

Voorneveld, Mark; Weibull, Jörgen W.

Working Paper

Prices and quality signals

SSE/EFI Working Paper Series in Economics and Finance, No. 551

Provided in Cooperation with:

EFI - The Economic Research Institute, Stockholm School of Economics

Suggested Citation: Voorneveld, Mark; Weibull, Jörgen W. (2004) : Prices and quality signals, SSE/EFI Working Paper Series in Economics and Finance, No. 551, Stockholm School of Economics, The Economic Research Institute (EFI), Stockholm

This Version is available at:

<https://hdl.handle.net/10419/56233>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Prices and Quality Signals*

Mark Voorneveld
Department of Economics
Stockholm School of Economics
P.O. Box 6501
113 83 Stockholm, Sweden

Jörgen W. Weibull
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215, USA

February 12, 2004. This revision July 1, 2004.

Abstract

We consider a market-for-lemons model where the seller is a price setter, and, in addition to observing the price, the buyer receives a private noisy signal of the product's quality, such as when a prospective buyer looks at a car or house for sale, or when an employer interviews a job candidate. Sufficient conditions are given for the existence of perfect Bayesian equilibria, and we analyze equilibrium prices, trading probabilities and gains of trade. In particular, we identify separating equilibria with partial and full adverse selection as well as pooling equilibria. We also study the role of the buyer's signal precision, from being completely uninformative (as in standard adverse-selection models) to being completely informative (as under symmetric information). The robustness of results for these two boundary cases is analyzed, and comparisons are made with established models of monopoly and perfect competition.

Keywords: lemons, adverse selection, noisy quality signals.

JEL codes: C72, D82.

SSE/EFI WORKING PAPER SERIES IN ECONOMICS AND FINANCE, No. 551.

*We would like to thank Lars-Göran Mattsson for many inspiring discussions at the beginning of this research project, and Carlos Alos-Ferrer, Tore Ellingsen, Jianjun Miao and Zvika Neeman for comments to an earlier version. Voorneveld thanks the Jan Wallander and Tom Hedelius Foundation for financial support.

1 Introduction

In many market contexts, prospective buyers are imperfectly informed about the utility, value or quality of the item for sale. Sellers, on the other hand, often know the quality of their product. In his classical “market for lemons” paper, George Akerlof has shown that such informational asymmetry can lead to adverse selection. Although there is room for beneficial trade in all qualities if only buyers were able to distinguish these from each other, the complete lack of such discriminatory power may imply a crowding-out of good items, leaving only the “lemons” in the market (Akerlof, 1970). This was one of the starting points for the surge of economic analyses of markets with asymmetric information, where typically one side of the market has better information than the other: the owner of a used car knowing more about the quality of his car than the prospective buyer, insurance clients knowing more about their risks than the insurance company, job applicants knowing more about their abilities than employers.¹

In the current paper, we formulate a simple — and we think natural — generalization of Akerlof’s model: in addition to observing the price posted by the seller, the buyer receives a private noisy signal about the product’s quality. We believe this is a realistic and relevant extension. Consumers typically have the opportunity to take a look at the objects for sale, thus getting an impression of their qualities, and employers interview job candidates, thus obtaining an impression of the candidates’ abilities, etc. In somewhat metaphorical language: while buyers have perfect vision in the classical model of symmetric information and are totally blind in Akerlof’s model, the present model covers a continuum range of intermediate degrees of vision, spanning from one of these extremes to the other. Moreover, instead of being a price taker, as in Akerlof’s model, the seller is here a strategic price setter. Thus, while Akerlof treated the case of perfect competition, we here treat the case of monopoly.

