frac{\alpha_A-1}{\alpha_A-\alpha_B},$$

we get

$$\begin{aligned}
\rho_1+(1-\rho_1)(p_1^{\bar{d}}+(1-p_1)^{\bar{d}}) &= \rho_1+(1-\rho_1)p_1^{\bar{d}-1}\\
(1-\rho_1)(p_1^{\bar{d}}+(1-p_1)^{\bar{d}}) &= (1-\rho_1)p_1^{\bar{d}-1}\\
p_1^{\bar{d}}+(1-p_1^{\bar{d}}) &= p_1^{\bar{d}-1}\\
p_1^{\bar{d}}+(1-p_1^{\bar{d}})-p_1^{\bar{d}-1} &= 0
\end{aligned}$$

where the condition is satisfied if $p_1 \in \{0.5, 1\}$. Since we are assuming positive probabilities for every signal in the favorable state, we have that $p_1 = p_0 = 0.5$.

Now, $p_1 = p_0$ implies that the conditions $\mathbb{P}(x|0) = \alpha_A \mathbb{P}(x|1)$ for $x \in \{1, 2, ...\bar{d}-1\}$ simply reduce to the condition defining $\rho_0 = 1-\alpha_A(1-\rho_1)$. This means that whatever the B-degree of node $i \in A$, he will always be persuaded upon seeing $s_i \in \{1, 2, ...\bar{d}-1\}$. On the other hand, every node $j \in B$ with degree $d_i < \bar{d}$ will never be persuaded upon seeing either $s_i \in \{0, d_i\}$ since conditions (9) and (11) are satisfied with equality only if $d_i = \bar{d}$ and with strict inequality if $d_i > \bar{d}$ since p_1 and p_0 are lower than one. Moreover, notice that the same conditions are always satisfied with strict inequality when substituting α_B with α_A, which implies that no A nodes are dissuaded when they observe $s_i \in \{0, d_i\}$.

---

[9] Notice that substituting into (9) would have delivered the same result.

Therefore, the value for the sender for the network-specific-NS strategy is equal to

$$\begin{aligned} V_{ns}^{SN} = & g_A + g_B\mu_S \sum_{d_i=\bar{d}}^{n_B} \delta_{BB}(d_i)\left[\rho_1 + (1-\rho_1)(p_1^{d_i} + (1-p_1)^{d_i})\right] + \\ & + g_B(1-\mu_S)\sum_{d_i=\bar{d}}^{n_B} \delta_{BB}(d_i)\left[\rho_0 + (1-\rho_0)(p_1^{d_i} + (1-p_1)^{d_i})\right] \end{aligned}$$

(iii) Since we assumed $PS_i = \{0, 1, 2, ..., \bar{d} - 1\}$ for targeted $i \in A$, the signal-specific conditions derived for the NS-strategy would apply here as well. This means that $p_0 = p_1$ and $\rho_0 = 1 - \alpha_A(1 - \rho_1)$. However, substituting for ρ_1 in the condition for $s_i = 0$, $\rho_0(1-p_0)+(1-\rho_0)(1-p_0)^{\bar{d}} = \alpha_B\left[\rho_1(1-p_1) + (1-\rho_1)(1-p_1)^{\bar{d}}\right]$, shows that the condition is never satisfied with equality unless $p_1 = 1$. This would imply that $\mathbb{P}(s_i|1) = 0$ for any $s_i \in \{1, 2, ...\bar{d}_i - 1\}$ which is impossible since, by definition of soft-news strategy, any $s_i \in S_i$ should persuade one of the targeted nodes in the favorable state. □

Notice that for both strategies, the same trade-offs underlined in the analysis of soft news strategies in an unconnected world still apply in the network case. In fact, the expected probability of persuading is increasing in the size of the group of believers, g_A, and the probability of the favorable state μ_S. Instead, it is decreasing in the polarization of the two groups $\alpha_A - \alpha_B$.

We conclude this section by stating two conjectures that have not been formally proved yet, but that are worth being mentioned given our analysis.

Conjecture 1. *The set of all possible soft-news strategies is exhausted by the unconnected-SN strategy and the network-specific-SN strategy (up to mirroring strategies).*

Conjecture 2. *Regardless of the target A of a soft-news strategies, no A receiver is ever dissuaded from taking the favorite action, regardless of his degree.*

The preliminary analysis conducted seems to validate conjecture 1. In turn, given that we showed that both the unconnected-SN strategy and the network-specific-SN strategy ensure that all nodes are always convinced to take action A, the proof of conjecture 2 would be immediate. We leave the formalization of these statements and their proof for future research.

