Optimal Patent Examination

Dietmar Harhoff, Thomas Roende, Konrad Stahl and Roland Strausz*

February 27, 2011

Abstract

Intellectual property rights may be generated in registration or examination systems. In registration systems the right is held valid until it is revoked in an administrative or court challenge. In examination systems, the examination outcome determines validity. In the first system, validation operates as an inspection device, by which the potential licensee checks the quality of the patent. In the second, it is a signaling device, by which the patent holder signals the quality of his invention. we show that, under quite general conditions the examination system welfare dominates the registration system. The essential reason is that in the former, only valid inventions are entering, whilst in the latter, opportunism brings the inventor to register also patents of low validity that may undergo costly examination.

*Contact details: Dietmar Harhoff, University of Munich (harhoff@lmu.de), Thomas Roende, University of Copenhagen (thr.ino@cbs.dk), Konrad Stahl, University of Mannheim (kos@econ.uni-mannheim.de), Roland Strausz, Humboldt University Berlin (strauszr@wiwi.hu-berlin.de). This research was supported by the Deutsche Forschungsgemeinschaft (SFB-TR 15). We are grateful to seminar audiences at...., and in particular ... for very constructive comments.

1 Introduction

The patent system serves many purposes. A first, widely undisputed one is that a patent protects intellectual property from their exploitation without consent by the inventor, a protection needed to provide incentives to innovate. A second property is much more disputed, namely the certification of an invention's quality. In some countries such as the U.S. of France, obtaining a patent involves not much more than registering an invention. By contrast, in other countries such as Germany, the patent is supposed to indicate, if not certify the quality of that invention.

The validation of an invention is warranted by both its creator and its potential user. The *inventor* needs the validation if he wants to license or sell its use. The *user* needs the validation because he may want to acquire a license at an appropriate price. In fact, if he is a rival to the inventor, he may feel restrained in his business activity by his competitor's patent, and may seek its annulation when found of little worth in the validation process.

The underlying reason of all this is that the inventor typically has better information about its quality than a potential user, and that unless verified by a third party, he often cannot credibly convey the quality to that user. In the extreme, this precludes the emergence of markets for trading intellectual property rights, which is bad for the inventor. From a user's or licensee's, or a challenger's perspective, the validation of the invention resolves asymmetric information and allows the licensee to exploit the exclusion function of the patent. Thus, both the inventor and potential licensees have demand for the validation of the invention. The inventor asks for validation in order to obtain the unfettered right to exploit the invention at appropriately high returns – provided the invention is of high quality. The licensee or challenger asks for validation because she does not want to license in a low quality invention at a high price or wants to avoid costly inventing-around a patent of uncertain validity.

In principle, there are two institutions of independent expertise available to validate the invention: the patent office to which the inventor applies, or the courts to which challengers or licensees appeal to. We develop in parallel two models in which we consider separately the two forms of validation of the innovation.

In the first model involving, as we call it, a *registration regime*, the inventor's intellectual property right (IPR) becomes valid automatically by registration without examination. A potential licensee or rival has to decide whether, if a high license price is quoted by the inventor, she wants to take the patent to court for examination of its validity. She will buy the license at high price only if it is proven to be of high value.

In the second model involving an *examination regime*, the inventor files an application at a patent office that truthfully certifies the quality of his invention as "high" or "low". The IPR is held valid only in the case of high quality. The license for using an invention is bought by a potential licensee at a low price when the invention is not patented, and at a high price when it is. This is in line with the notion that the patent right generates a patent premium while the unpatented invention may have value in its own right (e.g., because the invention cannot be fully replicated even without IPRs due to some related trade secret).

[... further summary and comments here]

2 Model

To produce an invention, the inventor chooses between high effort c_h or low effort c_l. In the first case, the invention will have high quality q_h with probability τ and low quality $q_l > 0$ with probability $1 - \tau$. In the second case, the invention will always have low quality q_l. The quality of the invention, before its examination, is revealed only to the inventor.

Before the quality of the invention is revealed, the patent's quality is high, q_h, with probability λ, and low ,q_l, with probability $1 - \lambda$ from the rival's or licensee's point of view, where $\Delta q \equiv q_h - q_l$.[1] The user's willingness

[1]This implies that λ is bounded away from 1.

to pay is identified with the invention's quality. The user, assumed risk neutral, is therefore willing to pay up to a price that equals expected quality $\bar{q} \equiv \lambda q_h + (1-\lambda) q_l$. If not licensing in the invention, her reservation payoff is zero. If not producing and selling the inventor's reservation payoff is also zero.

We consider two extremely different regimes, within which the validation of inventions, and with it the market for inventions can be organized: First, a *registration regime*, in which the patent office simply registers any inventor's application for a patent. In this case, the patent may be validated via a challenge brought to court by the licensee or rival. Second, an *examination regime*, in which the patent officer validates the invention and conveys the quality assessment to the public. In the baseline version of our model, we assume that the cost $c_c \in [0, q_h - c_h)$ of validating the quality of the invention is the same at the patent office and at the court house, and that both examination processes are fully revealing.[2]

We assume that the high quality invention delivers higher economic rents: $q_h - c_h > q_l - c_l = q_l > 0$. Moreover, the cost of producing a high quality invention exceeds the average quality, $c_h > \bar{q}$. Hence the problem to the inventor: without validation by either the patent office or the court, he would not wage the effort to invent, and thus the market outcome with informational asymmetry preserved would be inefficient.

Without the informational asymmetry, however, the high quality inventor could profitably license his invention for the price $q_h > c_h$. Consequently, the high quality inventor has demand for validation that reveals the invention's true quality to the licensee. Clearly, the high quality inventor is willing to pay the patent office at most $q_h - c_h$.

Yet the user/licensee has also demand for the validation of the invention. Whenever the inventor quotes a license price higher than that appropriate for the low quality invention, the user has demand for validation that ascertains the invention to be indeed of high quality, so that a higher price is justified.

[2] In the extension section we look at more realistic alternatives in which the cost structures differ and, more importantly, the examination processes are not fully revealing.

Summarizing, both the inventor and the user/licensee have a demand for the validation of the invention. This opens the question as to how should the patent system be organized with respect to the validation services. In particular, should the patent office, rather than purely registering an innovation, offer the validation service to the inventor, or should the court offer that validation, if asked by the potential user?

Before we analyze the two alternatives, let us briefly consider realistic reinterpretations of the set up introduced so far. The inventor may be interested in using the invention himself, rather than licensing it out, and his counterpart - so far called the user - may be a competitor whose *freedom of operation is endangered* by the patent, so the jargon. Consider the first interpretation, and suppose that the inventor, if using the invention himself, can generate utility u_h if the invention is of high quality, and u_l if it is of low quality – but only if he holds the IPR for it. Otherwise the invention is considered of zero worth to him.[3] Then, as we will show later, the analysis applies in full.

Similarly, consider now the second interpretation. Suppose that the competitor is not willing to license in the patent at the price quoted by the inventor, and challenges the patent before the court. Two outcomes are possible: First, the patent is annulated. Then the challenger has paid p_c and, by virtue of the annulation, is free to pursue his economic activity worth some $\bar{u} \geq p_c$. Alternatively, if the courts certify the patent as of high value q_h, then there are again two alternatives: first, the challenger pays the price p quoted by the inventor; or second, he *invents around* the patent holder's invention, which generates an (expected) net return $\hat{u}$ that may be above or below p.[4]

or to *invent around* with some effort e the patented innovation in order to generate rents.

In the ensuing two sections, we model first the situation we call the *registration regime*, in which the patent office merely registers the invention,

[3]This could reflect a perfectly competitive market in which the invention would be propagated and imitated immediately.

[4]That net return may be larger than p because of the fact that inventing around removes the challenger's dependence on the licensor.on the paper.

and the would-be licensee, on the basis of the license price quoted by the inventor, decides whether or not to challenge the patent and to ask the court to validate the patent. We then model the situation we call the *examination regime*, in which the patent office validates the invention. Our research question is simple, yet important when it comes to policy: which of the two regimes leads to a higher social surplus, as measured by the sum of producer and consumer surpluses?

3 Inventor–Induced Validation

Here we consider validation when induced by the user/licensee via a court procedure. Before analyzing the formal model, it is helpful to provide an intuition on the role of validation in this setup.

User–induced validation enables the user to check the inventor's quality claim. In particular, the validation offers the user protection against paying a high license fee for a low quality invention of an inventor who pretends to own the IPR of a high quality invention. From the user's perspective, therefore, validation is an *inspection device* to detect low quality inventions.

