University of Mannheim / Department of Economics

Working Paper Series

Certification and Market Transparency

Konrad O. Stahl   Roland Strausz

Working Paper 14-26

September 2014

Certification and Market Transparency

Konrad Stahl and Roland Strausz*

September 1, 2014

Abstract

We provide elementary insights into the effectiveness of certification to increase market transparency. In a market with opaque product quality, sellers use certification as a signaling device, while buyers use it as an inspection device. This difference alone implies that seller-certification yields more transparency and higher social welfare. Under buyer-certification profit maximizing certifiers further limit transparency, but because seller-certification yields larger profits, active regulation concerning the mode of certification is not needed. These findings are robust and widely applicable to, for instance, patents, automotive parts, and financial products.

JEL Classification Numbers: D82, G24, L15.
Keywords: Market Transparency, Certification, Information and Product Quality, Asymmetric Information

*Contact details: Konrad Stahl, University of Mannheim, D-68131 Mannheim, kos@econ.uni-mannheim.de. Roland Strausz, Humboldt–Universität Berlin, strauszr@wiwi.hu-berlin.de; Research supported by the Deutsche Forschungsgemeinschaft (SFB-TR 15). We are grateful to seminar audiences at the Free University of Berlin, ESSET in Gerzensee, Hebrew University, University of Copenhagen, University of Oxford, Tel Aviv University, University College London, the Institute for Advanced Studies Vienna, and in particular Helmut Bester, Gary Biglaiser, Patrick Bolton, Eddie Dekel, Leonardo Felli, Hans Gersbach, Matthias Lang, Guy Laroque, Raphael Levy, Alessandro Lizzeri, Zvika Neeman, Andras Niedermayer, Andrea Prat, Patrick Rey, Frank Rosar, and Lucy White for very constructive comments.

1 Introduction

A market exhibits limited transparency when sellers lack the ability to convey credibly their private information about product quality. As a result, a market with opaque product quality obtains, inducing economic inefficiencies. These inefficiencies create a demand for independent experts — certifiers, who increase market transparency by verifying quality independently. Examples abound. Credit rating agencies certify modern financial products, commercial testing agencies certify the quality of final and intermediate goods, patent offices certify the patentability of inventions, and academic journals certify the quality of research articles.

The above examples all have in common that, in principle, there is demand for transparency through certification from both sides of the market. High quality sellers have a demand for certifiers in order to obtain an appropriately high price for their product, while buyers have a demand for certification to ensure that they do not overspend on low quality.

Given that demand arises from either side of the market, we ask to what extent differences between the two business models, *buyer pays* vs. *seller pays*, affect market transparency and subsequent economic outcomes.[1] At first sight one may expect that the question who pays for certification is immaterial if under both models certification is equally effective in reducing the informational asymmetries and thus ensuring trade. Our main insight is however that even though the basic role of certification — increasing market transparency — is the same under either model, their economic roles differ fundamentally. In particular, we argue that under the *seller pays* model, certification acts as a *signalling device*. In contrast, certification acts as an *inspection device* under the *buyer pays* model.

Due to this economic difference alone, we obtain, surprisingly univocally, the result that certification under the 'seller pays' model is more effective in raising market transparency than under the 'buyer pays' model. It follows that more gains from trade are exhausted under the 'seller pays' model, so that social welfare is higher, bringing us to our *normative* statement that, all other things equal, a certifier *should* offer its services to the seller rather than the buyer. Moreover, we show that a certifier also obtains larger profits when it offers its services to the seller, leading us to our *positive* statement that, all other things equal, a certifier indeed *does opt* for the 'seller pays' model. The result implies that the certifier's preference are in line with social welfare; active regulation concerning the mode of certification is not needed.

[1] In the financial sector, the two alternatives are discussed under the terms *investor pays* vs. *issuer pays*.

We obtain these results by first studying a most parsimonious market with opaque product quality: an adverse selection setup between one buyer and one privately informed seller who sells a good with only two potential qualities. In this setup we characterize and compare equilibrium outcomes under the 'seller pays' and the 'buyer pays' model, respectively. In the outcome under the former, only the seller of the high quality good demands certification, in order to obtain an appropriately high price. Thus, under the 'seller pays' model, certification signals high quality and attains full market transparency so that all gains of trade are realized.

In contrast, under the 'buyer pays' model, both the high quality and the low quality seller pick a high price with positive probability. Upon seeing this high price, the buyer demands certification with a positive probability to prevent herself from overspending on low quality. Certification is therefore used as an inspection device, and the outcome exhibits the typical logic underlying inspection games: only a mixed strategy equilibrium exists, where the buyer certifies with positive probability, and the low quality seller mixes between charging a low and a high price. As a result, certification under the 'buyer pays' model does not attain full market transparency; due to the mixed equilibrium the seller's information remains private with positive probability.

Beyond the in-transparency created this way, we show that the certifier strategically exacerbates market in-transparency under buyer-certification. The reason is that the buyer's demand for certification is high when she is unsure about product quality, and therefore the certifier benefits from *minimizing* market transparency by setting a price of certification that reduces this transparency.

Summarizing, the 'seller pays' model is more effective in promoting market transparency than the 'buyer pays' model, for two reasons. First, its economic role as a signalling device is more suited to the task than the respective role of an inspection device under the 'buyer pays' model. Second, a certifier has no incentive to limit transparency under the former, whereas it has such an incentive under the latter. As a result, the 'seller pays' model generates more social welfare.

We further argue that also the certifier's equilibrium profits are larger in the 'seller pays' model so that the certifier's incentives are aligned with social welfare. This result is not straightforward, because, as we explicitly show, the aforementioned ability of the certifier to manipulate actively the buyer's demand for certification results in a higher certification intensity in the 'buyer pays' model under a large set of parameter constellations. This higher intensity is however offset by a lower equilibrium price of certification so that the certifier's revenues are lower in the 'buyer pays' model, despite a larger demand for certification.

In our formal model, it is relatively straightforward to derive these results under the assumption that certification is costless. The reasoning is substantively more involved when certification is costly, however, as we then need to trade-off the higher benefits from transparency under the 'seller pays' model against potentially higher certification costs in order to obtain our welfare result.

Whereas our parsimonious setup implies that our formal results are derived under restrictive assumptions, we examine their robustness in much detail. In particular, we argue that our results hold under moral hazard, where the high quality seller can actively choose between producing the high or low quality good. Moreover, we show that our results do not depend on the absence of renegotiation, multiple buyers or certifiers, imperfect certification, or the buyer's private information about her preferences.

Since our insights are so elementary, they help to shed light on a diverse range of topical debates. We here single out two, and discuss further applications in more detail in Section 7. First of all, our results contribute to the continuing debate about certification in financial markets. In the aftermath of the 2008 financial crisis, a frequently issued claim is that, due to concerns of capture, credit rating agencies (CRA) should abandon their 'issuer-pays' business model, and return to the 'investor pays' model adopted in earlier years, by offering their certification services to the buyers of financial products rather than their sellers (See White 2010).

Our contribution to this debate is twofold. First, our results imply that even in the absence of capture, it is natural to expect that in equilibrium, CRAs employ the 'issuer pays' model, because it yields them larger profits than the 'investor pays' one. More importantly yet, we provide a clear ceteris-paribus benchmark, that without any further differences between the two models, the current 'issuer pays' model leads to more market transparency and higher social welfare. Hence, any deviation from this model should be motivated by a violation of our ceteris paribus assumption that is strong enough to overturn these results. In the context of capture, this means that the 'investor pays' model will only lead to more market transparency and higher welfare if the problem of capture is *significantly* more severe under the 'issuer pays' model.

The second application singled out here involves the debate about certifying inventions in the form of patents. A prominent and controversially discussed view on this is provided by Lemley (2001), who argues that patent offices should essentially register patent applications, as opposed to tightly examine them. This way, only valuable patents would be challenged by potential users, with their examination delegated to the courts. In terms of our model, this mode of patenting corresponds to the 'buyer pays' model, with the courts playing the role of the certifier. Lemley contrasts this to the procedure in which the patent office directly examines the inventor's patent application before its commercial use, which in our context corresponds to the 'seller pays' model, with the patent office certifying the patent application.

With this interpretation of the two different patent systems, our analysis implies that examination by the patent office is preferable from a welfare point of view, even if we disregard the higher certification costs typically associated with legal courts. We argue this primarily on the basis of one of our extensions involving moral hazard in Section 6, by which revelation of the true quality under the 'seller pays' model induces the inventor to exercise effort towards obtaining a high quality invention – the only one he is willing to spend money for having its quality certified.[2]

The remainder of this paper is organized as follows. In Section 2 we discuss the related literature. In Section 3, we describe the baseline model. In Section 4, we derive the results for seller-certification. In Section 5 we derive the results for buyer-certification and compare them to the case of seller-certification. In Section 6 we discuss extensions of our baseline model and show the results to be robust. In Section 7 we discuss a variety of applications of third party certification, which our model can address. We summarize and conclude with Section 8. All proofs are relegated to the Appendix.

2 Related Literature

Dranove and Jin (2010) provide an extensive survey of the literature on certification. They point out that in the theoretical literature, third party certification is viewed as a means for sellers to *credibly disclose information*, thereby increasing market transparency. Hence, earlier authors (e.g. Viscusi 1978, Grossman and Hart 1980, Grossman 1981, Jovanovic 1982) focus exclusively on the seller's incentive to engage in certification and its effects on market outcomes. The implicit assumption is that only the seller has the ability to disclose information through certification.

[2] We also note that our certification model provides for a straightforward explanation of the so-called patent puzzle in the U.S.: Whereas until recently in the U.S., every innovation was registered as a patent, no matter whether of high or low quality, patents in Europe were granted for, and therefore applied to only for relatively high quality innovations. For a recent discussion of the value of patenting in the U.S., see Boldrin and Levine (2013).

In a second, more recent strand of the literature, authors focus on the certifier's incentives to *manipulate information disclosure*. In particular, they investigate the strategic incentives for a partial disclosure of information (e.g. Lizzeri 1999, Albano and Lizzeri 2001), distorting information disclosure (e.g. Faure-Grimaud et al. 2009, Skreta and Veldkamp 2009, Bolton et al. 2012) or engaging in outright collusion (e.g. Strausz 2005, Mathis et al. 2009, Peyrache and Quesada 2011). Given our interest in providing a ceteris paribus benchmark, we abstract from these concerns and assume that certification is honest and leads to unbiased disclosure. We also do not investigate the incentives of economic agents to become certifiers (e.g. Biglaiser 1983), the interaction between the acquisition and disclosure of information (e.g. Shavell 1994), or the effect of certifiers on market structure (e.g. Board 2009, Guo and Zhao 2009).

In contrast, we compare sellers' versus buyers' incentives in demanding certification, and to study their implications for a profit maximizing certifier and social welfare. In an older working-paper Durbin (1999) has a focus similar to ours. In contrast to our approach, his seller cannot make any (non-verifiable) claims about the quality of his product under buyer-induced certification. As a result, buyer-certification does not result in an inspection game. We consider unnatural the exclusion of any non-verifiable claims by the seller, since such claims arise naturally in the form of initial price quotations. Focusing on rating agencies, Fasten and Hofmann (2010) discuss the provision of certification to a seller versus individual buyers. In their setup, the certification model, however, directly affects its informational content, because seller-induced certification leads to a public signal, whereas buyer-induced certification is private to the buyer who bought it. Thus, a superior transparency of seller-certification is already build in by assumption.

3 The Setup

In our baseline model, we consider certification in an Akerlof adverse selection setup between one seller (he) and one buyer (she).[3] The good's quality q represents the buyer's willingness to pay and can either be high, q_h, or low, q_l, where $\Delta q \equiv q_h - q_l > 0$ and $q_l > 0$. High quality has production costs $c_h > 0$, while low quality has costs $c_l = 0$. The exact quality level is known only to the seller, while all other participants expect high quality with probability λ and low quality with probability $1 - \lambda$. High quality delivers higher economic rents, $q_h - c_h > q_l$, but its production costs exceed

[3]In Section 6, we discuss multiple extensions of the baseline model, including moral hazard, multiple buyers, and the possibility of renegotiation.

average quality, $c_h > \bar{q} \equiv \lambda q_h + (1-\lambda)q_l$. Outside options are zero: the seller obtains zero if he does not produce the good, and the buyer obtains zero if she does not buy.

As is well-known, the assumption $c_h > \bar{q}$ creates Akerlof 's (1970) adverse selection problem. Without being able to observe quality, the uninformed buyer is willing to pay at most a price $p = \bar{q}$. Because this price does not cover the production costs of high quality, a high quality seller will not produce and the market outcome is inefficient. Viscusi (1978) shows that Akerlof's framework creates a demand for an external certifier, who raises market transparency. We assume that such a certifier (it) is available and can, at some fixed cost $c_c \in [0, q_h - c_h]$, reveal truthfully and publicly the seller's quality and charge a price p_c for its services.

With its choice of the *certification model*, the certifier determines whether it offers its service to the seller, which in short we call *seller-certification*; or to the buyer, which we call *buyer-certification*. As motivated in the introduction, our main interest is to understand how each certification model affects market transparency and equilibrium outcomes. We therefore study separately the subgames that arise conditional upon the certifier's choice of the model.

Starting point of our analysis is the observation that in our setup a demand for certification exists from the high quality seller, because only with certification he is able to sell his good at the price $p = q_h$ and achieve a profit of $q_h - c_h > 0$. Because without certification his profits are zero, the high quality seller is willing to certify at any price $p_c \leq q_h - c_h$. A monopolistic certifier can therefore earn a profit up to $\Pi^f = \lambda(q_h - c_h - c_c)$ under seller-certification.