Roughly, the model is as follows: an indivisible good is available in two qualities, low and high. Buyers are interested in buying at most one unit of the good. At the time of purchase, they cannot observe with certainty the quality of the unit for sale. Each seller has a single unit for sale, and knows its quality. Sellers’ reservation prices and buyers’ willingness to pay are such that mutually beneficial trade is possible in both qualities. When a seller and a buyer meet, the seller sets a take-it-or-leave-it price for his unit. The buyer is informed of this price, and makes a noisy observation of the quality, modelled as a convex combination of actual quality and a random noise term. Hence, there is *two-sided* asymmetric information in the sense that the buyer does not know the quality of the item for sale, and the seller does not know the impression the buyer has of the item’s quality. This paper thus stands somewhat outside the usual literature on adverse selection, where it is usually assumed that only one side of the market is uninformed.² Our modelling approach is instead close to that used when analyzing moral hazard, where it is typically

¹We cannot possibly do justice to this rich literature here. The interested reader is referred to Riley (2001) for a recent survey and to Mas-Colell *et al.* (1995) for a textbook treatment.

²In this respect, the present model is more similar to principal-agent models, where the principal makes a noisy observation of the agent’s effort.

assumed that one party observes the other party's action with some noise. It should also be noted that the model is quite distinct from Spence's (1973) signalling model, which is a model of one-sided asymmetric information — the seller knows the signal received by the buyer. Moreover, the signal in Spence's model is chosen by the seller, and the signal need not, and in most applications does not, affect the value to the buyer of the unit for sale (say, the job candidate's productivity). By contrast, here the seller cannot (by assumption), affect the signal, but the signal is statistically correlated with the value of the unit to the buyer.

The solution concept we apply is that of (weak) perfect Bayesian equilibrium, that is, a player's beliefs along the path of play are required to be consistent with other players' strategies and the probabilities for nature's moves, and, at every information set of a player, the player's strategy should be sequentially rational in the sense of maximizing the player's conditionally expected payoff, given the player's beliefs at the information set and others' continuation strategies. We analyze the set of prices and trades that can be sustained in pure-strategy equilibria of this kind. The main focus of the analysis is on intermediate degrees of signal precision, but we also study the two boundary cases of *symmetric information*, that is, when buyers observe the quality of the unit for sale without error, and of *(one-sided) asymmetric information*, that is, when buyers have no information at all about the quality of the unit at hand. Taking limits of buyers' signal precision towards these two bounds allows us to study the robustness of equilibrium predictions for the two boundary cases.

Our main results are as follows. First, we prove the existence of separating equilibria. There are two distinct classes of such equilibria. In one class, high-quality items are bought with a positive probability below one. Buyers buy only if the quality signal is sufficiently high, and this keeps sellers with low-quality units from deviating to the high price. Low quality units, by contrast, are sold with probability one. There is a whole continuum of such separating equilibria, also in the limit as the buyer's signal precision is taken down to zero. The upper bound then falls to a positive value, allowing for some trade in high-quality units even when buyers have no information at all about the quality of the unit at hand. Low-quality sellers are kept from deviating to the high price because of the low probability that buyers will bite at that price, while high-quality sellers make insufficient profits at the low price. The other class of separating equilibria results in total adverse selection. If high-quality sellers' reservation price is not lower than buyers' willingness to pay for low quality, then sellers of high-quality units outprice themselves, leading to no trade in high-quality units. These sellers, who cannot make a profit at the low price, are kept from deviating to intermediate prices by consumers' skepticism. Buyers interpret any non-equilibrium price as coming from a seller with a low-quality unit.

Second, we provide necessary and sufficient conditions for pooling equilibrium. Unlike in the standard model of one-sided asymmetric information, buyers condition their purchasing decision, at the going pooling price, on the quality signal they receive: they buy only if this signal is above a certain equilibrium threshold. Sellers with low-quality goods thus have a lower probability of selling than sellers with high-quality goods (their cus-

tomers' signals are stochastically dominated by those emanating from high-quality units). However, the equilibrium selling probability is still sufficiently high to deter these sellers from deviating to buyers' willingness to pay for low quality, a price at which they would be sure to sell. There is in general a whole continuum of such pooling equilibrium prices, and buyers (rationally) condition on their signal even if their signal has low precision. Indeed, for some parameter constellations this conditioning occurs even in the limit as the precision of buyers' signals is taken down to zero. This occurs when the going price equals the average quality (the usual pooling price in models of one-sided asymmetric information). The reason for this phenomenon is that consumers *a priori* are indifferent between buying and not buying at that price, and hence even a weak negative signal will deter a buyer from purchase at that price. If the noise is symmetric and unimodal, the limiting probability for purchase is one half.

