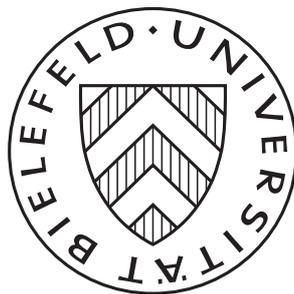


September 2014

Quality Uncertainty with Imperfect Information Acquisition

Christopher Gertz



Quality Uncertainty with Imperfect Information Acquisition

Christopher Gertz *

Center for Mathematical Economics, Bielefeld University

This version: September 3, 2014

[Abstract:] We analyze a monopolistic model of quality uncertainty but with the possibility of information acquisition on the consumer side. Information is costly and its amount is chosen by the consumer. The analysis of Bayesian equilibria shows the possibility of three equilibrium classes, only one of which leaves positive utility to the consumer. The classic adverse selection results of these markets are weakened in this situation. We show that cheaper information does not necessarily benefit the consumer but can instead rule out the buyer-friendly and welfare maximizing equilibria. Moreover, making quality search arbitrarily efficient does not lead to sure selling of the high quality product. A sustainable adverse selection effect, though weaker than in the classical model, remains even in the limit.

Keywords: Quality uncertainty, Price signaling, Adverse selection, Information acquisition, Two-sided incomplete information

JEL Classification Numbers: C72, D42, D82, D83

*This research was carried out within the International Research Training Group "Economic Behavior and Interaction Models" (EBIM) financed by the German Research Foundation (DFG) under contract GRK 1134/2. Financial support by the IMW is gratefully acknowledged.

1 Introduction

Along with the price, the quality of a product is one of the two major characteristics relevant to a purchase decision. While price information can be tricky to obtain in situations with negotiation possibilities or multiple retailers, it is usually even harder to get a good idea about the quality of a product. This phenomenon occurs not only in classical examples like second hand cars but is also present in almost every potential purchase for goods with which you do not have much experience. Whether you have not bought a TV for many years, want to buy wine in an unfamiliar supermarket, consider to buy an upgrade of the operating system on your computer or to invest in a financial product. You do not have full information about what you really get for your money if you are not an expert for these products or at least inform yourself prior to the purchase.

On the other hand, the seller of a product usually has much better information about the quality. A second hand car has been used by the owner for many years, giving him the knowledge of any accident, repair or defects that occurred over the last years and might still exist or be relevant. In the same way, a producer knows the characteristics and weaknesses of the product for sale. While positive qualities might be advertised or otherwise pointed out to the buyer, there is no incentive for the seller to do the same for bad qualities or missing features if not mandated by law.¹ If the consumer wants this information, she usually has to acquire it on her own. And even the advertised positive qualities may be exaggerated, untrue or not relevant to the consumer. In any case, it takes effort of the consumer to either gather information or check and filter the information provided.

This asymmetry of quality information has been present in the literature ever since George Akerlof's famous paper of 1970 and his previously mentioned second hand car example. The question of how such quality uncertainty influences the market outcome has since been discussed in many papers. There have been various attempts to mitigate between the extreme nature of Akerlof's model and the classical case of perfect information. For instance, Bagwell and Riordan (1991) introduce multiple periods in which firms can set different prices. High quality firms can then acquire customers by setting low prices in the first period and use this to charge highly afterwards. Milgrom and Roberts (1986) allow the firms to give an additional, costly advertising signal to distinguish oneself and keep low quality sellers from imitating the high quality sellers' behavior.

While these examples follow the idea of giving high quality firms certain additional possibilities to signal their excellence, it is only natural to also look at the other side of the market, namely the consumer. In the Akerlof model buyers do not have any other possibility than either to trust the firms or not, their posterior beliefs about quality depending on the price and being determined by Bayes' law in an equilibrium. Having the examples of the first paragraph in mind, the idea of the consumer having no further information is obviously a very strong assumption and not true in most purchasing situations. A

¹Dziuda (2011) shows that, in a game theoretic setting in which an expert might be biased or not, some bad properties may be disclosed even by a biased expert. In our market setting, however, there is no such ambiguity over the goal of the firm.

second hand car can be tested before buying, a bottle of wine can be bought and tried out before you decide to serve it at your dinner party and for most products you can find multiple tests, ratings and reviews online. Especially with the Internet, the amount of quality information available to consumers has dramatically increased in recent years. With more and more people owning smartphones, this information is available and can be looked up even inside the shop.

Few papers have so far considered to relax this part of Akerlof's model. Bester and Ritzberger (2001) let the consumer decide about buying a perfect quality signal and base their decision on the additional information. In Voorneveld and Weibull (2011) the buyer receives an additional, costless but noisy signal which is correlated to the true quality. One can interpret this as an independent, non-perfect test review that is observed by the consumer in any case. These two models of additional quality information do not quite capture the real life situations, as it is usually not costless (in terms of money or time) to acquire information and this information does not have to be perfect. Even if there are free tests and reviews on the Internet, one has to take the time and put effort in finding and reading these sources. These tests, on the other hand, may not contain all relevant information and can be incorrect, biased or based on a faulty product. The same holds for the information from friends and experts. Testing a TV in the store can not fully simulate the home environment and how the TV works together with other devices, etc. All this adds some unreliability to the information acquisition procedure and suggests some probability of false information.²

This chapter simultaneously covers both ideas of these two works; the consumer is able to choose how much effort or costs she wants to invest in acquiring quality information and this will result in a quality signal. This signal, on the other hand, will be more precise when exerting more search effort to such an extend that one might even reach perfect information.

Our analysis shows that in the market with imperfect information acquisition, different kinds of equilibria can occur. Under reasonable refinements, the most important two categories of equilibria are one in which the consumer does not spend effort on quality information and one in which she does. Only in the first type of these equilibria she has positive utility. Her ability to search, although not executed, lowers the price to below the expected quality. In equilibria with search, on the other hand, the price of the product is relatively high and all the consumer's possibility of acquiring information provides market power for the firm, not for the consumer.

Starting from these insights, we investigate how the existence and outcomes of these equilibria develop when information is available more easily. We find that a higher ability of acquiring quality information stops the existence of the consumer-friendly equilibrium and thus takes away all consumer utility.

At last, we investigate the limit behavior of the model in the case when quality infor-

²It is worth noting that the work of Kihlstrom (1974) was motivated by the same ideas as this paper. His analysis, however, solely focuses on the consumer side (the market for information) and does not give any indication about the implications for market equilibria.

mation is very expensive or very cheap. Surprisingly, making information acquisition arbitrarily easy does not lead to convergence to the full information case in that the error probability of the signal stays bounded away from zero and a non-vanishing share of high quality products is not traded. Although the possibility of imperfect information acquisition generally lowers adverse selection phenomena, it does not get rid of them even in the limit of perfect information. Making the information very expensive can lead to the outcomes of the classical models of quality uncertainty but often also, for a wide range of parameters, converges to an equilibrium that was previously disregarded in the literature.

The paper is structured as follows. The first section describes the model of quality information acquisition and the market participants. We then proceed by analyzing the consumer behavior. This is embedded in a formally defined monopolistic market model and the rational Bayesian equilibria are analyzed. Finally, we compare outcomes of different levels of search efficiency to investigate the market impact of cheaper or more expensive quality information.

2 The Model

We consider a monopolistic market with one product and one potential consumer (or “buyer”)³. The quality of the product can take two fixed values and is drawn by nature with a publicly known probability η of high quality.⁴ The realization is known only by the firm and will be denoted by its type $\theta \in \{H, L\}$.

A high quality firm faces production costs $c^H > 0$ while the low quality firm pays $c^L > 0$ for producing one unit of the product. This cost is only incurred when the good is actually sold. We assume $c^L < c^H$ so that high quality production is at least marginally more costly than for low quality. These costs can also be seen as losing an outside option. For example, the seller of a second hand car could also bring the vehicle to a professional dealer who would pay him the amount c^θ . This option is lost in case of a successful sale.

The seller (or firm) makes a take-it-or-leave-it offer and is a risk-neutral payoff maximizer. Selling one unit of the good for a price p with probability δ yields the payoff

$$\pi^\theta(p, \delta) = \delta \cdot (p - c^\theta), \quad \theta \in \{H, L\}.$$

Only observing the price p set by the firm, not the quality of the product, the risk-neutral consumer maximizes her expected gains from trade. After a purchase she learns her valuation; her utility from having bought a product of quality q for price p then is $q - p$.

We denote $q^H > q^L$ the consumer’s valuation for the high and the low quality product, respectively. To always ensure possible gains from trade, we assume $c^L < q^L$ and $c^H <$

³The analysis would not change if we assumed multiple, identical buyers. For simplicity, we only speak of one consumer.

⁴The assumption that quality is not a strategic decision of the firm is crucial in lemon markets. While quality uncertainty may also exist in equilibria under different models, they allow for separation even at the quality level. See Shaked and Sutton (1982) for an example of such a result.

q^H .⁵

The buyer always has a certain, endogenous belief of the quality after observing a price. Fix a price p and let $\hat{\mu} \in [0, 1]$ be the conceived probability of facing a high type firm. Then the consumer's expected utility from buying the good is

$$u_b(p, \hat{\mu}) := \hat{\mu}q^H + (1 - \hat{\mu})q^L - p$$

while the utility from not buying the good is $u_n := 0$.

She faces a third option, namely to pay a cost (or to exert effort) of a chosen level $k \geq 0$ to then obtain a binary signal $s \in \{s^H, s^L\}$ about the product quality. This signal might be incorrect with an error probability of $\varepsilon(k) \in [0, \frac{1}{2}]$. Mathematically this means⁶

$$\text{Prob}(s^H | \theta = L) = \text{Prob}(s^L | \theta = H) = \varepsilon(k).$$

The exogenously given error function ε satisfies the following assumptions.

- $\varepsilon : \mathbb{R}_+ \rightarrow [0, \frac{1}{2}]$ is continuous and non-increasing.
- $\varepsilon(0) = \frac{1}{2}$
- Denote $\bar{k} = \inf \{k \in \mathbb{R}_+ | \varepsilon(k) = 0\}$ the costs for a perfect signal where $\inf \emptyset := \infty$. Then ε is twice continuously differentiable on $(0, \bar{k})$.
- $\varepsilon'(k) < 0, \varepsilon''(k) > 0 \quad \forall k \in (0, \bar{k})$

While the first and third points are of technical nature, the second assumption says that the signal does not contain any information if the consumer exerts no effort. The last point ensures that higher effort always leads to a strictly higher signal precision while the marginal precision gain is diminishing. This accounts for the fact that information acquisition, such as reading reviews, will often give redundant information and thus the amount of new information gained via a certain increase of search effort is decreasing. Due to the second point we can assume that the consumer also receives the (non-informative) signal when she chooses $k = 0$.

Note that we allow for obtaining a *perfect signal*, i.e. there may be a finite cost \bar{k} for which the error probability is zero. Depending on the error function, this value might also be infinite so that perfect information would not be achievable. We do not restrict attention to any of these cases.

By the last assumption the expression

$$\varepsilon'(0) := \lim_{\substack{k \rightarrow 0 \\ k \in (0, \bar{k})}} \varepsilon'(k) = \inf_{k \in (0, \bar{k})} \varepsilon'(k) \in [-\infty, 0)$$

⁵Adriani and Deidda (2009) focus on a case in which trade would not always be beneficial under full information. They show that this leads to market breakdown in their setting under the D1-Refinement of Cho and Kreps (1987).

⁶The assumption of both error types being the same is certainly somehow restrictive but is not believed to have a qualitative impact on the results. See Martin (2012) for the use of a more complex information structure. In his analysis, however, the firm can choose only between two exogenously given prices.

is well-defined. This value is important in the analysis. For illustrating results, we use the simple function $\varepsilon(k) = \max\left\{\frac{1}{2} - \sqrt{k}, 0\right\}$ which satisfies the assumptions above.

Naturally, agents on this market do not act simultaneously. At the time when the consumer makes her decision, the price was already set by the firm and this requires the quality level to already be realized. Figure 1 outlines the timing of the market.

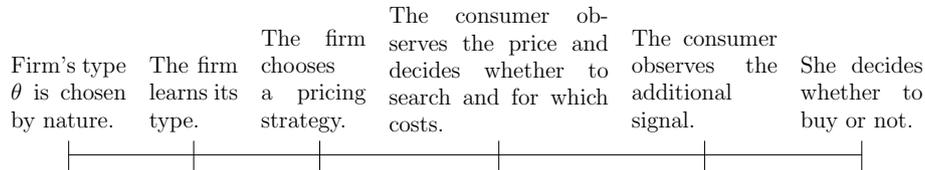


Figure 1: The timing of the market

The consumer holds a belief system $\mu : \mathbb{R}_+ \mapsto [0, 1]$, later determined by the equilibrium definition, which assigns to each possible price p a belief $\mu(p)$ about the probability that the product is of high quality conditional on the observed price. In the analysis it is sometimes useful to consider a fixed price p and a fixed corresponding posterior belief $\mu(p)$. In this case we abbreviate the latter by writing $\hat{\mu}$ instead of $\mu(p)$. We define the expected quality based on such beliefs by

$$\bar{q}_{\hat{\mu}} := \hat{\mu}q^H + (1 - \hat{\mu})q^L.$$

In the same way, to avoid imprecisions, single values of other functions are denoted similarly. Note that the true a priori probability of high quality is denoted by η while the letter μ is reserved for posterior belief values.

3 Consumer Behavior

Since we analyze a multi-stage game, we proceed by backward induction, thus first dealing with the buyer's decision problem. This problem itself has two stages. When observing the price $p \in \mathbb{R}_+$, she has to decide about the search amount $k \geq 0$. In the second step, she receives the signal and decides whether to buy the good or not. We allow for mixed strategies, so it is possible for the consumer to buy the good only with a certain probability. Remember that the two possible quality valuations q^L, q^H as well as the a priori probability η of facing a high type producer is known to the consumer.

3.1 After Receiving the Signal

Assume for now that k has been chosen. Let $\hat{\varepsilon} := \varepsilon(k)$ be the corresponding error probability and $\mu(p) \in (0, 1)$ the posterior belief that a product with price p has quality q^H . In this section, p and $\mu(p)$ can be viewed as fixed so that we write $\hat{\mu}$ for the posterior belief.

Conditional on observing the high signal s^H , the probability of the quality being high is

$$\frac{(1 - \hat{\varepsilon})\hat{\mu}}{(1 - \hat{\varepsilon})\hat{\mu} + \hat{\varepsilon}(1 - \hat{\mu})}$$

which follows from Bayes' law.

The expected utility from buying (not taking into account the sunk cost k), given this situation is then

$$\frac{(1 - \hat{\varepsilon})\hat{\mu}}{(1 - \hat{\varepsilon})\hat{\mu} + \hat{\varepsilon}(1 - \hat{\mu})} q^H + \frac{\hat{\varepsilon}(1 - \hat{\mu})}{(1 - \hat{\varepsilon})\hat{\mu} + \hat{\varepsilon}(1 - \hat{\mu})} q^L - p.$$

Note that with $\hat{\varepsilon} = \frac{1}{2}$ this is the original utility from buying without the additional signal.

The case on observing s^L is computed similarly. Clearly, the consumer will buy the good if this utility is above zero and not buy it if it is strictly below this value. The interesting insight here is that the signal is only relevant to her if not for all signals s^H and s^L the utilities lie both above or both below zero. Remember that she made a rational choice to pay an amount $k > 0$ and thus she can intuitively not be indifferent between the signal outcomes.

Lemma 3.1. *Let a price p with corresponding posterior belief $\hat{\mu}$ be given. If the consumer has optimally exerted positive search effort, she buys if and only if she receives the signal s^H .*

This result may not come as a surprise. If it was optimal to ignore a signal, it would be pointless to pay for its precision. The lemma is in the same spirit as the corresponding statement in Bester and Ritzberger (2001). It shows that the two pieces of information, namely the inherent information of the price given by the corresponding posterior belief $\hat{\mu}$ and the additional signal, are essentially not considered simultaneously. The former is used to decide about how much search effort to exert and if zero effort is chosen, it is used to determine whether to buy the good or not. Once the consumer decides to pay for signal precision, the buying decision only depends on the signal, not on the value $\hat{\mu}$ of the posterior belief. This, of course, does not occur in situations where an additional, informative signal is received regardless of the decision of the consumer as in Voorneveld and Weibull (2011).

While this effect also arises in Bester and Ritzberger (2001), consumers in their model observe a perfect signal and it is natural to dismiss prior information after learning the true state. In the situation at hand the reason is more subtle, basically lying in the backward induction argument. The probable implications of receiving various signals of a certain error probability are taken into account before the decision of costly acquiring the information is formed. Essentially, also the choice to buy only at a high signal is already made at that stage.

The proof of this lemma is straightforward. Like all others, it can be found in the appendix.

3.2 Choosing the Optimal Search Effort

We proceed by determining the optimal search costs k . Assume therefore that the consumer pays a cost $k > 0$ for search and that this level is optimal. We know by the previous lemma that the only possible behavior after receiving the signal is to buy if and only if the signal is s^H , i.e. if the quality is high and the signal is correct or if the quality is low and the signal is wrong. Then the expected utility, given price p and posterior beliefs $\hat{\mu}$, is

$$u_s(p, \hat{\mu}, k) := \underbrace{\hat{\mu}(1 - \varepsilon(k))(q^H - p)}_{\text{correct high signal}} + \underbrace{(1 - \hat{\mu})\varepsilon(k)(q^L - p)}_{\text{false high signal}} - k.$$

This formula consists of three terms. The (possibly subjective) probability of facing a high good is $\hat{\mu}$. The consumer then buys if she receives a correct signal which has the probability $1 - \varepsilon(k)$. This yields the utility $q^H - p$. The second term of the formula reflects the possibility and consequences of buying a low quality product because of a false high signal. The search costs k have to be paid regardless of the quality and the buying decision.