Yet, it seems we can just as well argue that also the buyer has a demand for certification for any price $p_c \leq q_h - c_h$. Indeed, the buyer's certification possibility allows a high quality seller to offer his good at a high price $p = q_h - p_c$, and allows the buyer to verify through certification that it is properly priced before buying it. Consequently, we can replicate the aforementioned certification outcome of seller-certification also under buyer-certification with a certification price $p_c = q_h - c_c$, a price $p = c_h$ for the high quality good, and a price $p = q_l$ for the low quality good. This reasoning suggests that the mode of certification does not matter.

A careful reader might however have noticed that the previous reasoning is misleading, because it disregards the buyer's lack of knowledge about the seller's quality. Indeed, the suggested replication of the outcome under seller-certification requires that the buyer demands certification only when facing a high quality seller. But exactly because of the informational asymmetry, the buyer does not know this. One may argue that even though she does not observe quality directly, she may deduce it indirectly from seeing a price exceeding q_l (because the low quality seller has no incentive to charge this price if the buyer certifies). But this would imply that, when observing this price, the buyer would have actually no incentive to certify at all (since it is costly and she does not expect to learn anything from it). However, if, upon seeing a high price, she would buy the good without certifying it, a low quality seller would also offer the good at this high price. Yet again, such behavior cannot be an equilibrium outcome. Thus, the equilibrium behavior under buyer-certification is not straightforward at all and requires a proper analysis if we want to understand how buyer-certification affects market outcomes.

In order to study the consequences from each certification model, we consider the following generic certification game:

t=1: The certifier sets a price p_c for certification.

t=2: Nature picks quality $q \in \{q_l, q_h\}$ and reveals it to the seller.

t=3: The seller decides whether to produce and sets a price p for the good.

t=4: The seller/buyer decides whether to certify.

t=5: The buyer decides whether or not to buy.

This setup captures the mode of certification in stage 4, where under seller-certification the seller decides whether to certify, whereas under buyer-certification the buyer decides. Moreover, the setup is sufficiently general to capture many different certification procedures in practice, as we discuss more extensively in Section 7.[4]

As argued, we are especially interested in the effectiveness of certification in both attaining market transparency and realizing potential gains of trade. For this reason, we say that a certification model is *information-effective* if it leads to an equilibrium outcome where the buyer perfectly learns the seller's quality before buying the good. When certification is information-effective, it achieves full market transparency. We, moreover, say that a certification model is *trade-effective* if it leads to an equilibrium outcome in which all potential gains of trade are realized, which in our setting means that the good is always produced and sold.

In our certification game, the certifier's price p_c set at $t = 1$ triggers a proper subgame, which is a Bayesian game in extensive form. Clearly, the equilibrium outcome of this subgame plays a crucial role in the determination of the certifier's optimal price

[4]It is straightforward to include the certifier's choice of certification model by an initial stage $t = 0$, where the certifier first decides whether in stage 4, the seller or buyer decides about certification.

p_c. For this reason, our approach is as follows. We first study, for a given p_c, the outcome of the seller-certification subgame $\Gamma^s(p_c)$, where at $t = 4$ the seller decides about certification. After characterizing this outcome, we solve for the monopolistic certifier's optimal price under seller-certification. We then contrast this analysis by studying the buyer-certification subgame $\Gamma^b(p_c)$, where at $t = 4$ the buyer rather than the seller decides about certification.

4 'Seller Pays' Model

We first consider seller-certification. As explained, we start with characterizing the equilibrium outcomes of the subgame $\Gamma^s(p_c)$. In this subgame, the seller picks a price p and decides to offer the good certified or uncertified. Observing the seller's decision and, possibly, the outcome of certification, the buyer decides whether to buy.

Allowing for mixed strategies, we denote the seller's strategy as a probability distribution over prices p and whether to certify the good. In particular, let $\sigma_i^c(p)$ denote the probability that a seller with quality q_i offers the good certified at a price p, and $\sigma_i^u(p)$ the probability that he offers the good uncertified at that price.[5] The seller's strategy σ_i is then a combination $(\sigma_i^c, \sigma_i^u), i \in \{l, h\}$ such that

$$\sum\nolimits_j \sigma_i^c(p_j) + \sum\nolimits_j \sigma_i^u(p_j) = 1.$$

After observing the seller's price and his decision to certify, the buyer forms a belief about the probability that the good has high quality. If the seller has his good certified, the buyer learns its true quality, and thus her beliefs after certification coincide with the true quality q_i. Consequently, she buys a certified good whenever $p \leq q_i$. If the good is uncertified, the buyer's belief $\mu(p)$ that it is of high quality is, in general, uncertain. It depends on the price p, since the buyer may interpret the price p as a signal of quality. In equilibrium, the belief must follow Bayes' rule whenever possible. Consequently, we say that *the belief $\mu(p)$ is consistent (with the seller's strategy (σ_l, σ_h))* if for any $\sigma_i^u(p) > 0$ it satisfies

$$\mu(p) = \frac{\lambda \sigma_h^u(p)}{\lambda \sigma_h^u(p) + (1 - \lambda)\sigma_l^u(p)}. \tag{1}$$

Facing an uncertified good at a price p, the buyer's belief equals $\mu(p)$, and it is optimal for her to buy when the expected quality $\mu(p)q_h + (1 - \mu(p))q_l$ exceeds the price p quoted by the seller. When that price exceeds expected quality, it is

[5] In order to circumvent measure-theoretical complications, we let the seller randomize over a finite but arbitrary number of prices.

optimal not to buy, and when expected quality coincides with the price, any random buying behavior is optimal. Let $\sigma(s_b|p,\mu)$ denote the probability that the buyer buys the good uncertified, i.e. takes the action s_b, given the seller has quoted the price p and the buyer's belief is μ. We say that $\sigma(s_b|p,\mu)$ *is optimal* if for any (p,μ), the corresponding buying behavior $\sigma(s_b|p,\mu)$ is optimal.

Let π_i^u denote the expected payoff of a seller with quality q_i, who offers the good uncertified. Given the buyer's belief $\mu(p)$ and her buying behavior $\sigma(s_b|p,\mu)$, a high quality seller and a low quality seller expect the following respective payoffs from offering the good uncertified at a price p:

$$\pi_h^u(p) = \sigma(s_b|p,\mu(p))p - c_h \quad \text{and} \quad \pi_l^u(p) = \sigma(s_b|p,\mu(p))p. \tag{2}$$

Hence, a strategy $\sigma_i = (\sigma_i^c, \sigma_i^u)$ yields the seller of quality q_i the expected payoff

$$\pi_i(\sigma_i) = \sum\nolimits_j \sigma_i^u(p_j)\pi_i^u(p_j) + \sum\nolimits_j \sigma_i^c(p_j)[p_j \mathbf{1}_i(p_j) - p_c - c_i],$$

where $\mathbf{1}_i(p)$ is an indicator function which equals 1 if $p \leq q_i$ and 0 otherwise. We say that *the seller strategy* σ_i^* *is optimal* if it maximizes $\pi_i(\sigma_i)$.

A perfect Bayesian equilibrium (PBE) of the subgame $\Gamma^s(p_c)$ is a combination $\{\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*\}$ for which the seller's strategies σ_l^* and σ_h^* are optimal, the belief μ^* is consistent, and the buyer's strategy σ^* is optimal. With this definition the following lemma characterizes the equilibrium outcomes corresponding to the subgame $\Gamma^s(p_c)$.

Lemma 1 *Consider the subgame* $\Gamma^s(p_c)$ *with seller-certification.*

i. For $p_c \leq q_h - c_h$*, a PBE exists for which the certifier obtains the payoff* $\lambda(p_c - c_c)$*, the good is always sold, the seller with quality* q_h *always certifies, whereas the seller with quality* q_l *does not. For* $p_c < q_h - c_h$*, this equilibrium outcome is unique.*

ii. For $p_c > q_h - c_h$*, the high and the low quality seller do not certify in any PBE and the outcome coincides with the market outcome without a certifier.*

The lemma shows that for a low enough price of certification, the high quality seller certifies to reveal his high quality. Hence, certification is used as a signaling device and the buyer interprets an uncertified good as revealing bad quality. For all certification prices different from $q_h - c_h$, the equilibrium outcome is unique. Note that this is in line with results about certification in competitive adverse selection markets (e.g. Viscusi 1978). Effectively, only certification serves as a credible signal, whereas the price p does not.

The lemma has the following direct implication for the effectiveness of seller-certification.

Corollary 1 *For $p_c < q_h - c_h$, seller-certification is information- and trade-effective.*

When choosing its price of certification, the certifier will take into account how it affects its demand as stated in the lemma. Let Π^s denote the certifier's payoff under seller-certification. The following proposition characterizes the outcome under seller-certification when we include the price setting decision of the certifier.

Proposition 1 *The game with seller-certification has a unique equilibrium outcome $\bar{p}_c^s = q_h - c_h$ with equilibrium expected payoffs $\Pi^s = \lambda(q_h - c_h - c_c)$ to the certifier, and $\pi_h^* = 0$ and $\pi_l^* = q_l$ to the seller. Moreover, the high quality seller certifies with certainty, the low quality seller does not certify, and the good is always traded.*

Unsurprisingly, the monopolistic certifier extracts all economic rents from certification. Consequently, the high quality seller is just as well off as without certification and obtains zero profits. Yet, in equilibrium all gains of trade are realized and the seller's quality is fully revealed. This yields the following corollary.

Corollary 2 *Monopolistic seller-certification is information- and trade-effective.*

5 'Buyer Pays' Model, Equilibrium, and Welfare Model Choice

We next study buyer-certification. Again, we first consider the subgame $\Gamma^b(p_c)$ for a given price of certification p_c. In this subgame, the seller first picks a price p and the buyer then decides whether to certify the good and to buy it.

Under buyer-certification, the seller's task is to pick a price. We let $\sigma_i(p_j)$ denote the probability that the seller with quality q_i sets a price p_j. Thus, for both $i \in \{l, h\}$,

$$\sum_j \sigma_i(p_j) = 1.$$

Observing the price p, the buyer forms a belief $\mu(p)$ about the probability that the good has high quality. Again, the buyer's belief follows Bayes' rule whenever possible. We, therefore, say that *the belief $\mu(p)$ is consistent (with the seller's strategy (σ_h, σ_l))* if for any $\sigma_i(p) > 0$ it satisfies

$$\mu(p) = \frac{\lambda\sigma_h(p)}{\lambda\sigma_h(p) + (1-\lambda)\sigma_l(p)}. \tag{3}$$

Given the price p and belief μ, the buyer has three relevant actions:

1. Action s_b: The buyer does not certify but buys the good. This yields payoff

$$U(s_b|p,\mu) = \mu q_h + (1-\mu)q_l - p.$$

2. Action s_n: The buyer does not certify, nor buy the good. This yields payoff

$$U(s_n|p,\mu) = 0.$$

3. Action s_h: The buyer certifies the good and buys only if certification reveals the high quality q_h. This yields payoff

$$U(s_h|p,\mu) = \mu(q_h - p) - p_c.$$

The other three actions open to the buyer — to certify and always buy, to certify but never buy, and to certify and buy only if quality is low — are clearly suboptimal. We therefore disregard them.

The action s_n is optimal whenever $U(s_n|p,\mu) \geq U(s_b|p,\mu)$ and $U(s_n|p,\mu) \geq U(s_h|p,\mu)$. Hence, the set of (p,μ) combinations for which s_n is optimal is

$$S(s_n|p_c) \equiv \{(p,\mu)|p \geq \mu q_h + (1-\mu)q_l \ \wedge \ p_c \geq \mu(q_h - p)\}.$$

Likewise, the action s_b is optimal whenever $U(s_b|p,\mu) \geq U(s_n|p,\mu)$ and $U(s_b|p,\mu) \geq U(s_h|p,\mu)$. Hence, the set of (p,μ) combinations for which s_b is optimal is

$$S(s_b|p_c) \equiv \{(p,\mu)|p \leq \mu q_h + (1-\mu)q_l \ \wedge \ p_c \geq (1-\mu)(p - q_l)\}.$$

Finally, the action s_h is optimal whenever $U(s_h|p,\mu) \geq U(s_n|p,\mu)$ and $U(s_h|p,\mu) \geq U(s_b|p,\mu)$. Hence, the set of (p,μ) combinations for which s_h is optimal is

$$S(s_h|p_c) \equiv \{(p,\mu)|p_c \leq \mu(q_h - p) \wedge p_c \leq (1-\mu)(p - q_l)\}.$$

Figure 1 illustrates the buyer's optimal actions. For low product prices p, the buyer buys the good uncertified, $(p,\mu) \in S(s_b)$, whereas for high prices p the buyer refrains from buying, $(p,\mu) \in S(s_n)$. It turns out that as long as $p_c < \Delta q/4$, there is an intermediate range of prices p and beliefs μ such that the buyer demands certification, i.e. $(p,\mu) \in S(s_h)$. In this case, the buyer only buys the product when certification reveals it to be of high quality. Note that apart from points on the thick, dividing lines, the buyer's optimal action is uniquely determined so that mixing over different actions is suboptimal.

For future reference we define

$$\tilde{p} \equiv \left(q_h + q_l + \sqrt{\Delta q(\Delta q - 4p_c)}\right)/2 \text{ and } \tilde{\mu} \equiv \left(1 + \sqrt{1 - 4p_c/\Delta q}\right)/2. \tag{4}$$

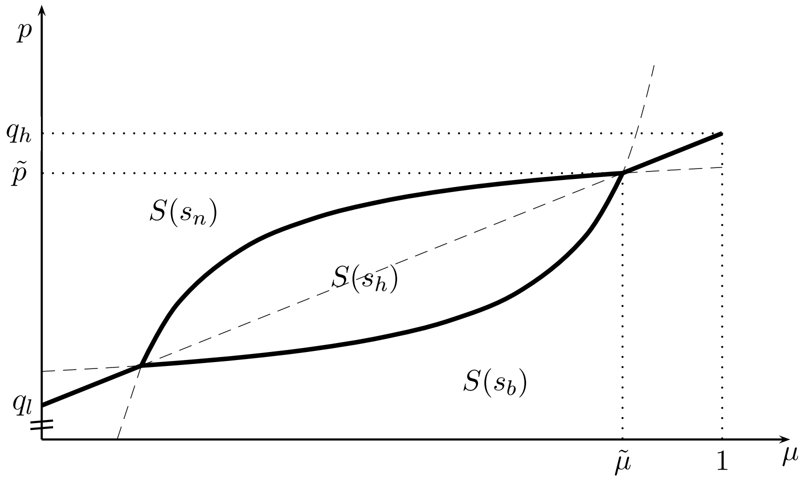

Figure 1: Buyer's buying behavior for given $p_c < \Delta q/4$.