The game underlying user–induced validation, therefore, is an inspection game. A mixed strategy equilibrium is typical for this type of game. Indeed, a pure strategy equilibrium in which the user never goes to court cannot exist, because it would give the low quality inventor an incentive to claim high quality – yet against this claim the user would have a strong incentive to have the invention validated. Likewise, an equilibrium in which the user always goes to court cannot exist either, because it keeps the low quality inventor from claiming high quality – yet against such behavior going to court is a costly and wasteful exercise for the user. Consequently, we typically have a mixed strategy equilibrium, where the inventor wants to license out a low quality invention at high price and thus cheats with some probability with the claim to offer high quality, and the user goes to court with some probability when the inventor indeed claims to have high quality.

Hence, user–induced validation plays the role of reducing cheating - yet not down to a zero probability event. The user's demand for validation will therefore be high when the problem of cheating is large. This reasoning suggests that a validation agent who targets his services towards the user will choose a price for its services that maximizes the user's cheating problem.

A closer look reveals that the user's cheating problem depends on two factors: the user's uncertainty and the inventor's price quotation. First, the user's cheating problem is the bigger the less certain she is about the true quality offered by the inventor. Second, checking true quality through the courts is especially worthwhile for intermediate license prices. Indeed, for a low price the user would not lose much from simply buying the invention without validation. By contrast, when the license price quoted by the inventor is high, the buyer would not lose much from not buying the invention at all, as the high price reduces her options to profitably use the license. Hence, the buyer's willingness to pay for a validation is largest for intermediate prices that are neither too low nor too high.

To sum up, our intuitive reasoning suggests that under user–induced validation the court will choose a price for the validation of the invention, p_c, so that it induces high uncertainty for the user and an intermediate license price. Below we show that this intuition is correct. Yet the formal analysis supporting it is not trivial at all.

With user–induced validation, the parties play the following game:

t=1 The court determines court costs p_c to handle the challenge, to be payable by the user, the plaintiff.

t=2 Nature selects the quality $q_i, i \in \{l, h\}$, of the invention, and conveys it to the inventor. The inventor automatically registers the invention at the patent office. It will be patented no matter its quality.

t=3 The inventor offering the invention of quality q_i at cost c_i decides about the price p at which he offers the invention to the potential licensee.

t=4 The user/licensee decides whether or not to demand validation of the invention, by challenging its quality in court.

t=5 Eventually on the basis of the court's verdict, the user decides whether or not to license the invention.

We focus on the Perfect Bayesian Equilibria (PBE) of the game described above. Note that after the court has set its price p_c, a proper sub-game, $\Gamma(p_c)$, starts with nature's decision about the quality of the invention. The sub-game $\Gamma(p_c)$ is a signalling game where the inventor's price p may or may not reveal his private information about the quality of the patent.

In the subsequent analysis, we first consider the PBE of the sub-game $\Gamma(p_c)$. A PBE specifies three components: First, the inventor's pricing strategy as a function of the quality q_i of the invention; second, the iser/licensee's belief $\mu(p)$ after observing the price p; and third, the user's behavior; in particular whether or not to challenge the patent, and whether or not to license it in.

We allow the inventor to randomize over prices. In order to circumvent measure–theoretical complications, we assume that the seller can randomize over an infinite but countable set. Consequently, we can express the inventor's pricing strategy for an invention of quality q_i by the function $\sigma_i : R_+ \to [0,1]$ with the interpretation that $\sigma_i(p_j)$ denotes the probability that the inventor endowed with quality q_i chooses the price p_j. Thus, for both $i \in \{h, l\}$,

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

The user's decisions are based on her belief specified as a function $\mu : \mathbb{R}_+ \to [0,1]$ with the interpretation that, after observing price p, the user believes that the invention is of type q_h with probability $\mu(p)$.

We can express the user's behavior after observing the price p and possessing some belief μ by the following six actions:

1. Action s_{nn}: The user does not challenge the patent nor license it. This action yields the payoff

$$U(s_{nn}|p, \mu) = 0.$$

2. Action s_{nb}: The user does not challenge the patent, but licenses it. This action yields the expected payoff

$$U(s_{nb}|p,\mu) = \mu q_h + (1-\mu)q_l - p.$$

3. Action s_{ch}: The user challenges the patent and licenses it only when the court reveals high quality. This action yields the expected payoff

$$U(s_{ch}|p,\mu) = \mu(q_h - p) - p_c.$$

4. Action s_{cb}: The user challenges the patent and licenses it irrespective of the outcome of certification. This action yields the expected payoff

$$U(s_{cb}|p,\mu) = \mu(q_h - p) + (1-\mu)(q_l - p) - p_c.$$

   Clearly, $U(s_{cb}|p,\mu) < U(s_{nb}|p,\mu)$ for any $p_c > 0$ so that the action s_{nb} dominates the action s_{cb}.

5. Action s_{cl}: The user challenges the patent and licenses it only when the court reveals low quality. This action yields the expected payoff

$$U(s_{cl}|p,\mu) = (1-\mu)(q_l - p) - p_c.$$

   Clearly, $U(s_{cl}|p,\mu) \leq U(s_{nb}|p,\mu)$ for $p \leq q_h$ and $U(s_{cl}|p,\mu) \leq U(s_{nn}|p,\mu)$ for $p > q_h$. Hence, also the action s_{cl} is weakly dominated.

6. Action s_{cn}: The buyer demands certification and does not buy the product. This action yields the expected payoff

$$U(s_{cn}|p,\mu) = -p_c.$$

   Clearly, $U(s_{cn}|p,\mu) < U(s_{nn}|p,\mu)$ for any $p_c > 0$ so that the action s_{cn} is dominated.

To summarize, only the first three actions s_{nn}, s_{nb}, s_{ch} are not (weakly) dominated for some combination (p,μ). The intuition is straightforward: the role of quality verification is to enable the would be licensee to discriminate between high and low quality. It is obviously only worthwhile to challenge the patent when the licensee uses it to screen out bad quality.[5]

In the following, we delete the weakly dominated actions from the user's action space. Consequently, we take the user's action space as $S \equiv \{s_{nn}, s_{nb}, s_{ch}\}$. Since we want to allow the user to use a mixed strategy, we let $\sigma(s|p, \mu) \in [0, 1]$ represent the probability that the user takes action $s \in S = \{s_{nn}, s_{nb}, s_{ch}\}$ given price p and belief μ. Thus

$$\sum_{s \in S} \sigma(s|p, \mu) = 1.$$

A PBE in our subgame $\Gamma(p_c)$ can now be described more specifically: it is a tuple of functions $\{\sigma_l, \sigma_h, \mu, \sigma\}$ satisfying the following three equilibrium conditions. First, inventor type i's pricing strategy σ_i must be optimal with respect to the user's strategy σ. Second, the user's belief μ must be consistent with the inventor's pricing strategy, whenever possible. Third, the user's strategy σ must be a best response given the observed price p and her beliefs μ.

We start our analysis of the Perfect Bayesian Equilibria of $\Gamma(p_c)$ by studying the third requirement: the optimality of the user's strategy given a price p and beliefs μ.

Fix a price $\bar{p}$ and a belief $\bar{\mu}$. Then the pure strategy s_{nn} is a best response whenever $U(s_{nn}|\bar{p}, \bar{\mu}) \geq U(s_{nb}|\bar{p}, \bar{\mu})$ and $U(s_{nn}|\bar{p}, \bar{\mu}) \geq U(s_{ch}|\bar{p}, \bar{\mu})$. It follows that the strategy s_{nn} is a best response whenever

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

Likewise, the pure strategy s_{nb} is (weakly) preferred whenever $U(s_{nb}|\bar{p}, \bar{\mu}, p_c) \geq U(s_{nn}|\bar{p}, \bar{\mu}, p_c)$ and $U(s_{nb}|\bar{p}, \bar{\mu}, p_c) \geq U(s_{ch}|\bar{p}, \bar{\mu}, p_c)$. It follows that the strategy s_{nb} is a best response whenever

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

[5]Observe that the strategy s_{ch} is not renegotiation proof, because even after certification has revealed low quality, gains could be realized by trading the low quality product. In Section 6, we will consider the simple extension to include renegotiation.

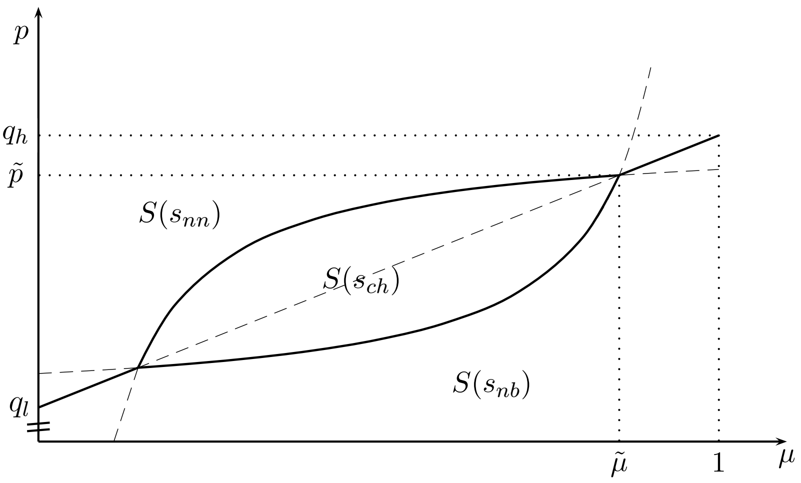

Figure 1: Buyer's buying behavior for given $p_c < \Delta q/4$.