2.5 Discussion and Future Research

This work provides an initial analysis of soft news strategy in networks and presents a thorough review of previous results that are instrumental for the analysis. Moreover, the definition of soft-new strategy proposed is consistent with Liporace (2021), whose model is taken as the basis for this work. Moreover, the characterization of the set of possible soft news strategies will help future research on the topic by limiting the number of strategies that need to be considered.

Future research will include the following steps. First, it will provide formal for the last conjectures stated. Consistent results would represent a substantial improvement in the analysis of soft news strategies in networks. In addition, it will be worthy to compare soft news strategies in networks to their hard news counterparts, in order to determine which characteristics of the network are better exploited by which strategy. To do that, it will be necessary to determine what are the optimal targets in each strategy for a given network. For now, it is possible to identify optimal targets only considering specific networks, but a general result might be found, in particular considering the *marginal value* added by nodes with different B-degree.

Finally, the model is based on a *rigid* signaling structure that limits the possible distribution of signals to be consistent with the distribution of correlated trials. It will be important to study when and how a *free* signaling structure strictly outperforms the network strategies considered in this paper. I leave this for future research.

2.6 References

Alonso, R., and O. Câmara, “Persuading voters”, American Economic Review, 106(11): 3590–3605, 2016.

Arieli, I., and Y. Babichenko, 2019, “Private Bayesian persuasion”, *Journal of Economic Theory*, 182, 185–217.

Egorov, G., and K. Sonin, 2020, “Persuasion on networks”, Technical report, National Bureau of Economic Research.

Innocenti, F., 2021, “Can media pluralism be harmful to news quality?”, Working Paper, University of Bonn and University of Mannheim.

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

Kerman, T., P. Herings, and Dominik Karos, 2020, “Persuading strategic voters”, Working Paper.

Kerman, T., and Tenev, A. P., 2021, “Persuading communicating voters”, GSBE Research Memoranda No. 003, Maastricht University.

Liporace, M., 2021, “Persuasion in Networks: a Model with Heterogeneous Agents”, Working Paper, Tilburg University.

Chapter 3

Influence of Polls on Elections: Strategic Responding and (Mis)reporting

3.1 Introduction

The influence of opinion polls on election outcomes has gained significant attention in the realm of social science, reflecting its critical implications for democratic processes. Despite extensive research on this topic, the empirical, experimental, and theoretical results remain inconclusive. This paper aims to develop a theoretical model to understand the conditions under which opinion polls influence electoral behavior and outcomes, provide predictions, and test them using data.

The "bandwagon effect" refers to the phenomenon where voters tend to support candidates perceived to be leading in the polls, thereby amplifying a front-runner's advantage (Grillo, 2017; Berent et al., 2015). Conversely, the "underdog effect" describes increased turnout or support for candidates trailing in the polls, motivated by a desire to correct perceived imbalances (Luxen, 2020; Blais et al., 2006). These definitions will guide the theoretical framework developed in this study.

This study stems from two well-established facts regarding pre-electoral polls and elections. First, empirical evidence has highlighted notable discrepancies between pre-electoral poll predictions and actual electoral outcomes. Recent noteworthy examples include the Brexit referendum, the US Presidential elections in 2016 (Kenet et al., 2018), and the Brazilian Presidential elections in 2022. Second, different studies have highlighted the increasing heterogeneity in poll forecasts among different pollsters, especially close to election dates (Shirani-Mehr et al., 2018; Dawson, 2022).

These differences have been attributed to several factors: misrepresentation problems due to declining respondent rates, social-desirability biases where respondents hide their true voting intentions for controversial candidates (Zhou et al., 2021), and the potential influence of polls on voter behavior, manifesting through the bandwagon and underdog effects. This paper focuses on the latter.

Whether pre-electoral polls can affect elections is the subject of more than a hundred empirical and experimental papers (Cancela and Geys, 2016), but the results are still mixed. Therefore, this research seeks to answer several critical questions using a game-theoretic approach:

- Can electoral polls affect voter behavior and, ultimately, election outcomes?

The mechanism we want to highlight is that the information provided by polls asymmetrically influences the turnout probability. Voter turnout decisions are often strategic and influenced by citizens' beliefs about other voters' preferences and turnout (Blais et al., 2006). Polls provide this information, allowing rational voters to update their beliefs. Thus, polls can affect elections by changing these beliefs and, consequently, the incentives of poll participants. Therefore, a second question arises:

- How should utility-maximizing partisan voters respond to pre-electoral polls?