If the seller quotes the price $\tilde{p}$ and the buyer has beliefs $\tilde{\mu}$, then the buyer is indifferent between all his three actions.

Let $\sigma(s|p, \mu)$ denote the probability that the buyer takes action $s \in S = \{s_b, s_n, s_h\}$ given price p and belief μ. We can then denote the buyer's (mixed) strategy by probabilities $\sigma(s|p, \mu)$ such that

$$\sigma(s_b|p, \mu) + \sigma(s_n|p, \mu) + \sigma(s_h|p, \mu) = 1.$$

We say that *the strategy σ^* is optimal* if it randomizes among those actions that are optimal: $\sigma^*(s|p, \mu) > 0$ implies that $(p, \mu) \in S(s|p_c)$.

Given buyer's belief μ and her strategy σ, a seller with quality q_h and a seller with quality q_l expect the following respective payoffs from offering the good at a price p:

$$\pi_h(p, \mu|\sigma) = [\sigma(s_b|p, \mu) + \sigma(s_h|p, \mu)]p - c_h \text{ and } \pi_l(p, \mu|\sigma) = \sigma(s_b|p, \mu)p.$$

Given that a price p leads to the belief $\mu(p)$, a seller with quality q_h and a seller with quality q_l expect the following respective payoffs from offering the good at a price p:

$$\pi_h^b(p) = \pi_h(p, \mu(p)|\sigma) \;\; \text{and} \;\; \pi_l^b(p) = \pi_l(p, \mu(p)|\sigma). \tag{5}$$

We say that *the seller's pricing strategy σ_i is optimal (with respect to the buyer's behavior (σ^*, μ^*))* if any price $\hat{p}$ such that $\sigma_i(\hat{p}) > 0$ maximizes $\pi_i^b(p)$:

$$\sigma_i(p) > 0 \;\Rightarrow\; \pi_i(p, \mu^*(p)|\sigma^*) \geq \pi_i(p', \mu^*(p')|\sigma^*), \; \forall p'. \tag{6}$$

A perfect Baysian equilibrium (PBE) of the subgame $\Gamma^b(p_c)$ is a combination $\{\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*\}$ for which the sellers' strategies σ_l^* and σ_h^* are optimal, the belief μ^* is consistent and the buyer's strategy σ^* is optimal.

It follows that in a PBE $(\sigma_h^*, \sigma_l^*, \mu^*, \sigma^*)$ the high quality seller's and the low quality seller's payoffs, respectively, are

$$\pi_h^* = \sum\nolimits_i \sigma_h^*(p_i)\pi_h(p_i, \mu^*(p_i)|\sigma^*) \text{ and } \pi_l^* = \sum\nolimits_i \sigma_l^*(p_i)\pi_l(p_i, \mu^*(p_i)|\sigma^*).$$

Corollary 1 showed that seller-certification at a price $p_c < c_h - q_h$ is both informational- and trade-effective. The following lemma reveals that buyer-certification does not lead to transparent markets if it is trade-effective.

Lemma 2 *If buyer-certification at some price p_c is trade-effective, then it is not information-effective.*

The lemma makes precise the sense in which buyer-certification is an imperfect tool for achieving market transparency, given that one ultimately wants to achieve efficient trade. The next lemma shows that, as a result, efficient trade is not attainable under buyer-certification.

Lemma 3 *Buyer-certification is not trade-effective.*

Lemma 3 indicates that under buyer-certification gains from trade are not fully realized, and some trading inefficiencies remain. This points to an important drawback of buyer-certification. When certification involves no costs ($c_c = 0$), this result allows us to conclude directly that welfare under seller-certification is higher than under buyer-certification, where welfare is defined simply as the sum of all the agents' surplus. Because in the seller pays model, the certifier is able to extract all the rents from certification, this also directly implies that the certifier can charge a higher price and its profits are larger. We therefore obtain the following corollary.

Corollary 3 *Suppose the certifier incurs no cost of certification ($c_c = 0$). Then the 'seller pays' model is welfare superior to the 'buyer pays' model. Moreover, the certifier charges a higher price and obtains larger profits under seller-certification so that its preferences concerning the certification model are in line with welfare.*

Lemma 3 is insufficient to make similar claims, however, when certification is costly ($c_c > 0$). In this case, overall welfare depends on both the indirect gains of trade from certification and the direct welfare costs of certification. Although the indirect gains are higher under seller-certification, we cannot exclude *a priori* that, due to a higher certification intensity, these higher gains are offset by larger certification costs. In order to address this question, we first need to characterize fully the equilibrium outcome in the subgame $\Gamma^b(p_c)$. This characterization will also enable us to show a further perverse effect of buyer-certification, that it induces certifiers to artificially limit market transparency.

We obtain a characterization of the market outcome in a series of three lemmata. The first of these makes precise the intuitive result that the seller's expected profits increase when the buyer is more optimistic about the good's quality.

Lemma 4 *In any PBE* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma^b(p_c)$ *with* $p_c > 0$ *the payoffs* $\pi_h(p, \mu|\sigma^*)$ *and* $\pi_l(p, \mu|\sigma^*)$ *are non-decreasing in* μ.

The next lemma shows the implications of the consistency requirement (3) on the belief $\mu(p)$. In particular, it shows that the seller, no matter his type, never sets a price below q_l, and the low quality seller never sets a price above q_h. The lemma also shows that, in equilibrium, the low quality seller never loses from the presence of asymmetric information, since he can always guarantee himself the payoff q_l that he obtains with observable quality. By contrast, the high quality seller loses from the presence of asymmetric information; his payoff is strictly smaller than $q_h - c_h$.

Lemma 5 *In any PBE* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma^b(p_c)$ *we have i)* $\sigma_l^*(p) = 0$ *for all* $p \notin [q_l, q_h]$ *and* $\sigma_h^*(p) = 0$ *for all* $p < q_l$*; ii)* $\pi_l^* \geq q_l$*; iii)* $\pi_h^* < q_h - c_h$.

As is well known, the concept of Perfect Bayesian Equilibrium places only weak restrictions on admissible beliefs. In particular, it does not place any restrictions on the buyer's beliefs, that are based on prices not played in equilibrium; any out-of-equilibrium belief is allowed. Hence, as is typical for signaling games, without any restrictions on out-of-equilibrium beliefs we cannot pin down behavior in the subgame $\Gamma^b(p_c)$ to a specific equilibrium. Especially by the use of pessimistic out-of-equilibrium beliefs, one can sustain many equilibrium pricing strategies.

In order to reduce the arbitrariness of equilibrium play, it is necessary to strengthen the solution concept of PBE by introducing more plausible restrictions on out-of-equilibrium beliefs. A standard belief restriction is the intuitive criterion of Cho-Kreps (1987), which in its standard formulation only has bite in an equilibrium where the signalling player reveals himself fully so that $\mu \in \{0, 1\}$ results. Bester and Ritzberger (2001) propose the following extension of the intuitive criterion to intermediate beliefs $\mu \notin \{0, 1\}$.

Belief Restriction (BR): A Perfect Bayesian Equilibrium $(\sigma_h^*, \sigma_l^*, \mu^*, \sigma^*)$ satisfies *BR* if, for any $\mu \in [0, 1]$ and any out-of-equilibrium price p, we have

$$\pi_l(p, \mu) < \pi_l^* \wedge \pi_h(p, \mu) > \pi_h^* \;\Rightarrow\; \mu^*(p) \geq \mu.$$

The belief restriction states intuitively that if a pessimistic belief μ gives only the high quality seller an incentive to deviate, then the restriction requires that the buyer's actual belief should not be even more pessimistic than μ. It extends the intuitive criterion of Cho-Kreps, which obtains for the special case $\mu = 1$. Indeed, the restriction extends the logic of the Cho-Kreps criterion to situations where the deviation to a price p is profitable only for the high quality seller when the buyer believes that the deviation originates from the high quality seller with probability μ. As we may have $\mu < 1$, the restriction considers more pessimistic beliefs than the Cho-Kreps criterion.

The next lemma characterizes equilibrium outcomes that satisfy the belief restriction (BR). In particular, the refinement implies that the high quality seller can sell his product at a price of at least $\tilde{p}$.

Lemma 6 *Any Perfect Bayesian Equilibrium* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma^b(p_c)$ *that satisfies BR exhibits i)* $\sigma_h^*(p) = 0$ *for all* $p < \tilde{p}$ *and ii)* $\pi_h^* \geq \tilde{p} - c_h$.

By combining the previous three lemmata, we are now able to characterize the equilibrium outcome.

Proposition 2 *Consider a PBE* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of* $\Gamma^b(p_c)$ *that satisfies BR. Then*

i. For $\tilde{\mu} > \lambda$ *and* $\tilde{p} > c_h$ *it is unique. The high quality seller sets the price* $\tilde{p}$ *with certainty,* $\sigma_h^*(\tilde{p}) = 1$*, while the low quality seller randomizes between price* $\tilde{p}$ *and* q_l *and the buyer randomizes between* s_b *and* s_h *upon observing the price* $\tilde{p}$*. The respective probabilities with which the low quality seller picks* $\tilde{p}$ *and the buyer certifies are*

$$\sigma_l^*(\tilde{p}) = \frac{\lambda(1-\tilde{\mu})}{\tilde{\mu}(1-\lambda)} \text{ and } \sigma^*(s_h|\tilde{p}, \tilde{\mu}) = \frac{\tilde{p} - q_l}{\tilde{p}}.$$

ii. For $\tilde{\mu} < \lambda$ *or* $\tilde{p} < c_h$*, certification does not take place in equilibrium.*

iii. For $\tilde{\mu} \geq \lambda$ *and* $\tilde{p} \geq c_h$*, an equilibrium outcome as described under i. exists.*

The proposition formalizes our insight that buyer-certification serves as an inspection device to discipline the low quality seller. Indeed, the high quality seller signals his quality by announcing $\tilde{p}$, while the buyer and the low quality seller play the mixed strategies typical of an inspection game: By choosing the low price q_l, the low quality seller provides an honest signal, whereas he cheats by picking the high price $\tilde{p}$. Whenever the buyer observes $\tilde{p}$, she cannot identify the good's quality. Therefore she certifies with positive probability.

A pure equilibrium does not exist. On one hand, if the buyer would always certify when seeing the high price, the low quality seller would not cheat by asking such a price; but without any cheating certification is suboptimal. On the other hand, if the buyer would never certify, then the low quality seller would have a strict incentive to cheat and to quote the high price; but with such cheating the buyer would want to certify. Hence, only a mixed equilibrium exists, where the buyer's certification probability keeps the low quality seller indifferent between cheating and honestly pricing his good, while at the same time the cheating probability of the low quality seller keeps the buyer indifferent between buying the good uncertified and asking for certification. In order to satisfy both indifference conditions, the high price must equal $\tilde{p}$ and the buyer's belief must equal $\tilde{\mu}$.

In Proposition 2 we characterize the equilibrium outcome under buyer-certification for a given price of certification p_c. The proposition allows us to derive the demand for buyer-certification by taking into account that $\tilde{\mu}$ and $\tilde{p}$ depend on p_c according to (4). We therefore write these dependencies explicitly as $\tilde{p}(p_c)$ and $\tilde{\mu}(p_c)$. Because the equilibrium probability of buyer-certification is the compounded probability that the seller picks the price $\tilde{p}$ and the buyer certifies, we can write demand as

$$x^b(p_c) = [\lambda + (1-\lambda)\sigma_l^*(\tilde{p}(p_c))]\sigma^*(s_h|\tilde{p}(p_c), \tilde{\mu}(p_c)),$$

whenever $\tilde{\mu}(p_c) \geq \lambda$ and $\tilde{p}(p_c) \geq c_h$, and as zero otherwise. Inserting $\sigma_l^*(\tilde{p})$ and $\sigma^*(s_h|\tilde{p}, \tilde{\mu})$ from Proposition 2, the certifier's profit under buyer-certification is

$$\Pi^b(p_c) = x^b(p_c)(p_c - c_c) = \frac{\lambda(\tilde{p}(p_c) - q_l)}{\tilde{\mu}(p_c)\tilde{p}(p_c)}(p_c - c_c), \tag{7}$$

whenever $\tilde{\mu}(p_c) \geq \lambda$ and $\tilde{p}(p_c) \geq c_h$, and zero otherwise. In the next proposition we derive the monopoly price of buyer-certification.

Proposition 3 *Consider the game with buyer-certification.*

i. For $c_h \leq (q_h + q_l)/2$, the certifier sets a price $\bar{p}_c^b = \Delta q/4$, which induces a subgame $\Gamma^b(\bar{p}_c^b)$ with $\tilde{\mu}(\bar{p}_c^b) = 1/2$ and a certification profit of

$$\Pi^b = \frac{\lambda \Delta q}{2(q_h + q_l)}(\Delta q - 4c_c).$$

ii. For $c_h > (q_h + q_l)/2$, the certifier sets the price $\bar{p}_c^b = (q_h - c_h)(c_h - q_l)/\Delta q$, which induces a subgame $\Gamma^b(\bar{p}_c^b)$ with $\tilde{p}(\bar{p}_c^b) = c_h$ and a certification profit of

$$\Pi^b = \frac{\lambda[(q_h - c_h)(c_h - q_l) - \Delta q c_c]}{c_h}.$$

The proposition reveals the perverse effect that buyer-certification induces the certifier to minimize market transparency artificially. According to Proposition 3i., the certifier picks a price $\bar{p}_c^b$ such that after observing the price $\tilde{p}$, the buyer has beliefs

$\tilde{\mu}(\bar{p}_c^b) = 1/2$. This maximizes her uncertainty about product quality and implies that market transparency is minimized.