Finally, the pure strategy s_{ch} is (weakly) preferred whenever $U(s_{ch}|\bar{p}, \bar{\mu}, p_c) \geq U(s_{nn}|\bar{p}, \bar{\mu}, p_c)$ and $U(s_{ch}|\bar{p}, \bar{\mu}, p_c) \geq U(s_{nb}|\bar{p}, \bar{\mu}, p_c)$. It follows that the strategy s_{ch} is a best response whenever

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

Since a mixed strategy is only optimal if it randomizes among those pure strategies that are a best response, we arrive at the following result:

Lemma 1 *In any Perfect Bayesian Equilibrium* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma(p_c)$ *we have for any* $s \in S = \{s_{nn}, s_{nb}, s_{ch}\}$,

$$\sigma^*(s|p, \mu) > 0 \Rightarrow (p, \mu^*(p)) \in S(s|p_c). \quad (1)$$

Figure 1 illustrates the user's behavior for a given price of the court case p_c. For low prices p the user licenses the patent uncertified, $(p, \mu) \in S(s_{nb})$, whereas for high prices p the user refrains from licensing, $(p, \mu) \in S(s_{nn})$. As long as $p_c < \Delta q/4$ there is an intermediate range of prices p and beliefs μ such that the user challenges the patent, i.e. $(p, \mu) \in S(s_{ch})$. In this case, the user only licenses the patent when the court reveals it to be of high quality. Intuitively, the user challenges the patent to ensure that it is worth the high price quoted by the inventor. Note that apart from points on the thick, dividing lines, the user's optimal behavior is uniquely determined, and mixing does not take place.

For future reference we define

$$\tilde{p} \equiv \left(q_h + q_l + \sqrt{\Delta q(\Delta q - 4p_c)}\right)/2$$

and

$$\tilde{\mu} \equiv \left(1 + \sqrt{1 - 4p_c/\Delta q}\right)/2.$$

Note that if the inventor prices at $\tilde{p}$ and the user has beliefs $\tilde{\mu}$, the user is indifferent between all three decisions namely not to license the patent, s_{nn}, to license the patent without challenging it, s_{nb}, or to license the patent only after it has been proven by the court as high quality, s_{ch}.

We previously argued that a profit maximizing court house benefits from high user uncertainty and an intermediate price of the patent. We now give precision to this. The user's willingness to pay for certification is the difference between her payoff from validation before the court and her next best alternative, namely either to license the patent unchallenged, or to not license it at all. More precisely, given her beliefs are μ, the difference in the user's expected payoffs between licensing the high quality patent when certified and licensing any good without quality assurance by the court is

$$\Delta U^1 \equiv \mu(q_h - p) - (\bar{q} - p).$$

Similarly, the difference in the user's expected payoffs between licensing the patent only when proven to be of high quality and not licensing the patent at all is

$$\Delta U^2 = \mu(q_h - p).$$

Hence, the user's willingness to pay for validation before the court is maximized for a price $\hat{p}$ and a belief $\hat{\mu}$ that solves

$$\max_{p,\mu} \min\{\Delta U^1, \Delta U^2\}.$$

The solution is $\hat{\mu} = 1/2$ and $\hat{p} = (q_h + q_l)/2$. We later demonstrate that under user–induced validation, the court chooses a price p_c for the validation to induce this outcome as closely as possible.

Next, we address the optimality of type i inventor's strategy $\sigma_i(p)$. For a given strategy σ of the user and a fixed belief μ, an inventor endowed with an invention of quality q_h expects the following payoff from setting a price p:

$$\Pi_h(p, \mu|\sigma) = [\sigma(s_{nb}|p, \mu) + \sigma(s_{ch}|p, \mu)]p - c_h.$$

Therefore, a specific strategy σ_h yields an inventor endowed with a high quality invention q_h an expected profit of

$$\bar{\Pi}_h(\sigma_h) = \sum_i \sigma_h(p_i)\Pi_h(p_i, \mu(p_i)|\sigma).$$

Likewise, a inventor endowed with an invention of quality q_l obtains the payoff

$$\Pi_l(p, \mu|\sigma) = \sigma(s_{nb}|p, \mu)p$$

and any strategy σ_l yields

$$\bar{\Pi}_l(\sigma_l) = \sum_i \sigma_l(p_i)\Pi_l(p_i, \mu(p_i)|\sigma).$$

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

$$\Pi_h^* = \sum_i \sigma_h^*(p_i)\Pi_h(p_i, \mu^*(p_i)|\sigma^*) \text{ and } \Pi_l^* = \sum_i \sigma_l^*(p_i)\Pi_l(p_i, \mu^*(p_i)|\sigma^*),$$

respectively.

The next lemma makes precise the intuitive result that the inventor's expected profits increase when the buyer has more positive beliefs about the good's quality.

Lemma 2 *In any PBE* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma(p_c)$ *with* $p_c > 0$ *the payoffs* $\Pi_h(p, \mu|\sigma^*)$ *and* $\Pi_l(p, \mu|\sigma^*)$ *are non–decreasing in* μ.

Inventor type i's pricing strategy σ_i is an optimal response to the user's behavior (σ^*, μ^*) exactly if, for any p', we have

$$\sigma_i^*(p) > 0 \Rightarrow \Pi_i(p, \mu^*(p)|\sigma^*) \geq \Pi_i(p', \mu^*(p')|\sigma^*). \tag{2}$$

Because the user's beliefs depend on the observed price p, it affects the user's behavior and, therefore, the belief function μ^* plays a role in condition (2).

Finally, a PBE demands that the user's beliefs μ^* have to be consistent with equilibrium play. In particular, they must follow Bayes' rule:

$$\sigma_i^*(p) > 0 \;\Rightarrow\; \mu^*(p) = \frac{\lambda \sigma_h^*(p)}{\lambda \sigma_h^*(p) + (1-\lambda)\sigma_l^*(p)}. \tag{3}$$

The next lemma shows some intuitive implications on PBEs that are due to Bayes' rule. In particular, it shows that the inventor, no matter his type, never sets a price below q_l, and the low quality inventor never sets a price above q_h. The lemma also shows that, in equilibrium, the low quality inventor 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 inventor loses from the presence of asymmetric information; his payoff is strictly smaller than $q_h - c_h$.

Lemma 3 *In any PBE* $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ *of the subgame* $\Gamma(p_c)$ *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 very weak restrictions on admissible beliefs. In particular, it does not place any restrictions on the patent user's beliefs for prices that are not played in equilibrium; in fact, any out–of–equilibrium belief is allowed. However, as is typical for signalling games, without any restrictions on out–of–equilibrium beliefs we cannot pin down behavior in the subgame $\Gamma(p_c)$ to a specific equilibrium. Especially by the use of pessimistic out–of–equilibrium beliefs, one can sustain many pricing strategies in equilibrium.

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. Bester and Ritzberger (2001) demonstrate that the following extension of the intuitive criterium of Cho–Kreps suffices to pin down equilibrium play.

Belief restriction (B.R.): A Perfect Bayesian Equilibrium $(\sigma_h^*, \sigma_l^*, \mu^*, \sigma^*)$ satisfies the Belief Restriction 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 contains the intuitive criterion of Cho–Kreps as the special case for $\mu = 1$. Indeed, the underlying idea of the restriction is to extend the idea behind the Cho–Kreps criterion to a situation where a deviation to p is profitable only for the q_h inventor when the user believes that the deviation originates from the q_h inventor with probability μ. As we may have $\mu < 1$, the restriction considers more pessimistic beliefs than the Cho–Kreps criterion. If such a pessimistic belief μ gives only the q_h seller an incentive to deviate, then the restriction requires that the user's actual belief should not be even more pessimistic than μ.

The next Lemma establishes characteristics of the equilibrium that are due to the belief refinement (B.R.). It shows that the belief restriction implies that the high quality seller can sell his product for a price of at least $\tilde{p}$.