We want to stress that this is the expected utility after observing the price and before receiving the signal, and only if the optimal search cost is positive. Lemma 3.1 allows us to ignore the updated beliefs after observing the additional quality information.

Maximizing this utility with respect to search costs, we get the first order condition

$$\begin{aligned} \hat{\mu}(-\varepsilon'(k))(q^H - p) + (1 - \hat{\mu})\varepsilon'(k)(q^L - p) &= 1 \\ \Leftrightarrow \varepsilon'(k) &= \frac{1}{-\hat{\mu}(q^H - p) + (1 - \hat{\mu})(q^L - p)} \\ &= \frac{-1}{\hat{\mu}(q^H - p) + (1 - \hat{\mu})(p - q^L)} =: d(p, \hat{\mu}) =: \hat{d}. \end{aligned}$$

The parameter \hat{d} depends both on the price p and the posterior belief $\hat{\mu}$ and is always negative in the relevant range of prices $[q^L, q^H]$ and when $\hat{\mu} \in (0, 1)$. Its value is roughly an indicator of whether the price fits the expected valuation given by the belief. If p and $\hat{\mu}$ are both high or both low, ε must have an extreme slope and thus the optimal k is low. If there is a discrepancy between p and $\hat{\mu}$, \hat{d} is closer to zero and thus k is higher. This shows that search is used more extensively if the consumer has reason not to trust the price. Figure 2 depicts this effect.

Note that we can rewrite the utility in the form

$$u_s(p, \hat{\mu}, k) = [\hat{\mu}(q^H - p) + (1 - \hat{\mu})(p - q^L)](-\varepsilon(k)) - k + \hat{\mu}(q^H - p)$$

so that the function is strictly concave in k in the range $(0, \bar{k})$ for all values $\hat{\mu} \in (0, 1)$ and $p \in [q^L, q^H]$. The first order condition thus provides the interior solution if and only if there is one.

It follows that the utility maximizing search cost for the consumer problem is

$$k^*(p, \hat{\mu}) := \begin{cases} 0 & \hat{d} \leq \varepsilon'(0) \\ (\varepsilon')^{-1}(\hat{d}) & \varepsilon'(0) < \hat{d} < \varepsilon'(\bar{k}) \\ \bar{k} & \hat{d} \geq \varepsilon'(\bar{k}). \end{cases} \quad (1)$$

This function is continuous and piecewise differentiable in both arguments. However, its form presents some problem for the analysis, namely that there is a saddle point at $(\frac{q^H - q^L}{2}, \frac{1}{2})$. Figure 2 shows an example of this function. It also shows the effect that search effort is high in the areas in which $\hat{\mu}$ and p do not correspond to each other.

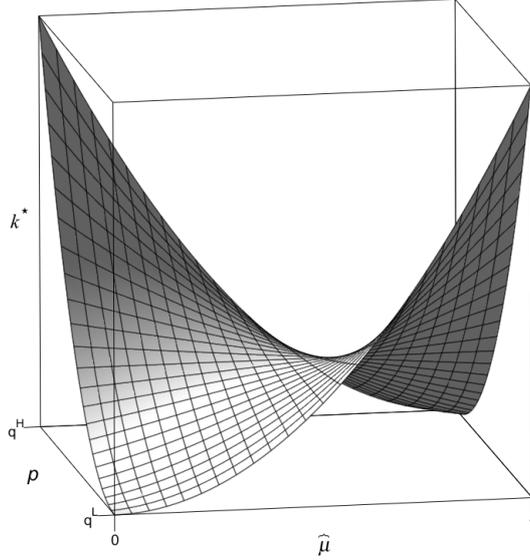


Figure 2: The function k^* for $\varepsilon(k) = \max \left\{ \frac{1}{2} - \sqrt{k}, 0 \right\}$.

As mentioned above, this analysis is based on Lemma 3.1 and thus gives a *necessary* condition. If the consumer optimally pays a positive cost, it has to be given by the function k^* . To ensure that paying this cost and then acting in accordance to the signal (provided k^* is positive) is optimal, the corresponding error probability must be low enough to yield positive utility when the signal is s^H and negative utility in case of receiving s^L . We thus have to test whether k^* meets this condition. In general, this is not the case for all pairs $(p, \hat{\mu}) \in [q^L, q^H] \times (0, 1)$. The following lemma, however, shows that this is never an issue when utility implied by the optimal search behavior exceeds the one from not buying or from buying without extra information.

Definition 3.2. *Let*

$$u_s^*(p, \hat{\mu}) := u_s(p, \hat{\mu}, k^*(p, \hat{\mu}))$$

denote the maximal achievable utility if the consumer was committed to buy if and only if she receives signal s^H .

Lemma 3.3. *Let $(p, \hat{\mu}) \in [q^L, q^H] \times (0, 1)$ be given and denote $\hat{k}^* := k^*(p, \hat{\mu})$. Moreover, assume*

$$u_s^*(p, \hat{\mu}) > \max\{0, \bar{q}_{\hat{\mu}} - p\} = \max\{u_n, u_b(p, \hat{\mu})\}. \quad (2)$$

Then we have $\hat{k}^ > 0$ and the error probability $\varepsilon(\hat{k}^*)$ is low enough so that the consumer buys the product if and only if she receives the signal s^H .*

We denote the optimal utility, given a price p and a corresponding posterior belief $\hat{\mu}$ by

$$u^*(p, \hat{\mu}) := \max\{u_b(p, \hat{\mu}), u_n, u_s^*(p, \hat{\mu})\}.$$

Having the three options of *searching*, *not buying* and *buying without search*, the consumer acquires information if the condition (2) of the previous lemma is met (while there can be mixed strategies in case of equality). We continue by investigating when this is the case and when the consumer prefers either of the two other options, depending on the observed price p and the corresponding posterior belief $\hat{\mu}$. Note that, due to the complicated behavior of the optimal search costs and hence the signal precision, the area in which positive search effort occurs is not trivially well-shaped.

Lemma 3.4. *For all $\hat{\mu} \in (0, 1)$ there are prices $\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}}$ such that*

$$q^L < \underline{p}_{\hat{\mu}} \leq \bar{q}_{\hat{\mu}} \leq \bar{p}_{\hat{\mu}} < q^H$$

and the consumer strictly prefers buying without search whenever the price p is below $\underline{p}_{\hat{\mu}}$, she strictly prefers searching whenever $p \in (\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}})$ and she strictly prefers not buying whenever $p > \bar{p}_{\hat{\mu}}$, provided that $\hat{\mu}$ is the corresponding posterior belief to p .

Figure 3 gives a graphical intuition for how the utility of each of the three options depends on p for a fixed value of $\hat{\mu}$. There is a counter-intuitive effect when $\hat{\mu} > \frac{1}{2}$. The optimal search effort k^* is then decreasing in the price and hence a price increase could have a positive effect for the consumer's utility. The proof of Lemma 3.4 shows that this effect is, however, negligible such that we indeed always have a decreasing behavior of the search payoff in the price variable. The thicker line in Figure 3 depicts the function u^* , the maximum utility value of all three options “search”, “buy” and “don't buy”. Note that u_s^* is not a linear function but the proof shows that its slope is always below zero and above the slope of u_b which leads to the result above.

Having this lemma, we are particularly interested in situations where the interval $(\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}})$ is not empty. As it turns out, this is always the case as long as the marginal gain of signal precision from search effort is sufficiently high at zero.

Lemma 3.5. *For all $\hat{\mu} \in (0, 1)$, the strict inequality $\underline{p}_{\hat{\mu}} < \bar{p}_{\hat{\mu}}$ holds if and only if*

$$\varepsilon'(0) < \frac{-1}{2\hat{\mu}(1-\hat{\mu})(q^H - q^L)}.$$

In this case, one even has $\underline{p}_{\hat{\mu}} < \bar{q}_{\hat{\mu}} < \bar{p}_{\hat{\mu}}$.

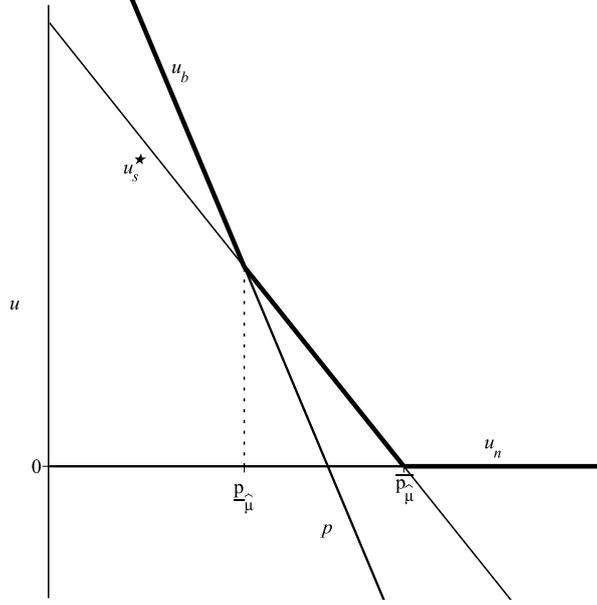


Figure 3: The utility development with p and determination of the prices $\underline{p}_{\hat{\mu}}$ and $\bar{p}_{\hat{\mu}}$ for fixed $\hat{\mu}$. The function u^* is given by the thick upper contour line.

In other words: Every non-degenerate posterior belief can lead to search behavior if the marginal benefit from search is sufficiently high.

Note that search and hence a positive probability of trade exists even with prices above expected quality $\bar{q}_{\hat{\mu}}$. It is of importance for the later analysis that the statement of this lemma is always true if we have $\varepsilon'(0) = -\infty$.

A similar result to Lemma 3.4 is true for the dependence of consumer's behavior on the posterior belief $\hat{\mu}$. This follows from the following, stronger observation.

Lemma 3.6. *The values $\underline{p}_{\hat{\mu}}$ and $\bar{p}_{\hat{\mu}}$ are continuous and piecewise differentiable in $\hat{\mu}$. Moreover, we have*

$$\frac{\partial}{\partial \hat{\mu}} \underline{p}_{\hat{\mu}} > 0 \text{ and } \frac{\partial}{\partial \hat{\mu}} \bar{p}_{\hat{\mu}} > 0$$

for each point in which the respective function is differentiable and

$$\lim_{\hat{\mu} \rightarrow 0^+} \underline{p}_{\hat{\mu}} = \lim_{\hat{\mu} \rightarrow 0^+} \bar{p}_{\hat{\mu}} = q^L \quad \lim_{\hat{\mu} \rightarrow 1^-} \underline{p}_{\hat{\mu}} = \lim_{\hat{\mu} \rightarrow 1^-} \bar{p}_{\hat{\mu}} = q^H.$$

This relatively nice behavior of the lower and upper bound for prices for which search is optimal comes as a surprise considering the shape of the optimal search effort function. It is needless to say that these properties facilitate the following equilibrium analysis.

To give a better feeling for how the three options of “search”, “buy” and “don't buy” are

distributed, we give a graphical example. Figure 4 shows the various areas for $q^H = 1, q^L = .5, \varepsilon(k) = \max\left\{\frac{1}{2} - \frac{3}{2}\sqrt{k}, 0\right\}$. Note that this error function satisfies $\varepsilon'(0) = -\infty$ and hence for every non-degenerate value of $\hat{\mu}$ there is a price for which search is strictly optimal.

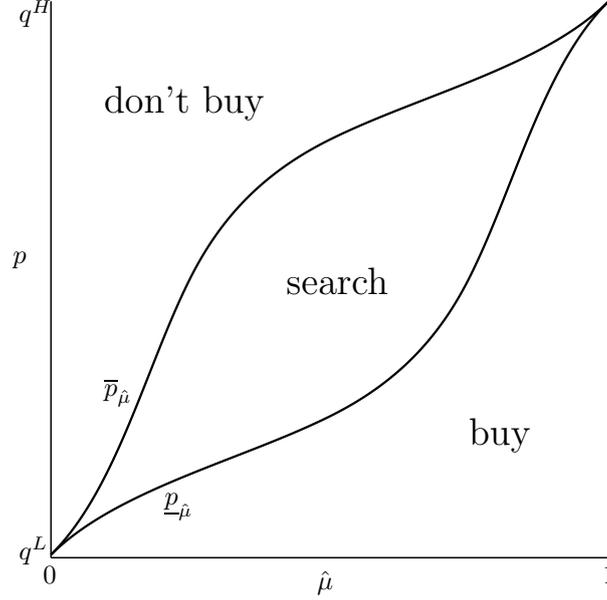


Figure 4: The areas of consumer behavior and the development of $\underline{p}_{\hat{\mu}}$ and $\bar{p}_{\hat{\mu}}$.

4 The Market and Equilibrium Behavior

Having determined the behavior of the consumer, we investigate how this leads to various equilibria. We first need to formally define the game, i.e. the strategies and the equilibrium concept.

Definition 4.1. A consumer strategy is a function $b : \mathbb{R}_+ \rightarrow \mathbb{R}_+ \times [0, 1]^2$ where, for every price p , $b(p) = (k, \gamma^H, \gamma^L)$ denotes the amount of search effort k and the probabilities γ^H, γ^L of buying the product conditional on receiving the high or low signal.

A firm's strategy $a : \{H, L\} \rightarrow \Delta(\mathbb{R}_+)$ is a mapping that maps each type to a probability distribution over the price space \mathbb{R}_+ .

We write a_H and a_L instead of $a(H)$ or $a(L)$. Using Lemma 3.1 of the previous section, we know that the consumer optimally either pays a positive search cost and then buys if and only if a positive signal arises or she does not search and buys with a certain probability $\gamma \in [0, 1]$ independent of the signal that does not convey any information⁷.

⁷Of course, still having different probabilities for each (meaningless) signal is possible. It is clear,

Based on this behavior, it is convenient to narrow down the set of possible consumer strategies.

Definition 4.2. A consistent consumer strategy is a strategy where for all $p \in \mathbb{R}_+$ we have $b(p) = (k, 1, 0)$ or $b(p) = (0, \gamma, \gamma)$ with $k > 0$ and $\gamma \in [0, 1]$.

Having this, we give the formal definition of an equilibrium in this setting.

Definition 4.3. Let (a, μ, b) be a tuple where a is the firm's strategy, $\mu : \mathbb{R}_+ \rightarrow [0, 1]$ is a posterior belief system of the consumer and b is a consistent consumer strategy.

This tuple is an equilibrium if

- Every price in the support of a_H and a_L maximizes the profit of the respective type
- μ is determined by Bayes' law whenever possible⁸
- b maximizes the consumer's utility with respect to $\mu(p)$ for each prize p .

Note that this is similar to the classical weak Perfect Bayesian Equilibrium as used in the text book by Mas-Colell et al. (1995) but adapted to the general strategy space of this model. An equilibrium in which $a_H = a_L$ is called a *pooling equilibrium* while a *separating equilibrium* is one in which the supports of a_H and a_L have an empty intersection. We call any other equilibrium a *hybrid equilibrium*.

For an equilibrium $EQ = (a, \mu, b)$, any price p that is in at least one of the supports of a_L or a_H is called an *equilibrium price of EQ*. If additionally $b(p)$ has the form $(k, 1, 0)$, we call p a *search price* of EQ , otherwise a *no-search price*. Abusing notation, we denote $u(p, \mu(p), b(p))$ the consumer's expected utility when observing a price p with corresponding posterior belief $\mu(p)$ and playing strategy $b(p)$. For each firm type θ we define the equilibrium profit $\pi^\theta(EQ)$ as the expected profit when setting a price in the support of a_θ . This value is well-defined by the first point in the equilibrium definition.

It is trivial to see that there can be a separating equilibrium in which the high type always sets price q^H , the low type sets the price q^L , the consumer has the belief system $\mu(p) = \mathbf{1}_{\{p=q^H\}}$ and only buys for prices smaller than or equal to q^L . For this to actually be an equilibrium, one must have $c^H \geq q^L$ so the high quality firm has no strict incentive to set the price q^L . This equilibrium would also occur if one did not allow for information acquisition and is present in many other models of markets with quality uncertainty. Note that high quality is not traded at all in this setting. We thus refer to this constellation as the *total adverse selection* (TAS) equilibrium.

however, that playing a strategy $b(p) = (0, \alpha, \beta)$ is equivalent to playing $b(p) = (0, \gamma, \gamma)$ with $\gamma = \frac{1}{2}\alpha + \frac{1}{2}\beta$.

⁸This point is often not precisely formulated in the literature. Formally, we apply the classic version of Bayes' law for every price p where $a_H(\{p\}) + a_L(\{p\}) > 0$. For prices that are in the support of exactly one of the two distributions, we assume that the posterior belief is either 1 or 0, according to the type that uses p . No restriction is made for prices that are in both supports but have probability 0.

To emphasize the relation to the classical model of quality uncertainty and the perfect information case, we briefly discuss these two cases.

With full information, the situation is quite obvious. Since the consumer always knows the type, the firm can always demand the true value q^θ and the buyer buys with probability one. Otherwise, any slightly lower price would lead to sure buying and thus causes the firm to deviate. There are no other equilibria.

If the consumer had no possibility of obtaining information about the product quality, the described situation corresponds to a lemon market model in the spirit of Akerlof that is similar but not quite equal to the analysis of Ellingsen (1997)⁹. It appears as a boundary case of our model if we set $\varepsilon(k) = \frac{1}{2}$ for all k (which, of course, would not satisfy the assumptions). In that setting, if $c^H \leq \bar{q}_\eta$ ¹⁰, pooling equilibria exist for a price in $[\max\{q^L, c^H\}, \bar{q}_\eta]$ while separating equilibria with prices q^L and q^H always exist in which the low quality firm sells with probability one and the high quality firm with a probability in

$$\left[\max \left\{ 0, \frac{q^L - c^H}{q^H - c^H} \right\}, \frac{q^L - c^L}{q^H - c^L} \right].$$

In particular, the total adverse selection equilibrium exists if and only if $q^L \leq c^H$ as was already observed in the setting of this paper. There are other, hybrid equilibria in Ellingsen's setting. While they are disregarded due to his refinements and although his analysis is not completely applicable to this setting, such equilibria also appear here.