The assumptions regarding the information transmitted through polls, given that respondents can answer untruthfully or strategically, are key in this setting. We aim for our model to address this question without resorting to any bounded-rationality assumptions whenever possible. Clearly, under the assumption that polls can effectively influence electoral outcomes, we would like to extend our analysis by considering pollsters as actual agents in the model. These agents derive their utility from the election results and thus have incentives to induce favorable poll results. For example, Dawson (2022) reports evidence suggesting that polling numbers were manipulated for favorable outcomes by Recep Tayyip Erdogan in 2013, and reportedly by Donald Trump's former lawyer Michael Cohen before the 2016 presidential election. Whether pollsters act for their benefit or under the influence of a political party participating in the election, these considerations lead to a third question:

- Can a pollster manipulate poll results to influence electoral outcomes?

The assumptions here are again key: if pollsters can simply misreport poll answers and voters are fully rational, the message conveyed by publishing the poll is completely uninformative.[1] However, a polling organization might have other ways to convey favorable signals to influence voter decisions, such as partially manipulated poll results and the power to select the sample of voters to whom the poll is submitted, under constraints due to the usual requirements of the sample's representativity compared to the actual population within an accepted margin of error. Moreover, given the above-mentioned evidence of heterogeneity in poll forecasts, another question seems worth exploring:

- What are the consequences of competition among pollsters?

3.1.1 Literature Review

The relationship between pre-electoral polls and electoral outcomes has been extensively studied in empirical, experimental, and theoretical literature. This review focuses on the most recent theoretical studies that model how polls can effectively alter election results, aligning closely with the aim of this study.

Luxen (2020) examines how polls impact the participation decisions of citizens in a two-candidate election with costly voting. The study reveals that if all poll participants answer truthfully, the underdog effect leads supporters of the trailing candidate to turn out at higher rates. However, strategic misrepresentation by poll respondents can negate the informative value of the polls, resulting in an almost certain victory for the majority candidate regardless of voters' posterior beliefs.

Similarly, Taylor and Yildirim (2010) explore the role of public information in electoral bias, finding that public information about the distribution of voters' preferences can

---

[1] Unless deciding to convey an uninformative message represents an informative message itself.

reduce the likelihood of the majority-preferred candidate winning the election. Their results indicate that polls mobilize minority groups, which may lead to less desirable outcomes for the majority.

Strategic behavior in responding to polls can significantly alter the expected outcomes of elections. Goeree and Großer (2007) present a model showing that public polls can be self-defeating, as they stimulate the "wrong" group to participate, leading to welfare-reducing outcomes. They argue that public release of information through polls can increase turnout but reduce overall welfare due to the adverse selection of participants.

Durazzo and Turchick (2023) also discuss the potential welfare improvements from misreported polls, suggesting that strategic misrepresentation in polls can sometimes lead to better election outcomes by correcting biases in voter turnout. This is consistent with the findings of Jo (2023), who demonstrate that pre-election polls can help elections aggregate information more successfully, even if the poll results are not always accurate.

The welfare implications of pre-election polls are mixed. While some studies suggest that polls can reduce welfare by promoting strategic misrepresentation and turnout among less-informed voters, others indicate that polls can enhance welfare by providing valuable information to voters and candidates.

For instance, Grillo (2017) analyzes the risk aversion and bandwagon effects in the pivotal voter model, concluding that polls can sometimes reduce welfare by encouraging strategic voting and turnout among voters who might otherwise abstain. Conversely, Luxen (2020) argues that under certain conditions, truthful reporting in polls can be welfare-enhancing by discouraging minority participation and ensuring that the majority's preferences are accurately reflected in the election outcome.

3.2 Baseline Model and Extensions

This section provides an overview of the baseline model, based on Luxen (2020), who built upon the canonical model of costly voting by Palfrey and Rosenthal (1983), highlighting its limitation for the scope of the paper, and suggesting the intended characteristics and modifications necessary to enrich the analysis. This will be available in the following versions of the paper.

Baseline Model Following Luxen (2020), we will build on a model considering electoral competition involving two candidates, A and B. The model is characterized by independent private values, so that each voter i receives a value $v > 0$ if their preferred candidate wins the election. There are two possible states of the world, $\omega \in \Omega = \alpha, \beta$, that determine the probabilities associated with each candidate winning the election given by $Pr(A|\alpha) = q > \frac{1}{2}$ and $Pr(A|\beta) = 1 - q$.