Lemma 4 *Any Perfect Bayesian Equilibrium $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ of the subgame $\Gamma(p_c)$ that satisfies B.R. exhibits i) $\sigma_h^*(p) = 0$ for all $p < \tilde{p}$ and ii) $\Pi_h^* \geq \tilde{p} - c_h$.*

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

Proposition 1 *Consider a PBE $(\sigma_l^*, \sigma_h^*, \mu^*, \sigma^*)$ of the subgame $\Gamma(p_c)$ that satisfies B.R. Then*

i) for $\lambda < \tilde{\mu}$ and $c_h < \tilde{p}$ it exhibits unique pricing behavior by the inventor and unique buying behavior by the user. In particular, the high quality inventor sets the price $\tilde{p}$ with certainty, and the low quality inventor randomizes between the price $\tilde{p}$ and q_l. Observing the price $\tilde{p}$ the user buys certification with positive probability. The court's equilibrium profit equals

$$\Pi_c(p_c) = \frac{\lambda(\tilde{p} - q_l)}{\tilde{\mu}\tilde{p}}(p_c - c_c). \tag{4}$$

ii) For $\lambda > \tilde{\mu}$ *or* $c_h > \tilde{p}$ *we have* $\Pi_c(p_c) = 0$ *in any equilibrium.*

iii) For $\lambda \leq \tilde{\mu}$ *and* $c_h \leq \tilde{p}$ *there exists an equilibrium outcome, in which the court's profits equal expression* (4).

The Proposition shows that the user and the low quality inventor play the mixed strategies that reflect the typical outcome of an inspection game. Indeed, by choosing the low price q_l a low quality inventor honestly signals his low quality.

In contrast, we may interpret a low quality inventor, who sets a high price $\tilde{p}$, as trying to cheat. Hence, whenever the user observes the price $\tilde{p}$, she is uncertain whether the patent is supplied by the high quality or the low quality inventor. She therefore wants the patent inspected by going to court with positive probability. Through inspection, the user tries to dissuade the low quality inventor to set the "cheating" price $\tilde{p}$. Yet, as in an inspection game, the user has only an incentive to buy validation and inspect when the low quality inventor cheats "often enough".

This gives rise to the use of mixed strategies. As in an inspection game the user's probability of going to court is such that the low quality inventor is indifferent between cheating, i.e., setting the high price $\tilde{p}$, and honestly signaling his low quality by setting the price q_l. On the other hand, the probability with which the low quality inventor chooses the high price $\tilde{p}$ is such that the user is indifferent between licensing the patent uncertified, and asking the court for validation.

Proposition 1 also describes the court's profits in the subgame $\Gamma(p_c)$. The court anticipates this outcome when choosing its price p_c for validating the patent's quality. When the court maximizes its profits Π_c with respect to the validation price p_c, it must take into account that $\tilde{\mu}$ depends on p_c itself and the court therefore anticipates that the very case distinction $\lambda \lesseqgtr \tilde{\mu}$ and $c_h \gtreqless \tilde{p}$ depends on its choice of p_c. The following proposition shows that expression (4) is increasing in p_c. Hence, the court picks the largest price such that $\lambda \leq \tilde{\mu}$ and $c_h \leq \tilde{p}$.

Proposition 2 *Consider the full game with user–induced validation.*

i.) Suppose that $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$. Then the court sets a price $p_c^b = \Delta q/4$ and obtains a profit of

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

ii.) Suppose that $\lambda > 1/2$ or $c_h > (q_h + q_l)/2$. Then the court sets the price $p_c^b = (q_h - c_h)(c_h - q_l)/\Delta q$ and obtains a profit of

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

We argued that the court benefits from a relatively high uncertainty for the user and an intermediate price of the patent, in the sense that this maximizes the user's interest in the validation of the patent by the court; we also showed that the user's willingness to pay for the validation is maximized for $\hat{\mu} = 1/2$ and $\hat{p} = (q_h + q_l)/2$. A comparison demonstrates that, for the parameter constellation $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$, the equilibrium induces exactly this outcome. Indeed, the optimal price $p_c = \Delta q/4$ charged by the court leads to a price $p = (q_h + q_l)/2$ and a belief $\mu = 1/2$ and maximizes the expression

$$\min\{\Delta U^1, \Delta U^2\}.$$

For $c_h > (q_h + q_l)/2$, the price $p = (q_h + q_l)/2$ would imply a loss to the high quality inventor and, intuitively, the court cannot induce this maximum degree of uncertainty. For $\lambda > 1/2$, the ex ante belief of the user about the quality of the invention exceeds 1/2. Consequently, the court is unable to induce the belief $\mu = 1/2$. Instead, the court is restricted and maximizes the expression $\min\{\Delta U^1, \Delta U^2\}$ under a feasibility constraint. That is, the court's price maximizes the user's uncertainty about the inventor's quality and, thereby, her willingness to pay.

4 Inventor Induced Certification

In this section we consider the case in which the patent office perfectly validates any incoming patent application, and the inventor instead of the user/licensee buys validation via the patent application. Here validation plays a different role. Rather than giving the user/licensee the possibility to protect herself from bad quality, it enables a high quality inventor to ascertain the quality of his patent to the buyer. Although the distinction seems small, it has a major impact on the equilibrium outcome, primarily because *only the high quality inventor is prepared to demand certification.*

Under inventor–induced validation the parties play the following game:

t=1 The patent office sets a price p_c.

t=2 Nature selects the quality $q_i, i \in \{l, h\}$ of the patent offered by the inventor.

t=3 The inventor offering the patent at quality q_i and cost c_i decides about the price p at which he offers the invention.

t=4 The inventor decides whether or not apply for a patent.

t=5 The user decides whether or not to adopt the innovation.

When comparing to the model described in the previous section, we only change stage four by letting the inventor, rather than the user, decide about the validation of the invention. Note that the sequence of stages 3 and 4 is immaterial. Our setting where the inventor first chooses his price and then decides about validation is strategically equivalent to the situation where he simultaneously takes both decisions, or reverses their order.

We again focus on Perfect Bayesian Equilibria of this game. Note again that after the patent office has set its price p_c a proper subgame, $\Gamma(p_c)$, starts with nature's decision about the quality of the invention offered by the inventor. The subgame $\Gamma(p_c)$ is a pure signalling game if the inventor does not apply for a patent in stage 4. In contrast, if the inventor does decide to apply for a patent, its quality is revealed to the user, and there is no asymmetric information. In the subsequent subgame, the q_h inventor sells his patent at price $p = q_h$, whence the low quality inventor sells his patent at a price $p = q_l$.

In order to capture the inventor's option to validate, we expand the actions open to the inventor by an action c that represents the seller's option to validate and to charge the maximum price q_i. Hence, the inventor's payoffs associated with the action c are $\Pi_h(c) = q_h - c_h$ and $\Pi_l(c) = q_l$ for a high and low quality inventor, respectively. Let $\sigma_i(c)$ denote the probability that the q_i inventor applies for a patent. We further adopt the notation of the previous section. Then we may express a mixed strategy of the seller q_i over validation and a, possibly, infinite but countable number of prices by probabilities $\sigma_i(p_j)$ such that

$$\sigma_i(c) + \sum_j \sigma_i(p_j) = 1. \tag{5}$$

In contrast to the previous section, the user can no longer decide to buy validation, so that her actions are now constrained to s_{nn} and s_{nb}. As before let $\mu(p)$ represent the user's belief upon observing a non–validated invention priced at p. Consequently, s_{nb} is individually rational whenever

$$\mu(p)\Delta q + q_l \geq p$$

and s_{nn} is individually rational whenever

$$\mu(p)\Delta q + q_l \leq p.$$

Proposition 3 *For any price of patent application and validation $p_c < q_h - c_h$, the equilibrium outcome in the subgame $\Gamma(p_c)$ is unique. The high quality inventor applies with probability 1 and obtains the profit $\Pi_h^* = q_h - c_h - p_c > 0$, whereas the low quality inventor does not apply and obtains the payoff $\Pi_l^* = q_l$. For any price $p_c > q_h - c_h$, any equilibrium outcome of the subgame $\Gamma(p_c)$ involves no validation. For $p_c = q_h - c_h$, the subgame $\Gamma(p_c)$ has an equilibrium in which high quality inventor applies with probability 1 and obtains the profit $\Pi_h^* = 0$, whereas the low quality inventor does not apply and obtains the payoff $\Pi_l^* = q_l$.*

The proposition characterizes the equilibrium outcome of the subgame $\Gamma(p_c)$. From this characterization, we can derive the equilibrium of the overall game of inventor–induced validation.

Proposition 4 *The full game with inventor–induced validation has the unique equilibrium outcome* $p_c = q_h - c_h$ *with equilibrium payoffs* $\Pi_c^s = \lambda(q_h - c_h - c_c)$, $\Pi_h^* = 0$, *and* $\Pi_l^* = q_l$.

Comparing the outcome of inventor–induced validation by the patent office with the outcome under user–induced validation by the court, we get the following result.

Proposition 5 *The patent office obtains a higher profit under inventor–induced validation, than the court does under user–induced validation:* $\Pi_c^s > \Pi_c^b$.

The proposition shows that profits are higher when the patent office offers the validation of inventions and sells it to the inventor. The intuition behind this result is that if the user decides whether or not to have the validation conducted by the court, her decision cannot be made contingent on the actual quality of the invention. This is different from when the inventor decides about the validation of his invention by the patent office. Clearly, an inventor with low quality q_l will never apply for a patent. In contrast, we showed that, in any equilibrium, the inventor q_h always applies. The intuition is that if inventor q_h does not have the invention validated at a price p_c quoted by the patent office, then the patent office gets zero profits from the inventor. It, therefore, does strictly better by lowering the price for validating the invention to a level where it is worthwhile for the inventor to have the invention validated by a patent.