Also in comparison with the classical symmetric-information model do we note two discontinuities. First, suppose that the sellers' reservation price for high-quality units is at least as high as buyers' willingness to pay for low quality. Then there are separating equilibria in our model where high-quality sellers withdraw from the market by outpricing themselves, as mentioned above, even when buyers' signals are arbitrarily close to being perfect. By contrast, no such separating equilibria exist if the buyer can observe quality perfectly. When buyers' willingness to pay for low quality is at least as high as the sellers' reservation price for high-quality units, then there exists a pooling equilibrium in our model with a price equal to the buyer's willingness to pay for low quality. By contrast, the standard symmetric information model has no equilibria where low and high-quality sellers charge the same price.

Although there certainly is a related literature, to the best of our knowledge the present model is new. For example, in Milgrom and Roberts (1986), sellers signal their quality by way of a posted price and expenditures on advertisement. The seller knows the combined price-advertisement signal received by buyers, so that is a model of one-sided asymmetric information. Bagwell and Riordan (1991) develop a model with informed and uninformed consumers, where the latter update their initial beliefs about quality on the basis of the price posted by the seller. Hence, also here the seller knows the signal — his price — received by buyers. Schlee (1996) considers a model in which there is quality uncertainty on both sides of the market, but where buyers and sellers are symmetrically informed, that is, available information about quality is public. Hence, in none of these models is there two-sided asymmetric information, and in none of the models do consumers have some discriminatory power concerning quality. Ellingsen (1997) considers a version of Akerlof's (1970) model of one-sided asymmetric information, where the seller, like here, is a strategic price setter, rather than a price-taker. Bester and Ritzberger (2001) analyze interactions of the type modelled in Ellingsen, but where the buyer may undertake a test and find out the exact quality. The relation to these two papers will be highlighted below.

The paper is organized as follows. The model is formalized in section 2. Section 3 prepares the ground for the equilibrium analysis by way of providing certain necessary conditions on prices to be consistent with pure-strategy perfect Bayesian equilibrium. Separating equilibria are studied in section 4 and pooling equilibria in section 5. In

section 6, we consider an equilibrium refinement based on the behavioral assumption that buyers also off the equilibrium path use Bayesian updating to form beliefs. In section 7 we compare the equilibria under one-sided asymmetric information with those that remain in the limit as the signal precision is brought down towards zero, and, similarly, we compare the equilibria under symmetric information with those that remain in the limit as the signal precision approaches one. Section 8 discusses non-pure behavior strategies, and section 9 concludes with a discussion of directions for further research.

2 Model

An indivisible good is available in two qualities, low and high, denoted q_L and q_H , where $0 < q_L < q_H$. Buyers are *ex ante* identical. Without loss of generality, we take the *buyers'* valuation of (or *willingness to pay* for) the two qualities to be $v_L = q_L$ and $v_H = q_H$, respectively. Each seller has one unit of the good for sale. The seller knows the quality of his or her unit. His or her *valuation* (or *reservation price*) is $w_L < v_L$ for a low-quality item and $w_H < v_H$ for a high-quality item, where $w_H \geq w_L$. The probability that the seller's unit is of low quality is $\lambda \in (0, 1)$. Let $q \in \{q_L, q_H\}$ be the quality of his product, and let \bar{q} denote the average quality: $\bar{q} = \lambda q_L + (1 - \lambda) q_H$. This is thus also the buyer's average valuation, $\bar{v} = \bar{q}$.³