The number of voters N is finite but uncertain and follows a Poisson distribution with mean n, so that the probability of there being k citizens is given by the following expression:

$$Pr(N = k) = e^{-n}\frac{n^k}{k!}.$$

Each voter i has a strategy set $S_i = a, b, nv$, where a and b represent votes for candidates A and B respectively, and nv represents not voting. The cost of voting, c_i, is independently and identically distributed (i.i.d.) according to a distribution $F[0, \bar{c}]$.

Before the election takes place, an opinion poll is conducted involving m citizens, yielding results $\tau = (\tau_A, \tau_B)$, where τ_i is the vote gathered by candidate $i \in A, B$. This poll is costless to the participants and provides insight into the potential outcome of the election. In the election, the candidate with the majority of votes wins, with ties broken by a fair coin toss.

The game unfolds as follows:

1. Nature determines the number of voters N and the state ω.
2. Voters' preferences for candidates A or B and their voting costs c_i are determined based on the state ω.
3. The pre-electoral poll is conducted and the results are published.
4. Voters cast their votes and the candidate with the most votes wins the election.

The model induces an extended Poisson game (Myerson, 1998), and the solution concept used in this model is symmetric Perfect Bayesian Equilibria (PBE), denoted as p_A and p_B, representing the probabilities that a voter will vote for candidates A and B respectively.

Results and Limitations First, Luxen (2020) proves that a voting equilibrium with strictly positive participation rates by both groups of voters always exists. Moreover, as the mean number of voters $n \to \infty$, the participation rate for the minority candidate increases, consistent with the empirically observed underdog effect. The reason is that an informative signal from the poll results induces a relatively higher turnout in the minority candidate's voters since it indicates a higher probability of being pivotal in the election. However, this effect is never complete, so the majority candidate wins almost surely. This is the main limitation of the model: empirical evidence suggests that many times the underdog effect is strong enough to flip the election in favor of the candidate behind in the polls right before an election.

Another limitation of the model is that the *incomplete* underdog effect is present both in the equilibrium where voters respond truthfully to the opinion poll (and the poll is informative) and in the equilibrium where voters respond strategically to the poll, which becomes completely uninformative. In both cases, the majority candidate wins almost surely.

Additional Features and Extensions To address the limitations and allow for a complete underdog effect, where the minority candidate can win with a probability greater than zero, and the bandwagon effect, where the margin between the two candidates increases, the model requires additional features. In particular, consistent with empirical evidence, I propose to enhance the model with measures of polarization for the candidates and their supporters. Although not formalized yet, we expect that including such measures in the utility function for the voters will plausibly induce the existence of a threshold above which the equilibrium outcome exhibits an underdog effect, and below which a bandwagon effect takes place. I will formalize and explore this conjecture in the following version of this project.

3.2.1 Extended Model without Competition

We extend the baseline framework by introducing a representative pollster who reports an aggregate forecast $\hat{\tau}_A$ for party A. The pollster chooses weights w_i on observations to maximize utility:

$$U_p = \lambda \cdot \text{Accuracy}(\hat{\tau}, \tau) + (1 - \lambda) \cdot \text{Partisan}(\hat{\tau}),$$

where $\lambda \in [0, 1]$ captures the reputation motive (minimizing deviation from the true voting share) and $1 - \lambda$ reflects the partisan motive favoring party A (Bullock, 2011; Leonard and Druckman, 2018). If forecast accuracy falls below a threshold ϵ, the pollster exits the market, modeling a reputation constraint.

Voters observe the published forecast $\hat{\tau}_A$ and update their beliefs $Pr(A)$ via Bayes' rule, treating $\hat{\tau}_A$ as a noisy signal of the true share. Alternatively, voters may form a weighted belief

$$\pi \hat{\tau}_A + (1 - \pi)\pi_0,$$

where π evolves with past forecast accuracy (Kelly, 2022). Election outcomes then materialize, and both pollster and voters observe the realized vote share, allowing voters to update π and pollster type beliefs in subsequent elections. Competition between multiple pollsters introduces additional signals and strategic interactions among pollsters, which we will analyze later.

3.3 Empirical Analysis

In addition to the theoretical analysis, this study plans to test the main predictions of the model using empirical data. Fetzer and Yotzov (2023) investigate the impact of surprise election outcomes on economic cycles using a newly constructed dataset. This dataset combines election results with voting intentions from opinion polls conducted before each election, covering 233 elections across 51 countries from 1980 to 2020. It includes data from approximately 13,600 opinion polls and over 100,000 voting intention estimates, along with details about election characteristics, voter turnout, and the names of incumbents and winners. Data on pre-electoral polls in the US can also be found at the 538 repository.[2]

Additional sources will include the Parliaments and Governments (Parlgov) Database[3],

---

[2] https://data.fivethirtyeight.com

[3] https://www.parlgov.org/data-info

which encodes the political orientation of parties on a 0-10 scale, with scores above five indicating right-wing orientation and scores below five indicating left-wing orientation. However, identifying potential measures of the political orientation of pollsters, if any, will require additional effort.