5 Welfare

The validation of his invention enables the high quality (and high cost) inventor to sell his patent at the appropriate high price. Otherwise he could not profitably sell his invention. This obvious increase in social efficiency obtains both under user-induced validation by the courts, and under inventor-induced validation by the patent office. From an efficiency perspective, the difference between the two regimes relates to the difference in the probability at which the low quality invention is licensed, and the difference in the cases in which costly validation arises.

First, under inventor–induced validation by the patent office the low quality invention, if offered at all, is always sold. This is different under user–induced validation by the courts, where the invention is not sold when the low quality inventor picks the high price $\tilde{p}$ and the buyer goes to court for its validation. This happens with probability

$$\omega = \sigma_l^*(\tilde{p})\sigma^*(s_{ch}|\tilde{p}, \mu^*(\tilde{p})).$$

Thus, under user–induced validation an efficiency loss of q_l occurs with probability $(1-\lambda)\omega$.[6]

Second, the different regimes may lead to different intensities of validation and therefore differences in expected validation costs. In particular, the probability of the user going to court user is

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

Remember that the user demands validation only if the inventor quotes a high price. Now, the cornered bracket contains the probabilities at which the inventor quotes that high price, which include the probability λ at which he sells the high quality invention, and the probability $(1-\lambda)\sigma_l^*(\tilde{p})$ by which he has a low quality invention but quotes the high price.

By comparison, under inventor–induced validation by the patent office the probability of the inventor applying for a patent is

$$x^s = \lambda.$$

Let $WF^i, i = b, s$ denote social welfare under user and inventor–induced validation, respectively. As usual, it is defined as the sum of consumer and producer surplus. Then, social welfare under user–induced validation by the courts is

$$WF^b = \lambda(q_h - c_h) + (1-\lambda)(1-\omega)q_l - x^b c_c,$$

whereas social welfare under inventor–induced validation equals

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

[6]While in the baseline model, we assume that the low quality invention then is not sold, we show in the extension section that price renegotiation towards the price of the low quality invention does not change our result.

Consequently, the difference in social welfare between the two regimes is

$$\Delta WF = WF^s - WF^b = (1-\lambda)\omega q_l + (x^b - x^s)c_c,$$

In Proposition 5 we have established that the profits of the patent office validating patent applications are larger than those by the courts validating a ex post registered patent. Therefore, from a purely institutional profit point of view, the validation of patent applications by the patent office should be preferred. We now want to check whether this preference is aligned with social efficiency. Clearly, if validation costs were zero, this would follow immediately. The more interesting case is therefore when the cost of validation, c_c, is strictly positive. In this case, the preference given from a profit point of view is still in line with social efficiency, when the probability of certification is smaller under inventor–induced validation. In the next lemma we compare the probability of validation in both regimes.

Lemma 5 *For $\lambda > 1/2$ or $c_h > (q_h + q_l)/2$ the probability x^s of patent application and thus validation by the inventor is lower than the probability x^b of the user going to court towards the validation of the patent. For $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$ the probability x^s of patent application by the inventor is higher than the probability x^b of the user going to court towards the validation of the patent, if and only if $q_h < 3q_l$.*

The lemma identifies a case where the probability of certification is higher under inventor–induced validation than under user–induced validation. This leaves open the possibility that validation of the patent office on the inventor's deman is not in the interest of social efficiency. In particular, if validation costs, c_c, are large, validation by the patent office may be suboptimal. Yet, the following proposition shows that this possibility does not arise. Whenever the the court's profit under user–induced certification is non–negative, social welfare is larger under inventor–induced validation, in spite of possibly a higher probability of validation of the invention via the patent application.

Proposition 6 *Social welfare is higher under inventor–induced validation by the patent office than under user–induced validation by the courts.*

BIS DAHIN HABE ICH DAS PAPIER UMGEARBEITET.

6 Extensions

Our central result that the certifier is better off selling its services to the better informed party, and that its decision is socially efficient is derived within a very stylized model. In this section, we informally discuss extensions in order to show that our result is robust.

To begin, we assumed that, because of the high price, the buyer does not purchase the good in spite of gains from trade, if certification reveals low quality. This assumption is realistic as long as the costs of renegotiating the price after certification are sufficiently high. Yet our results do not depend on the absence of renegotiation. To see this, suppose renegotiation is costless so that, after certification, the buyer and a low quality seller always renegotiate to trade the low quality good at the price $p = q_l$. In this case, the low quality seller always has an incentive to quote the higher price for the low quality good before certification, because he is ensured the low quality price even when the buyer demands certification. Hence, ex post renegotiation actually worsens the outcome of the inspection game by raising the seller's cheating incentives - yet it does not change the outcome of the signalling game.

Our results are also robust to the introduction of imperfect certification technologies. Consider a certification technology that reveals the correct quality only with probability $\pi > 1/2$, whereas it identifies the wrong quality with probability $1 - \pi$. Although the imperfect certification technology reduces the profitability of buyer–induced certification, it does not qualitatively change the equilibrium. Intuitively, a less informative certification technology shrinks the intermediate area in Figure 1, where $S(s_{ch})$ is optimal, in a continuous way. Imperfect certification also does not change the nature of the equilibrium outcome with seller–induced certification. In particular, an equilibrium exists where the certifier charges the certification price $p_c = \pi q_h - c_h$, the high quality seller certifies and charges the price q_h, and the low quality seller sells the good uncertified at a price q_l. The equilibrium is sustained by a buyer who buys the good at the price q_h only if it is certified as of high quality and, consistent with equilibrium play, only believes that the good has high quality when it is certified and the price is q_h. Hence, as shown in Strausz (2010) and in contrast to De and Nabar (1991), the equilibrium outcome remains separating also with imperfect certification. Consequently, the equilibrium outcomes under buyer– and seller–induced certification are continuous in π. As a result, our results are robust to imperfect certification technologies that are not completely uninformative.

Starting from an industrial organization perspective, we assumed that the buyer, seller, and certifier can only use unconditional prices rather than sophisticated contracts to coordinate their exchange. This raises the question whether more complicated contracts, such as prices that condition on the certification outcome, can change our ranking between seller–induced and buyer–induced certification. As one can formally show with optimal mechanism design, this is not the case. The intuition is that with seller–induced certification, the certifier extracts all the rents from certification, and hence, the certifier cannot do better than in our context with seller–induced certification. Stated more formally, the equilibrium payoffs under the optimal mechanism coincide with the equilibrium payoffs in our certification game with seller–induced certification.

In the baseline model, the seller can produce only one fixed quality. Suppose alternatively that a high quality producer actually has the choice to produce alternatively high or low quality, whence a low quality producer can produce only low quality. In this case, the high quality seller's next best alternative to producing high quality and having this certified is to sell low quality without certification. This changes the outside option of the high quality seller from zero to q_l and limits the certifier's possibility to exploit him. Nevertheless, all our qualitative results are upheld. In particular, the certifier obtains the higher profits from seller–induced certification, because, as explained in the previous paragraph, it enables it to extract all rents from certification – even though the rents from certification are now smaller. Similarly, welfare is higher under seller–induced certification.

We finally emphasize that the bilateral seller–buyer framework, within which we have developed our argument, is not crucial. As a particular example, consider a setting which applies particularly well to the financial market, where one seller can sell n units of the good to n identical buyers. Essentially, there are two possible information structures. A first one in which buyers cannot share the certification result but each individually must buy certification. Under buyer certification, our formal results carry through and, hence, the certifier's profits are simply multiplied by n. Under seller certification, Proposition 3 is changed so that the profits from selling the product are also multiplied by n, and $p_c = n[q_h - c_h]$. Because the certifier's profits from selling to buyers and sellers are both multiplied by n, both the ranking of seller–induced vs. buyer–induced certification by the certifier and from a welfare point of view are as in our baseline model.

The second information structure is one in which buyers collude to collectively initiate certification. Under buyer certification, the market structure remains as in the baseline model, yet with n times the buyer's benefit that can be exploited by the certifier. Under seller certification, the same change of Proposition 3 takes place as above. Again, the results remain unchanged.

7 Empirical Examples

Our model and results apply one–to–one to situations in which certification is both product and customer specific. This is the case, for example, in the automotive industry. We first argue that this industry motivates particularly well our theoretical model used.[7] We then move on to other examples – in particular to certification in the financial market.