To see that this perverse effect results directly from the role of buyer-certification as an inspection device, observe that the value of an inspection device is typically higher when the underlying uncertainty is larger. Hence, the buyer's willingness to pay for certification and her demand are highest when, conditional upon observing the price $\tilde{p}$, market transparency is minimized. The certifier's most preferred price p_c is, therefore, such that $\tilde{\mu}(p_c) = 1/2$. The certifier must however ensure that at this price the high quality seller does not drop out of the market. In the case specified in Proposition 3ii, this limits the certifier's ability to fully minimize transparency.

In Corollary 3 we showed that, for zero certification costs, seller-certification outperforms buyer-certification both from a social welfare and an equilibrium perspective. By contrasting the equilibrium outcomes under seller- and buyer-certification as derived in Proposition 1 and 3, we now show that these two results also obtain when certification costs are positive. We first show this for the certifier's profits:

Proposition 4 *For any cost of certification $c_c \in [0, q_h - c_h]$, the certifier obtains a higher profit and charges higher prices under seller-certification than under buyer-certification, $\Pi^s > \Pi^b$ and $\bar{p}_c^s > \bar{p}_c^b$. Hence, the certifier prefers the 'seller pays' to the 'buyer pays' model. The certification intensity in the 'buyer pays' model exceeds the certification intensity in the 'seller pays' model, whenever $c_h > (q_h + q_l)/2$ or $q_h < 3q_l$.*

Next we show that Corollary 3 extends to positive certification costs also for social welfare. We thereby ideally want to establish that social welfare is higher not only for the respective monopoly prices $\bar{p}_c^s$ and $\bar{p}_c^b$ but also for lower price combinations. If this were the case, then our welfare result would also hold when certification markets are more competitive in that they exhibit equilibrium prices below monopoly. Under perfect competition we expect certification prices to equal marginal costs c_c. For intermediate forms of competition, where certifiers have some market power, we expect prices to exceed marginal costs but not reach monopoly levels. Hence, we want to show that for any combination of certification prices, (p_c^s, p_c^b), in between marginal costs and the respective monopoly price, social welfare is larger under seller-certification.

For any price of certification p_c^s that lies in between marginal cost c_c and the monopoly price under seller-certification $\bar{p}_c^s$, the high quality seller certifies and the good is always traded. Hence, welfare under seller-certification is

$$W^s = \lambda(q_h - c_h) + (1 - \lambda)q_l - \lambda c_c.$$

As long as the price of certification, p_c^s, does not exceed the monopoly price $\bar{p}_c^s$, welfare under seller-certification is independent of the actual price, because for such prices demand is inelastic so that the price represents a pure welfare transfer.

This is different under buyer-certification, because here the certification price affects directly the gains of trade. This is because buyer-certification is not trade-effective; the good is not sold when the low quality seller picks a price exceeding q_l and the buyer certifies. According to Proposition 2, this happens with probability

$$\omega(p_c^b) = \sigma_l^*(\tilde{p}(p_c^b))\sigma^*(s_h|\tilde{p}(p_c^b), \tilde{\mu}(p_c^b)),$$

which depends explicitly on the price of certification p_c^b. For any certification price that does not exceed the monopoly price under buyer certification $\bar{p}_c^b$, the high quality good is always sold so that social welfare under buyer-certification is

$$W^b(p_c^b) = \lambda(q_h - c_h) + (1-\lambda)(1-\omega(p_c^b))q_l - x^b(p_c^b)c_c.$$

The difference in welfare is therefore

$$\Delta W(p_c^b) \equiv W^s - W^b(p_c^b) = (1-\lambda)\omega(p_c^b)q_l - [\lambda - x^b(p_c^b)]c_c. \tag{8}$$

The expression illustrates the trade-off between differences in trade-effectiveness — represented by the first, positive term $(1-\lambda)\omega(p_c^b)q_l$ — and the cost of certification — represented by the second, possibly negative term $[\lambda - x^b(p_c^b)]c_c$. For zero certification costs, the second term disappears and the expression is strictly positive. This confirms Corollary 3. When certification costs are positive, then we cannot directly draw a conclusion, because when the certification intensity under buyer-certification, $x^b(p_c^b)$, is substantially lower than the certification intensity λ under seller-certification, the second term outweighs the first term and renders $\Delta W(p_c^b)$ negative.

The next proposition shows, however, that this is not the case for any buyer-certification price p_c^b in between marginal costs c_c and the monopoly price $\bar{p}_c^b$.

Proposition 5 *For any cost of certification $c_c \in [0, q_h - c_h]$ and any combination of seller-certification and buyer-certification prices such that each price lie in between marginal costs and the respective monopoly price, $(p_c^s, p_c^b) \in [c_c, \bar{p}_c^s] \times [c_c, \bar{p}_c^b]$, welfare under seller-certification exceeds welfare under buyer-certification.*

6 Extensions

We derived our formal results in a highly stylized Akerlof-style model. In this section, we argue that they are robust to many extensions. Specifically, we argue the results remain unchanged with the introduction of seller moral hazard concerning the endogenous choice of the good's quality; of multiple quality levels; of many buyers; of imperfect certification; of competition between certifiers; and of *ex post* price renegotiations. We further show that our results do not crucially depend on the absence of seller uncertainty about product quality and the absence of the buyer's private information about her preferences.

6.1 Moral hazard

We analyzed certification in a pure adverse selection setting, where quality is determined exogenously in that a high quality seller cannot produce at low quality. If we consider a model in which the high quality seller has this choice, moral hazard results. Here we show that such moral hazard does not affect our insights. We do so by extending our certification game with and intermediate stage in between $t = 2$ and $t = 3$:

t=2.5: Seller type q_h decides whether to produce quality q_h or q_l.

The introduction of moral hazard improves the outside option of the high quality seller, because in addition to not producing, the seller now can also decide to produce at low quality. Thus, the certifier can extract less rents and its equilibrium profits decrease. As we now argue, this however does not affect our insights.

To make this precise, note that with moral hazard, type q_h's relevant outside option is to produce q_l (leading to profit q_l) rather than not sell at all (leading to profit 0). As a result, the certifier's profits from seller-certification reduce as follows:

Proposition 6 *Under seller-certification and moral hazard the certification game has the unique equilibrium outcome* $\bar{p}_c^s = \Delta q - c_h$ *with equilibrium payoffs* $\Pi^s = \lambda(\Delta q - c_h - c_c)$*,* $\pi_h^* = q_l$*, and* $\pi_l^* = q_l$*.*

The outside option changes similarly under buyer-certification, where rather than ensuring that $\tilde{p} - c_h \geq 0$, the certifier now has to ensure that $\tilde{p} - c_h \geq q_l$. The next proposition makes precise how Proposition 3 changes in the presence of moral hazard.

Proposition 7 *Consider buyer-certification with moral hazard.*
i. For $\lambda \leq 1/2$ *and* $c_h \leq \Delta q/2$*, the certifier sets a price* $\bar{p}_c^b = \Delta q/4$ *and obtains*

$$\Pi^b = \frac{\lambda \Delta q}{2(q_h + q_l)}(\Delta q - 4c_c).$$

ii. For $\lambda > 1/2$ *or* $c_h > \Delta q/2$*, the certifier sets a price* $\bar{p}_c^b = c_h(1 - c_h/\Delta q)$ *and obtains*

$$\Pi^b = \frac{\lambda[c_h(\Delta q - c_h) - \Delta q c_c]}{c_h + q_l}.$$

Although under seller-certification the certifier's profit declines with moral hazard, the next proposition shows that the certifier's profits remain higher under seller-certification.

Proposition 8 *With moral hazard the certifier obtains a higher profit under seller- than under buyer-certification:* $\Pi^s > \Pi^b$.

Hence, including moral hazard does not affect our result that the certifier prefers the 'seller pays' model. Moreover, our welfare result remains also unchanged, because the equilibrium including moral hazard involves no change in the allocation, but only a redistribution of rents away from the certifier towards the seller. We therefore conclude that our results readily extend to the presence of moral hazard.

We point out that this extension is particularly relevant for the application of our results to the alternative patents examination schemes introduced in the introduction: under the 'seller pays' model, inventors take extra effort to provide high quality, and have only high quality examined by the patent office, whereas under the 'buyer pays' model, low quality patents go through with a positive probability, and are wastefully examined by the courts.

6.2 More than two qualities

In our formal analysis, we assumed that sellers have only two qualities. This assumption allowed us to fully characterize the equilibrium outcome under buyer-certification. Yet, even though the problem of characterizing the equilibrium outcome becomes intractable with more than two types, it should be clear that our main results do not depend on its explicit characterization. In particular, Corollary 3 extends, because the result that seller-certification is information- and trade-effective (Lemma 1) whereas buyer-certification is not (Lemma 2-3), is valid for any number of quality levels.

Also the logic underlying the perverse effect of buyer-certification that a certifier wants to minimize market transparency, is independent of the number of seller types. This effect is directly linked to our insight that buyer-certification plays the role of an inspection device, because it implies that the demand for buyer-certification is larger when the market is less transparent. This demand effect gives a strategic certifier an incentive to make markets more opaque. Effectively, our explicit characterization of the equilibrium outcome was only needed to show this effect formally and to be able to compute explicitly the certifier's payoffs under buyer-certification.

6.3 Multiple buyers

We can also extend our results to multiple buyers. In particular, consider a setting that applies particularly well to the financial market, where one seller can sell n units to n identical buyers. Essentially, there are two possible information structures.

In the first one, buyers cannot share the certification result but they each must buy certification individually. Under seller-certification, Proposition 1 is changed so that the profits from selling the product are multiplied by n. Under buyer-certification, our formal results carry through, implying that the certifier's profits are also multiplied by n. Because the certifier's profits from selling to buyers and sellers are both multiplied by n, the equilibrium ranking of seller- vs. buyer-certification from the perspective of both a monopolistic certifier and social welfare remains unchanged.

The second information structure is one in which buyers collude to collectively initiate certification. Again both under seller- and buyer-certification, the certifier's profits are multiplied by n. Again, the results remain unchanged.

6.4 Competitive certification

Due to both technical and reputational economies of scale, certification markets tend to be highly concentrated. For this reason, our assumption of a monopolistic certifier represents a good baseline. We want to point out however that our two results — buyer-certification leads to both less transparent markets and lower social welfare than seller-certification — holds independent of the market structure in the certification market. In particular, the lemmata 1,2, and 3, which imply our transparency results, do not depend on the market structure, while Proposition 5, which states our welfare ordering for any cost of certification, holds for any price charged by the certifier that lies in between marginal costs and monopoly prices, and therefore for any mode of competition in the certification markets.

Moreover, the logic behind the perverse effect that under buyer-certification the certifier has an incentive to minimize market transparency artificially implies that it arises whenever certifiers can set their prices strategically to influence demand. Hence, as soon as certifiers have market power the perverse effect crops up and its degree depends on the certifiers' effective market power.

Finally, it is relatively straightforward to see that even in the extreme case of perfect Bertrand competition our positive result that the 'seller pays' model yields certifiers higher profits is not overturned. Bertrand competition leads the certifiers to offer their service at a price $p_c = c_c$ under both seller- and buyer-certification, thus making certifiers indifferent between the two modes of certification.

6.5 Imperfect certification

Consider a certification technology is informative, but not perfectly so. Assume specifically that the certifier can reveal the correct quality only with probability $\pi > 1/2$, and identifies the wrong quality with corresponding probability $(1 - \pi) > 0$. Although this imperfection reduces the profitability of both seller- and buyer-certification, it does not qualitatively change the equilibrium. In the case of seller-certification, the equilibrium outcome remains separating also with imperfect certification and is continuous in π.[6] Imperfect certification also does not change the nature of the equilibrium outcome with buyer-certification. Intuitively, a less informative certification technology shrinks the intermediate area in Figure 1, where $S(s_h)$ is optimal. Hence, the equilibrium outcomes under buyer- and seller-certification are continuous in π so that at least for π close to 1 our results remain unchanged. Our results are therefore robust to the introduction of imperfections in the certification technology.

6.6 Ex post renegotiation

In our pure adverse selection setting, we assumed that the buyer does not purchase the good if she finds out that the seller has quoted an inappropriately high price. Implicit in this assumption is the idea that the seller cannot renegotiate and lower that price; for instance because renegotiation is too costly or requires too much time.

To see that both our normative and positive result do not depend on the absence of such renegotiation, consider the other extreme where renegotiation is costless, so that after certifying a low quality good, renegotiation leads to trade at the price $p = q_l$. In this case, the low quality seller has always an incentive to initially quote a high price, because renegotiation ensures him that he can trade at the lower price even when the buyer certifies. Hence, ex post renegotiation raises the seller's cheating incentives under buyer-certification. This induces the buyer to increase her certification intensity and, in comparison to seller-certification, this leads to excessive certification and lower social welfare. Our normative welfare result is therefore robust to renegotiation. To see that this implies that also our positive result of higher certification profits under seller-certification is robust, recall that seller-certification is trade-effective. The certifier is therefore able to appropriate the entire increase in aggregate surplus from certification, which as we just argued is larger under seller-certification.

[6] See Strausz (2010) for a formal derivation of this result.