The seller sets a take-it-or-leave-it price p for his unit, conditional on its quality. The buyer is informed of the price p set by the seller, but cannot observe the quality of the item at hand. Instead, the buyer makes a noisy (or perturbed) observation of the quality: he or she observes a mean-preserving mix of the unit's true quality and a noise term. More exactly, the buyer observes

$$\tilde{q} = \alpha q + (1 - \alpha) \varepsilon, \tag{1}$$

where $\alpha \in [0, 1]$, and where ε is a random variable with mean \bar{q} , distributed according to a cumulative probability distribution function $F : \mathbb{R} \rightarrow [0, 1]$ with continuous and everywhere positive density $f : \mathbb{R} \rightarrow \mathbb{R}_+$.⁴ The draws of the true quality q and the noise term ε are statistically independent. Given F , the parameter α thus determines the *precision* of the buyer's information about the quality of the unit at hand: the higher α is, the more precise is the buyer's signal \tilde{q} about the unit's quality. In particular, when $\alpha = 1$, then we have the classical case of *symmetric information* when buyers observe quality without error, while the opposite extreme case, $\alpha = 0$, when buyers have no information at all about the quality of the unit at hand, corresponds to the usual

³Mathematically, our normalization to set buyers' valuation v^θ equal to quality q^θ allows us to use the notation v and q interchangeably. Conceptually, however, the two are different: we think of quality as something objective while buyers' valuations thereof are subjective. Our reason for introducing the two different notations is that in models with heterogeneous buyers — not considered in this paper — the same quality may well imply different valuations for different buyers. We have strived to use the notation q when referring to the physical aspect of quality and v when referring to its subjective aspect, its value to buyers.

⁴To be more precise, we assume F to be absolutely continuous: $F(z) = \int_{-\infty}^z f(z) dz$ for all $z \in \mathbb{R}$.

model of one-sided *asymmetric information*, pioneered by Akerlof (1970).⁵ In all cases between these two extremes, buyers make imperfect observations of quality, and sellers know this. In a random matching between a seller and a buyer, the seller does not know the quality signal \tilde{q} observed by the buyer. We will refer to these intermediate cases, $\alpha \in (0, 1)$, as interactions with *two-sided asymmetric information*. This parametrization of the buyer’s signal is made in order to include the traditional cases of symmetric and one-sided asymmetric information as special cases.⁶

The sellers are expected-profit maximizers, where $\pi = p - w$ is the profit to a seller with valuation $w \in \{w_L, w_H\}$ who sells a unit at price p . Likewise, the buyers are risk-neutral utility maximizers, where $u = v - p$ is the utility to a buyer when paying p for a unit of value $v \in \{v_L, v_H\}$ to the buyer. The buyer’s utility and the seller’s profit are both normalized to zero if no trade takes place. Hence, in force of the assumptions made above, there are positive potential gains of trade in both cases: $v_L - w_L > 0$ and $v_H - w_H > 0$. Expected *potential* gains of trade are thus

$$\hat{W} = \lambda(v_L - w_L) + (1 - \lambda)(v_H - w_H) > 0. \quad (2)$$

One of the main concerns from a welfare viewpoint is to what extent these potential gains will be realized in equilibrium, and how this depends on the parameters in the model, in particular, on the precision α of buyers’ information.

2.1 The seller-buyer game

Formally, we model the interaction outlined above as an extensive-form game with incomplete information between a seller, player 1, and a buyer, player 2, as follows:

1. Nature chooses $\theta \in \{L, H\}$, where L has probability λ and H probability $1 - \lambda$. Nature also chooses $\varepsilon \in \mathbb{R}$, according to the distribution F , and these two moves by nature are statistically independent.
2. The seller observes θ and chooses a price $p_\theta \in \mathbb{R}$. A pure strategy for the seller is thus a pair $s = (p_L, p_H) \in \mathbb{R}^2$ of prices, where p_L is the price when $\theta = L$ and p_H the price when $\theta = H$.
3. The buyer observes the price-signal pair (p_θ, \tilde{q}) , where $\tilde{q} = \alpha q_\theta + (1 - \alpha)\varepsilon$ (see equation (1)), and decides whether to buy the item or not. A pure strategy for

⁵In the case $\alpha = 0$, there is a subtle difference, though: while the buyer observes no signal at all in the traditional model of one-sided asymmetric information, here the buyer observes a signal that is statistically independent from the true quality of the item at hand. See section 7 for a discussion of this aspect.