4.1 Equilibrium Analysis

We start with observing some rather obvious and intuitive features that are quite standard and can be found in similar form in other models. They are nevertheless important for the analysis of equilibria.

Lemma 4.4. *In every equilibrium, the following statements hold.*

- i) The support of a_L is a subset of $[q^L, q^H]$, the support of a_H is a subset of $[q^L, \infty)$.*
- ii) The low type does not set the price q^H with positive probability.*
- iii) The low type's profit is weakly larger than $q^L - c^L$.*
- iv) Every price in (q^L, q^H) is either in both supports of a_L and a_H or in neither.*

These points are not surprising considering the nature of an equilibrium. Any price below q^L would induce sure buying and thus always yield a lower profit than a higher price with the same property. The low type thus always has the option to deviate to a price arbitrarily close to q^L and to receive a profit close to $q^L - c^L$ which shows iii).

⁹Ellingsen assumes equal differences between valuation and production costs for each type, thus corresponding to the case $q^H - c^H = q^L - c^L$.

¹⁰Although Ellingsen excludes this case in his paper, the set of pooling equilibria is easy to derive. The separating equilibria are the same.

For the low type, setting a price q^H or higher with positive probability would lead to a posterior belief below 1 and hence the consumer does not buy. The resulting profit is zero and contradicts iii).

If a price is set by one type but not by the other, the equilibrium definition implies that the consumer knows the true quality. If it was a low type and the price was above q^L , this would result in not buying at all, making it irrational for the low type to set this price. On the other hand, a price below q^H set by only the high firm would result in sure buying and this would attract the low quality firm to imitate that behavior. The formal versions of these arguments can be found in the appendix.

Since the consumer never buys a product for a price higher than q^H , every such strategy is at least weakly dominated by any price in $(c^H, q^H]$. We thus assume that also the high type does not set a price above q^H .

We now know that, apart from the prices q^L and q^H , every price is either in both types' support or in neither of them. However, there could in principle still be a large number of such prices, making further analysis even more complicated by adding measure theoretic obstacles. We show that this is in fact not the case and that there cannot be more than two such non-boundary prices played in equilibrium.

Lemma 4.5. *In an equilibrium, there are no two prices that are in both supports of a_L and a_H and for which the consumer searches.*

Lemma 4.6. *In an equilibrium, there cannot exist two different prices that are in both supports and for which the consumer does not search.*

For both of these lemmas, the first property of the equilibrium definition implies that both types must be indifferent between the prices in the support of their price distribution. In the proof, we show that this can not be the case for two search prices or two no-search prices. It is, however, possible that both types are indifferent between a search price and a no-search price.

These observations already significantly reduce the set of possible equilibrium strategies. Although we put no a priori restrictions on the firm's price-setting behavior, in equilibrium, each type does not play more than two prices in the set (q^L, q^H) .

If the error function ε satisfies an additional, Inada-like condition, we can rule out even more equilibria. As seen in Lemma 3.5, the value of $\varepsilon'(0)$ is of importance when it comes to determining the consumer reaction. It has to be low enough to ensure the existence of a search price for any given posterior belief $\hat{\mu} \in (0, 1)$. The bound itself depends on this belief and hence may vary between different equilibria or even between different equilibrium prices. It is hence convenient to define the following property.

Definition 4.7. *An error function ε satisfies the assumption (I) if*

$$\varepsilon'(0) = -\infty.$$

Having this, we can even go further in narrowing down the set of equilibria.

Lemma 4.8. *Assume that ε satisfies assumption (I) and let $p \in [q^L, q^H)$ be a no-search equilibrium price. Then $b(p) = (0, 1, 1)$ so that the consumer buys with probability one.*

This statement follows from Lemma 3.5. If the consumer buys with a probability in $(0, 1)$, she is indifferent between buying and not buying, hence $\bar{q}_{\mu(p)} = p$. This price, however, leads to search when $\varepsilon'(0)$ is low enough. Having $b(p) = (0, 0, 0)$ would give zero profit to both firms and thus violates Lemma 4.4 iii).

The previous lemmas now allow us to define quite precisely the form of possible equilibria in the model.

Proposition 4.9. *If assumption (I) is satisfied, in every equilibrium the inclusions*

$$\begin{aligned} \text{supp}(a_L) &\subset \{q^L, p_s\} & \text{supp}(a_H) &\subset \{p_s, q^H\} \\ \text{or } \text{supp}(a_L) &= \{p_1\} & \text{supp}(a_H) &\subset \{p_1, q^H\} \end{aligned}$$

hold where p_s is a search price, q^L and p_1 induce sure buying and if q^H is played, we have $b(q^H) = (0, \gamma, \gamma)$ with γ low enough to not attract the low type firm.

Summarized, these are the different types of potential equilibria in the model

- Separating adverse selection equilibria
- Pooling equilibria without search in which both types set the same price $p_1 \leq \underline{p}_\eta$
- Pooling equilibria with search and a price $p_s \in [\underline{p}_\eta, \bar{p}_\eta]$
- Hybrid equilibria in which the high type firm demands a high search price p_s and the low type plays $a_L(p_s) = \alpha$, $a_L(q^L) = 1 - \alpha$ for some $\alpha \in (0, 1)$.
- Other equilibria with $q^H \in \text{supp}(a_H)$ and $b(q^H) = (0, \gamma, \gamma)$, $\gamma > 0$.

All these equilibria exist provided the buyer and the high type firm make non-negative profit and the low type earns at least $q^L - c^L$. We denote the pooling search equilibrium with the highest possible price \bar{p}_η as PE_s and the pooling no-search equilibrium with the price \underline{p}_η as PE_b . If at least one hybrid equilibrium exists, the one with the highest search price p_s is denoted as HE . These are the important equilibria due to the following robustness check.

The set of potential equilibria is significantly narrowed down but still too large to draw qualitative conclusions from the model. In what follows, we argue in which way some of these equilibria, and in particular the belief systems by which they are supported, can be disregarded.

4.2 Selection of Equilibria

There are various, well established refinements to eliminate implausible equilibria in signaling games. Bester and Ritzberger (2001) use a modification of the well-known Intuitive Criterion introduced by Cho and Kreps (1987). In this model, as well as in

theirs, the original version of the Intuitive Criterion is not sufficient. The modification used by Bester and Ritzberger, however, is not well defined in our setting since the firms' profit functions are not monotone in beliefs. We thus follow another approach of arguing which consumer beliefs are unconvincing and hence rule out the equilibria supported by these beliefs.

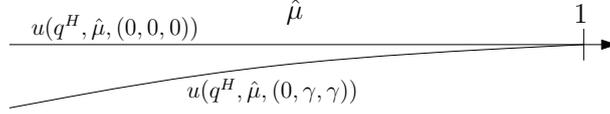


Figure 5: Buying for the price q^H leads to negative utility if $\hat{\mu} < 1$

To illustrate the idea of the following refinement, consider an adverse selection equilibrium in which the high type firm makes positive profit, i.e. a separating equilibrium in which the low type sets price q^L , the high type price q^H and the consumer buys the high quality product with some probability $\gamma > 0$. From the equilibrium property we must have $\mu(q^H) = 1$ so the consumer knows the quality when she sees the high price. Note that she is then indifferent between buying and not buying since the price matches her valuation. If she had any doubts about her posterior belief $\mu(q^H)$, i.e. if she admits that there is even the smallest possibility to be wrong about her belief, “not buying” would be strictly better than her strategy $b(q^H) = (0, \gamma, \gamma)$. Since “not buying” is optimal even for her rational belief $\hat{\mu} = 1$, her strategy is dominated in a certain sense. This idea is depicted in Figure 5 and formally written down in the following refinement.

Definition 4.10. *Let $p, \mu(p)$ be given. The action $b \in \mathbb{R}_+ \times [0, 1]^2$ is locally dominated in beliefs if there exists another action $b^* \in \mathbb{R}_+ \times [0, 1]^2$ and a $\delta > 0$ such that*

$$u(p, \hat{\mu}, b^*) \geq u(p, \hat{\mu}, b) \quad \forall \hat{\mu} \in (\mu(p) - \delta, \mu(p) + \delta) \cap [0, 1]$$

and the inequality is strict for $\hat{\mu} \neq \mu(p)$.

An equilibrium (a, μ, b) has belief-robust responses if for no equilibrium price p and corresponding belief $\mu(p)$ the action $b(p)$ is dominated in beliefs.

This condition reflects some doubts about the posterior beliefs. A best response b which violates this criterion is not a strict one, meaning that there is another best response b^* to $(p, \mu(p))$ that yields the same payoff. Moreover, choosing b over b^* is not a robust behavior and only rational if the buyer is absolutely confident about the firm's strategy.

The criterion is one of local robustness of the strategy. Other criteria in the same spirit can be found in the literature, for example the *robust best reply* definition in Okada (1983).

Note that this condition does not in general rule out mixed strategies of the consumer. In this case, however, it leads to eliminating all equilibria in which the buyer plays a mixed strategy for the highest possible price q^H , including the classic adverse selection equilibria, mentioned above, in which the high type makes positive profits.

Lemma 4.11. *Let assumption (I) be satisfied. For an equilibrium, the following is equivalent.*

- i) The equilibrium has belief-robust responses.*
- ii) The price q^H is not an equilibrium price or the equilibrium is the separating equilibrium with total adverse selection.*

The reason why most adverse selection equilibria are ruled out is not specific to this setting. In fact, a similar refinement excludes these equilibria e.g. in Ellingsen (1997).¹¹ In that paper, he uses another refinement under which only the separating equilibrium with the highest possible high type trade probability survives. While this is a legitimate approach, the richness of equilibria in our setting allows us to exclude these equilibria and still obtain interesting results.¹²

As one can see, a lot of the equilibria survive this refinement. This gives us the opportunity to address another issue of implausible consumer behavior, namely the possibility of extreme belief changes.

Imagine two situations in which the consumer observes a price p or a similar price that is very close to p . It does not seem intuitive that the posterior beliefs should differ too much, especially if we let the difference of the two prices be arbitrarily small. Even if one admits that real prices usually can not differ by less than one cent, posterior beliefs that assign $\mu(p) = 1$ and $\mu(p + 0.01) = 0$ seem quite extreme. In fact, marginal price changes are often due to retailer behavior and may not even be perfectly perceived by consumers.¹³ It is thus more realistic that the consumer acknowledges the closeness of the prices by assigning a similar posterior belief. Formally, we postulate continuity of beliefs in those prices that actually occur in equilibrium.

Definition 4.12. *An equilibrium (a, μ, b) satisfies the locally continuous beliefs condition if for every equilibrium price p the function μ is continuous in p .*

Local continuity is not a very strong assumption considering that it just excludes jumps in beliefs but still allows for arbitrarily strongly increasing or decreasing posteriors. The described behavior for the one-cent difference in the motivating example would actually still be possible under locally continuous beliefs. However, this slight step has a big impact on the number of equilibria.

¹¹Compare Proposition 5 of Ellingsen (1997). Note that elimination of strategies that are locally dominated in beliefs could be substituted by elimination of weakly dominated strategies in this paper without changing the results.

¹²Interestingly enough, Ellingsen justifies using his other refinement by saying “in reality, a seller will typically not know exactly the buyer’s valuation” which is true. In the same spirit, however, the idea that the consumer might not be perfectly confident about her posterior beliefs should not be ignored. Ellingsen’s idea of “elastic demand” is incorporated in the next refinement and thus our approach covers both aspects of imperfections to some extent.

¹³See Zeithaml (1988) for an overview on perception of price and other product characteristics by consumers.

Before we determine the consequences of this refinement, note that it usually¹⁴ rules out the pooling equilibrium without search (PE_b) if we have $\varepsilon(k^*(\underline{p}_\eta, \eta)) = 0$ so that there is perfect search for the border case of a pooling equilibrium price in which “buy” and “search” yield the same outcome to the consumer. The reason for this is that, with continuous beliefs, the high quality firm would want to deviate to a slightly higher price than \underline{p}_η which, because the consumer receives a perfect signal, also yields a selling probability of one for high quality products.

Proposition 4.13. *Let assumption (I) be satisfied. The strategies (a, b) can form an equilibrium with a posterior belief system that satisfies locally continuous beliefs and such that it has belief-robust responses if and only if they are the strategies of one of the following equilibria:*

- *the pooling no-search equilibrium PE_b with price \underline{p}_η . This equilibrium exists if and only if $\underline{p}_\eta \geq c^H$ and $\varepsilon(k^*(\underline{p}_\eta, \eta)) > 0$*
- *the pooling search equilibrium PE_s with price \bar{p}_η . It exists if and only if $\bar{p}_\eta \geq c^H$ and $\pi^L \geq q^L - c^L$*
- *hybrid equilibria in which the high quality firm sets a price $p = \bar{p}_{\hat{\mu}}$ and the low quality firm sets this price with probability $\alpha \in (0, 1)$ while setting q^L with probability $1 - \alpha$ and we have $\hat{\mu} = \mu(p) = \frac{\eta}{\eta + \alpha(1 - \eta)} > \eta$. This equilibrium exists if and only if $\bar{p}_{\hat{\mu}} \geq c^H$ and $\pi^L = q^L - c^L$.*
- *the total adverse selection equilibrium (TAS). It exists if $c^H \geq q^L$.*

While three of these equilibria are unique within their class if they exist, there may be multiple hybrid equilibria. Every value $\hat{\mu} > \eta$ for which the equation

$$\varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu})) \cdot (\bar{p}_{\hat{\mu}} - c^L) = q^L - c^L$$

holds yields such an equilibrium if c^H does not exceed the price $\bar{p}_{\hat{\mu}}$. The reason is that a low quality firm must be indifferent between the prices $\bar{p}_{\hat{\mu}}$ and q^L . Figure 6 shows such a constellation in which not only multiple hybrid equilibria but also the pooling search equilibrium PE_s exist at the same time. It is useful to note that the existence of a hybrid equilibrium implies

$$\begin{aligned} & \frac{1}{2}(q^H - c^L) > \varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu})) \cdot (\bar{p}_{\hat{\mu}} - c^L) = q^L - c^L \\ \Rightarrow & \quad \quad \quad q^L - c^L < q^H - q^L. \end{aligned}$$

¹⁴There can be cases in which the pair $(\underline{p}_\eta, \eta)$ is exactly on the border defined by (1) so that there might be a continuous “path” $\mu(p)$ of posterior beliefs under which the PE_b equilibrium can be sustained. Since this is a non-generic case, we omit the detailed analysis and just write $\varepsilon(k^*(\underline{p}_\eta, \eta)) > 0$ as condition for the existence of PE_b .

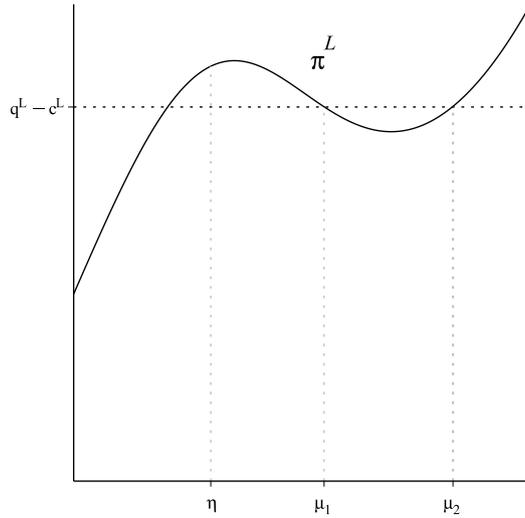


Figure 6: The coexistence of PE_s and multiple hybrid equilibria. The function depicts the low type profit for each $\hat{\mu}$ when setting the price $\bar{p}_{\hat{\mu}}$. The values μ_1 and μ_2 are the posterior beliefs of search prices in hybrid equilibria.

It is a common result in lemon markets that all equilibria are not efficient (so that some goods are not traded with full probability) or the consumer has a chance of buying a good for a higher price than his valuation. We also observe this, here. Note that, although we focus on take-it-or-leave-it offers, the famous result of Myerson and Satterthwaite (1984) suggests that this can not be overcome when using a different mechanism.¹⁵

To give an overview over the qualitative implications of these equilibria, their properties are summarized in the following table. The “ \gtrsim ” symbol indicates generic strict inequalities, i.e. the set of parameters for which equality occurs is a Lebesgue null set in the parameter space.¹⁶

¹⁵Their formal result does not apply here. To give the connection, production costs c can be seen as the seller’s valuation, q as the buyer’s value. In contrast to the original result, they are not independent and not drawn from an interval $[c^L, c^H]$, $[q^L, q^H]$. The only efficient (unrefined) equilibrium that guarantees non-negative profits and consumer surplus in every outcome is the pooling equilibrium on the price q^L . It exists if and only if $c^H \leq q^L$ which directly translated to a violation of the assumption of Myerson and Satterthwaite that the intervals $[c^L, c^H]$ and $[q^L, q^H]$ overlap.

¹⁶For example, in the PE_s equilibrium if \bar{p}_{η} happens to be exactly c^H , the high type makes no profit. The value of \bar{p}_{η} does not depend on c^H so this is a Lebesgue null set.