In the automotive industry, most of the development and production of a complex part for a premium automobile is done by only one supplier — the seller, whom the automotive producer — the buyer — selects explicitly. Because the part is customer specific, the buyer–seller relationship is a bilateral monopoly. Moreover, before the so called null–series production, information between the buyer and the seller about the quality of the part is asymmetric. The automotive industry provides independent certifiers, whose role is

[7] The evidence is taken from Mueller et al. (2008), and from a large scale study conducted in 2007/08 by Konrad 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.

to mediate these information asymmetries.[8] Due to significant economies of scope involving the analytical instruments, the certification industry is highly concentrated. One of the key test criteria is the fulfilment of safety norms. It turns out that the testing of car modules and systems is predominantly performed on the request of the upstream supplier rather than the buyer. Moreover, the buyer conditions his actual purchase on the quality certification. Our model, therefore, captures the procurement relationships in the automotive industry and our results are consistent with the observations in this industry.

While our model applies particularly well to cases in which certification is both product and customer specific, the results also help us understanding purely product specific certification. Examples of purely product specific certification range from the certification of foodstuff for production without herbicides or pesticides; to the certification of toys for production without aggressive chemicals, to the certification of building materials, or of fire–proof safes.

A particularly timely and controversially discussed example is certification in the financial industry. Before the financial crises was triggered, financial products were certified by a heavily concentrated rating industry. The fact that many actual buyers now admit that they poorly understood the products' complexities underscores the large degree of asymmetric information in this market and the rating agencies' central role in reducing it. Before the crisis and consistent with our result, certification was initiated by the issuers – the sellers, who paid rating agencies. A controversial claim is that seller–induced certification led to capture of the certifier and inflated ratings, which precipitated the financial crisis. Proponents of this claim, therefore, argue for a regulatory response to transfer the rating decision from sellers to buyers.

Yet by our results, certification should continue to be initiated by the

[8]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

sellers, since this has robust welfare superior properties. Given these welfare properties, we caution against regulatory pressure in favor of buyer– rather than seller–induced certification. Since capture is the issue, regulatory initiative should concentrate on directly preventing this, by designing a certification system in which capture is minimized or excluded. A particularly successful example of this is the German "Stiftung Warentest" originally founded by the Federal Government to prevent capture, and later privatized. Yet the design of an efficient, capture–proof regulatory mechanism addressing certification in financial markets lies beyond the scope of this paper.

8 Conclusion

Under asymmetric quality information, a demand for certification may arise from both buyers and sellers. Buyers do not want to be cheated if offered a good of unknown quality at a high price. In turn, sellers want to offer the good at a high price – especially if it is of high quality. So to whom does, and, from a welfare point of view, to whom should a credible certifier sell his services, to the buyer or to the seller? Within a parsimonious model, we give straightforward answers to these questions: a certifier does, and should sell to the better informed party.

While the answers to these questions appear deceptively simple, their justification needs an elaborate argument. In particular, we show that certification to the buyer and certification to the seller play very different economic roles and lead to different games, namely an inspection game with the typical mixed strategy equilibrium (which is semi-separating in our context), and a signalling game with a separating equilibrium, respectively.

Our result is consistent with certification in real life – in particular in the automobile industry and financial markets. As to the current discussion about certification in the latter markets, it leads to a clear policy implication. In contrast to much of the current discussion about transferring the initiation to certify to buyers, we provide an argument in favor of seller–initiated certification. This may caution policy makers to think of other means to prevent capture rather than simply reverting from seller– to buyer–induced certification.

We also demonstrated the robustness of our results by considering many extensions. Clearly, further extensions and refinements of the approach are possible. In order to focus on our central point, we have purposively excluded seller reactions to certification, such as adapting quality, as this is discussed in other papers. For the same reason, we also have excluded certifier capture by the seller. Finally, we excluded competition between many sellers, or many certifiers. Arguably, the latter is less important, in view of the technical economies of scale and reputation effects associated with certification. The former, competition between sellers, enhances sellers's demand for certification, but tends not to qualitatively change our insights.

9 Appendix

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

Proof of Lemma 1: Follows directly from the text. Q.E.D.

Proof of Lemma 2: To show that $\Pi_h(p,\mu|\sigma^*)$ is non–decreasing in μ we first establish that, in any PBE, $\sigma^*(s_{nn}|p,\mu)$ is weakly decreasing in μ. Suppose not, then we may find $\mu_1 < \mu_2$ such that $0 \leq \sigma^*(s_{nn}|p,\mu_1) < \sigma^*(s_{nn}|p,\mu_2) \leq 1$. Lemma 1 implies that $(p,\mu_2) \in S(s_{nn}|p_c)$. That is,

$$p \geq \mu_2 q_h + (1-\mu_2)q_l \tag{6}$$

and

$$p_c \geq \mu_2(q_h - p). \tag{7}$$

Now since $\sigma^*(s_{nn}|p,\mu_1) < 1$ we have either $\sigma^*(s_{nb}|p,\mu_1) > 0$ or $\sigma^*(s_{ch}|p,\mu_1) > 0$. Suppose first $\sigma^*(s_{nb}|p,\mu_1) > 0$, then by Lemma 1 we have $p \leq \mu_1 q_h + (1-\mu_1)q_l$. 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 (6). Suppose therefore that $\sigma^*(s_{ch}|p,\mu_1) > 0$, then by Lemma 1 we have $\mu_1(q_h - p) \geq p_c > 0$. This requires $q_h > p$. But then, due to $\mu_2 > \mu_1$, we get $\mu_2(q_h - p) > p_c$, which contradicts (7).

Hence, we establish that $\sigma^*(s_{nn}|p,\mu)$ is weakly decreasing in μ and therefore $\sigma^*(s_{nb}|p,\mu) + \sigma^*(s_{ch}|p,\mu)$ must be weakly increasing in μ. Consequently, $\Pi_h(p,\mu|\sigma^*)$ is weakly increasing in μ.

Next we show that in any PBE $\sigma^*(s_{nb}|p,\mu)$ is weakly increasing in μ. Suppose not, then we may find $\mu_1 < \mu_2$ such that $1 \geq \sigma^*(s_{nb}|p,\mu_1) > \sigma^*(s_{nb}|p,\mu_2) \geq 0$. Since $\sigma^*(s_{nb}|p,\mu_1) > 0$, Lemma 1 implies that $(p,\mu_1) \in S(s_{nb}|p_c)$. That is,

$$p \leq \mu_1 q_h + (1-\mu_1) q_l \tag{8}$$

and

$$p_c \geq (1-\mu_1)(p-q_l). \tag{9}$$

Now since $\sigma^*(s_{nb}|p,\mu_2) < 1$ we have either $\sigma^*(s_{nn}|p,\mu_2) > 0$ or $\sigma^*(s_{ch}|p,\mu_2) > 0$. Suppose first $\sigma^*(s_{nn}|p,\mu_2) > 0$, then by Lemma 1 this 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 (8). Suppose therefore that $\sigma^*(s_{ch}|p,\mu_2) > 0$, then by Lemma 1 we have $(1-\mu_2)(p-q_l) \geq p_c > 0$. 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 (9). Hence, $\sigma^*(s_{nb}|p,\mu)$ must be weakly increasing in μ. Consequently, $\Pi_l(p,\mu|\sigma^*)$ is weakly increasing in μ. Q.E.D.

Proof of Lemma 3: i) For any $\bar{p} < q_l$, $\mu \in [0,1]$ we have $(\bar{p},\mu) \notin S(s_{nn})$, $(\bar{p},\mu) \notin S(s_{ch})$ and $(\bar{p},\mu) \in S(s_{nb})$. Hence, $\sigma^*(s_{nb}|\bar{p},\mu^*(\bar{p})) = 1$. Now suppose for some $\bar{p} < q_l$ we have $\sigma_i^*(\bar{p}) > 0$. This would violate (2), because instead of charging $\bar{p}$ seller q_i could have raised profits by $\varepsilon\sigma_i(\bar{p})$ 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, $\sigma^*(s_{nb}|\bar{p}+\varepsilon,\mu) = 1$ for all μ and in particular for $\mu = \mu^*(\bar{p}+\varepsilon)$.