6.7 Certification informative to the seller

In line with the literature, we modelled certification as a technology that does not generate new, but only verifies the seller's private information. We can extend our results also to settings where certification generates additional information.

To make this claim more precise, let the value of the good be still either q_h or q_l, but assume that the seller in stage $t = 2$ receives only an imperfect private signal $\theta \in \{h, l\}$ about the good's actual value in that the signal is only correct with probability $\pi \geq 1/2$. Bayes' rule then implies that upon receiving the signal $i \in \{h, l\}$, the seller learns that the good has value q_i with probability λ_i, where

$$\lambda_h = \Pr\{q_h|\theta = h\} = \frac{\lambda\pi}{\lambda\pi + (1-\lambda)(1-\pi)}; \; \lambda_l = \Pr\{q_h|\theta = l\} = \frac{\lambda(1-\pi)}{\lambda(1-\pi) + (1-\lambda)\pi}.$$

In other words, upon receiving the signal $i \in \{h, l\}$ the seller privately learns that the buyer's expected value is $v_i = \lambda_i q_h + (1 - \lambda_i) q_l$ rather than $\lambda q_h + (1 - \lambda) q_l$. Now let certification be informative in that it reveals publicly the true value of the good, q_i, rather than only the seller's private information, v_i. Since our framework obtains for $\pi = 1$, this setup extends it.

One may show that the equilibrium outcome under seller-certification changes as follows: the certifier sets a price $p_c = \lambda_h(q_h - c_h)$, the θ_h seller certifies and sets a price $p = q_h$, whereas the θ_l seller sells the good uncertified at a price v_l. The equilibrium outcome is therefore continuous in π.

One may further show that for buyer-certification our reasoning determining $\tilde{p}$ and $\tilde{\mu}$ does not change. Propositions 2 and 3 then extend naturally. In particular, for the relevant case $\lambda < \tilde{\mu}$ and $c_h < \tilde{p}$, the seller θ_h picks the price $\tilde{p}$ with probability 1, seller θ_l picks the price $\tilde{p}$ with some probability $\sigma_l > 0$ and v_l with $1 - \sigma_l > 0$, while the buyer always buys when observing the price $p = v_l$ and certifies with probability $\sigma^* > 0$ when observing the price $p = \tilde{p}$. The equilibrium outcome under buyer-certification is therefore also continuous in π.

Hence, the equilibrium outcomes of seller- and buyer-certification do not change qualitatively and are, in particular, continuous in π. Consequently, our results are robust to the consideration of informative certification, because by continuity they are guaranteed to hold for $\pi < 1$ close to 1.

6.8 Buyer's private information

Our results also do not depend crucially on the absence of private information on buyer's side. To make this claim more precise, suppose that the buyer has private information about her willingness to pay for the product. In particular, let buyer type $\theta \in \{\theta_l, \theta_h\}$ value the quality q_i with θq_i, where $0 \leq \theta_l \leq \theta_h = 1$. Let θ be private information of the buyer, whereas the seller and certifier share the common belief $\nu \in (0,1)$ that the buyer has the higher valuation $\theta_h = 1$.

In this formulation θ_l measures the degree of private information. For the case $\theta_l = \theta_h$ private information is meaningless and the equilibrium coincides with the model without asymmetric information. The extension is therefore especially interesting with θ_l small.

Yet for small $\theta_l \geq 0$, one can formally show that there exists a threshold $\nu(\theta_l)$ below and bounded away from 1 such that for any $\nu \in [\nu(\theta_l), 1]$, the equilibrium outcome under seller-certification is such that the low quality seller charges a price q_l and sells only to the θ_h-buyer, whereas the high quality seller certifies at a price $p_c = \nu(q_h - c_h)$ and sells only to the θ_h-buyer at a price q_h.

Similarly, the equilibrium outcome under buyer-certification extends straightforwardly for any $\nu \in [\nu(\theta_l), 1]$ in the following sense. The seller and the θ_h-buyer behave in the same way as in the equilibrium of our model without private information, whereas in this equilibrium the θ_l-buyer does not certify nor buy the product.

Comparing outcomes, we once again obtain that seller-certification yields a higher social welfare than buyer-certification, because welfare under both modes of certification decreases proportionally to ν. Moreover, the certifier's profits are necessarily higher under seller-certification for $\nu < 1$ but close enough to 1. Hence, for $\nu < 1$ close to 1 our results are robust to the introduction of asymmetric information on part of the buyer, even when the difference $\theta_h - \theta_l$ in buyer valuation is large.

6.9 More general contracts

Starting from an industrial organization perspective, we naturally assumed that the buyer and the seller, as well as the certifier can only use prices rather than more sophisticated contracts to coordinate their exchange. Indeed, in many contexts explicit regulation limits the ability of certifiers to use general contracts. For instance, the European Securities and Markets Authority (ESMA) explicitly forbids credit rating agencies to charge discriminatory fees or base them on any form of contingency.[7] We may nevertheless investigate the theoretical question whether more complicated

[7] EU Regulation No 462/2013 (31.5.2013) states "A credit rating agency shall ensure that fees charged to its clients for the provision of credit rating and ancillary services are not discriminatory and are based on actual costs. Fees charged for credit rating services shall not depend on the level of the credit rating issued by the credit rating agency or on any other result or outcome of the work performed." Such regulatory restrictions are often justified by the claim that uniform certification prices give the certifier less incentives to manipulate certification outcomes.

contracts, such as prices that condition on the certification outcome, can change our ranking between seller- and buyer-certification.

An extreme way to investigate this is to allow the certifier to offer a general profit maximizing mechanism towards seller or buyer. From this perspective, the two modes of certification models studied here are but two possible mechanisms which the certifier may then use. The question is then whether there exist better mechanisms than the seller-certification one.

The answer to this question depends crucially on the commitment power underlying the mechanism. In the extreme where commitment to any possible mechanism and, in particular, random mechanisms, is possible, the certifier can achieve the first best arbitrarily closely. By these mechanisms, the seller is asked to report his quality, while the certifier commits to certify with arbitrary small probability and punish the seller arbitrarily harshly when certification does not confirm the announced quality.

The applicability of these mechanisms, however, is subject to the critique that the randomization is not sequentially rational: In equilibrium the seller always reports honestly, and thus there is an incentive to manipulate the randomization outcome so as not to certify, in order to save on the costs of certification. The mechanism above therefore depends crucially on the ability of the parties to commit to an honest randomization, even though the certification is suboptimal in equilibrium.

If the certifier cannot commit to a random mechanism, then it is relatively straightforward to show that the certifier cannot improve on seller-certification. Hence, seller-certification implements indirectly the optimal deterministic mechanism.

6.10 The ceteris paribus assumption

After arguing that our results are robust to all these extensions, we wish to emphasize that they were obtained under the assumption that, apart from the question who initiates certification, all other things are equal. Clearly, our results can be overturned if this condition is not met. For instance, if for some reason the cost of certification c_c is substantially higher under seller-certification than under buyer-certification, then buyer-certification is optimal both from a profit maximization and a welfare perspective. Moreover, concerns of certifier capture or more certifier credibility may depend on the mode of certification. Depending on their direction, such additional differences may strengthen or dampen the effects we identified here, and thereby explain why in some cases we may see the 'buyer pays' model implemented.

7 Applications

In addition to the two examples given in the introduction, a further class of economic transactions for which our model applies particularly well involves situations in which certification is both product and customer specific.[8] A specific example is parts-procurement in the automotive industry.[9] The development and production of a complex part for a premium automobile is typically done by only one supplier — our seller, whom the automotive producer — our buyer — selects explicitly. Because the part supplied is customer specific, the buyer-seller relationship after the buyer's selection is a bilateral monopoly.

Due to significant economies of scope involving the analytical instruments, the certification industry is highly concentrated.[10] Key test criteria are the functionality of parts (measured in failure of parts per million) and safety norms, characteristics about which the seller as the producer will typically possess private information. As it turns out, in about 80 per cent of all cases the testing of car modules and systems is performed on the request of the upstream supplier rather than the buyer — and if performed on request of the buyer it is paid by the seller. Moreover, the buyer conditions her purchase decision on the outcome of the certification process. Our model, therefore, captures the typical procurement relationship in the automotive industry, and our equilibrium result is consistent with the observations in this industry.

Similarly interesting and important examples belonging to this class are parts procurement in the aerospace industry, or the provision of building construction services, where again functionality and safety considerations figure prominently.

Whereas our model applies directly to cases in which certification is both product and customer specific, the results also help understanding purely product specific certification. Examples range from the certification of foodstuff in terms of production without herbicides or pesticides; to the certification of toys in terms of production without aggressive chemicals, to the certification of building materials, or ecologically

[8]Headwaters (2012) values this Testing, Inspection and Certification (TIC) sector at 100 billion euro (125 billion dollar) for 2012.

[9]The evidence is taken from Mueller et al. (2008), and from a large scale study conducted in 2007/08 by Stahl et al. for the German Association of Automotive Manufacturers (VDA) on Upstream Relationships in the Automotive Industry. Survey participants were car producers and their upstream suppliers. All German car producers and 13 first tier counterparts were questioned as to their procurement relationships. A description of the data base is found in Koenen et al.(2011)

[10]An example is EDAG, an engineering company centering on the development and prototype-construction of cars, as well as on independent certification of car modules and systems. In this function it serves all major car producers world wide. See http://www.edag.de/produkte/prueftechnik/automotive/index_html

correct inputs into the production of particular products, or of the fire-resistance of safes.[11] An example close to our academic activity is the certification process induced by the editors of academic journals, on request of the producers of research papers.

There are counterexamples as well, however, with mixed evidence between 'seller pays' and 'buyer pays' certification. Both models are used in real estate and in used car markets. One reason seems to be that certifiers in these markets are not all trustworthy, which violates our credibility assumption on certifier behavior.

Once again, a particularly timely and controversially discussed application is the certification of financial products. Certification is produced in a heavily concentrated rating industry.[12] Many buyers of financial products have admitted that they poorly understood the products' complexities. This underscores the opaqueness of these markets, and the importance of the rating agencies' role in promoting transparency. Before the crisis, and consistent with our result on the superiority of seller-certification, rating agencies used the 'issuer pays' model. A controversial claim is that such seller-certification led to certifiers' capture, and inflated ratings precipitated the financial crisis. Proponents of this claim, therefore, argue for a regulatory response to transfer the rating decision from sellers to buyers. Due to the superior welfare properties of seller-certification, however, our results caution against regulatory pressure in favor of buyer-certification.

8 Conclusion

In a market with opaque product quality, demand for certification to raise market transparency arises from both buyers and sellers. We provide new, elementary insights into the economic role of such third party certification by examining how the certification model affects transparency and market outcomes. In particular, we show that sellers use certification as a device to signal their quality. In contrast, buyers use certification as an inspection device to safeguard themselves against low quality sellers. Due to these differences, the 'seller pays' model is more effective in raising market transparency than the 'buyer pays' model. As a result, it also generates larger gains of trade, more social welfare, and higher profits to the certifier.

Our analysis leads to a clear policy implication, that is relevant especially in the ongoing discussion about patent certification, and certification in financial markets.

[11] Our certification setup captures both product and process certification. The distinction being that product certification regards inspecting the final good, while under process certification, the certifier inspects the quality of the production process.

[12] That concentration is reinforced by the regulatory decree. See White (2010).

Against a current argument about changing to the 'buyer pays' model, we point out that, based on elementary but fundamental differences in their economic use, the 'seller pays' model has a natural advantage over the 'buyer pays' one.

Appendix

The appendix contains all formal proofs to our Lemmata and Propositions.

Proof of Lemma 1: Consider the subgame $\Gamma^s(p_c)$ with $p_c \leq q_h - c_h$. Let the q_h-seller's strategy be the pure strategy $\sigma_h^c(q_h) = 1$, and the q_l-seller be the pure strategy $\sigma_l^u(q_l) = 1$. Moreover, let the Bayes' consistent buyer's belief satisfy $\mu(p) = 0$ for all p and let $\sigma(s_b|p, \mu)$ equal 1 if $p \leq q_l$ and zero otherwise. These strategies and beliefs describe a perfect Bayesian equilibrium of the game $\Gamma^s(p_c)$ with an outcome as described in the lemma.

To show uniqueness for $p_c < q_h - c_h$, note first that by certifying and charging the price $p = q_h$, the q_h-seller can guarantee himself a payoff $\pi_h^c \equiv q_h - c_h - p_c > 0$. Hence, in any equilibrium of the subgame $\Gamma^s(p_c)$ the q_h-seller must obtain a payoff of at least $\pi_h^c > 0$. Moreover, if the q_h-seller always certifies, he obtains the payoff π_h^c only if charging a price $p = q_h$. Hence, given that the q_h-seller always certifies, the equilibrium outcome is unique. We next show that there does not exist a perfect Bayesian equilibrium where the q_h-seller certifies with a probability less than 1. For suppose such an equilibrium would exist, then there exist prices $\tilde{p}$ such that the high quality seller does offer the good uncertified with positive probability, i.e. $\sigma_h^u(\tilde{p}) > 0$. For $\tilde{p}$ to be an equilibrium price, the associated profits to the q_h-seller, $\pi_h^u(\tilde{p})$, must at least match $\pi_h^c > 0$. Hence, at any such price $\tilde{p}$, the buyer must buy with positive probability: $\sigma(s_b|\tilde{p}, \mu(\tilde{p})) > 0$. This however requires that $\mu(\tilde{p})q_h + (1 - \mu(\tilde{p}))q_l \geq \tilde{p}$. This implies that $\mu(\tilde{p}) > \lambda$, for if not, then $\tilde{p} \leq \mu(\tilde{p})q_h + (1-\mu(\tilde{p}))q_l < \lambda q_h + (1-\lambda)q_l < c_h$, so that the high quality seller would not want to offer his product at price $\tilde{p}$. Hence, by (1), it must hold that $\sigma_h^u(\tilde{p}) > \sigma_l^u(\tilde{p})$ for each price $\tilde{p}$ such that $\sigma_h^u(\tilde{p}) > 0$. Adding over all such prices, we get the contradiction

$$1 \geq \sum\nolimits_{\tilde{p}:\sigma_h^u(\tilde{p})>0} \sigma_h^u(\tilde{p}) > \sum\nolimits_{\tilde{p}:\sigma_h^u(\tilde{p})>0} \sigma_l^u(\tilde{p}) = 1,$$

where the last equality follows, because if the q_l-seller picks a price $\bar{p}$ with $\sigma_h^u(\bar{p}) = 0$, then by (1), $\mu(\bar{p}) = 0$ so that either $\sigma(s_b|\bar{p}, \mu(\bar{p})) = 0$ or $\bar{p} \leq q_l$. In either case, the profits to the q_l-seller are less than from a price $\tilde{p}$ such that $\sigma_h^u(\tilde{p}) > 0$, because for such a $\tilde{p}$, $\pi_l^u(\tilde{p}) = \pi_h^u(\tilde{p}) + c_h \geq \pi_h^c + c_h \geq c_h > \bar{q} > q_l$.