	π^H	π^L	consumer utility	existence condition
PE_b	$\gtrsim 0$	$> q^L - c^L$	> 0	$\underline{p}_\eta \geq c^H$ and $\varepsilon(\underline{p}_\eta) > 0$
PE_s	$\gtrsim 0$	$\gtrsim q^L - c^L$	0	$\bar{p}_\eta \geq c^H$ and $\pi^L \geq q^L - c^L$
hybrid	> 0	$q^L - c^L$	0	$\exists \hat{\mu} > \eta : \bar{p}_{\hat{\mu}} \geq c^H$ and $\pi^L = q^L - c^L$
TAS	0	$q^L - c^L$	0	$c^H \geq q^L$

Table 1: The properties of equilibria surviving the refinements

This table shows an interesting aspect especially about the consumer utility. There is only one equilibrium in which she has positive utility and this does not involve search. The possibility of search does not allow the PE_b equilibrium to have a higher price than \underline{p}_η . Remember that in the classic lemon market this price would be equal to \bar{q}_η so that the consumer had zero expected utility if we apply the same refinements to the pooling equilibria of the classical case. Introducing search can thus benefit the consumer but only if she does not use this new “ability”. Naturally, this consumer friendly equilibrium only exists if the price is still high enough for a high quality firm to make positive profit. It also shows, however, that if the optimal search effort on the pooling price leads to perfect information, this equilibrium fails the refinements. In this case, the consumer’s ability to search destroys her only chance of having positive utility. We elaborate on this effect in the next section.

Quality uncertainty situations being famous for their adverse selection effects, we can now investigate how the model behaves in this regard. The following shows that introducing search, as one would assume, indeed reduces the advantage of low quality goods over high quality goods in terms of traded amount.

Observation 4.14. In PE_b and PE_s , a high quality firm has a weakly higher probability of selling the good than the low type. In any hybrid equilibrium of Proposition 4.13, the probability for a high firm of selling the good is higher than in any separating adverse selection equilibrium.

Note that in a hybrid equilibrium, the low type firm can have a higher chance of selling its good than the other type. This value is $1 - \alpha + \alpha\hat{\varepsilon}$ where α is the share with which it sets the high search price and $\hat{\varepsilon}$ is the error probability of that price.

To go even further, observe that there is a partial ranking in Pareto dominance between the existing equilibria.

Definition 4.15. An equilibrium (a, μ, b) Pareto dominates another equilibrium $(\tilde{a}, \tilde{\mu}, \tilde{b})$ if the equilibrium payoffs satisfy

$$\pi^H \geq \tilde{\pi}^H, \pi^L \geq \tilde{\pi}^L \text{ and } u^* \geq \tilde{u}^*$$

and at least one of these inequalities is strict.

This definition of Pareto dominance is taken after the quality of the firm is revealed, thus taking each type’s profit into account separately. This gives a stricter version than

an a priori Pareto dominance in which one would only consider the expected profit before the firm learns its type. However, an interesting dominance ranking holds even with this stronger condition.

Lemma 4.16. *Ignoring non-generic cases, the following items reflect the full Pareto dominance ranking between the equilibria of Proposition 4.13.*

- *If multiple hybrid equilibria exist, the one with the highest search price (HE) dominates the others.*
- *TAS is dominated by PE_b , PE_s and HE whenever one of these equilibria exists.*
- *PE_s and HE are dominated by PE_b if and only if $\pi^H(PE_s) \leq \pi^H(PE_b)$ or $\pi^H(HE) \leq \pi^H(PE_b)$, respectively.*

It is quite natural to observe that the equilibria PE_s and HE are somehow similar. In both equilibria, there is a search price on the upper border of the search area and the consumer has zero utility. Indeed, the coexistence of these equilibria is rare and does never occur if the probability of having high quality is sufficiently high.

Lemma 4.17. *Let $q^L - c^L < q^H - q^L$. There is a lower bound $\underline{\eta} < \frac{1}{2}$ such that whenever $\eta > \underline{\eta}$, there exists either PE_s or HE provided that the search price of one of these equilibria exceeds c^H .*

The reason for having this lower bound lies in the profit of the low quality type. It strictly increases when the posterior $\hat{\mu}$ goes from $\frac{1}{2}$ to 1 and the price is $\bar{p}_{\hat{\mu}}$. Thus, the PE_s condition $\pi^L \geq q^L - c^L$ implies that for all higher beliefs the low type's profit is even larger. In HE , however, the profit must exactly attain this bound. The situation in Figure 6 corresponds to a case in which $\eta < \underline{\eta}$. In this figure, $\underline{\eta}$ can be chosen to be μ_1 .

Notice also that the actual value of η might be zero so that the negative profit effect of losing customers never outweighs the positive effect of a higher price for the low type firm. The condition $q^L - c^L < q^H - q^L$ of the lemma follows from the existence condition of a hybrid equilibrium. If this is violated, a hybrid equilibrium can never exist.

The so far established results already shed some light on how the market outcome is influenced by introducing information acquisition costs in the classical model of quality uncertainty. It shows that if the cost for high quality production is low, a pooling equilibrium without search exists. While this is also true in the classical model, there are qualitatively different aspects, namely that the actual price to pay in the pooling equilibrium is strictly below the average quality valuation and hence the consumer has strictly positive utility. Of course, this effect is caused by the same issue that rules out these equilibria for high quality costs between \underline{p}_{η} and \bar{q}_{η} . In these cases, introducing the possibility of information acquisition leads to search behavior but does not help the consumer.

A different phenomenon can be observed in the PE_s and HE equilibria. They exist whenever \bar{p} is high enough. Since these equilibria contain search prices, they do not

occur in the classical model but dominate and thus eliminate the otherwise existing separating equilibria. Although also the payoffs are different, the main contribution of these equilibria is the weakening or complete elimination of adverse selection phenomena.

5 Search Efficiency

The previous section investigates a market in which the search possibility for the consumer is fixed by the function ε . As mentioned in the introduction, we are also interested in comparing situations in which consumer might have higher or lower costs for searching. Since the actual costs k are endogenously chosen by the consumer, we have to clarify what “lower search costs” means in this setting. It is rather to be viewed as higher “search efficiency” which means that the consumer gets a more precise signal for the same search effort. Think about someone who wants to buy a TV in 2013 or someone in the 1980s. Getting information about a certain product is much easier now than it was back then, due to the Internet, multiple test magazines and websites. It is safe to say that it is both less time consuming and cheaper to get the same amount of information now than it was back then.

To capture this efficiency in the model we introduce a parameter a to the function ε . The extended function satisfies the following properties.

- $\varepsilon : \mathbb{R}_+ \times \mathbb{R}_{++} \rightarrow [0, \frac{1}{2}]$, $(k, a) \mapsto \varepsilon(k, a)$ is continuous.
- For every $a > 0$, $\varepsilon(\cdot, a)$ satisfies the assumptions from section 2 on page 4 and assumption (I). Denote $\bar{k}(a)$ the perfect information cost for parameter a .
- The function is twice continuously differentiable on

$$K := \{(k, a) | 0 < k < \bar{k}(a)\} = \{(k, a) | 0 < \varepsilon(k, a) < \frac{1}{2}\},$$

the area in which information is neither perfect nor meaningless.

- $\varepsilon_a(k, a) := \frac{\partial}{\partial a} \varepsilon(k, a) < 0$ for all $(k, a) \in K$.
- For all $k > 0$ we have $\lim_{a \rightarrow \infty} \varepsilon(k, a) = 0$ and $\lim_{a \rightarrow 0} \varepsilon(k, a) = \frac{1}{2}$.

Having $\varepsilon(\cdot, a)$ satisfy the same conditions as before, one can use the results of the previous model and perform comparative statics by varying parameter a . The fourth point ensures that for increasing a , the signal precision for the same search effort becomes higher. Finally, the last item ensures that in the pointwise limit, the error function reflects perfect information (for $a \rightarrow \infty$) or the classical lemon market without information acquisition (for $a \rightarrow 0$). Hence, it allows us to use the parameter as a mediator between these two widely acknowledged models. One simple example for such a function is

$$\varepsilon(k, a) = \max \left\{ \frac{1}{2} - a\sqrt{k}, 0 \right\}. \quad (3)$$

Most expressions of the previous sections now depend on the new parameter. We denote them in the same way but adding the value a as the last argument of every function.

The aim of this section is to compare the various types of equilibria and their level of price, consumer utility, average quality etc. under a change of search efficiency. Is it also interesting to see whether the limit behavior of the error probability function, when taking $a \rightarrow \infty$ or $a \rightarrow 0$, also leads to market behavior that converges to the equilibria of the classical lemon market or the perfect information case as discussed above.

This analysis is necessarily different from the one of Bester and Ritzberger (2001) since their cost has an exogenously given value and could thus just be directly increased or decreased. Here, the effort level is chosen by a rational consumer. A direct change of the costs can thus not be done. We rather facilitate the access to information by giving more signal precision for the same effort.

5.1 Analytical Results

The first result analyzes the price behavior under pooling equilibria where no search occurs. Remember that the consumer is indifferent between *searching* and *buying without search* in the pooling price \underline{p}_η of these equilibria.

Proposition 5.1. *The price $\underline{p}_\eta(a)$ of the PE_b equilibrium is continuous, piecewise differentiable and non-increasing in a . Moreover,*

$$\lim_{a \rightarrow 0} \underline{p}_\eta(a) = \bar{q}_\eta \text{ and } \lim_{a \rightarrow \infty} \underline{p}_\eta(a) = q^L$$

holds. If $q^L < c^H < \bar{q}_\eta$, there are values $0 < \underline{a} \leq \bar{a}$ such that the pooling no-search equilibrium PE_b exists if $a \leq \underline{a}$ and it does not exist for $a \geq \bar{a}$.

The reason why $\underline{a} \neq \bar{a}$ can not be excluded despite the monotonicity of the price in PE_b is that the condition $\hat{\varepsilon} > 0$ for PE_b to be an equilibrium might be violated for some lower a but be true for a higher search efficiency. This, however, appears for rather special parameters and is not further investigated. A direct consequence of this proposition is that the profit of both firms decreases with increasing search efficiency while the consumer's utility rises in the pooling equilibrium without search. This can be seen by just observing that neither the average quality nor the amount of trade is different between each of these equilibria.

Remember that in the end of the last section we conclude that the existence of PE_b is due to moderate production costs of high quality goods. The proposition provides a similar statement in terms of search costs. Only when search costs are high, pooling no-search equilibria can exist. However, as long as $a < \underline{a}$, making search more efficient lets the equilibrium price decrease and thus gives a higher utility to the consumer. This supports the first intuition that a more efficient way of searching should increase the consumer's power and thus increases her surplus. Note, however, that no search occurs in these equilibria. Instead, all products are sold for a price that decreases with better search efficiency. Here, the possibility of search is rather used as a threat than as a tool. If search gets too efficient, there is no low-price equilibrium and thus only those equilibrium can exist which provide zero consumer utility. In a sense, quality information is too cheap

from the view of the consumer, the producer of a high quality product benefits from the higher information level.¹⁷

One can deduce a similar proposition for the upper bound \bar{p}_η of the consumer's search area.

Proposition 5.2. *For every $\hat{\mu} \in (0, 1)$ the function $\bar{p}_{\hat{\mu}}(a)$ is continuous, piecewise differentiable and non-decreasing in a . Moreover, we have*

$$\lim_{a \rightarrow 0} \bar{p}_{\hat{\mu}}(a) = \bar{q}_\eta \text{ and } \lim_{a \rightarrow \infty} \bar{p}_{\hat{\mu}}(a) = q^H.$$

Inferring to equilibrium behavior from this proposition is not as easy as it was before, since both PE_s and HE make use of search prices. It is not obvious which of these equilibria exist for a given a . However, noticing that the high price of HE is always higher than \bar{p}_η , it follows that the equilibrium search price behavior for $a \rightarrow \infty$ is not influenced by this question.

Corollary 5.3. *Let $c^H > q^L$ and $q^L - c^L < q^H - q^L$. For $a \rightarrow \infty$, all undominated equilibria converge to a separating state in which the high type sells with probability $1 - \frac{q^L - c^L}{q^H - c^L}$.*

This result is quite interesting in the background of the Grossman-Stiglitz paradox which stems from their 1980 paper. Essentially they argue that in situations in which all arbitrage opportunities are eliminated because all available information is reflected in the prices, there is no incentive for any market participant to obtain this information. Hence the assumption of freely available information would not be justified anymore. While their reasoning originally applied to financial markets, it is in the same way questionable how (in the perfect information case) quality information can be public knowledge when prices perfectly signal the quality and thus the consumer does not need to obtain this information. Here, we see that even with arbitrarily easily accessible information, this paradox does not occur since there is always an incentive for the consumer to obtain information. However, this comes with the cost of the high type not being able to sell with full probability.

If the condition $c^H > q^L$ was violated, the pooling equilibrium without search PE_b would exist for all a . The second condition ensures the existence of HE even for high values of a . If this inequality is not true, neither of the equilibria PE_b , PE_s and HE exist for high values of a . It follows that the TAS equilibrium is the only equilibrium and hence efficient information would lead to an even worse outcome for the high quality firm. Since every equilibrium in this setting has only a finite number of prices and we are not interested in the out-of-equilibrium beliefs, convergence of equilibria is just taken as convergence of the firms' actions.

This result is striking in that even though the information costs approach zero, the consumer still makes errors and does not get perfect information. The firm's behavior

¹⁷In a related but weaker result, Bar-Isaak, Caruana and Cu nat (2012)[?] show, in a setting with quality investments, fixed prices and an information acquisition procedure close to Bester and Ritzberger (2001), that consumer utility can be non-monotone in the acquisition costs. The reason for this effects are, however, quite different from our setting.

converges to the separating state that would appear in the perfect information case but the probability of selling does not approach one. Note that there is an important difference to the similar result by Bester and Ritzberger (2001). They show that in their setting, letting the information costs k approach zero, a non-vanishing share of consumers still acquires information but the high type still sells to all buyers that would buy the product under perfect information. In contrast, the limit behavior at hand shows a sustainable loss in sales for the high quality type.

In the setting of Voorneveld and Weibull (2011), they show that for the limit of perfect information there exists a continuum of limit equilibria between the perfect information equilibrium and the one given in Corollary 5.3. Their result is thus similar but weaker. Moreover, they argue that between these equilibria, the perfect information equilibrium with full sales of high quality Pareto dominates all the other limit situations.

The analysis of equilibria in the direction of high search costs is quite tricky. When search becomes inefficient, the prices of PE_b and PE_s necessarily converge to \bar{q}_η and the question if this is an equilibrium depends on whether the high quality production costs exceed this value. If they do, PE_b and PE_s do not exist for low values of a but TAS does. In other words, these equilibria converge to the two possible classic lemon market equilibria with the exact same existence conditions. One might expect that this price convergence is also true for HE since $\bar{p}_\mu(a)$ is the high quality price. The issue here is that in HE , the posterior $\hat{\mu}$ for the search price itself depends on a so that the convergence result from Proposition 5.2 does not apply and may, in fact, not be true for the equilibrium price.

Proposition 5.4. *Let $\eta(q^H - q^L) < q^L - c^L < q^H - q^L$ and $c^H < 2q^L - c^L$. Then, for each value of a which is close enough to 0, HE exists and for $a \rightarrow 0$ these equilibria converge to a semi-separating state with the high type setting $\hat{p} := 2q^L - c^L$ while the low type mixes between this price and q^L .*

The limit state of the HE equilibria involves a high price $\hat{p} = 2q^L - c^L$ and a posterior belief $\mu(\hat{p}) = \frac{q^L - c^L}{q^H - q^L}$ while the probability of selling is $\frac{1}{2}$. If $q^H - c^H = q^L - c^L$, this situation corresponds to one of the *semi-mixed equilibria* computed by Ellingsen (1997) in his model without information acquisition. There, this equilibrium exists with many others of a similar type and fails the D1 refinement by Cho and Kreps (1987). However, the behavior here shows that this equilibrium, if it exists, is robust with respect to costly information acquisition.

5.2 A Numerical Example

To increase the understanding of what happens with various levels of search efficiency, we continue with a concrete example. Even with a relatively simple error function satisfying the assumptions, the model is too complex to solve for explicit expressions of the various equilibria. We thus rely on numerical calculations to illustrate the results on the development of the model outcomes.

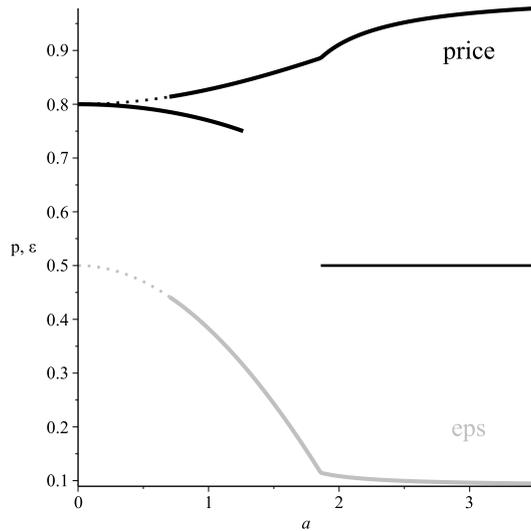


Figure 7: Price and error probability development for increasing search efficiency

For these calculations we choose the error function (3) and set $\eta = .6$, $q^H = 1$, $q^L = \frac{1}{2}$, $c^H = \frac{3}{4}$ and $c^L = .45$. Note that since $c^H < \bar{q}_\eta = .8$, the existence of pooling no-search equilibria is possible for small values of a . Choosing $\eta \geq \frac{1}{2}$ ensures, using Lemma 4.17, that PE_s and HE do not exist at the same time to obtain clearer pictures. Corollary 5.3 is applicable to this setting but Proposition 5.4 is not. We give a second example to illustrate its result.