For any $\bar{p} > q_h$, $\mu \in [0,1]$ we have $(\bar{p},\mu) \in S(s_{nn})$, $(\bar{p},\mu) \notin S(s_{ch})$ and $(\bar{p},\mu) \notin S(s_{nb})$. Hence, $\sigma^*(s_{nn}|\bar{p},\mu^*(\bar{p})) = 1$. Now suppose we have $\sigma_l(\bar{p}) > 0$. This would violate (2), because instead of charging $\bar{p}$ seller q_l could have raised profits by $(q_l-\varepsilon)\sigma_l(\bar{p})$ by charging the price $q_l-\varepsilon$.

ii) Suppose $q_l - \Pi_l^* = \delta > 0$. Now consider a price $p' = q_l - \varepsilon$ with $\varepsilon \in (0,\delta)$ then for any $\mu' \in [0,1]$ we have $(p',\mu') \in S(s_{nb})$ and $(p',\mu') \notin S(s_{nn}) \cup S(s_{ch})$ so that, by (1), we have $\sigma^*(s_{nb}|p'^*,\mu^*(p')) = 1$ and, therefore, $\Pi_l(p'^*,\mu^*(p')|\sigma^*) = p' > \Pi_l^*$. This contradicts (2).

iii) For any p such that $\sigma_h^*(p) > 0$, we have $\Pi_h^* = \Pi_h(p,\mu^*(p)|\sigma^*) = [\sigma^*(s_{nb}|p,\mu^*(p)) + \sigma^*(s_{ch}|p,\mu^*(p))]p - c_h$. As argued in i), we have $\sigma^*(s_{nn}|p,\mu) =$

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 $\sigma^*(s_{nb}|q_h,\mu^*(q_h)) + \sigma^*(s_{ch}|q_h,\mu^*(q_h)) = 1$. But, due to $\mu^*(q_h)(q_h - q_h) = 0 < p_c$, we have $(q_h,\mu^*(q_h)) \notin S(s_{ch}|q_h)$ so that $\sigma^*(s_{ch}|q_h,\mu^*(q_h)) = 0$. Hence, we must have $\sigma^*(s_{nb}|q_h,\mu^*(q_h)) = 1$. This requires $(q_h,\mu^*(q_h)) \in S(s_{nb}|p_c)$ so that we must have $\mu^*(q_h) = 1$. By (3), this requires $\sigma_l^*(q_h) = 0$. But since $\Pi_l(q_h,1|\sigma^*) = \sigma^*(s_{nb}|q_h,\mu^*(q_h))q_h = q_h$ we must, by (2), 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 4: 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 $\mu' \in (p_c/(q_h - p'), 1 - p_c/(p' - q_l)) \neq \emptyset$[9] we have $(p',\mu') \in S(\sigma_{ch})$ and $(p',\mu') \notin S(\sigma_{nn}) \cup S(\sigma_{nb})$. Consequently, $\sigma^*(s_{ch}|p',\mu') = 1$. Hence, $\Pi_h(p',\mu') = p' - c_h = \tilde{p} - c_h - \varepsilon > \tilde{p} - c_h - \delta = \Pi_h^*$ and $\Pi_l(p',\mu') = 0 < q_l \leq \Pi_l^*$. Therefore, by B.R. the buyer's equilibrium belief must satisfy $\mu^*(p') \geq \mu'$. By Lemma 2 it follows $\Pi_h(p',\mu^*(p')) \geq \Pi_h(p',\mu') = \tilde{p} - c_h - \varepsilon > \Pi_h^*$. This contradicts (2). 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 (2). Q.E.D.

Proof of Proposition 1: i): First we show that for $\lambda < \tilde{\mu}$ and $c_h < \tilde{p}$ 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 4.i, we have $\bar{p} \geq \tilde{p}$ and, by Lemma 3.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, $\sigma^*(s_{nn}|\bar{p},\mu^*(\bar{p})) = 1$ and $\Pi_h(\bar{p},\mu^*(\bar{p})) = 0$. As a result, $\sigma_h^*(\bar{p}) > 0$ contradicts (2), because, by Lemma 4.ii, $\Pi_h^* \geq \tilde{p} - c_h > 0 = \Pi_h(\bar{p},\mu^*(\bar{p}))$.

Second, suppose that for some $\bar{p} > \tilde{p}$ we have $\sigma_h^*(\bar{p}) > 0$ then, by definition of $\tilde{p}$, we have $(\bar{p},\mu) \notin S(s_{ch})$ for any $\mu \in [0,1]$. Hence, $\sigma^*(s_{ch}|\bar{p},\mu^*(\bar{p})) = 0$

[9]Let $l(p) \equiv p_c/(q_h - p)$ 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 $l'(p) = p_c/(q_h - p)^2 > h'(p) = p_c/(p - q_l)^2 > 0$. Hence, $l(\tilde{p} - \varepsilon) < h(\tilde{p} - \varepsilon)$ for $\varepsilon > 0$ so that $\tilde{p} - \varepsilon > q_l$ and, therefore, $l(p') < h(p')$.

so that $\Pi_l(\bar{p}, \mu^*(\bar{p})) = \Pi_h(\bar{p}, \mu^*(\bar{p})) + c_h$. From Lemma 4.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 take a $\bar{p} > \tilde{p}$ with $\sigma_l(\bar{p}) > 0$ then, by Lemma 3.ii and (2), $0 < q_l \leq \Pi_l^* = \Pi_l(\bar{p}, \mu^*(\bar{p})|\sigma^*) = \sigma(s_{nb}|\bar{p}, \mu^*(\bar{p}))\bar{p}$. This requires $\sigma(s_{nb}|\bar{p}, \mu^*(\bar{p})) > 0$ and therefore $(\bar{p}, \mu^*(\bar{p})) \in S(s_{nb}|p_c)$ 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 $\bar{p} > \tilde{p}$. Hence, if an equilibrium for $\lambda < \tilde{\mu}$ and $\tilde{p} > c_h$ exists then, by Lemma 4, it exhibits $\sigma_h^*(\tilde{p}) = 1$, $\Pi_h^* = \tilde{p} - c_h$ and $\sigma^*(s_{ch}|\tilde{p}, \tilde{\mu}) + \sigma^*(s_{nb}|\tilde{p}, \tilde{\mu}) = 1$.

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

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

which is smaller than one exactly when $\lambda < \tilde{\mu}$. By definition, $(\tilde{p}, \tilde{\mu}) \in S(s_{ch}) \cap S(s_{nb})$ so that any buying behavior with $\sigma^*(s_{ch}|\tilde{p}, \tilde{\mu}) + \sigma^*(s_{nb}|\tilde{p}, \tilde{\mu}) = 1$ is consistent in equilibrium. In particular, $\sigma^*(s_{nb}|\tilde{p}, \tilde{\mu}) = q_l/\tilde{p} < 1$ 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 $\sigma^*(s_{nb}|q_l, \mu^*(q_l)) = 1$ and $\mu^*(p) = 0$ for any price p larger than q_l and unequal to $\tilde{p}$. This out–of–equilibrium beliefs satisfies B.R.. Hence, the expected profit to the certifier is

$$\Pi_c(p_c) = (\lambda + (1-\lambda)\sigma_l^*(\tilde{p}))\,\sigma^*(s_{ch}|\tilde{p}, \tilde{\mu})(p_c - c_c) = \frac{\lambda(\tilde{p} - q_l)}{\tilde{\mu}\tilde{p}}(p_c - c_c).$$

ii) In order to show that, in any equilibrium of $\Gamma(p_c)$, we have $\Pi_c(p_c) = 0$ whenever $\lambda > \tilde{\mu}$, we prove that for any $\bar{p}$ such that $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$, it must hold $\sigma_h^*(\bar{p}) = \sigma_l^*(\bar{p}) = 0$. Suppose we have $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$, then $(\bar{p}, \mu^*(\bar{p})) \in S(s_{ch})$ and, necessarily, $\bar{p} \leq \tilde{p}$. But by Lemma 4.i, $\sigma_h^*(\bar{p}) > 0$ also implies $\bar{p} \geq \tilde{p}$. Therefore, we must have $\bar{p} = \tilde{p}$. But $(\tilde{p}, \mu) \in S(s_{ch})$ 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 $(p', \mu) \notin S(s_{ch}|p_c)$ 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 $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$ 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 (2), $\sigma_l^*(\bar{p}) = 0$. Since $\bar{p} = \tilde{p}$, this violates $\sigma_l^*(\tilde{p}) > \sigma_h^*(\tilde{p}) \geq 0$. As a result, $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$ implies $\sigma_h^*(\bar{p}) = 0$.

In order to show that we must also have $\sigma_l^*(\bar{p}) = 0$, assume again that $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$. 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 $(\bar{p}, \mu^*(\bar{p})) \notin S(s_{ch})$, which contradicts $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$.

In order to show that $\tilde{p} < c_h$ implies $\Pi_c(p_c) = 0$ suppose, on the contrary that, $\Pi_c(p_c) > 0$. This requires that there exists some $\bar{p}$ such that $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$ and $\sigma_i^*(\bar{p}) > 0$ for some $i \in \{l, h\}$. First note that $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) > 0$ implies $\bar{p} \leq \tilde{p}$. Now suppose $\sigma_h^*(\bar{p}) > 0$ then $\Pi_h(\bar{p}, \mu^*(\bar{p})|\sigma^*) = (\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) + \sigma^*(s_{nb}|\bar{p}, \mu^*(\bar{p})))\bar{p} - c_h < 0$ so that the high quality seller would make a loss and, thus, violates (2). Therefore, we have $\sigma_h^*(\bar{p}) = 0$. Now if $\sigma_l^*(\bar{p}) > 0$ then (3) implies $\mu^*(\bar{p}) = 0$ so that $\sigma^*(s_{ch}|\bar{p}, \mu^*(\bar{p})) = 0$, which contradicts $\Pi_c(p_c) > 0$. Q.E.D.