For a subgame with $p_c > q_h - c_h$, the q_h-seller cannot obtain a profit from certification, because after certification, he can sell the good at a price of at most q_h, which yields the negative payoff, since $q_h - p_c - c_h < 0$. Consequently, an equilibrium in which the q_h-seller certifies with positive probability does not exist, because he is better off not offering his good to the market at all. Due to the lemon problem, an equilibrium where the q_h-seller offers his good uncertified does not exist. Such an equilibrium would have a price of at most $\bar{q}$, which exceeds the seller's production costs. Q.E.D.

Proof of Proposition 1: An equilibrium in which the certifier obtains a profit strictly less than $\lambda(q_h - c_h - c_c)$ does not exist, because, by Lemma 1, the certifier can guarantee itself a payoff arbitrarily close to $\lambda(q_h - c_h - c_c)$ by setting a price p_c slightly below $q_h - c_h$. Hence, if an equilibrium exists, it must exhibit $\Pi_c^s = \lambda(q_h - c_h - c_c)$. This profit is attainable only if the certifier sets a price of certification $p_c = q_h - c_h$ and the q_h-seller always certifies. According to Lemma 1 this is indeed an equilibrium outcome of the subgame $\Gamma^s(q_h - c_h)$. Q.E.D.

Proof of Lemma 2: Let buyer-certification be trade-effective. Then for any $\sigma_l(p) > 0$ it must hold $\sigma(s_b|p, \mu(p)) = 1$; for any price which the q_l-seller picks with positive probability, the buyer must buy with probability 1. Moreover, for any $\sigma_h(p) > 0$ it must hold $\sigma(s_n|p, \mu(p)) = 0$; for any price which the q_h-seller picks with positive probability, the buyer may not refrain from buying.

Suppose to the contrary that buyer-certification is information-effective. This means that for any $\sigma_i(p) > 0$ it holds $\mu(p) \in \{0, 1\}$; for any price that is picked with positive probability, the buyer learns the good's quality perfectly. It follows that the buyer will never certify, so that $\sigma(s_h|p, \mu(p)) = 0$. Hence, if buyer-certification is both information- and trade-effective, then for any $\sigma_h(p) > 0$ chosen by the seller, the buyer chooses $\sigma(s_b|p, \mu(p)) = 1$. But then the seller's strategy is suboptimal if there exist $p_1 \neq p_2$ such that $\sigma_h(p_1) > 0$ and $\sigma_l(p_2) > 0$ (because the higher price yields either seller type the highest profit). Hence for any $\sigma_h(p) > 0$ it must hold $\sigma_l(p) > 0$. But then (3) implies $\mu(p) \in (0, 1)$, which contradicts $\mu(p) \in \{0, 1\}$. Hence, buyer-certification cannot be both trade- and information-effective. Q.E.D.

Proof of Lemma 3 Consider any $p_c \geq 0$ and suppose to the contrary that buyer-certification is trade-effective so that $\Gamma^b(p_c)$ has an equilibrium where trade takes place with probability 1. Let $P_l = \{p|\sigma_l(p) > 0\}$ denote the set of prices that the q_l-seller charges with positive probability in this equilibrium. Trade-effectiveness implies that for any $p \in P_l$, we must have $\sigma(s_b|p, \mu(p)) = 1$. Let $P_h = \{p|\sigma_h(p) > 0\}$ denote the set of prices that the q_h-seller charges with positive probability in equilibrium. For any $p \in P_h$, we must have $p \geq c_h$, because otherwise the q_h-seller makes a loss from offering this p. Now suppose $P_l \cap P_h \neq \emptyset$ and let $\hat{p}$ denote the highest price in $P_l \cap P_h$. Then trade-effectiveness implies $\sigma(s_b|\hat{p}, \mu(\hat{p})) = 1$ and $\hat{p} \geq c_h$. Hence, the q_l-seller obtains an equilibrium profit of at least $\sigma(s_b|\hat{p}, \mu(\hat{p}))\hat{p} = \hat{p}$. The set P_l, therefore, cannot contain a price below the highest price $\hat{p}$. Hence, if $P_l \cap P_h \neq \emptyset$, the set P_l contains only one element. But then, $\sigma_l(\hat{p}) = 1 \geq \sigma_h(\hat{p})$ so that (3) implies that $\mu(\hat{p}) \leq \lambda$. But then $\mu(\hat{p})q_h + (1 - \mu(\hat{p}))q_l < c_h \leq \hat{p}$ so that we get the contradiction

$\sigma(s_b|\hat{p}, \mu(\hat{p})) = 0$. Hence, if a trade-effective equilibrium exists then $P_l \cap P_h = \emptyset$. But it then follows that for any $p \in P_h$ we have $\sigma_l(p) = 0$ so that (3) implies $\mu(p) = 1$ and, hence, $\sigma(s_b|p, \mu(p)) = 1$. Moreover, since $P_l \cap P_h = \emptyset$ we have for any $p_l \in P_l$ and $p_h \in P_h$ either $p_l < p_h$ or $p_l > p_h$. If $p_l < p_h$ then p_l yields the q_l-seller less than p_h (because as established $\sigma(s_b|p, \mu(p)) = 1$) so that we obtain the contradiction that $\sigma_l(p_l) > 0$ is not part of an optimal strategy. Likewise, if $p_l > p_h$ the price p_l yields the q_h-seller strictly more than p_h and, hence, we obtain the contradiction that $\sigma_h(p_h) > 0$ is not part of an optimal strategy. Q.E.D.

Proof of Lemma 4: To show that $\pi_h(p, \mu|\sigma^*)$ is non-decreasing in μ we first establish that, in any PBE, $\sigma^*(s_n|p, \mu)$ is weakly decreasing in μ. Suppose not, then we may find $\mu_1 < \mu_2$ such that $0 \leq \sigma^*(s_n|p, \mu_1) < \sigma^*(s_n|p, \mu_2) \leq 1$. $\sigma^*(s_n|p, \mu_2) > 0$ implies hat $(p, \mu_2) \in S(s_n|p_c)$ and, consequently,

$$p \geq \mu_2 q_h + (1 - \mu_2) q_l \tag{9}$$

and

$$p_c \geq \mu_2 (q_h - p). \tag{10}$$

Now since $\sigma^*(s_n|p, \mu_1) < 1$ we have either $\sigma^*(s_b|p, \mu_1) > 0$ or $\sigma^*(s_h|p, \mu_1) > 0$. Suppose first $\sigma^*(s_b|p, \mu_1) > 0$, then $(p, \mu_1) \in S(s_b|p_c)$, which implies $p \leq \mu_1 q_h + (1 - \mu_1) q_l$. But from $\mu_2 > \mu_1$ and $q_h > q_l$ it then follows that $\mu_2 q_h + (1 - \mu_2) q_l > p$, which contradicts (9). Suppose therefore that $\sigma^*(s_h|p, \mu_1) > 0$, then $(p, \mu_1) \in S(s_h|p_c)$, which implies $\mu_1(q_h - p) \geq p_c > 0$. This requires $q_h > p$. But then, due to $\mu_2 > \mu_1$, we get $\mu_2(q_h - p) > p_c$, which contradicts (10). Hence, we establish that $\sigma^*(s_n|p, \mu)$ is weakly decreasing in μ and therefore $\sigma^*(s_b|p, \mu) + \sigma^*(s_h|p, \mu)$ must be weakly increasing in μ. This directly implies that $\pi_h(p, \mu|\sigma^*)$ is weakly increasing in μ.

Next we show that in any PBE $\sigma^*(s_b|p, \mu)$ is weakly increasing in μ. Suppose not, then we may find $\mu_1 < \mu_2$ such that $1 \geq \sigma^*(s_b|p, \mu_1) > \sigma^*(s_b|p, \mu_2) \geq 0$. Since $\sigma^*(s_b|p, \mu_1) > 0$, it holds $(p, \mu_1) \in S(s_b|p_c)$ and, consequently,

$$p \leq \mu_1 q_h + (1 - \mu_1) q_l \tag{11}$$

and

$$p_c \geq (1 - \mu_1)(p - q_l). \tag{12}$$

Now since $\sigma^*(s_b|p, \mu_2) < 1$ we have $\sigma^*(s_n|p, \mu_2) > 0$ or $\sigma^*(s_h|p, \mu_2) > 0$. Suppose first $\sigma^*(s_n|p, \mu_2) > 0$, then $(p, \mu_2) \in S(s_n|p_c)$, which implies $p \geq \mu_2 q_h + (1 - \mu_2) q_l$.

But due to $\mu_2 > \mu_1$ and $q_h > q_l$ we get $p > \mu_1 q_h + (1-\mu_1)q_l$. This contradicts (11). Suppose therefore that $\sigma^*(s_h|p,\mu_2) > 0$, then $(p,\mu_2) \in S(s_h|p_c)$, which implies $(1-\mu_2)(p-q_l) \geq p_c > 0$. This requires $p > q_l$. But then, due to $\mu_2 > \mu_1$, we get $(1-\mu_1)(p-q_l) > p_c$. This contradicts (12). Hence, $\sigma^*(s_b|p,\mu)$ must be weakly increasing in μ. This directly implies that $\pi_l(p,\mu|\sigma^*)$ is weakly increasing in μ. Q.E.D.

Proof of Lemma 5: i) For any $\bar{p} < q_l$, $\mu \in [0,1]$ we have $(\bar{p},\mu) \notin S(s_n|p_c)$, $(\bar{p},\mu) \notin S(s_h|p_c)$ and $(\bar{p},\mu) \in S(s_b|p_c)$. Hence, $\sigma^*(s_b|\bar{p},\mu^*(\bar{p})) = 1$. Now suppose for some $\bar{p} < q_l$ we have $\sigma_i^*(\bar{p}) > 0$. This would violate (6), because instead of charging $\bar{p}$ seller q_i could have raised profits by ε by charging the higher price $\bar{p}+\varepsilon < q_l$ with $\varepsilon \in (0, q_l - \bar{p})$. At $\bar{p}+\varepsilon < q_l$ the buyer always buys, because, as established, $\sigma^*(s_b|\bar{p}+\varepsilon,\mu) = 1$ for all μ and in particular for $\mu = \mu^*(\bar{p}+\varepsilon)$.

For any $\bar{p} > q_h$, $\mu \in [0,1]$ we have $(\bar{p},\mu) \in S(s_n|p_c)$, $(\bar{p},\mu) \notin S(s_h|p_c)$ and $(\bar{p},\mu) \notin S(s_b|p_c)$. Hence, $\sigma^*(s_n|\bar{p},\mu^*(\bar{p})) = 1$. Now suppose we have $\sigma_l(\bar{p}) > 0$. This would violate (6), because instead of charging $\bar{p}$, which due to $\sigma^*(s_n|\bar{p},\mu^*(\bar{p})) = 1$ leads to zero profits, seller q_l could have obtained strictly positive profits by charging the price $q_l - \varepsilon$, where $\varepsilon \in (0, q_l)$.

ii) Suppose to the contrary that $\delta = q_l - \pi_l^* > 0$. Now consider a price $p' = q_l - \varepsilon$ with $\varepsilon \in (0,\delta)$ then for any $\mu' \in [0,1]$ we have $(p',\mu') \in S(s_b|p_c)$ and $(p',\mu') \notin S(s_n|p_c) \cup S(s_h|p_c)$ so that we have $\sigma^*(s_b|p',\mu^*(p')) = 1$ and, therefore, $\pi_l(p',\mu^*(p')|\sigma^*) = p' > \pi_l^*$. This contradicts (6).

iii) For any p such that $\sigma_h^*(p) > 0$, we have $\pi_h^* = \pi_h(p,\mu^*(p)|\sigma^*) = [\sigma^*(s_b|p,\mu^*(p)) + \sigma^*(s_h|p,\mu^*(p))]p - c_h$. As argued in i), we have $\sigma^*(s_n|p,\mu) = 1$ for all $p > q_h$ and $\mu \in [0,1]$. Hence, $\pi_h(p,\mu|\sigma^*) = 0$ whenever $p > q_h$. But for any price $p \leq q_h$ we have $\pi_h(p,\mu|\sigma^*) \leq q_h - c_h$. Hence, it follows that $\pi_h^* \leq q_h - c_h$. Now suppose $\pi_h^* = q_h - c_h$. Then we must have $\sigma_h^*(q_h) = 1$ and $\sigma^*(s_b|q_h,\mu^*(q_h)) + \sigma^*(s_h|q_h,\mu^*(q_h)) = 1$. But, due to $\mu^*(q_h)(q_h - q_h) = 0 < p_c$, we have $(q_h,\mu^*(q_h)) \notin S(s_h|q_h)$ so that $\sigma^*(s_h|q_h,\mu^*(q_h)) = 0$. Hence, we must have $\sigma^*(s_b|q_h,\mu^*(q_h)) = 1$. This requires $(q_h,\mu^*(q_h)) \in S(s_b|p_c)$ so that we must have $\mu^*(q_h) = 1$. By (3), this requires $\sigma_l^*(q_h) = 0$. But since $\pi_l(q_h,1|\sigma^*) = \sigma^*(s_b|q_h,\mu^*(q_h))q_h = q_h$ we must, by (6), have $\pi_l^* \geq q_h$. Together with $\sigma_l^*(q_h) = 0$, it would require $\sigma_l^*(p) > 0$ for some $p > q_h$ and leads to a contradiction with i). Q.E.D.