Figure 7 shows the development of the price and the corresponding error probability for the various search prices. The thin black line is the price q^L for the low quality in HE . To show the convergence even of the dominated pooling search equilibria, their values are displayed as dotted lines. One observes that for low values of a , the pooling no-search equilibrium exists and that its price is decreasing. While at first the other equilibrium is dominated, there is an interval of values of a in which both pooling equilibria exist at the same time until the lower price falls under the production costs of the high type. As search gets more efficient, the error probability decreases and converges to $\frac{q^L - c^L}{q^H - c^L} \approx 0,091$ as predicted by Corollary 5.3.

Figure 8 illustrates the profit development of the different types in these equilibria.

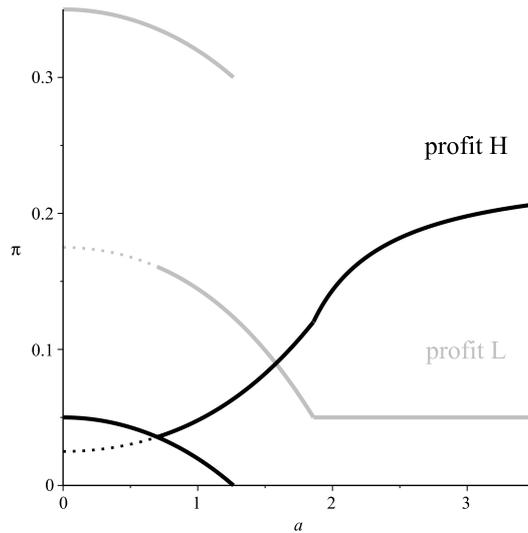


Figure 8: The profits of both types, depending on the search efficiency and the equilibrium. Dotted Lines show Pareto dominated equilibria.

One can see very well how profits decrease in PE_b and how PE_s stops being Pareto dominated when its high type's profit catches up with the one from PE_b . The low type's profit never falls below $q^L - c^L = .05$ and attains this value in the hybrid equilibrium.

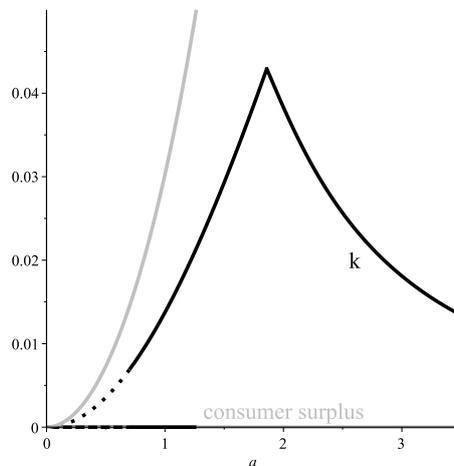


Figure 9: Utility and search effort of the consumer

We look at the consumer side of the market in Figure 9. The utility of the consumer behaves exactly as predicted, dropping to zero when PE_b does not exist anymore. An interesting effect is observable for the search costs. Although search gets more efficient

with increasing a , the absolute effort increases in the pooling search equilibrium. This effect has already been observed by Bester and Ritzberger (2001) and is reproduced here under a different model of search behavior. It shows how the power given to the consumer by allowing for information acquisition can be exploited by the high type to increase prices, search behavior and thus also the probability of selling the high quality product.

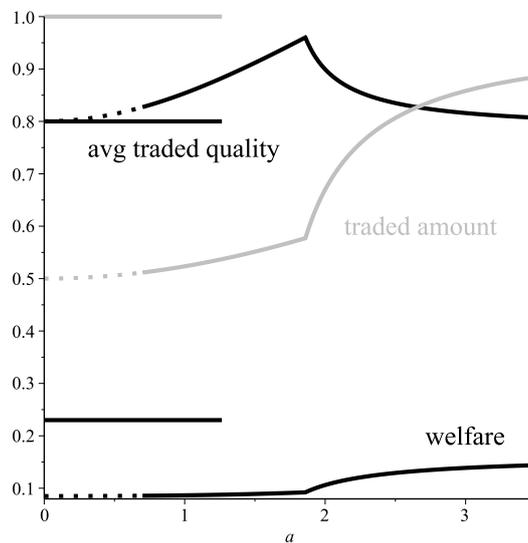


Figure 10: The development of various market characteristics

Finally, Figure 10 depicts the development of some market figures. Welfare here is simply computed as a sum of the firm's expected profit $\eta\pi^H + (1-\eta)\pi^L$ and the consumer utility. The high welfare value in PE_b stems not only from the higher consumer utility but mostly from the fact that all products are traded with probability one and no utility is "wasted" on search effort. All possible gains from trade are thus exploited and distributed among the market participants.

In contrast to markets with quality uncertainty being famous for their adverse selection phenomena, average traded quality in this setting is even higher than the offered one. This is of course due to the higher trade probabilities for the high type on search prices and thus occurs in PE_s and HE (except for very high search efficiency). As search gets more efficient, the low type shifts its price distribution more to q^L and hence sells with an overall higher probability which causes average traded quality to go down and, because of the not vanishing error probability, to go even below the a priori expected quality.

To also give an example of Proposition 5.4, we consider the case where $q^H = 1, q^L = 0.5, c^H = 0.7, c^L = 0.2$ and $\eta = .35$. Note that this also implies that PE_b and PE_s do not exist if a is close to 0 because of their low price close to $\bar{q}_\eta = .675$. The values are explicitly chosen so that they fit the model of Ellingsen (1997) in which we have

$q^H - c^H = q^L - c^L$. As mentioned before, his refinements forecast a separating equilibrium for this case without information acquisition. This result may serve as a hint that these equilibria might not be disregarded, after all.

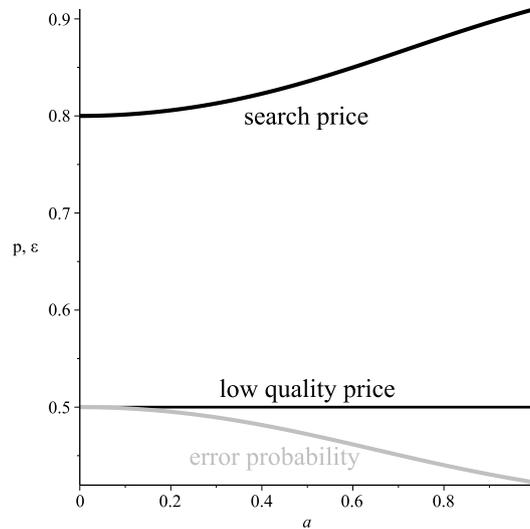


Figure 11: The convergence of HE for vanishing search efficiency.

6 Conclusion and Possible Extensions

The paper shows the outcome of a monopolistic market with quality uncertainty in which the consumer has the possibility to costly acquire information about the product quality. This information could be perfect or imperfect, the exact precision depending on the endogenous search effort exerted by the consumer.

Given the optimal consumer behavior, the market offers many possible equilibria, some of which are already present in the classical model without information acquisition. After eliminating implausible and Pareto dominated equilibria, we are left with three main categories of market behavior. In pooling equilibria without search, the consumer has a positive profit and the highest possible welfare was reached. Equilibria which comprise search leave no utility to the consumer but, except for when information is extremely cheap, have an average traded quality that is above the actual average and thus show the opposite effect to the classical adverse selection results on lemon markets. The third category of total adverse selection occurs when high quality production costs are high and search is very efficient. The analysis shows that information acquisition possibilities only benefit the consumer if she does not acquire any information. If she can use her search abilities as a threat rather than actually acquire costly information, she forces prices to be lower than the average quality and thus have positive utility. In contrast, actual quality search in equilibrium leads to a higher market power of the high type and

thus to higher prices. The consumer welfare is zero in these equilibria.

An important contribution of this paper is the comparison of situations with different search efficiencies. We show that an increase in efficiency can benefit or hurt the consumer and that the consumer's utility will with certainty drop to zero after a certain threshold of search efficiency.

At last, the analysis shows that the case of perfect information is not the limit case of high search efficiency. Even when making information acquisition arbitrarily cheap, the probability of consumers receiving a false signal does not vanish. Moreover, this limit error level does not depend on the error function. Making information acquisition inefficient can lead to the same behavior as predicted in the classical models but for a substantial range of parameters the limit equilibrium is one that was previously disregarded.

Starting from this model, certain extensions come to mind to enrich the analysis and lead to a more realistic behavior.

As most other papers in this field, this work does not incorporate competition between multiple firms. The consumer is always confronted with exactly one good and her only choices are on how much to search and if to buy the product. In the same way the firm does not have to worry about actual competition. The type of "rival" it faces exists only theoretically in the head of the consumer who has to figure out which type she is facing. Extending this setting to an oligopolistic market will certainly prove to be interesting.

We only look at one consumer. As we have pointed out, this could be extended to multiple identical consumers without changing the analysis and thus the outcome. In reality, this is not realistic. Besides different valuations or outside options that people might have, the relative efficiency of today's quality information is mainly due to new technologies which in return can not be assumed to be accessible for everyone with the same efficiency, even when restricted to single countries. Some people are more adapted or have better access to these new technologies than others and this difference can be quite severe from one person to another. It is hence important to account for this in a more realistic model.

Other restrictive aspects of the model might be generalized such as the amount of quality levels and signals. Especially this former aspect is important to obtain meaningful results about the relationship between prices and quality in lemon markets with information acquisition.

7 Appendix

Proof of Lemma 3.1. Fix a price p and a corresponding posterior belief $\hat{\mu} \in (0, 1)$, assume that the consumer has paid a cost $k \geq 0$ for the signal precision and denote $\hat{\varepsilon} := \varepsilon(k)$ the error probability. Receiving the high signal s^H , the updated posterior belief is

$$\hat{\mu}_H := \text{Prob}(q^H | s^H) = \frac{\hat{\mu}(1 - \hat{\varepsilon})}{\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}}$$

which is just Bayes' law applied.

The expected quality with respect to this information is then

$$\bar{q}_{\hat{\mu}_H} := \hat{\mu}_H q^H + (1 - \hat{\mu}_H) q^L.$$

With similar calculations, let $\bar{q}_{\hat{\mu}_L}$ be the expected quality on receiving a low signal. We have

$$\begin{aligned} \hat{\mu}_H - \hat{\mu}_L &= \frac{\hat{\mu}(1 - \hat{\varepsilon})}{\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}} - \frac{\hat{\mu}\hat{\varepsilon}}{\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon})} \\ &= \frac{\hat{\mu}(1 - \hat{\varepsilon})(\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon})) - \hat{\mu}\hat{\varepsilon}(\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon})}{(\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon}))(\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon})} \\ &= \frac{\hat{\mu}(1 - \hat{\mu})(1 - 2\hat{\varepsilon})}{(\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon}))(\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon})} \geq 0. \end{aligned}$$

The inequality follows from $\hat{\varepsilon} \in [0, \frac{1}{2}]$. Thus

$$\bar{q}_{\hat{\mu}_L} \leq \bar{q}_{\hat{\mu}_H}$$

where equality holds if and only if $\hat{\mu}_H = \hat{\mu}_L$ which is equivalent to $\hat{\varepsilon} = \frac{1}{2}$.

There are now three cases that can occur, regarding the level of the price p .

First case: $\bar{q}_{\hat{\mu}_L} < p < \bar{q}_{\hat{\mu}_H}$

This implies that $\hat{\varepsilon} < \frac{1}{2}$, $k > 0$ and that the consumer only buys if she receives the high signal.

Second case: $p \leq \bar{q}_{\hat{\mu}_L}$

The consumer would either buy with each signal or mix between “buying” and “not buying” in the case where $p = \bar{q}_{\hat{\mu}_L}$ and the low signal appears. In both cases, the payoff (with search costs and before observing s) is

$$\hat{\mu}q^H + (1 - \hat{\mu})q^L - p - k$$

which clearly has a maximum at $k = 0$.

Third case: $p \geq \bar{q}_{\hat{\mu}_H}$

The consumer would not buy on any signal (while with equality she may buy on s^H but gets utility $\bar{q}_{\hat{\mu}_H} - p - k = -k$) so also here optimality implies $k = 0$.

\Rightarrow If $k > 0$, the consumer buys if and only if the signal is s^H . \square

Proof of Lemma 3.3. We write $\hat{k}^* := k^*(p, \hat{\mu})$ and $\hat{\varepsilon} := \varepsilon(\hat{k}^*)$. Note first that with $\hat{k}^* = 0$ we would have $\hat{\varepsilon} = \frac{1}{2}$ and thus

$$u_s^*(p, \hat{\mu}) = \frac{1}{2}\hat{\mu}(q^H - p) + \frac{1}{2}(1 - \hat{\mu})(q^L - p) = \frac{1}{2}u_b(p, \hat{\mu}) \leq \max\{u_b(p, \hat{\mu}), u_n\}$$

which contradicts the assumptions. Define

$$E_H := \frac{\hat{\mu}(1 - \hat{\varepsilon})}{\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}} q^H + \frac{(1 - \hat{\mu})\hat{\varepsilon}}{\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}} q^L - p \quad (4)$$

$$E_L := \frac{\hat{\mu}\hat{\varepsilon}}{\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon})} q^H + \frac{(1 - \hat{\mu})(1 - \hat{\varepsilon})}{\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon})} q^L - p \quad (5)$$

the expected utility from buying, disregarding the sunk search costs, when receiving a high signal or a low signal, respectively.

The following two important relations are immediate from these formulas.

$$\begin{aligned}
u_s^*(p, \hat{\mu}) &= \hat{\mu}(1 - \hat{\varepsilon})(q^H - p) + (1 - \hat{\mu})\hat{\varepsilon}(q^L - p) - \hat{k}^* \\
&= \hat{\mu}(1 - \hat{\varepsilon})q^H + (1 - \hat{\mu})\hat{\varepsilon}q^L - (\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon})p - \hat{k}^* \\
&\stackrel{(4)}{=} (\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon})E_H - \hat{k}^*
\end{aligned} \tag{6}$$

$$\begin{aligned}
u_s^*(p, \hat{\mu}) &= \hat{\mu}(1 - \hat{\varepsilon})(q^H - p) + (1 - \hat{\mu})\hat{\varepsilon}(q^L - p) - \hat{k}^* \\
&= \hat{\mu}(q^H - p) + (1 - \hat{\mu})(q^L - p) - \hat{\mu}\hat{\varepsilon}(q^H - p) - (1 - \hat{\mu})(1 - \hat{\varepsilon})(q^L - p) - \hat{k}^* \\
&= \bar{q}_{\hat{\mu}} - p - [\hat{\mu}\hat{\varepsilon}q^H + (1 - \hat{\mu})(1 - \hat{\varepsilon})q^L - (\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon}))p] - \hat{k}^* \\
&\stackrel{(5)}{=} \bar{q}_{\hat{\mu}} - p - (\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon}))E_L - \hat{k}^*
\end{aligned} \tag{7}$$

Assume that the signal is not precise enough in the sense of the lemma. This can have two reasons. Either the expected value from buying is below zero even when receiving a high signal ($E_H \leq 0$) or it is above zero even on receiving a low signal ($E_L \geq 0$).

In the first case, equation (6) implies $u_s^*(p, \hat{\mu}) < 0 = u_n$ which contradicts the conditions of the lemma.

In the second case, equation (7) implies

$$u_s^*(p, \hat{\mu}) = \bar{q}_{\hat{\mu}} - p - \underbrace{\left((\hat{\mu}\hat{\varepsilon} + (1 - \hat{\mu})(1 - \hat{\varepsilon}))E_L + \hat{k}^* \right)}_{>0} < \bar{q}_{\hat{\mu}} - p = u_b(p, \hat{\mu})$$

which is a contradiction for the same reason.