⁶For intermediate cases of two-sided asymmetric information, it would be sufficient to use the simpler formulation $\tilde{q} = q + \varepsilon'$, where ε' is a random variable. Indeed, for $\alpha \in (0, 1)$ fixed, the two representations are mathematically equivalent. Comparative statics with respect to buyer information in the simpler formulation can be done by varying the “dispersion” of the error term. For example, with $\varepsilon' \sim N(0, \sigma^2)$ one could study the role of σ^2 very much along the same lines as we here study the role of α .

the buyer is thus a (Borel-measurable) function $b : \mathbb{R}^2 \rightarrow \{0, 1\}$, where $b(p, \tilde{q}) = 1$ means “buy” and $b(p, \tilde{q}) = 0$ means “don’t buy,” given any price p and quality signal \tilde{q} .

4. The game ends and the players receive their payoffs,

$$\pi = (p_\theta - w_\theta) b(p_\theta, \tilde{q}) \quad \text{and} \quad u = (v_\theta - p_\theta) b(p_\theta, \tilde{q}). \quad (3)$$

A decision node for the seller is thus a pair (θ, ε) , but the seller does not observe ε , the noise term in the buyer’s perception of the quality of the item at hand. The seller thus has two information sets (or “seller types”), namely $\{(L, \varepsilon) : \varepsilon \in \mathbb{R}\}$ and $\{(H, \varepsilon) : \varepsilon \in \mathbb{R}\}$, depending on whether his item is of low or high quality. Likewise, a decision node for the buyer is a triplet (θ, ε, p) , but the buyer observes only the seller’s price p and the signal \tilde{q} . The buyer thus has infinitely many information sets (or “buyer types”). Each of the buyer’s information sets is characterized by a pair $(p, \tilde{q}) \in \mathbb{R}^2$, specifying the seller’s price p and the received quality signal \tilde{q} . For $\alpha < 1$, each such information set contains exactly two nodes of the game tree, namely (L, ε_L, p) and (H, ε_H, p) , where

$$\varepsilon_\theta = \frac{\tilde{q} - \alpha q_\theta}{1 - \alpha} \quad \text{for } \theta \in \{L, H\}. \quad (4)$$

By contrast, if $\alpha = 1$, then, by (1), the quality signal \tilde{q} coincides with the true quality: $\tilde{q} \in \{q_L, q_H\}$. Hence, for any $\theta \in \{L, H\}$, the buyer information set (p, q_θ) contains infinitely many nodes of the game tree, namely (θ, ε, p) , for all $\varepsilon \in \mathbb{R}$. In other words, the observation of \tilde{q} then does not give the buyer any clue about ε . However, in this boundary case of perfect precision, the random variable ε is irrelevant to the buyer’s decision making, since we then have $\tilde{q} = q_\theta$: the buyer observes the true quality without error. As noted above: what makes this game different from the standard “lemons model” (see, for example, chapter 13 in Mas-Colell, Whinston, and Green, 1995), is that the seller is a price setter rather than price taker and that there is incomplete information on both sides of the market: the buyer does not know the quality of the item for sale, and the seller does not know what signal the buyer has about the unit’s quality. Note also that the “types” of the seller and the buyer are correlated: according to (1) the buyer observes the random variable $\tilde{q} = \alpha q_\theta + (1 - \alpha)\varepsilon$, where $\theta \in \{L, H\}$ is the random variable observed by the seller.