Proof of Proposition 2: 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 (4) 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 (4) 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 $\lambda \leq 1/2$, it 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 $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$, the maximum obtains for $p_c = \Delta q/4$ with

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

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

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

Second, 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

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

Q.E.D.

Proof of Proposition 3 Fix some $p_c < q_h - c_h$. By certifying, seller q_h guarantees himself the payoff $\Pi_h(c) = q_h - c_h - p_c > 0$. Hence, in any equilibrium of the subgame $\Gamma(p_c)$ seller q_h must obtain a payoff of at least $\Pi_h(c) > 0$.

Now suppose that there exists some equilibrium in which $\sigma_h(c) < 1$. Then, by (5) there exists some price $\tilde{p}$ such that $\sigma_h(\tilde{p}) > 0$. For $\tilde{p}$ to be optimal, it is required that $\Pi_h(\tilde{p}, \mu^*(\tilde{p})|\sigma^*) = \tilde{p}\sigma(s_{nb}|\tilde{p}, \mu^*(\tilde{p})) - c_h \geq \Pi_h(c) > 0$. This implies $\Pi_l(\tilde{p}, \mu^*(\tilde{p})|\sigma^*) = \tilde{p}\sigma(s_{nb}|\tilde{p}, \mu^*(\tilde{p})) > c_h$ so that the equilibrium payoff of seller q_l is $\Pi_l^* > c_h > \bar{q}$. Consequently, $\sigma_l^*(p) = 0$ for any $p < \bar{q}$ and therefore

$$\sum_{p \geq \bar{q}} \sigma_l^*(p) = 1. \tag{10}$$

But if $\sigma_l^*(p) > 0$ then we must have $p\sigma(s_{nb}|p, \mu^*(p)) > c_h$. This requires $\sigma(s_{nb}|p, \mu^*(p)) > 0$. Therefore, s_{nb} must be an optimal response given price p and belief $\mu^*(p)$. Hence, $\mu^*(p)\Delta q + q_l \geq p > c_h > \lambda\Delta q + q_l$. As a result, $\mu^*(p) > \lambda$ and, due to (3), it holds $\sigma_h^*(p) > \sigma_l^*(p)$ for any $\sigma_l^*(p) > 0$. Together with (10) we arrive at the contradiction

$$\sum_{p \geq \bar{q}} \sigma_h^*(p) > \sum_{p \geq \bar{q}} \sigma_l^*(p) = 1. \tag{11}$$

It is straightforward to verify that for $p_c \leq q_h - c_h$, the strategies $\sigma_h(c) = 1$, $\sigma_l(q_l) = 1$, $\sigma^*(s_{nn}|p, \mu) = 1$ whenever $\mu\Delta q + q_l \geq p$ and zero otherwise together with $\mu^*(p) = q_l$ constitute an equilibrium that sustains the equilibrium outcome.

For $p_c > q_h - c_h$, certification would yield seller q_h a negative payoff: $\Pi_h(c) = q_h - c_h - p_c < 0$. Certification would yield seller q_l a payoff $\Pi_l(c) = q_l - p_c < q_l$, whereas seller q_l could guarantee himself the payoff q_l by not certifying. Q.E.D.

Proof of Proposition 4: First, suppose there exists an equilibrium in which the payoff of the certifier, Π_c^*, is strictly smaller than $\lambda(q_h - c_h - c_c)$. That is, $\delta = \lambda(q_h - c_h - c_c) - \Pi_c^* > 0$. Now note that the price $p_c = q_h - c_h - \delta/2 < q_h - c_h$ yields the certifier a payoff $\lambda(q_h - c_h + \delta/2) > \Pi_c^*$, because Proposition 3 shows that its subgame $\Gamma(p_c)$ has the unique outcome that seller q_h always certifies and seller q_l does not. Second, note that the certifier cannot obtain a profit that exceeds $\lambda(q_h - c_h - c_c)$, because it would require that the price of certification exceeds $q_h - c_h$ or that the low quality seller certifies with a strictly positive probability. Hence, in any equilibrium the certifier obtains the payoff $\lambda(q_h - c_h - c_c)$. According to Proposition 3 the certifier may become this payoff only for $p_c = q_h - c_h$ with $\sigma_h(c) = 1$. Q.E.D.

Proof of Proposition 5: For $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$ we have $\Pi_c^s = \lambda(q_h - c_h - c_c) \geq \lambda(q_h - c_h - c_c)\frac{q_h - q_l}{q_h + q_l} \geq \lambda(q_h - (q_h + q_l)/2 - c_c)\frac{q_h - q_l}{q_h + q_l} = \lambda(q_h - q_l - 2c_c)\frac{q_h - q_l}{2(q_h + q_l)} \geq \lambda(q_h - q_l - 4c_c)\frac{q_h - q_l}{2(q_h + q_l)} = \Pi_s^b$, where the second inequality uses $c_h \leq (q_h + q_l)/2$.

For $\lambda > 1/2$ or $c_h > (q_h + q_l)/2$ it follows that $\Pi_c^b = \frac{\lambda[(q_h - c_h)(c_h - q_l) - \Delta q c_c]}{c_h} < \frac{\lambda[(q_h - c_h)(c_h - q_l) - (c_h - q_l)c_c]}{c_h} = \lambda(q_h - c_h - c_c)\frac{c_h - q_l}{c_h} \leq \lambda(q_h - c_h - c_c) = \Pi_b^s$, where the first inequality uses $q_h > c_h$. Q.E.D.

Proof of Lemma 5: For $\lambda > 1/2$ or $c_h > (q_h + q_l)/2$, it follows

$$x_c^b = (\lambda + (1 - \lambda)\sigma_l^*(\tilde{p}))\sigma(s_{ch}|\tilde{p}, \mu_h^*) = \lambda\frac{\Delta q}{c_h} \leq \lambda = x_c^s,$$

where the inequality obtains from $q_h - c_h - c_c > q_l \;\Rightarrow\; \Delta q < c_h + c_c < c_h$.

For $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$, it follows

$$x_c^b = (\lambda + (1 - \lambda)\sigma_l^*(\tilde{p}))\sigma(s_{ch}|\tilde{p}, \mu_h^*) = \lambda\frac{2\Delta q}{q_h + q_l}.$$

Hence, $x_c^b < x_c^s$ if and only if $2\Delta q < q_h + q_l$. This yields the condition $q_h < 3q_l$. Q.E.D.

Proof of Proposition 6: Due to Lemma 5 we need only check for the case $\lambda \leq 1/2$ and $c_h \leq (q_h + q_l)/2$ and $q_h < 3q_l$. According to Proposition 2 the certifier in this case makes non–negative profits exactly when $p_c^b = \Delta q/4 \geq c_c$. The differences in social welfare for this case is

$$\begin{aligned}
\Delta WF &= \lambda\frac{\Delta q}{q_h + q_l}q_l + \lambda\left(\frac{2\Delta q}{q_h + q_l} - 1\right)c_c && (12)\\
&= \frac{\lambda}{q_h + q_l}\left(\Delta q q_l - (3q_l - q_h)c_c\right) && (13)\\
&\geq \frac{\lambda}{q_h + q_l}\left(\Delta q q_l - (3q_l - q_h)\Delta q/4\right) && (14)\\
&= \lambda\Delta q/4 > 0. && (15)
\end{aligned}$$

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., 1993, Middlemen as Experts, *RAND Journal of Economics* 24, 212–223.
Bouvard, M. and R. Levy (2009), Humouring Both Parties: a Model of Two Sided Reputation, *mimeo, Universite de Toulouse I*
De, S. and P. Nabar (1991), Economic Implications of Imperfect Quality Certification, *Economics Letters*, 37, 333-337
Emons, W. (1997), Credence Goods and Fraudulent Experts, *The RAND Journal of Economics*, 28, 107-119
Fasten, E. and D. Hofmann (2010), Two-sided Certification: The Market for Rating Agencies, *mimeo, Humboldt University of Berlin*
Faulhaber, G. and D. Yao (1989), "Fly-by-Night Firms and the Market for Product Review", *The Journal of Industrial Economics*, 38, 65-77
Guerra, G. (2001), Certification Disclosure and Informational Efficiency: A Case for Ordered Ranking of Levels, *mimeo, University of Oxford*
Hvide, H.(2004), A Theory of Certification with an Application to the Market for Auditing Services, *mimeo, Norvegian School of Economics, Bergen*
Lizzeri, A., 1999, Information Revelation and Certification Intermediaries, *RAND Journal of Economics* 30, 214–231.
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 (2004), Strategic Certification, *mimeo, University of Wisconsin-Madison*

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 (1978) A Note on "Lemons" Markets with Quality Certification, *Bell Journal of Economics* 9, 277-279
Wolinsky (1993) Competition in a Market for Informed Experts, *RAND Journal of Economics* 24, 380-398 277-279