These contradictions prove that $E_L < 0 < E_H$. Hence the consumer buys if and only if she receives the high signal. \square

Proof of Lemma 3.4. We begin by finding values $\underline{p}_{\hat{\mu}} \leq \bar{p}_{\hat{\mu}}$ such that the strict inequalities hold.

$$\begin{aligned}
u_b(p_1, \hat{\mu}) &> \max\{u_n, u_s^*(p_1, \hat{\mu})\} & \forall p_1 \in [q^L, \underline{p}_{\hat{\mu}}] \\
u_s^*(p_2, \hat{\mu}) &> \max\{u_b(p_2, \hat{\mu}), u_n\} & \forall p_2 \in (\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}}) \\
u_n &> \max\{u_b(p_3, \hat{\mu}), u_s^*(p_3, \hat{\mu})\} & \forall p_3 \in (\bar{p}_{\hat{\mu}}, q^H].
\end{aligned} \tag{8}$$

It is obvious from $\hat{\mu} \in (0, 1)$ that the first inequality is satisfied for $p_1 = q^L$ and the last inequality for $p_3 = q^H$. It is thus only left to show that there is monotonic behavior in p in the pairwise differences between $u_b(p, \hat{\mu})$, $u_s^*(p, \hat{\mu})$ and u_n . We show that the inequalities

$$\frac{\partial}{\partial p} u_b(p, \hat{\mu}) < \frac{\partial}{\partial p} u_s^*(p, \hat{\mu}) < \frac{\partial}{\partial p} u_n$$

hold on the interval (q^L, q^H) wherever k^* and thus u_s^* is differentiable in p . The left and right component of this expression are obvious from their definitions.

$$\frac{\partial}{\partial p} u_b(p, \hat{\mu}) = \frac{\partial}{\partial p} (\hat{\mu} q^H + (1 - \hat{\mu}) q^L - p) = -1 \quad \frac{\partial}{\partial p} u_n = 0$$

By the shape of the function k^* (on page 8), we know that $u_s^*(p, \hat{\mu})$ is continuous and piecewise differentiable in p on some (possibly empty) intervals (q^L, p') , (p', p'') , (p'', q^H) and we have

$$u_s^*(p, \hat{\mu}) = \begin{cases} u_s(p, \hat{\mu}, 0) = \hat{\mu} \frac{1}{2} (q^H - p) + (1 - \hat{\mu}) \frac{1}{2} (q^L - p) & \text{if } p \in I_1 \\ u_s(p, \hat{\mu}, (\varepsilon')^{-1}(\hat{d})) & \text{if } p \in I_2 \\ u_s(p, \hat{\mu}, \bar{k}) = \hat{\mu} (q^H - p) - \bar{k} & \text{if } p \in I_3. \end{cases}$$

In this expression we have $I_2 = (p', p'')$ and I_1, I_3 are the sets $[q^L, p')$ and $(p'', q^H]$.¹⁸

While the first and third components are easy to differentiate, the middle one becomes

$$\begin{aligned} & \frac{\partial}{\partial p} u_s(p, \hat{\mu}, (\varepsilon')^{-1}(\hat{d})) \\ &= \frac{\partial}{\partial p} \left(\hat{\mu} (1 - \varepsilon(k^*)) (q^H - p) + (1 - \hat{\mu}) \varepsilon(k^*) (q^L - p) - k^* \right) \\ &= -\hat{\mu} (1 - \hat{\varepsilon}) - \hat{\mu} (q^H - p) \frac{\partial}{\partial p} \varepsilon(k^*) - (1 - \hat{\mu}) \hat{\varepsilon} + (1 - \hat{\mu}) (q^L - p) \frac{\partial}{\partial p} \varepsilon(k^*) - \frac{\partial}{\partial p} k^* \\ &= -\hat{\mu} (1 - \hat{\varepsilon}) - (1 - \hat{\mu}) \hat{\varepsilon} + \frac{1}{d} \frac{\partial}{\partial p} \varepsilon(k^*) - \frac{\partial}{\partial p} k^* \\ &= -\hat{\mu} (1 - \hat{\varepsilon}) - (1 - \hat{\mu}) \hat{\varepsilon} + \frac{1}{d} \varepsilon'(k^*) \frac{\partial}{\partial p} k^* - \frac{\partial}{\partial p} k^* \\ &= -\hat{\mu} (1 - \hat{\varepsilon}) - (1 - \hat{\mu}) \hat{\varepsilon} + \frac{1}{d} \hat{d} \frac{\partial}{\partial p} k^* - \frac{\partial}{\partial p} k^* \\ &= -\hat{\mu} (1 - \hat{\varepsilon}) - (1 - \hat{\mu}) \hat{\varepsilon} \in (\hat{\varepsilon} - 1, -\hat{\varepsilon}), \end{aligned}$$

using the chain rule, omitting the arguments of k^* and writing $\hat{\varepsilon} := \varepsilon(k^*)$.

Summarized, we end up with the following expression

$$\frac{\partial}{\partial p} u_s^*(p, \hat{\mu}) = \begin{cases} -\frac{1}{2} & \text{if } p \in I_1 \\ -\hat{\mu} (1 - \hat{\varepsilon}) - (1 - \hat{\mu}) \hat{\varepsilon} & \text{if } p \in I_2 \\ -\hat{\mu} & \text{if } p \in I_3. \end{cases}$$

which is always strictly between -1 and 0 and even continuous. This proves the inequalities (8).

It follows from Lemma 3.3 that the consumer really searches in the region $(\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}})$. Finally, the price $p = \bar{q}_{\hat{\mu}}$ implies

$$u_n = u_b(p, \hat{\mu})$$

so that we must have $\underline{p}_{\hat{\mu}} \leq \bar{q}_{\hat{\mu}} \leq \bar{p}_{\hat{\mu}}$ for (8) to be true. \square

¹⁸Their order depends on the value of $\hat{\mu}$ which determines whether search effort increases or decreases in p . For $\hat{\mu} = \frac{1}{2}$, search effort is constant and two of the intervals are empty.

Proof of Lemma 3.5. Set $\hat{p} = \bar{q}_{\hat{\mu}}$. Clearly $u_b(\hat{p}, \hat{\mu}) := \bar{q}_{\hat{\mu}} - \hat{p} = 0 = u_s(\hat{p}, \hat{\mu}, 0)$ holds. Moreover, we have

$$\frac{\partial}{\partial k} u_s(\hat{p}, \hat{\mu}, k) = -\varepsilon'(k)\hat{\mu}(q^H - \hat{p}) + \varepsilon'(k)(1 - \hat{\mu})(q^L - \hat{p}) - 1$$

for all $k \in (0, \bar{k})$ and thus, taking the limit $k \rightarrow 0$

$$\begin{aligned} \lim_{k \rightarrow 0} \frac{\partial}{\partial k} u_s(\hat{p}, \hat{\mu}, k) &= -\varepsilon'(0)[\hat{\mu}(q^H - \hat{p}) + (1 - \hat{\mu})(\hat{p} - q^L)] - 1 \\ &\stackrel{\hat{p}=\bar{q}_{\hat{\mu}}}{=} -\varepsilon'(0)[\hat{\mu}(q^H - (\hat{\mu}q^H + (1 - \hat{\mu})q^L)) \\ &\quad + (1 - \hat{\mu})((\hat{\mu}q^H + (1 - \hat{\mu})q^L) - q^L)] - 1 \\ &= -\varepsilon'(0)(2\hat{\mu}(1 - \hat{\mu})(q^H - q^L)) - 1. \end{aligned}$$

This is positive by the assumption. Hence, $u_s^*(\hat{p}, \hat{\mu}) > u_s(\hat{p}, \hat{\mu}, 0) = u_b(\hat{p}, \hat{\mu}) = u_n$ which proofs $q_{\hat{\mu}} = \hat{p} \in (\underline{p}_{\hat{\mu}}, \bar{p}_{\hat{\mu}})$.

Assume now that the inequality stated in the lemma is not true. Then, by the calculations above and the strict concavity of ε , $u_s(\hat{p}, \hat{\mu}, k)$ is decreasing in k so that $k = 0$ is the optimal choice of search effort. Thus

$$u_s^*(\hat{p}, \hat{\mu}) = u_s(\hat{p}, \hat{\mu}, 0) = u_b(\hat{p}, \hat{\mu}) = u_n.$$

From the proof of Lemma 3.4, we know that

$$\frac{\partial}{\partial p} u_b(p, \hat{\mu}) < \frac{\partial}{\partial p} u_s^*(p, \hat{\mu}) < \frac{\partial}{\partial p} u_n$$

for all $p \in (q^L, q^H)$ and thus $u_b(p, \hat{\mu}) > u_s^*(p, \hat{\mu})$ for all $p < \bar{q}_{\hat{\mu}}$ and $u_n > u_s^*(p, \hat{\mu})$ for all $p > \bar{q}_{\hat{\mu}}$. Hence we have $\underline{p}_{\hat{\mu}} = \bar{p}_{\hat{\mu}} = \bar{q}_{\hat{\mu}}$. \square

Proof of Lemma 3.6. We only show the claim for $\underline{p}_{\hat{\mu}}$ since the other part is basically the same proof with even simpler arguments. Note that, given $\hat{\mu} \in (0, 1)$, $\underline{p}_{\hat{\mu}}$ is uniquely determined¹⁹ by solving

$$\begin{aligned} u_b(p, \hat{\mu}) &= u_s^*(p, \hat{\mu}) \\ \Leftrightarrow \hat{\mu}q^H + (1 - \hat{\mu})q^L - p &= \hat{\mu}(1 - \varepsilon)(q^H - p) + (1 - \hat{\mu})\varepsilon(q^L - p) - k^* \end{aligned} \quad (9)$$

and since these expressions are continuous and piecewise differentiable, the function $\underline{p}_{\hat{\mu}}$ also has these properties.²⁰

¹⁹This is also true if $\underline{p}_{\hat{\mu}} = \bar{q}_{\hat{\mu}}$. From Lemma 3.4 it follows that in this case $k^*(p, \hat{\mu}) = 0$, $\varepsilon(k^*(p, \hat{\mu})) = \frac{1}{2}$ and the equation (9) holds.

²⁰ For the differentiability, the only problem occurs on the set

$$\{\hat{\mu} \in [0, 1] \mid d(p_{\hat{\mu}}, \hat{\mu}) = \varepsilon'(0) \text{ or } d(p_{\hat{\mu}}, \hat{\mu}) = \varepsilon'(\bar{k})\}$$

in which the possible non-differentiable points of k^* are touched. By continuity, this set is closed and hence compact. It can thus be written as the union of finitely many open intervals and finitely many singletons. Within these intervals, the differentiation for $k^*(p, \hat{\mu}) = \bar{k}$ or $k^*(p, \hat{\mu}) = 0$ applies. The singletons are the only candidates in which $p_{\hat{\mu}}$ may not be differentiable in $\hat{\mu}$.

In the areas of differentiability, we either have $k^*(p, \hat{\mu}) = 0$ (implying $\varepsilon = \frac{1}{2}$), $k^*(p, \hat{\mu}) = \bar{k}$ (with $\varepsilon = 0$) or $k^*(p, \hat{\mu}) = (\varepsilon')^{-1}(d(p, \hat{\mu}))$. In the first two cases, (9) yields

$$\underline{p}_{\hat{\mu}} = \bar{q}_{\hat{\mu}} = \hat{\mu}q^H + (1 - \hat{\mu})q^L \quad \text{or} \quad \underline{p}_{\hat{\mu}} = q^L + \frac{\bar{k}}{1 - \hat{\mu}}$$

which both induce a strictly positive derivative in $\hat{\mu}$.

In the third case, differentiating (9) with respect to $\hat{\mu}$ yields

$$\begin{aligned} q^H - q^L - \underline{p}'_{\hat{\mu}} &= (1 - \hat{\varepsilon})(q^H - \hat{p}) + \hat{\varepsilon}(\hat{p} - q^L) - \hat{\mu}\left(\frac{\partial}{\partial \hat{\mu}}\varepsilon\right)(q^H - \hat{p}) \\ &\quad + (1 - \hat{\mu})\left(\frac{\partial}{\partial \hat{\mu}}\varepsilon\right)(q^L - \hat{p}) - \underline{p}'_{\hat{\mu}}(\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}) - \frac{\partial}{\partial \hat{\mu}}k^* \\ &= (1 - \hat{\varepsilon})(q^H - \hat{p}) + \hat{\varepsilon}(\hat{p} - q^L) + \frac{1}{d}\left(\frac{\partial}{\partial \hat{\mu}}\varepsilon\right) \\ &\quad - \underline{p}'_{\hat{\mu}}(\hat{\mu}(1 - \hat{\varepsilon}) + (1 - \hat{\mu})\hat{\varepsilon}) - \frac{\partial}{\partial \hat{\mu}}k^* \end{aligned}$$

where we left out the arguments for d, ε and k^* and wrote $\hat{p} = \underline{p}_{\hat{\mu}}$, $\hat{\varepsilon} = \varepsilon(k^*(\hat{p}, \hat{\mu}))$.

Reordering this equation and using $\frac{\partial}{\partial \hat{\mu}}\varepsilon = \varepsilon'(\hat{k}^*)\frac{\partial}{\partial \hat{\mu}}k^* = \hat{d} \cdot \frac{\partial}{\partial \hat{\mu}}k^*$, we get

$$\underline{p}'_{\hat{\mu}} = \frac{\hat{\varepsilon}(q^H - \hat{p}) + (1 - \hat{\varepsilon})(\hat{p} - q^L)}{1 - \hat{\mu}(1 - \hat{\varepsilon}) - (1 - \hat{\mu})\hat{\varepsilon}} > 0.$$

The limit behavior $\lim_{\hat{\mu} \rightarrow 0} \underline{p}_{\hat{\mu}} = q^L$ is clear, since we have $q^L < \underline{p}_{\hat{\mu}} \leq \bar{q}_{\hat{\mu}}$ for all values $\hat{\mu} \in (0, 1)$.

For $\hat{\mu}$ going to one, note that the convergence of $\underline{p}_{\hat{\mu}}$ is guaranteed by the strict monotonicity. Since k^* can not be higher than q^H for the equation (9) to be true, it is bounded and hence there is an increasing sequence $(\hat{\mu}_n)$, converging to 1, for which $k^*(\underline{p}_{\hat{\mu}_n}, \hat{\mu}_n)$ converges to a value $\kappa \geq 0$. Take such a sequence and the limit $n \rightarrow \infty$ in (9). We then obtain

$$\begin{aligned} q^H - \lim_{\hat{\mu} \rightarrow 1} \underline{p}_{\hat{\mu}} &= (1 - \varepsilon(\kappa)) \left(q^H - \lim_{\hat{\mu} \rightarrow 1} \underline{p}_{\hat{\mu}} \right) - \kappa \\ \Leftrightarrow \quad \varepsilon(\kappa) \left(q^H - \lim_{\hat{\mu} \rightarrow 1} \underline{p}_{\hat{\mu}} \right) &= -\kappa \end{aligned}$$

which, since the left hand side is weakly positive, implies $\kappa = 0$, thus $\varepsilon(\kappa) = \frac{1}{2}$ and finally

$$\lim_{\hat{\mu} \rightarrow 1} \underline{p}_{\hat{\mu}} = q^H.$$

□

Proof of Lemma 4.4. As seen from section 3, the consumer would buy for any price below q^L . This shows that no such price $p < q^L$ can be part of an equilibrium because a deviation to any price in (p, q^L) would yield a higher payoff. This shows the lower bound of i).

Assume that the equilibrium profit π^L for the low type is strictly below $q^L - c^L$. A deviation to the price $q^L - \frac{q^L - c^L - \pi^L}{2}$ would then yield the profit

$$q^L - \frac{q^L - c^L - \pi^L}{2} - c^L = \frac{q^L - c^L}{2} + \frac{\pi^L}{2} > 2\frac{\pi^L}{2} = \pi^L.$$

This concludes the proof of statement iii). Having the low type set the price q^H with positive probability, the equilibrium definition implies $\mu(q^H) < 1$ and thus it is optimal for the consumer to not buy the product, implying zero profit for the low type. This contradicts iii). Using an obvious similar argument, we conclude that the low type can not have any price above q^H in its support. This shows ii) and the rest of i).

It is left to show statement iv). Let $p \in (q^L, q^H)$ be a price that is in the support of a_L but not of a_H . By the equilibrium definition we must have $\mu(p) = 0$ so the consumer would know the true quality when observing price p . Hence, he would not buy the product and the low type would make no profit which contradicts the previous point.

Assume now that $p \in (q^L, q^H)$ is a price in the support of a_H but not of a_L . Since $\mu(p) = 1$ and $p < q^H$, the consumer buys with probability 1. So there must be at least one price $p_L > p, p_L \in \text{supp}(a_L)$, otherwise the low type would deviate from any price to p . Since only consistent strategies are played in equilibrium, we are left with two cases.

First case: $b(p_L) = (0, \gamma, \gamma), \gamma \in [0, 1]$

By optimality of the low type's strategy, this price satisfies

$$p - c^L \leq (p_L - c^L)\gamma. \quad (10)$$

That is, the low type must make at least as much profit with setting price p_L than with price p . By the previous part of the proof, p_L is in the support of a_H . Hence the high type must be indifferent between setting prices p or p_L .

$$\begin{aligned} p - c^H &= (p_L - c^H)\gamma \\ \stackrel{(10)}{\Rightarrow} & (p_L - c^L)\gamma + c^L - c^H \geq (p_L - c^H)\gamma \\ \Leftrightarrow & (1 - \gamma)c^L \geq (1 - \gamma)c^H \end{aligned}$$

which, since the first equation also implies $\gamma \neq 1$, is not compatible with the assumption $c^L < c^H$.

Second case: $b(p_L) = (k, 1, 0), k > 0$

Let $\hat{\varepsilon} := \varepsilon(k)$ be the probability of a false signal. This equals the chance that the low type will sell her product for the price p_L . Again by optimality of the low type's choice the following inequality holds.

$$p - c^L \leq (p_L - c^L)\hat{\varepsilon} \quad (11)$$

As before, the price p_L must also be in the support of the high type. This yields

$$\begin{aligned} p - c^H &= (p_L - c^H)(1 - \hat{\varepsilon}) \\ \stackrel{(11)}{\Rightarrow} & (p_L - c^L)\hat{\varepsilon} + c^L - c^H \geq (p_L - c^H)(1 - \hat{\varepsilon}) \\ \Leftrightarrow & \hat{\varepsilon}(p_L - c^H) \geq (1 - \hat{\varepsilon})(p_L - c^L) \end{aligned}$$

which gives a contradiction for the same reason as before and using $\hat{\varepsilon} < 1 - \hat{\varepsilon}$. \square

Proof of Lemma 4.5. Let $p < p'$ be two such prices, $\hat{\varepsilon}, \hat{\varepsilon}'$ the corresponding error probabilities. For search to be possible, the prices have to be in both supports and hence, by the indifference principle for both types,

$$\begin{aligned}(p - c^H)(1 - \hat{\varepsilon}) &= (p' - c^H)(1 - \hat{\varepsilon}') \Rightarrow (1 - \hat{\varepsilon}) > (1 - \hat{\varepsilon}') \Leftrightarrow \hat{\varepsilon} < \hat{\varepsilon}' \\ (p - c^L)\hat{\varepsilon} &= (p' - c^L)\hat{\varepsilon}' \Rightarrow \hat{\varepsilon} > \hat{\varepsilon}'\end{aligned}$$

which gives a contradiction. \square

Proof of Lemma 4.6. Without search, there are two probabilities $\gamma, \gamma' \in [0, 1]$ of the consumer buying for the prices $p < p'$. By the optimality of the firm's strategy we have

$$\begin{aligned}(p - c^H)\gamma &= (p' - c^H)\gamma' \\ (p - c^L)\gamma &= (p' - c^L)\gamma'.\end{aligned}$$

Note that, since the low type always has positive profit, all of these factors must be strictly above zero. Reassembling these equations gives two different values for the ratio $\frac{\gamma}{\gamma'}$ since $p \neq p'$ and $c^L \neq c^H$.²¹ This is a contradiction. \square

Proof of Lemma 4.8. Assume such a price $p \in (q^L, q^H)$ exists for which the consumer buys with probability lower than one. Then, because of Lemma 4.4 iv) and iii), *not buying* is not possible so she buys with probability $\gamma \in (0, 1)$. She is thus indifferent between buying without search and not buying. In this case $p = \bar{q}_{\mu(p)}$ holds. Since $\varepsilon'(0) = -\infty$, she would search on that price by Lemma 3.5.