2.2 Equilibrium

We solve this game by adapting the usual (weak) perfect Bayesian equilibrium concept to the present setting (see, *e.g.*, Mas-Colell *et al.* (1995)). Such an equilibrium specifies a strategy for each player, and beliefs for the players over the nodes in their information sets, such that (a) beliefs are consistent with all players’ strategies and the probabilities for nature’s moves, along the induced path of play, and (b) at each information set, the strategy of the concerned player is optimal, given the player’s belief at the information

set and the strategies of all other players (there is no future move by nature in our game). Requiring best replies also off the path of play, this concept is more stringent than (Bayesian) Nash equilibrium, but it imposes no constraints on beliefs at information sets off the path of play; in this sense it is weaker than sequential equilibrium.⁷ We will only consider pure strategies. The present section concerns the case of incompletely informed buyers ($\alpha < 1$); the case of completely informed buyers is discussed in section 3.2. Formally (and following Kreps and Wilson, 1982), a *belief system* is a function μ that assigns to each information set a probability distribution over the nodes in that information set. A *pure-strategy profile* in the present game is a pair (s, b) , where $s = (p_L, p_H) \in \mathbb{R}^2$, a pure strategy for the seller, and $b : \mathbb{R}^2 \rightarrow \{0, 1\}$ is a (Borel-measurable) function, a pure strategy for the buyer. We will call a triplet (μ, s, b) an *equilibrium* if it meets conditions [B1]-[B3] and [S] below.

[B1] At each seller information set $\{(\theta, \varepsilon) : \varepsilon \in \mathbb{R}\}$, and for every $x \in \mathbb{R}$, μ assigns probability $F(x)$ to the subset of nodes $\{(\theta, \varepsilon) : \varepsilon \leq x\}$.

[B2] If $p_L \neq p_H$, then at each buyer information set (p_θ, \tilde{q}) , μ assigns probability 1 to the node $(\theta, \varepsilon_\theta, p_\theta)$, where ε_θ is defined in (4).

[B3] If $p_L = p_H$, then at each buyer information set (p, \tilde{q}) where $p = p_L$, μ assigns probability

$$\mu_L(\tilde{q}) = \frac{\lambda f(\varepsilon_L)}{\lambda f(\varepsilon_L) + (1 - \lambda)f(\varepsilon_H)} \quad (5)$$

to the node (L, ε_L, p) , and thus probability $\mu_H(\tilde{q}) = 1 - \mu_L(\tilde{q})$ to the node (H, ε_H, p) .⁸

[S] At all information sets: the concerned player's strategy is a best response to the other player's strategy, under the belief induced by μ at that information set.

Consistency condition [B1] is met if the seller knows the quality of his item and knows the c.d.f. F of the buyer's noise term ε . At each of his two information sets, $\{(L, \varepsilon) : \varepsilon \in \mathbb{R}\}$ and $\{(H, \varepsilon) : \varepsilon \in \mathbb{R}\}$, the seller's beliefs are then given by the distribution function F for ε . Conditions [B2] and [B3] are met if, at every information set of the buyer that can be reached under the strategy profile (s, b) , the buyer's beliefs are consistent with (s, b) , by way of Bayesian updating. More exactly, condition [B2] requires that if the seller charges different prices, p_L and p_H , depending on the quality of his good, an incompletely informed buyer should believe that the item is of low (high) quality if he price is p_L (p_H). Similarly, condition [B3] requires that if the seller charges the same price for high and low-quality items, the buyer's belief about which of the two nodes was reached should be consistent with Bayes' rule applied to the information set in question.⁹ Finally, condition

⁷Note also that the present game has infinitely many information sets, and thus departs from the usual setting for sequential equilibrium. For exact definitions, see Fudenberg and Tirole (1991) and Mas-Colell *et al.* (1995). For additional remarks on belief constraints, see section 9.

⁸In the boundary case $\alpha = 0$, $\varepsilon_\theta = \tilde{q}$ for both $\theta = L$ and $\theta = H$, and thus $\mu_L(\tilde{q}) = \lambda$ for all \tilde{q} , by (5).

⁹Expressed differently, [B3] requires the probability $\mu_L(\tilde{q})$ to be consistent with the conditional likelihood of the true quality, given the observed quality \tilde{q} .