3.4 Conclusion and future steps

This research proposal aims to develop a theoretical model to better understand the influence of opinion polls on voter behavior and election outcomes. The outcome of the theoretical analyses should be some testable predictions on the electoral outcome depending on specific characteristics of interest. Future empirical validation using extensive datasets on pre-electoral polls and elections will test the model's predictions, ensuring alignment with observed electoral phenomena. This step is crucial for enhancing the model's robustness and applicability. The ultimate goal is to provide policy recommendations on the rules governing the implementation and dissemination of polls.

3.5 References

Blais, André, Gidengil, Elisabeth, and Neil Nevitte. 2006. “Do polls influence the vote?” In H. E. Brady and J. Richard (Eds.), Capturing Campaign Effects, Chapter 11, pp. 263–279. Ann Arbor: University of Michigan Press.

Berent, Matthew K., Robin S. Cowan, and Kyle E. Browning. 2015. “A Theory of Bandwagon Voting.” *Journal of Theoretical Politics.*

Bullock, John G. 2011. “Elite Strategic Weighting and Pollster Reputation.” *Public Opinion Quarterly* 75(2): 271–294.

Cancela, João, and Benny Geys. 2016. “Explaining Voter Turnout: A Meta-Analysis of National and Subnational Elections.” *Electoral Studies* 42: 264-275.

Dawson, Stephen. 2024. “Poll Wars: Perceptions of Poll Credibility and Voting Behaviour.” *The International Journal of Press/Politics* 29(1): 206-226.

Durazzo, Felipe R., and David Turchick. 2023. “Welfare-Improving Misreported Polls.” *Economic Theory* 75(2): 523-565.

Fetzer, Thiemo, and Ivan Yotzov. 2023. “(How) Do Electoral Surprises Drive Business Cycles? Evidence from a New Dataset.” CEPR Discussion Paper 18306.

Frankovic, Kathy, Johnson, Timothy and Marina Stavrakantonaki. 2018. “Freedom to conduct opinion polls: A 2017 worldwide update.” Technical report, ESOMAR and WAPOR.

Goeree, Jacob K., and Jens Großer. 2007. “Welfare Reducing Polls.” *Economic Theory* 31 (1): 51-68.

Grillo, Alberto. 2017. “Risk Aversion and Bandwagon Effect in the Pivotal Voter Model.” *Public Choice* 172 (3-4): 465-482.

Jo, J. 2023. “Informational Roles of Pre-Election Polls.” *Journal of Public Economic Theory*, 25, 441–458.

Kenett, Ron S., Pfeffermann, Denny, and David M. Steinberg. 2018. "Election Polls—A Survey, A Critique, and Proposals." *Annual Review of Statistics and Its Application* 5: 1-26.

Kelly, Nathan J. 2022. "Voter Belief Updating and Poll Accuracy." *Electoral Studies* 74: 102482.

Leonard, Mary, and James N. Druckman. 2018. "Pollster Motivations and Forecast Weighting." *Public Opinion Quarterly* 82(3): 512–536.

Luxen, Christina. 2020. "Polls and Elections: Strategic Respondents and Turnout Implications." CRC TR 224 Discussion Paper Series, University of Bonn and University of Mannheim, Germany.

Myerson, Roger B. 1998. "Extended Poisson games and the Condorcet jury theorem." *Games and Economic Behavior* 25 (1): 111–131.

Palfrey, T. R., and H. Rosenthal. 1985. "Voter participation and strategic uncertainty." *The American Political Science Review* 79 (1): 62–78.

Shirani-Mehr, Houshmand, Rothschild, David, Goela, Sharad, and Andrew Gelman. 2018. "Disentangling Bias and Variance in Election Polls." *Journal of the American Statistical Association* 113: 607–614.

Taylor, Curtis R., and Huseyin Yildirim. 2010. "Public Information and Electoral Bias." *Games and Economic Behavior* 68 (1): 353-375.

Zhou, Zhenkun, Serafino, Matteo, Cohan, Luciano, Caldarelli, Guido, and Hernán A. Makse. 2021. "Why polls fail to predict elections." *Journal of Big Data* 8:137.