The case $q = q^L$ follows from the same argument as the proof of Lemma 4.4 iii). If sales had less than full probability, a slightly lower price would yield full sales and hence a higher profit. \square

Proof of Proposition 4.9. The previous lemmas already show that not more than two equilibrium prices can exist in (q^L, q^H) . We then show that there can not be a search price p_s and a no-search price p_1 played by both types. If this was the case, we must have $p_1 < p_s$ since otherwise there would be incentives to deviate from p_s to p_1 . For this, remember that Lemma 4.8 shows that the price p_1 induces sure buying.

Let now be p_1 and p_s be played in an equilibrium by both types. Applying the indifference principle for both firms we get

$$\begin{aligned}p_1 - c^H &= (1 - \hat{\varepsilon})(p_s - c^H) \\ p_1 - c^L &= \hat{\varepsilon}(p_s - c^L)\end{aligned}$$

and thus

$$1 - \hat{\varepsilon} = \frac{p_1 - c^H}{p_s - c^H} \begin{matrix} p_1 < p_s \\ c^H > c^L \\ < \end{matrix} \frac{p_1 - c^L}{p_s - c^L} = \hat{\varepsilon}$$

²¹Here, we use that the fraction $\frac{p'-c}{p-c}$ for $p \neq p'$ is strictly monotone in c for $c < p, p'$. This statement is easy to check via differentiation.

which contradicts $\hat{\varepsilon} \in [0, \frac{1}{2}]$.

Finally, note that if the price p_1 is played, the consumer buys with certainty and p_1 is thus the lowest price in the equilibrium. Thus, $q^L \notin \text{supp}(a_L)$.

The statement about the low value of γ for $b(q^H) = (0, \gamma, \gamma)$ is obvious. If we denote the low type's equilibrium profit by π^L , one upper bound for γ is $\frac{\pi^L}{q^H - c^L}$ which is strictly smaller than one. Note that there might also be a lower bound for this value, e.g. if $c^H < q^L$. See for example the existence condition for the total adverse selection equilibrium. \square

Proof of Lemma 4.11. Let (a, μ, b) be an equilibrium in which q^H is played by the high firm and the consumer has strategy $b(q^H) = (0, \gamma, \gamma)$, $\gamma > 0$. This is a best response because of $\hat{\mu} := \mu(q^H) = 1$ and thus the consumer's utility

$$\gamma \cdot u_b(q^H, \hat{\mu}) + (1 - \gamma) \cdot \underbrace{u_n}_{=0} = \gamma(\hat{\mu}q^H + (1 - \hat{\mu})q^L - q^H) = \gamma(1 - \hat{\mu})(q^L - q^H)$$

attains its maximum in all values of γ . For slightly lower $\hat{\mu}$, however, this value has a unique maximum in $\gamma = 0$ and thus the original strategy $b(q^H)$ is locally dominated by $(0, 0, 0)$. On the other hand, the total adverse selection equilibrium does not have this problem since $(0, 0, 0)$ is the equilibrium strategy for q^H .

Every other equilibrium price is a search price p_s or a non-search price p_1 being in both supports of the firm's strategy or the price q^L , set by only the low type. We show that non of these prices is locally dominated in beliefs.

Let p_s be an equilibrium search price (implying that $\mu(p_s) \in (0, 1)$) with consumer behavior $b(p_s) = (k, 1, 0)$, $k > 0$. By the analysis of section 3, this strategy is the unique maximum over all search behaviors. The remaining candidates for domination are thus $(0, 0, 0)$ ("don't buy") and $(0, 1, 1)$ ("buy")²². If "don't buy" had the same utility, we had

$$u_s(p_s, \mu(p_s), k) = \mu(p_s)(1 - \varepsilon(k))(q^H - p_s) + (1 - \mu(p_s))\varepsilon(k)(q^L - p_s) - k = 0 = u_n.$$

Differentiating this with respect to the posterior belief, one sees that

$$\frac{\partial}{\partial m} u_s(p_s, m, k) = (1 - \varepsilon(k))(q^H - p_s) + \varepsilon(k)(p_s - q^L) > 0$$

such that we have $u_s(p_s, m, k) > u_n$ for all $m > \mu(p_s)$. The strategy $b(p_s)$ is thus not locally dominated by the strategy $(0, 0, 0)$.

If "buy" had the same utility as b , we had

$$\begin{aligned} u_s(p_s, \mu(p_s), k) &= \mu(p_s)(1 - \varepsilon(k))(q^H - p_s) + (1 - \mu(p_s))\varepsilon(k)(q^L - p_s) - k \\ &= \mu(p_s)q^H + (1 - \mu(p_s))q^L - p_s = u_b(p_s, \mu(p_s)). \end{aligned}$$

²²Note that the proof shows that also their convex combinations can not be local best responses in this case and thus are no candidates for dominating strategies.

and the derivatives

$$\begin{aligned} \frac{\partial}{\partial m} u_s(p_s, m, k) &= (1 - \varepsilon(k))(q^H - p_s) + \varepsilon(k)(p_s - q^L) \\ &\leq \max\{q^H - p_s, p_s - q^L\} < q^H - q^L = \frac{\partial}{\partial m} u_b(p_s, m). \end{aligned} \quad (12)$$

This shows that $b(p_s)$ is strictly better than $(0, 1, 1)$ for any $m < \mu(p_s)$. The strategy $b(p_s)$ is thus not locally dominated.

Now let $p_1 < q^H$ be an equilibrium price. If $b(p_1) = (0, 1, 1)$ is not a unique best response, there is a search strategy $(k, 1, 0)$ with the same payoff²³. This also implies $\mu(p_1) \in (0, 1)$ since search being optimal is not possible for degenerate posteriors. Using (12), we know that for a marginally higher posterior belief, this search strategy is worse than “buy”. Strategy $b(p_1)$ is thus not locally dominated.

The last price to check is q^L for which the equilibrium behavior $(0, 1, 1)$ is clearly the unique best response for any posterior belief $m > 0$ so that local domination is also excluded here. This also concludes the proof for showing that the total adverse selection equilibrium has belief-robust responses. \square

Proof of Proposition 4.13. Let (a, μ, b) be an equilibrium. We first show that there are no equilibrium prices p which differ from $q^L, q^H, \underline{p}_{\hat{\mu}}$ and $\bar{p}_{\hat{\mu}}$ where $\hat{\mu}$ is determined according to Bayes’ law.

First, assume that there is an equilibrium no-search price p in (q^L, q^H) which is not equal to $\underline{p}_{\mu(p)}$. In this case, Lemma 3.4 shows that $u_b(p, \mu(p)) > u_s^*(p, \mu(p))$. Assuming that μ is continuous in p , the continuity of u_b and u_s^* implies that “buying” will still be better than “searching” for a marginal increase of the price. This is an incentive for both types to deviate which contradicts the equilibrium property. If $p = q^L$ but played by both types (thus $\mu(p) > 0$), the non-emptiness of $(q^L, \underline{p}_{\mu(p)})$ shows that the same argument holds.

Second, assume the existence of an equilibrium search price $p \neq \bar{p}_{\mu(p)}$. It follows again from Lemma 3.4 that we must have $u_s^*(p, \mu(p)) > u_n$. Again the continuity of these expressions implies that the consumer will also search for a marginal higher price, although the search effort and thus the error probability $\hat{\varepsilon}$ might change. Note that the profit of the firm, depending on the type, is $(p - c^H)(1 - \hat{\varepsilon})$ or $(p - c^L)\hat{\varepsilon}$ such that, for a higher price, at least one of these values will increase. This gives an incentive for at least one type to deviate.

These two arguments together with Lemma 4.11 rule out all the equilibria from Proposition 4.9 that are not mentioned in Proposition 4.13. We thus only have to show that the rest of the equilibria can be supported by a locally continuous belief system.

Since all equilibria have a finite number of prices, these prices can be considered independently from each other by looking at non-intersecting environments of them. The price q^L , if only played by the low quality firm, can obviously be supported by $\mu(p) = 0$ in any environment of q^L . For the price q^H note that for all values of $\hat{\mu} = \mu(q^H)$ the

²³Lemma 3.5 rules out no-search strategies giving the same payoff.

interval $(\bar{p}_{\hat{\mu}}, q^H)$ is non-empty and the lower bound is continuously increasing in $\hat{\mu}$ (see Lemma 3.6) so that there is an invertible, continuous function $p(\hat{\mu})$ with $p(\hat{\mu}) \in (\bar{p}_{\hat{\mu}}, q^H)$ for all $\hat{\mu} \in (0, 1)$. By the definition of $\bar{p}_{\hat{\mu}}$, the inverse $\mu(p)$ of this function satisfies

$$u_n > \max\{u_b(p, \mu(p)), u_s^*(p, \mu(p))\} \quad \forall p \in (p(\frac{1}{2}), q^H)$$

so that the consumer would not buy with that belief system for any other price and hence there is no incentive for a deviation by any firm.

Let now p_s be an equilibrium search price on the upper line of the search area. That is

$$p_s = \bar{p}_{\mu(p_s)}.$$

Keeping the function μ constant above p_s leads to the consumer not buying for higher prices. This was shown in the proof of Lemma 4.11. For lower prices, we can use the same argument as before of $\bar{p}_{\hat{\mu}}$ being strictly increasing in $\hat{\mu}$ to show that there exists a continuous and even increasing belief system $\mu(p)$ for $p < p_s$ for which the consumer does not buy on lower prices.

At last, assume $p_1 = \underline{p}_{\mu(p)}$ for a no-search price, implying that the consumer buys with probability one. Moreover, assume $\varepsilon(k^*(p_1, \mu(p_1))) < 1$. Independent of the belief system, no firm would deviate to a lower price since the probability of selling can not grow (p_1 induces sure buying). Keeping the posterior belief constant to $\mu(p_1)$ for higher prices leads to search behavior for these prices. By continuity of $\varepsilon(k^*)$, we have

$$\lim_{p \downarrow p_1} (1 - \varepsilon(k^*(p, \mu(p))))(p - c^H) = (1 - \varepsilon(k^*(p_1, \mu(p_1))))(p_1 - c^H) \stackrel{\text{ass.}}{<} p_1 - c^H$$

so that a high quality firm has a lower profit for slightly higher prices. It is straightforward to show the same for the low type.

The existence conditions are obvious from the proof, the minimum payoffs of both types and the indifference principle for all prices in a type's price support. \square

Proof of Observation 4.14. In PE_b , both types trade for sure while in PE_s the high type firm has some trading probability $1 - \hat{\varepsilon} > \hat{\varepsilon}$. It suffices to show the observation for any hybrid equilibrium.

Remember that the existence of a hybrid equilibrium implies $q^L - c^L < q^H - q^L$. In this case, the selling probability for the price q^H in a separating equilibrium can not exceed

$$\frac{q^L - c^L}{q^H - c^L} = \frac{q^L - c^L}{q^H - q^L + q^L - c^L} < \frac{q^L - c^L}{q^L - c^L + q^L - c^L} = \frac{1}{2}.$$

to not make the low firm deviate to the higher price. The selling probability for a high quality firm in a hybrid equilibrium, however, is $1 - \hat{\varepsilon} > \frac{1}{2}$ which proves the observation. \square

Proof of Lemma 4.16. If two or more hybrid equilibria exist, the consumer surplus and the low type's profit are both at their minimum. Observe that the latter implies that the

search precision $\varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu}))$ in these equilibria is higher if the price $\bar{p}_{\hat{\mu}}$ is higher. This also implies a higher chance of selling high quality goods and thus an overall strictly higher profit for the high type. Hence, HE dominates all other hybrid equilibria.

Let either of PE_b , PE_s or HE exist. We know that in all these equilibria we have $\pi^L \geq q^L - c^L$, $\pi^H \geq 0$ and $u^* \geq 0$ which are the payoffs of the TAS equilibrium. In PE_b , the consumer has positive utility while in PE_s and HE , the high type firm has positive profit. Hence TAS is dominated.

Now let EQ either denote PE_s or HE and denote p_s and $\hat{\varepsilon}$ the search price and its corresponding signal imprecision in EQ . We know that the consumer is strictly better off in PE_b . Assume $\pi^H(PE_b) \geq \pi^H(EQ)$. It then follows, since $p_s \geq \underline{p}_{\eta}$

$$\begin{aligned} \frac{\underline{p}_{\eta} - c^L}{p_s - c^L} &\geq \frac{\underline{p}_{\eta} - c^H}{p_s - c^H} = \frac{\pi^H(PE_b)}{\frac{1}{1-\hat{\varepsilon}}\pi^H(EQ)} \geq 1 - \hat{\varepsilon} > \hat{\varepsilon} \\ \Rightarrow \quad \pi^L(PE_b) &= \underline{p}_{\eta} - c^L > \hat{\varepsilon} \cdot (p_s - c^L) = \pi^L(EQ). \end{aligned}$$

This shows that PE_b dominates EQ .

It is straightforward to show that HE and PE_s do not dominate each other. The low type profit in PE_s is higher than in HE but this, together with the higher search price in HE , implies a lower sale probability for the low type and hence a higher probability for the high type. Hence the profit of the high type is higher in HE . \square

Proof of Lemma 4.17. Note first that the existence of PE_s implies that the low type's profit $\hat{\varepsilon} \cdot (\bar{p}_{\eta} - c^L)$ is at least $q^L - c^L$. In the hybrid equilibrium, this bound must be attained for the low type to justify playing both prices. We show that this condition can not be met in both equilibria.

Taking the derivative of $\pi_L(\bar{p}_{\hat{\mu}}, \hat{\mu})$ with respect to the posterior belief $\hat{\mu}$ (for values in which this is differentiable) yields

$$\begin{aligned} &\frac{\partial}{\partial \hat{\mu}} \pi_L(\bar{p}_{\hat{\mu}}, \hat{\mu}) \\ &= \frac{\partial}{\partial \hat{\mu}} \varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu})) \cdot (\bar{p}_{\hat{\mu}} - c^L) \\ &= \left(\frac{\partial}{\partial \hat{\mu}} \varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu})) \right) (\bar{p}_{\hat{\mu}} - c^L) + \varepsilon \bar{p}'_{\hat{\mu}} \\ &= \begin{cases} 0 & \text{if } \varepsilon = 0 \\ \varepsilon'(k^*) \left([(\varepsilon')^{-1}]' (\hat{d}) \frac{\partial}{\partial \hat{\mu}} d(\bar{p}_{\hat{\mu}}, \hat{\mu}) \right) (\bar{p}_{\hat{\mu}} - c^L) + \varepsilon \bar{p}'_{\hat{\mu}} & \text{if } \varepsilon > 0 \end{cases} \\ &= \begin{cases} 0 & \text{if } \varepsilon = 0 \\ \varepsilon'(k^*) \frac{1}{\varepsilon''(k^*)} \hat{d}^2 \left(q^H - q^L - 2\bar{p}_{\hat{\mu}} + \bar{p}'_{\hat{\mu}}(1 - 2\hat{\mu}) \right) (\bar{p}_{\hat{\mu}} - c^L) + \varepsilon \bar{p}'_{\hat{\mu}} & \text{if } \varepsilon > 0 \end{cases} \end{aligned}$$

For this, we had to use standard results for the derivative of the inverse function $[(\varepsilon')^{-1}]' = \frac{1}{\varepsilon''}$ and the quotient differentiation theorem for the derivative of d .

While the right hand side of the lower term is always positive, the left hand side is positive if $\hat{\mu} \geq \frac{1}{2}$ and $\bar{p}_{\hat{\mu}} \geq \frac{q^H + q^L}{2}$. Note that the former condition implies the latter as we always have $\bar{p}_{\hat{\mu}} \geq \bar{q}_{\hat{\mu}}$. It is thus sufficient to have $\hat{\mu} \geq \frac{1}{2}$. Moreover, in an equilibrium where the high type always sets the search price, Bayes' law implies $\hat{\mu} \geq \eta$ on that price. The derivative of the low type's profit is thus always non-negative for $\hat{\mu} \in [\frac{1}{2}, 1)$. It is even strictly positive for all beliefs in which $\varepsilon(\bar{p}_{\hat{\mu}}, \hat{\mu}) > 0$ which must be the case in equilibrium. If $\eta \geq \frac{1}{2}$, this monotonicity results hold for all posterior beliefs that can occur on equilibrium search prices.

By Lemma 3.6, this profit converges to $\frac{1}{2}(q^H - c^L)$ when taking $\hat{\mu} \rightarrow 1$. Using the condition $q^L - c^L < q^H - q^L$, we know

$$\frac{1}{2}(q^H - c^L) = \frac{1}{2}(q^H - q^L + q^L - c^L) > q^L - c^L.$$

Let $\eta \geq \frac{1}{2}$ and $\varepsilon(k^*(\bar{p}_{\eta}, \eta))(\bar{p}_{\eta} - c^L) < q^L - c^L$. It then follows from the strict monotonicity and the convergence that there exists exactly one $\hat{\mu} \in (\eta, 1)$ with $\varepsilon(k^*(\bar{p}_{\hat{\mu}}, \hat{\mu}))(\bar{p}_{\hat{\mu}} - c^L) = q^L - c^L$. On the other hand, if we had $\varepsilon(k^*(\bar{p}_{\eta}, \eta))(\bar{p}_{\eta} - c^L) \geq q^L - c^L$, we have no such value for $\hat{\mu}$.