[S] requires sequential rationality as defined in Kreps and Wilson (1982). The subsequent analysis is focused on pure-strategy equilibria $(\mu, (p_L, p_H), b)$ as described above. Such an equilibrium is *separating* if $p_L \neq p_H$ and *pooling* if $p_L = p_H$. Let ρ_L and ρ_H denote the equilibrium *trading probabilities* for low and high quality units:

$$\rho_\theta = \Pr[b(p_\theta, \tilde{q}) = 1 \mid \theta] \quad (6)$$

for $\theta = L, H$. Expected equilibrium gains of trade can then be written

$$W^* = \lambda \rho_L (v_L - w_L) + (1 - \lambda) \rho_H (v_H - w_H). \quad (7)$$

By an *equilibrium outcome* we mean a quadruple $(p_L, p_H, \rho_L, \rho_H)$ such that $(\mu, (p_L, p_H), b)$ is a perfect Bayesian equilibrium for some buyer strategy b satisfying (6).

3 Preliminaries

Before embarking on the equilibrium analysis, we briefly consider certain necessary conditions on price pairs (p_L, p_H) for these to be sustained in equilibrium, and consider the boundary case of symmetric information.¹⁰

3.1 Necessary conditions for equilibrium

The necessary conditions for equilibrium given below follow more or less immediately from first principles. For instance, it is sequentially rational for the buyer to buy if the seller would ask a price below the buyer's willingness to pay for a low-quality unit, and it is always sequentially rational not to buy if the seller would ask a price above the buyer's willingness to pay for a high-quality unit. In sum:

$$\forall p < v_L, \forall \tilde{q} \in \mathbb{R} : b(p, \tilde{q}) = 1. \quad (8)$$

and

$$\forall p > v_H, \forall \tilde{q} \in \mathbb{R} : b(p, \tilde{q}) = 0. \quad (9)$$

Using this simple observation, a wide range of price pairs can be ruled out from equilibrium *a priori*:

Lemma 3.1 *For $\alpha < 1$ there are no equilibria $(\mu, (p_L, p_H), b)$ where*

(a) $\min \{p_L, p_H\} < v_L$,

(b) $p_L \neq p_H$ and $p_L > v_L$,

¹⁰Recall that our model exhibits a form of “discontinuity” at that boundary. At $\alpha = 1$, the buyer's information set contains infinitely many nodes, while for any $\alpha < 1$ it contains only two nodes.

(c) $p_L = v_L$ and $p_H \in (v_L, v_H)$,

(d) $p_L = p_H \geq v_H$.

Proof: (a) If a seller of type $\theta \in \{L, H\}$ sets a price below v_L , then, by (8), his expected payoff can be increased from $p - w_\theta$ to $p' - w_\theta$ by deviating to a price $p' \in (p, v_L)$.

(b) Suppose $p_L \neq p_H$ and $p_L > v_L$. Consider an information set (p_L, \tilde{q}) of the buyer. By [B2], he believes that the item is of low quality, with probability one. Since $p_L > v_L$, the optimal decision by [S] is not to buy. So $b(p_L, \tilde{q}) = 0$ for all $\tilde{q} \in \mathbb{R}$. Hence, the expected payoff to a seller with a low-quality unit equals zero, while (8) implies that the seller can deviate to a price $p \in (w_L, v_L)$ and obtain a positive payoff, $p - w_L$.

(c) Suppose $p_L = v_L$ and $p_H \in (v_L, v_H)$. Consider an information set (p_H, \tilde{q}) of the buyer. By [B2], he believes with probability one that the item is of high quality. Since $p_H < v_H$, the optimal decision by [S] is to buy. So $b(p_H, \tilde{q}) = 1$ for all $\tilde{q} \in \mathbb{R}$. The low-quality seller can then increase his expected payoff from $(q_L - w_L) \cdot \Pr[\tilde{q} \in \mathbb{R} \mid b(q_L, \tilde{q}) = 1]$ to $p_H - w_L$ by charging price p_H instead of q_L .