Proof of Lemma 6: We first prove ii): Suppose to the contrary that $\delta \equiv \tilde{p} - c_h - \pi_h^* > 0$. Then, due to the countable number of equilibrium prices, we can find an out-of-equilibrium price $p' = \tilde{p} - \varepsilon$ for some $\varepsilon \in (0,\delta)$. Then for any belief $\mu' \in (p_c/(q_h - p'), 1 - p_c/(p' - q_l))$[13] we have $(p',\mu') \in S(s_h)$ and $(p',\mu') \notin S(s_n) \cup S(s_b)$.

[13]To see that $p_c/(q_h - p') < 1 - p_c/(p' - q_l)$ define $l(p) \equiv p_c/(q_h - p)$ and $h(p) \equiv 1 - p_c/(p - q_l)$. Then by the definition of $\tilde{p}$ we have $l(\tilde{p}) = h(\tilde{p})$. Moreover, for $q_l < p < q_h$ we have $l'(p) =$

Consequently, $\sigma^*(s_h|p', \mu') = 1$. Hence, $\pi_h(p', \mu'|\sigma^*) = p' - c_h = \tilde{p} - c_h - \varepsilon > \tilde{p} - c_h - \delta = \pi_h^*$ and $\pi_l(p', \mu'|\sigma^*) = 0 < q_l \leq \pi_l^*$. Therefore, by BR the buyer's equilibrium belief must satisfy $\mu^*(p') \geq \mu'$. By Lemma 4 it follows $\pi_h(p', \mu^*(p')|\sigma^*) \geq \pi_h(p', \mu'|\sigma^*) = \tilde{p} - c_h - \varepsilon > \pi_h^*$. This contradicts (6). Consequently, we must have $\pi_h^* \geq \tilde{p} - c_h$. To show i) note that for all $p < \tilde{p}$ and $\mu \in [0, 1]$ we have $\pi_h(p, \mu|\sigma) \leq p - c_h < \tilde{p} - c_h \leq \pi_h^*$ so that $\sigma_h(p) > 0$ would violate (6). Q.E.D.

Proof of Proposition 2: i): First we show that for $\tilde{\mu} > \lambda$ and $\tilde{p} > c_h$ there exists no pooling, i.e., there exists no price $\bar{p}$ such that $\sigma_h^*(\bar{p}) = \sigma_l^*(\bar{p}) > 0$. For suppose there does. Then, by Lemma 6.i, we have $\bar{p} \geq \tilde{p}$ and, by Lemma 5.i, we have $\bar{p} \leq q_h$. Yet, due to (3) we have $\mu^*(\bar{p}) = \lambda < \tilde{\mu}$ so that $q_l + \mu^*(\bar{p})\Delta q - \bar{p} < q_l + \tilde{\mu}\Delta q - \tilde{p} = 0$. Moreover, $\mu^*(\bar{p})(q_h - \bar{p}) < \tilde{\mu}(q_h - \tilde{p}) = p_c$. Therefore, $\sigma^*(s_n|\bar{p}, \mu^*(\bar{p})) = 1$ and, hence, $\pi_h(\bar{p}, \mu^*(\bar{p})) = 0$. As a result, $\sigma_h^*(\bar{p}) > 0$ contradicts (6), because, by Lemma 6.ii, $\pi_h^* \geq \tilde{p} - c_h > 0 = \pi_h(\bar{p}, \mu^*(\bar{p}))$.

Second, we show that for $\tilde{\mu} > \lambda$, we cannot have $\sigma_h^*(\bar{p}) > 0$ for some $\bar{p} > \tilde{p}$. Suppose to the contrary we find such a $\bar{p}$ then, by definition of $\tilde{p}$, we have $(\bar{p}, \mu) \notin S(s_h)$ for any $\mu \in [0, 1]$. Hence, $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) = 0$ so that $\pi_l(\bar{p}, \mu^*(\bar{p})) = \pi_h(\bar{p}, \mu^*(\bar{p})) + c_h$. From Lemma 6.ii it then follows $\pi_l(\bar{p}, \mu^*(\bar{p})) \geq \tilde{p}$ and, therefore, $\sum_{p \geq \tilde{p}} \sigma_l^*(p) = 1$. From $\bar{p} > \tilde{p}$ and $\tilde{\mu} > \lambda$ it follows $\lambda\Delta q + q_l - \bar{p} < \tilde{\mu}\Delta q + q_l - \tilde{p} = 0$ so that $\lambda\Delta q + q_l < \bar{p}$. Now suppose it also holds $\sigma_l^*(\bar{p}) > 0$ then, by Lemma 5.ii and (6), $0 < q_l \leq \pi_l^* = \pi_l(\bar{p}, \mu^*(\bar{p})|\sigma^*) = \sigma^*(s_b|\bar{p}, \mu^*(\bar{p}))\bar{p}$. This requires $\sigma^*(s_b|\bar{p}, \mu^*(\bar{p})) > 0$ and therefore $(\bar{p}, \mu^*(\bar{p})) \in S(s_b|p_c)$ and, hence, $\mu^*(\bar{p})\Delta q + q_l \geq \bar{p}$. Combining the latter inequality with our observation that $\lambda\Delta q + q_l < \bar{p}$ and using (3), it follows

$$\lambda\Delta q + q_l < \frac{\lambda\sigma_h^*(\bar{p})}{\lambda\sigma_h^*(\bar{p}) + (1 - \lambda)\sigma_l^*(\bar{p})}\Delta q + q_l,$$

which is equivalent to $\sigma_h^*(\bar{p}) > \sigma_l^*(\bar{p})$. Summing over all $p \geq \tilde{p}$ and using $\sum_{p \geq \tilde{p}} \sigma_l^*(p) = 1$ yields the contradiction $\sum_{p \geq \tilde{p}} \sigma_h^*(p) > 1$. Hence, we must have $\sigma_l^*(\bar{p}) = 0$ for any $\bar{p} > \tilde{p}$. But this contradicts $\sum_{p \geq \tilde{p}} \sigma_l^*(p) = 1$ and, therefore, we must have $\sigma_h^*(\bar{p}) = 0$ for all $\bar{p} > \tilde{p}$.

This second observation implies that if an equilibrium for $\tilde{\mu} > \lambda$ and $\tilde{p} > c_h$ exists then, by Lemma 6, it exhibits $\sigma_h^*(\tilde{p}) = 1$, $\pi_h^* = \tilde{p} - c_h$ and $\sigma^*(s_h|\tilde{p}, \tilde{\mu}) + \sigma^*(s_b|\tilde{p}, \tilde{\mu}) = 1$. We now show existence of such an equilibrium and demonstrate that any such equilibrium has a unique equilibrium outcome. If $\sigma_h^*(\tilde{p}) = 1$ then (3) implies that $\mu^*(\tilde{p}) = \tilde{\mu}$ whenever

$$\sigma_l^*(\tilde{p}) = \frac{\lambda(1 - \tilde{\mu})}{\tilde{\mu}(1 - \lambda)},$$

$p_c/(q_h - p)^2 > h'(p) = p_c/(p - q_l)^2 > 0$. Hence, $l(\tilde{p} - \varepsilon) < h(\tilde{p} - \varepsilon)$ for $\varepsilon > 0$ small. Since $p' = \tilde{p} - \varepsilon$ we have $l(p') < h(p')$.

which is smaller than one exactly when $\lambda < \tilde{\mu}$. By definition, $(\tilde{p}, \tilde{\mu}) \in S(s_h) \cap S(s_b)$ so that any buying behavior with $\sigma^*(s_h|\tilde{p}, \tilde{\mu}) + \sigma^*(s_b|\tilde{p}, \tilde{\mu}) = 1$ is consistent in equilibrium. In particular, $\sigma^*(s_b|\tilde{p}, \tilde{\mu}) = q_l/\tilde{p} < 1$ is consistent in equilibrium. Only for this buying behavior we have $\pi_l(q_l, 0) = q_l = \pi_l(\tilde{p}, \tilde{\mu})$ so that seller q_l is indifferent between price $\tilde{p}$ and q_l. The equilibrium therefore prescribes $\sigma_l^*(q_l) = 1 - \sigma_l^*(\tilde{p})$. Finally, let $\mu^*(q_l) = 0$ and $\sigma^*(s_b|q_l, \mu^*(q_l)) = 1$ and $\mu^*(p) = 0$ for any price p larger than q_l and unequal to $\tilde{p}$. This out-of-equilibrium beliefs satisfies BR.

ii) In order to show that, in any equilibrium of $\Gamma^b(p_c)$, we have $\Pi^b(p_c) = 0$ whenever $\lambda > \tilde{\mu}$, we prove that for any $\bar{p}$ such that $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$, it must hold $\sigma_h^*(\bar{p}) = \sigma_l^*(\bar{p}) = 0$. Suppose we have $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$, then $(\bar{p}, \mu^*(\bar{p})) \in S(s_h)$ and, necessarily, $\bar{p} \leq \tilde{p}$. But by Lemma 6.i, $\sigma_h^*(\bar{p}) > 0$ also implies $\bar{p} \geq \tilde{p}$. Therefore, we must have $\bar{p} = \tilde{p}$. But $(\tilde{p}, \mu) \in S(s_h)$ only if $\mu = \tilde{\mu}$. Hence, we must have $\mu^*(\tilde{p}) = \tilde{\mu}$. By (3) it therefore must hold

$$\tilde{\mu} = \mu^*(\tilde{p}) = \frac{\lambda \sigma_h^*(\tilde{p})}{\lambda \sigma_h^*(\tilde{p}) + (1-\lambda)\sigma_l^*(\tilde{p})}.$$

For $\lambda > \tilde{\mu}$ this requires $\sigma_h^*(\tilde{p}) < \sigma_l^*(\tilde{p}) \leq 1$ and therefore there is some other $p' > \tilde{p}$ such that $\sigma_h^*(p') > 0$. But if also p' is an equilibrium price, then $\pi_h(\tilde{p}, \mu^*(\tilde{p})|\sigma^*) = \pi_h(p', \mu^*(p')|\sigma^*)$. Yet, for any $p' > \tilde{p}$ it holds $(p', \mu) \notin S(s_h|p_c)$ for any $\mu \in [0, 1]$ so that $\pi_l(p', \mu|\sigma^*) = \pi_h(p', \mu|\sigma^*) + c_h$ and, together with our assumption $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$ yields $\pi_l(\bar{p}, \mu^*(\bar{p})|\sigma^*) < \pi_h(\bar{p}, \mu^*(\bar{p})|\sigma^*) + c_h = \pi_h(p', \mu^*(p')|\sigma^*) + c_h = \pi_l(p', \mu^*(p')|\sigma^*)$ so that, by (6), $\sigma_l^*(\bar{p}) = 0$. Since $\bar{p} = \tilde{p}$, this violates $\sigma_l^*(\tilde{p}) > \sigma_h^*(\tilde{p}) \geq 0$. As a result, $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$ implies $\sigma_h^*(\bar{p}) = 0$.

In order to show that we must also have $\sigma_l^*(\bar{p}) = 0$, assume again that $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$. We have shown that his implies $\sigma_h^*(\bar{p}) = 0$. Now if $\sigma_l^*(\bar{p}) > 0$ then, by (3), it follows $\mu^*(\bar{p}) = 0$. But then $q_l + \mu^*(\bar{p})\Delta q - \bar{p} - p_c = q_l - \bar{p} - p_c < q_l - \bar{p}$ so that $(\bar{p}, \mu^*(\bar{p})) \notin S(s_h)$, which contradicts $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$.

In order to show that $\tilde{p} < c_h$ implies $\Pi^b(p_c) = 0$ suppose, on the contrary that, $\Pi^b(p_c) > 0$. This requires that there exists some $\bar{p}$ such that $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$ and $\sigma_i^*(\bar{p}) > 0$ for some $i \in \{l, h\}$. First note that $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) > 0$ implies $\bar{p} \leq \tilde{p}$. Now suppose $\sigma_h^*(\bar{p}) > 0$ then $\pi_h(\bar{p}, \mu^*(\bar{p})|\sigma^*) = (\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) + \sigma^*(s_b|\bar{p}, \mu^*(\bar{p})))\bar{p} - c_h < 0$ so that the high quality seller would make a loss and, thus, violates (6). Therefore, we have $\sigma_h^*(\bar{p}) = 0$. Now if $\sigma_l^*(\bar{p}) > 0$ then (3) implies $\mu^*(\bar{p}) = 0$ so that $\sigma^*(s_h|\bar{p}, \mu^*(\bar{p})) = 0$, which contradicts $\Pi^b(p_c) > 0$. Q.E.D.