We see that the conditions of the low type having a higher profit than $q^L - c^L$ in one equilibrium and exactly this profit in the other are mutually exclusive. By continuity, this extends to an open interval of values of η below $\frac{1}{2}$ which proves the existence of $\underline{\eta}$. \square

Proof of Proposition 5.1. In this proof, we often write $p(a)$ instead of $\underline{p}_{\hat{\mu}}(a)$ for expositional reasons. Define

$$P := \{(p, a) | d(p, \eta) < \varepsilon_k(\bar{k}(a), a)\}$$

the open area for which $\varepsilon(k^*(p, \eta, a), a)$ is positive.

Note that $k^*(p, \eta, a)$ as defined by (1) is continuous²⁴ and piecewise differentiable in (p, a) in P and its complement P^c . It thus also holds for the composition $\varepsilon(k^*, a)$. Hence $\underline{p}_{\eta}(a)$ has the same properties²⁵, being the unique implicit solution of the equation

$$\begin{aligned} u_b(p(a), \eta, a) &= u_s^*(p(a), \eta, a) \\ \bar{q}_{\eta} - p(a) &= \eta(1 - \varepsilon(k^*, a))(q^H - p(a)) + (1 - \eta)\varepsilon(k^*, a)(q^L - p(a)) - k^* \end{aligned} \quad (13)$$

²⁴While \bar{k} can take the value ∞ , k^* can not. Since \bar{k} is continuous when restricted to the open set on which it is finite, k^* is continuous.

²⁵A similar differentiability argument as in footnote 20 on page 34 applies here. The set

$$\{a | d(p(a), \eta) = \varepsilon_k(\bar{k}(a), a)\}$$

might not be bounded and hence not compact. But every intersection with $[0, A]$, $A > 0$ is compact. The previous argument applies to these sets and hence there are countably many, ordered potential discontinuities $a_1 < a_2 < \dots$.

where we left out the arguments of k^* . Differentiating this expression with respect to a , in the areas where this is differentiable, and writing $p_a := \frac{\partial}{\partial a}p(a)$ yields

$$\begin{aligned} -p_a &= \eta \left[\left(-\frac{\partial}{\partial a}\varepsilon \right) (q^H - p) + (1 - \varepsilon)(-p_a) \right] + (1 - \eta) \left[\left(\frac{\partial}{\partial a}\varepsilon \right) (q^L - p) + \varepsilon(-p_a) \right] - \frac{\partial}{\partial a}k^* \\ &= -\left(\eta(1 - \varepsilon) + (1 - \eta)\varepsilon \right) p_a - \left(\eta(q^H - p) + (1 - \eta)(p - q^L) \right) \frac{\partial}{\partial a}\varepsilon - \frac{\partial}{\partial a}k^* \\ &= -\left(\eta(1 - \varepsilon) + (1 - \eta)\varepsilon \right) p_a + \frac{1}{d(p, \eta)} \frac{\partial}{\partial a}\varepsilon - \frac{\partial}{\partial a}k^* \end{aligned}$$

and

$$\begin{aligned} \frac{\partial}{\partial a}\varepsilon &= \frac{\partial}{\partial a}\varepsilon(k^*, a) = \varepsilon_k(k^*, a) \frac{\partial}{\partial a}k^*(p(a), a) + \varepsilon_a(k^*, a) \\ &= \varepsilon_k \frac{\partial}{\partial a}k^* + \varepsilon_a \\ &= d(p, \eta) \frac{\partial}{\partial a}k^* + \varepsilon_a. \end{aligned}$$

whenever $(p(a), a) \in P$ and

$$\frac{\partial}{\partial a}\varepsilon = \frac{\partial}{\partial a}0 = 0$$

in every open subset of P^c . Combining these expressions, we either get

$$\begin{aligned} -p_a &= -\left(\eta(1 - \varepsilon) + (1 - \eta)\varepsilon \right) p_a + \frac{\varepsilon_a}{d} \\ \Leftrightarrow p_a &= \frac{\eta(q^H - p) + (1 - \eta)(p - q^L)}{1 - \eta(1 - \varepsilon) - (1 - \eta)\varepsilon} \varepsilon_a < 0 \end{aligned}$$

or

$$p_a = \frac{1}{1 - \eta(1 - \varepsilon) - (1 - \eta)\varepsilon} \frac{\partial}{\partial a}\bar{k}(a) \leq 0.$$

It is left to show the limit of $p(a)$ when we let a go to 0 or ∞ . We begin with the latter case.

First claim: $\lim_{a \rightarrow \infty} \varepsilon(k^*(\underline{p}_\eta(a), \eta, a), a) = 0$

We write $k^*(a) := k^*(\underline{p}_\eta(a), \eta, a)$. Assume that $\limsup_{a \rightarrow \infty} \varepsilon(k^*(a), a) =: e > 0$. This also implies

$$\varepsilon_k(k^*(a_n), a_n) = d(\underline{p}_\eta(a_n), \eta) \text{ and } \varepsilon(k^*(a_n), a_n) > \frac{e}{2} \quad \forall n \in \mathbb{N}$$

for some sequence (a_n) going to infinity and having $\lim_{n \rightarrow \infty} \varepsilon(k^*(a_n), a_n) = e$. Since $\underline{p}_\eta(a_n)$ converges (due to the monotonicity), so does $d(\underline{p}_\eta(a_n), \eta)$ and we have

$$\lim_{n \rightarrow \infty} \varepsilon_k(k^*(a_n), a_n) = \lim_{n \rightarrow \infty} d(\underline{p}_\eta(a_n), \eta) =: \delta \leq \frac{-1}{\max\{\eta, 1 - \eta\}(q^H - q^L)} < 0$$

which implies $\varepsilon_k(k^*(a_n), a_n) > 2\delta$ for large n . Choose $k = -\frac{e}{8\delta} > 0$ and n large enough such that this inequality holds. We then have for all these n either

$$\varepsilon(k, a_n) \geq \varepsilon(k^*(a_n), a_n) > \frac{e}{2}$$

if $k \leq k^*(a_n)$ or

$$\begin{aligned}\varepsilon(k, a_n) &= \varepsilon(k^*(a_n), a_n) + \int_{k^*(a_n)}^k \underbrace{\varepsilon_k(l, a_n)}_{\geq \varepsilon_k(k^*(a_n), a_n) > 2\delta} dl > \frac{e}{2} + (k - k^*(a_n))2\delta \\ &\geq \frac{e}{2} + k2\delta = \frac{e}{2} - \frac{e}{4} = \frac{e}{4}\end{aligned}$$

otherwise. This is a contradiction to $\lim_{n \rightarrow \infty} \varepsilon(k, a_n) = 0$.

Second claim: $\lim_{a \rightarrow \infty} k^*(\underline{p}_\eta(a), \eta, a) = 0$

The argument here is almost the same. Assume $\limsup_{a \rightarrow \infty} k^*(\underline{p}_\eta(a), \eta, a) =: \kappa > 0$. Take a sequence a_n , $\lim_{n \rightarrow \infty} a_n = \infty$ with $\lim_{n \rightarrow \infty} k^*(a_n) = \kappa$ and $k^*(a_n) > \frac{\kappa}{2}$ for all n . Let δ be as before, and let n be large enough so that the inequality $\varepsilon_k(k^*(a_n), a_n) > 2\delta$ holds. We then have

$$\begin{aligned}\varepsilon\left(\frac{\kappa}{2}, a_n\right) &= \varepsilon(k^*(a_n), a_n) - \int_{\frac{\kappa}{2}}^{k^*(a_n)} \varepsilon_k(l, a_n) dl \geq \varepsilon(k^*(a_n), a_n) - \left(k^*(a_n) - \frac{\kappa}{2}\right) \frac{\delta}{2} \\ &\geq -\left(k^*(a_n) - \frac{\kappa}{2}\right) \frac{\delta}{2} \rightarrow -\frac{\kappa \delta}{2} > 0, \quad n \rightarrow \infty,\end{aligned}$$

contradicting $\lim_{n \rightarrow \infty} \varepsilon\left(\frac{\kappa}{2}, a_n\right) = 0$.

From the two claims, it now follows easily by equation (13) that

$$\begin{aligned}\bar{q}_\eta - \lim_{a \rightarrow \infty} \underline{p}_\eta(a) &= \eta(q^H - \lim_{a \rightarrow \infty} \underline{p}_\eta(a)) \\ \Rightarrow \lim_{a \rightarrow \infty} \underline{p}_\eta(a) &= q^L.\end{aligned}$$

This concludes the proof for the case $a \rightarrow \infty$.

The proof for $a \rightarrow 0$ is quite similar. Using basically the same arguments, we show that $\lim_{a \rightarrow 0} \varepsilon(k^*(p(a), a), a) = \frac{1}{2}$ and that $\lim_{a \rightarrow 0} k^*(p(a), a) = 0$. Hence the limit of equation (13) yields

$$\begin{aligned}\bar{q}_\eta - \lim_{a \rightarrow 0} \underline{p}_\eta(a) &= \frac{1}{2} \eta(q^H - \lim_{a \rightarrow 0} \underline{p}_\eta(a)) + \frac{1}{2} (1 - \eta)(q^L - \lim_{a \rightarrow 0} \underline{p}_\eta(a)) \\ \Rightarrow \bar{q}_\eta - \lim_{a \rightarrow 0} \underline{p}_\eta(a) &= \frac{1}{2} \bar{q}_\eta - \frac{1}{2} \lim_{a \rightarrow 0} \underline{p}_\eta(a) \\ \Rightarrow \lim_{a \rightarrow 0} \underline{p}_\eta(a) &= \bar{q}_\eta.\end{aligned}$$

The last part of the proposition follows from this convergence of \underline{p}_η (so that for low values, the price of PE_b is above c^H and for high a it is not) and the previously shown

$$\lim_{a \rightarrow 0} \varepsilon(k^*(p(a), a), a) = \frac{1}{2} > 0$$

which ensures that $\varepsilon > 0$ for low values of a . This is part of the existence condition for PE_b . \square

Proof of Proposition 5.2. The arguments here are basically the same as in the previous proof, using that $\bar{p}_\eta(a)$ is implicitly defined by the equation

$$u_n = u_s^*(p(a), \eta, a)$$

$$0 = \eta(1 - \varepsilon(k^*, a))(q^H - p(a)) + (1 - \eta)\varepsilon(k^*, a)(q^L - p(a)) - k^*.$$

The derivative is thus either

$$\frac{\eta(q^H - p) + (1 - \eta)(p - q^L)}{-\eta(1 - \varepsilon) - (1 - \eta)\varepsilon} \varepsilon_a > 0 \text{ or } \frac{1}{-\eta(1 - \varepsilon) - (1 - \eta)\varepsilon} \frac{\partial}{\partial a} \bar{k}(a) \geq 0.$$

The arguments for the convergence to \bar{q}_η and q^H are again very similar to the previous proof and are thus omitted. \square

Proof of Corollary 5.3. It follows from the convergence of $\underline{p}_\eta(a)$ that PE_b does not exist for high values of a , since its price would be lower than the high quality production costs c^H from some point on. The search prices of both, PE_s and HE , converge to q^H . The proof of Proposition 5.2 shows that for PE_s , the corresponding signal error $\varepsilon(k^*(\bar{p}_\eta(a), \eta), a)$ converges to 0 when a goes to infinity. Thus for high values of a we have

$$\varepsilon(k^*(\bar{p}_\eta(a), \eta), a) (\bar{p}_\eta(a) - c^L) < q^L - c^L \quad (14)$$

so that PE_s does not exist. The convergence of $\bar{p}_\eta(a)$ to q^H also shows that for high values of a we must have $\bar{p}_\eta(a) > c^H$. For each of these values of a , since the left hand side of (14) converges to $\frac{1}{2}(q^H - c^L) > q^L - c^L$ when η goes to one, there exists $\hat{\mu} \in (\eta, 1)$ so that

$$\varepsilon(k^*(\bar{p}_{\hat{\mu}}(a), \hat{\mu}), a) (\bar{p}_{\hat{\mu}}(a) - c^L) = q^L - c^L.$$

This constitutes the existence of a hybrid equilibrium and thus of HE .

The above equality combined with the limit behavior

$$q^H \geq \bar{p}_{\hat{\mu}}(a) \geq \bar{p}_\eta(a) \rightarrow q^H, \quad a \rightarrow \infty$$

implies the convergence of $\varepsilon(k^*(\bar{p}_{\hat{\mu}}(a), \hat{\mu}), a)$ to $\frac{q^L - c^L}{q^H - c^L}$. \square

Proof of Proposition 5.4. From the proof of Proposition 5.2 we know that

$$\lim_{a \rightarrow 0} \bar{p}_\eta(a) = \bar{q}_\eta, \quad \lim_{a \rightarrow 0} k^*(\bar{p}_\eta(a), \eta, a) = 0 \text{ and } \lim_{a \rightarrow 0} \varepsilon(k^*(\bar{p}_\eta(a), \eta, a), a) = \frac{1}{2}.$$

The low type profit in PE_s thus converges to

$$\begin{aligned} \lim_{a \rightarrow 0} \varepsilon(k^*(\underline{p}_\eta(a), \eta, a), a) (\bar{p}_\eta(a) - c^L) &= \frac{1}{2}(\bar{q}_\eta - c^L) \\ &= \frac{1}{2}(\eta q^H + (1 - \eta)q^L - c^L) \\ &= \frac{1}{2}(\underbrace{\eta(q^H - q^L)}_{< q^L - c^L} + q^L - c^L) \\ &< q^L - c^L \end{aligned}$$

which shows that for low a this equilibrium type does not exist. It follows from the proof of Lemma 4.17 that for each such a there is at least one belief $\hat{\mu}(a) > \eta$ for which the low type exactly attains the profit $q^L - c^L$ at price $\hat{p}(a) := \bar{p}_{\hat{\mu}(a)}(a)$ if the consumer behaves optimally. Writing $\hat{k}^*(a) := k^*(\hat{p}(a), \hat{\mu}(a), a)$, this means

$$\underbrace{\varepsilon(\hat{k}^*(a), a)}_{=:\hat{\varepsilon}(a)} \cdot (\hat{p}(a) - c^L) = q^L - c^L.$$

We can use a similar argument to the one in the proof of Proposition 5.1 to show that

$$\lim_{a \rightarrow 0} \hat{k}^*(a) = 0 \text{ and } \lim_{a \rightarrow 0} \hat{\varepsilon}(a) = \frac{1}{2}.$$

The above equality then dictates that $\lim_{a \rightarrow 0} \hat{p}(a) = 2q^L - c^L$ and, since the posterior belief satisfies

$$\hat{\mu}(a)(1 - \hat{\varepsilon}(a))(q^H - \hat{p}(a)) + (1 - \hat{\mu}(a))\hat{\varepsilon}(a)(q^L - \hat{p}(a)) - \hat{k}^*(a) = 0$$

for all a , taking the limit and applying the result yields $\lim_{a \rightarrow 0} \hat{\mu}(a) = \frac{q^L - c^L}{q^H - q^L}$.

Finally, the condition $c^H < 2q^L - c^L = \lim_{a \rightarrow 0} \hat{p}(a)$ ensures that these prices indeed form hybrid equilibria for low values of a . \square

References

- [1] ADRIANI, FABRIZIO and LUCA G. DEIDDA: *Price signaling and the strategic benefits of price rigidities*. Games and Economic Behavior, 67(2):335–350, November 2009.
- [2] AKERLOF, GEORGE A.: *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*. The quarterly journal of economics, 84(3):488–500, 1970.
- [3] BAGWELL, KYLE and MICHAEL H. RIORDAN: *High and Declining Prices Signal Product Quality*. The American Economic Review, 81(1):224–239, 1991.
- [4] BESTER, HELMUT and KLAUS RITZBERGER: *Strategic pricing, signalling, and costly information acquisition*. International Journal of Industrial Organization, 19(9):1347–1361, 2001.
- [5] CHO, IN-KOO and DAVID M KREPS: *Signaling Games and Stable Equilibria*. The Quarterly Journal of Economics, 102(2):179–221, January 1987.
- [6] ELLINGSEN, TORE: *Price signals quality: The case of perfectly inelastic demand*. International Journal of Industrial Organization, 16(1):43–61, 1997.
- [7] GROSSMAN, SANFORD J and JOSEPH E STIGLITZ: *On the impossibility of informationally efficient markets*. The American Economic Review, 70(3):393–408, 1980.
- [8] KIHILSTROM, RICHARD: *A Bayesian model of demand for information about product quality*. International Economic Review, 15(1):99–118, 1974.

-
- [9] MARTIN, DANIEL: *Strategic Pricing with Rational Inattention to Quality*. Job Market Paper, New York University, October 21, 2012.
- [10] MAS-COLELL, ANDREU, MICHAEL DENNIS WHINSTON, JERRY R GREEN et al.: *Microeconomic theory*, volume 1. Oxford university press New York, 1995.
- [11] MILGROM, PAUL and JOHN ROBERTS: *Price and Advertising Signals of Product Quality*. *Journal of Political Economy*, 94(4):796–821, 1986.
- [12] OKADA, AKIRA: *Robustness of equilibrium points in strategic games*. Inst. of Technology, Department of Information Sciences, 1983.
- [13] VOORNEVELD, MARK and JÖRGEN W WEIBULL: *A Scent of Lemon - Seller Meets Buyer with a Noisy Quality Observation*. *Games*, 2(1):163–186, 2011.
- [14] ZEITHAML, VALARIE A.: *Consumer Perceptions of Price, Quality, and Value: A Means-End Model and Synthesis of Evidence*. *Journal of Marketing*, 52(3):2–22, 1988.