Proof of Proposition 3: In order to express the dependence of $\tilde{\mu}$ and $\tilde{p}$ on p_c explicitly, we write $\tilde{\mu}(p_c)$ and $\tilde{p}(p_c)$, respectively. We maximize expression (7) with respect to p_c over the relevant domain

$$P = \{p_c | p_c \leq \Delta q/4 \wedge \tilde{\mu}(p_c) \geq \lambda \wedge \tilde{p}(p_c) \geq c_h\}.$$

First, we show that (7) is increasing in p_c. Define

$$\alpha(p_c) \equiv \frac{\lambda(\tilde{p}(p_c) - q_l)}{\tilde{\mu}(p_c)\tilde{p}(p_c)}$$

so that $\Pi_c(p_c) = \alpha(p_c)(p_c - c_c)$. We have

$$\alpha'(p_c) = \frac{4\lambda\Delta q^2}{\sqrt{\Delta q(\Delta q - 4p_c)}\left(q_h + q_l + \sqrt{\Delta q(\Delta q - 4p_c)}\right)^2} > 0$$

so that $\alpha(p_c)$ is increasing in p_c and, hence, $\Pi_c(p_c)$ is increasing in p_c and maximized for $\max P$.

We distinguish two cases. First, for $c_h \leq (q_h + q_l)/2$ it follows $1/2 = \Delta q/(2\Delta q) \geq (c_h - q_l)/\Delta q > \lambda$, where the last inequality follows from $c_h > \bar{q}$. From $\lambda < 1/2$, it then follows $\tilde{\mu}(p_c) \geq 1/2 \geq \lambda$. Therefore,

$$P = \{p_c | p_c \leq \Delta q/4 \wedge \tilde{p}(p_c) \geq c_h\}.$$

Hence, $\max P$ is either $p_c = \Delta q/4$ or such that $\tilde{p}(p_c) = c_h$. Because $\tilde{p}(\Delta q/4) = (q_h + q_l)/2$, it follows that for $c_h \leq (q_h + q_l)/2$, the maximum obtains at $p_c = \Delta q/4$ with

$$\Pi^b = \frac{\lambda \Delta q}{2(q_h + q_l)}(\Delta q - 4c_c).$$

Second, for $c_h > (q_h + q_l)/2$ the maximum obtains for p_c such that $\tilde{p}(p_c) = c_h$ in case $\lambda \leq 1/2$. This yields $p_c = (q_h - c_h)(c_h - q_l)/\Delta q$ with

$$\Pi^b = \frac{\lambda[(q_h - c_h)(c_h - q_l) - \Delta q c_c]}{c_h};$$

while for $\lambda > 1/2$ we have

$$\tilde{\mu}(p_c) \geq \lambda \Leftrightarrow p_c \leq \lambda(1 - \lambda)\Delta q.$$

Since $\lambda(1 - \lambda) \leq 1/4$ the requirement $p_c < \lambda(1 - \lambda)\Delta q$ automatically implies $p_c \leq \Delta q/4$. Hence for $\lambda > 1/2$ we have

$$P = \{p_c | p_c \leq \lambda(1 - \lambda)\Delta q \wedge \tilde{p}(p_c) \geq c_h\}.$$

Because, $\tilde{p}(\lambda(1 - \lambda)\Delta q) = \lambda q_h + (1 - \lambda)q_l$, which by assumption is smaller than c_h, we have $\max P = (q_h - c_h)(c_h - q_l)/\Delta q$. Note that $c_h > \lambda q_h + (1 - \lambda)q_l$ and $\lambda > 1/2$ implies that $c_h > (q_h + q_l)/2$. It follows $\tilde{\mu} = (c_h - q_l)/\Delta q$ and

$$\Pi^b = \frac{\lambda[(q_h - c_h)(c_h - q_l) - \Delta q c_c]}{c_h};$$

Q.E.D.

Proof of Proposition 4: For $c_h \leq (q_h + q_l)/2$ we have $\Pi^s = \lambda(q_h - c_h - c_c) \geq \lambda(q_h - c_h - c_c)\frac{q_h - q_l}{q_h + q_l} \geq \lambda(q_h - (q_h + q_l)/2 - c_c)\frac{q_h - q_l}{q_h + q_l} = \lambda(q_h - q_l - 2c_c)\frac{q_h - q_l}{2(q_h + q_l)} \geq \lambda(q_h - q_l - 4c_c)\frac{q_h - q_l}{2(q_h + q_l)} = \Pi^b$, where the second inequality uses $c_h \leq (q_h + q_l)/2$. Moreover, the certification intensity in the 'buyer pays' model is $x^b(\bar{p}_c^b) = x^b(\Delta q/4) = \lambda \Delta q / c_h$, which exceeds the certification intensity in the 'seller pays' model, λ, because due to $q_h - c_h - c_c > q_l$ it holds $\Delta q < c_h + c_c < c_h$.

For $c_h > (q_h + q_l)/2$ it follows $\Pi^b = \frac{\lambda[(q_h - c_h)(c_h - q_l) - \Delta q c_c]}{c_h} < \frac{\lambda[(q_h - c_h)(c_h - q_l) - (c_h - q_l)c_c]}{c_h} = \lambda(q_h - c_h - c_c)\frac{c_h - q_l}{c_h} \leq \lambda(q_h - c_h - c_c) = \Pi^s$, where the first inequality uses $q_h > c_h$. Moreover, $x^b(\bar{p}_c^b) = \lambda \frac{2\Delta q}{q_h + q_l}$, which is smaller than λ if and only if $q_h < 3q_l$. Q.E.D.

Proof of Proposition 5: For a combination of certification prices $(p_c^s, p_c^b) \in [c_c, \bar{p}_c^s] \times [c_c, \bar{p}_c^b]$, it follows

$$\begin{aligned}
\Delta W(p_c^b) &\equiv W^s - W^b(p_c^b) = (1 - \lambda)\omega(p_c^b)q_l + (x^b(p_c^b) - \lambda)c_c && (13)\\
&= \frac{\lambda}{\tilde{\mu}(p_c^b)\tilde{p}(p_c^b)}[(1 - \tilde{\mu}(p_c^b))(\tilde{p}(p_c^b) - q_l)(q_l + c_c) - \tilde{\mu}(p_c^b)q_l c_c] && (14)\\
&= \frac{\lambda}{\tilde{\mu}(p_c^b)\tilde{p}(p_c^b)}[(1 - \tilde{\mu}(p_c^b))\tilde{\mu}(p_c^b)\Delta q(q_l + c_c) - \tilde{\mu}(p_c^b)q_l c_c] && (15)\\
&= \frac{\lambda}{\tilde{p}(p_c^b)}[(1 - \tilde{\mu}(p_c^b))\Delta q(q_l + c_c) - q_l c_c] && (16)\\
&\geq \frac{\lambda}{\tilde{p}(p_c^b)}[(1 - \tilde{\mu}(c_c))\Delta q(q_l + c_c) - q_l c_c] && (17)\\
&= \frac{\lambda}{2\tilde{p}(p_c^b)}[(1 - \sqrt{1 - 4c_c/\Delta q})\Delta q(q_l + c_c) - 2q_l c_c] && (18)\\
&= \frac{\lambda}{2\tilde{p}(p_c^b)}\left[\Delta q(q_l + c_c) - 2q_l c_c - \sqrt{1 - 4c_c/\Delta q}\Delta q(q_l + c_c)\right], && (19)
\end{aligned}$$

where the inequality holds because $\tilde{\mu}$ is decreasing in p_c^b and $p_c^b \geq c_c$ if the certifier is not to make a loss. It remains to show that the term in the squared bracket is positive for any $c_c \in [0, q_h - c_c]$. That is, we need to show

$$\Delta q(q_l + c_c) - 2q_l c_c > \sqrt{1 - 4c_c/\Delta q}\Delta q(q_l + c_c).$$

To see this first note that the left hand side is indeed positive, since $\Delta q \geq 4c_c$ implies $\Delta q(c_c + q_l) > \Delta q q_l/2 \geq 2c_c q_l$. Squaring both sides yields

$$\Delta q^2(q_l + c_c)^2 - 4\Delta q(q_l + c_c)q_l c_c + 4q_l^2 c_c^2 > (1 - 4c_c/\Delta q)\Delta q^2(q_l + c_c)^2,$$

which is equivalent to $c_c \Delta q(q_l + c_c)c_c + q_l^2 c_c^2 > 0$, which is evidently true. As a result, $\Delta W(p_c^b) > 0$. Q.E.D.

Proof of Proposition 6: Follows from applying the same arguments as in the proof of Proposition 1 but with the high quality seller's outside option of $\Pi_h = q_l$ instead of $\Pi_h = 0$. The certifier therefore can at most ask for $\bar{p}_c^s = \Delta q - c_h$. Q.E.D.

Proof of Proposition 7: Mimics the arguments in the proof of Proposition 3 where the critical threshold for c_h is $\tilde{p} - q_l$ rather than $\tilde{p}$. Q.E.D.

Proof of Proposition 8: For $\lambda \leq 1/2$ and $c_h \leq \Delta q/2$, it follows

$$\begin{aligned}\Pi^b &= \frac{\lambda \Delta q}{2(q_h + q_l)}(\Delta q - 4c_c) < \lambda(\Delta q/2 - 2c_c) = \lambda(\Delta q - \Delta q/2 - 2c_c) \\ &\leq \lambda(\Delta q - c_h - 2c_c) < \lambda(\Delta q - c_h - c_c) = \Pi^s,\end{aligned}$$

where the first inequality uses $\Delta q < q_h + q_l$ and the second uses $c_h \leq \Delta q/2$.

For $\lambda > 1/2$ or $c_h > \Delta q/2$, it follows

$$\begin{aligned}\Pi^b &= \frac{\lambda[c_h(\Delta q - c_h) - \Delta q c_c]}{c_h + q_l} = \frac{\lambda[c_h(\Delta q - c_h - c_c) - (\Delta q - c_h)c_c]}{c_h + q_l} \\ &\leq \frac{\lambda c_h(\Delta q - c_h - c_c)}{c_h + q_l} \leq \lambda(\Delta q - c_h - c_c) = \Pi^s,\end{aligned}$$

where the first inequality follows because $\Delta q \geq c_h$. Q.E.D.

Literature

Akerlof, G., 1970, The Market for Lemons: Quality Uncertainty and the Market Mechanism, *Quarterly Journal of Economics*, 84, 488-500.

Albano, G. and A. Lizzeri (2001), Strategic Certification and Provision of Quality, *International Economic Review*, 42, 267-284.

Bester, H. and K. Ritzberger (2001), Strategic Pricing, Signalling, and Costly Information Acquisition, *International Journal of Industrial Organization*, 19, 1347-1361

Biglaiser, G. (1983), Middlemen as Experts, *The RAND Journal of Economics*, 24, 212-223.

Board, O. (2009), Competition and disclosure, *Journal of Industrial Organization* 57, 197-213

Boldrin, M. and D. K. Levine (2013), The Case Against Patents, *The Journal of Economic Perspectives* 27, 3-22

Bolton, P., X. Freixas and J. Shapiro (2012), The Credit Ratings Game, *The Journal of Finance*, 67, 85-111.

Cho, I-K. and D. Kreps (1987), Signaling games and stable equilibria, *Quarterly Journal of Economics* 102, 179-221.

Durbin, E. (1999), McDonald's or the Michelin Guide? Revealing Quality through Private-sector Certification, mimeo, Olin School of Business, Washington University

Dranove, D. and G. Jin (2010), Quality Disclosure and Certification, *Journal of Economic Literature*, 48, 935-963

Fasten, E. and D. Hofmann (2010), Two-sided Certification: The Market for Rating Agencies, mimeo, Humboldt University of Berlin

Faure-Grimaud, A., E. Peyrache, and L. Quesada (2009), The ownership of ratings, *The RAND Journal of Economics* 40, 234-257.

Felli, L., J. Koenen and K. Stahl (2011) Competition and Trust: Evidence from German Car Manufacturers, CEPR Discussion Paper 8265

Guo, L. and Y. Zhao (2009), Voluntary Quality Disclosure and Market Interaction, *Marketing Science* 28, 488-501.

Grossman, S. (1981), The informational role of warranties and private disclosure about product quality, *Journal of Law and Economics* 24, 461-83.

Grossman, S. and O. Hart (1980), Disclosure Laws and Takeover Bids, *Journal of Finance* 35, 323-334

Jovanovic, B. (1982), Truthful Disclosure of Information, *Bell Journal of Economics* 13, 36-44.

Koenen, J., L. Felli and K. Stahl (2011), Competition and Trust: Evidence from German Car Manufacturers, *CEPR DP 8265*

Lemley, M. (2001), Rational Ignorance at the Patent Office, *Northwestern University Law Review* 95, 1-34

Lizzeri, A., 1999, Information Revelation and Certification Intermediaries, *The RAND Journal of Economics* 30, 214-231.

Mathis, J., J. McAndrews, J.C Rochet, 2009, Rating the raters: Are reputation concerns powerful enough to discipline rating agencies?, *Journal of Monetary Economics* 56, 657-674.

Headwaters, 2012, Global Testing, Inspection and Certification M&A update. http://headwatersmb.com/rsrc/research/Global_Testing_Inspection_&_Certification.pdf (last retrieved 6/9/2014)

Mueller, F., Stahl, K. and F. Wachtler (2008), Upstream Relationships in the Automotive Industry: A Contractual Perspective, mimeo, University of Mannheim

Peyrache, E. and L. Quesada (2011), Intermediaries, Credibility and Incentives to Collude, *Journal of Economics and Management Strategy* 20, 1099-1133

Shavell, S. (1994), Acquisition and Disclosure of Information Prior to Sale, *The RAND Journal of Economics*, 25, 20-36

Skreta, V. and L. Veldkamp (2009), Ratings shopping and asset complexity: A theory of ratings inflation, *Journal of Monetary Economics* 56, 678-695.

Strausz, R. (2005), Honest Certification and the Threat of Capture, *International Journal of Industrial Organization*, 23, 45-62

Strausz, R. (2010), Separating Equilibria with Imperfect Certification, mimeo Humboldt Universität Berlin.

Viscusi, W. (1978), A Note on "Lemons" Markets with Quality Certification, *Bell Journal of Economics* 9, 277-279

White, L. (2010), Markets: The Credit Rating Agencies, *Journal of Economic Perspectives* 24, 211-226