(d) Suppose first $p_L = p_H > v_H$. By (9), the good is not sold at this price, so the expected payoff to the low-quality seller equals zero, while (8) implies that he can deviate to a price $p \in (w_L, v_L)$ and make profit $p - w_L > 0$. Suppose now that $p_L = p_H = v_H$, and consider an information set (v_H, \tilde{q}) of the buyer. By [B3], he believes with probability (5) that the item is of low quality, in which case a purchase would give negative utility ($v_L - v_H < 0$), and with the residual probability that the item is of high quality, in which case a purchase would give zero utility ($v_H - v_H = 0$). The expected payoff thus being negative, the optimal decision by [S] is not to buy. The expected payoff of the seller with a low-quality unit hence equals zero, while (8) implies that such a seller can deviate to a price $p \in (w_L, v_L)$ and guarantee himself a positive payoff $p - w_L$. **End of proof.**

We note that, by (a) and (b), in separating equilibria the price of low-quality units equals buyers' willingness to pay for such units, and, by (c), the price of high-quality units is never below the buyers' willingness to pay for such units. When trade occurs in separating equilibria, the seller thus reaps all the gains of trade. This is not surprising, since by assumption the sellers by assumption can commit to take-it or leave-it offers. By (a) and (d): in pooling equilibria the price is never below buyers' willingness to pay for low-quality units but always lower than buyers' willingness to pay for high-quality units. In sum:

Remark 3.2 It follows from the lemma that the only remaining candidate pure-strategy equilibria, when $\alpha < 1$, are: (i) separating equilibria with $p_L = v_L$ and $p_H \geq v_H$, (ii) pooling equilibria with $p_L = p_H \in [v_L, v_H)$. \triangleleft

Before embarking on an analysis of separating and pooling equilibria in cases when buyers have less than perfect information, we briefly consider the limiting case of symmetric information ($\alpha = 1$).

3.2 Symmetric information

The classical case of symmetric information, that is, when buyers know exactly the quality of the goods they buy, corresponds to $\alpha = 1$ in the present model. Pure-strategy perfect Bayesian equilibrium then requires the buyer to make a purchase if and only if the posted price does not exceed the received quality signal $\tilde{q} = q_\theta$, where q_θ is the quality of the unit at hand.¹¹ The seller's unique best reply to this buyer strategy is to set the price equal to the buyer's willingness to pay, $(p_L, p_H) = (v_L, v_H)$. Hence, both low and high quality items are traded with probability one in equilibrium: $\rho_L = \rho_H = 1$. Moreover, consumer surplus is zero, $\mathbb{E}[u] = 0$ (see (3)), while the seller's expected profit equals the potential gains of trade: $\mathbb{E}[\pi] = \hat{W}$ (see (2)). Hence, all potential gains of trade are realized in equilibrium in this boundary case: $W^* = \hat{W}$.

4 Separating equilibria

The aim of this section is to establish conditions for the existence of separating equilibria, and to investigate their nature. Having already dealt with the boundary case $\alpha = 1$, we from now on focus on $\alpha \in [0, 1)$. We will show that low-quality items are sold with probability one ($\rho_L = 1$) in all separating (pure-strategy) equilibria and identify two varieties of separating equilibria. In one variety, high-quality units are not sold ($\rho_H = 0$), while in the other variety they are sold with positive probability below one ($0 < \rho_H < 1$). We will refer to the first variety as “total” adverse selection. These are equilibria in which high-quality sellers outprice themselves, and thereby “leave the market,” just as in “classical” adverse selection. The second variety will be referred to as “partial” adverse selection.

4.1 Total adverse selection

We noted above that the lower bound on the equilibrium probability for trade in high-quality units is zero if $v_L \leq w_H$. Indeed, in this case, and only then, do there exist separating equilibria in which the price of high-quality items exceeds v_H . Since it is suboptimal to buy at such high prices, there is no trade in high-quality goods. The reason why sellers with high-quality units do not lower their price is that, in these equilibria, buyers believe that such a seller may well hold a low-quality item, and therefore buyers do not “bite,” unless, of course, the seller would set the price as low as v_L , resulting in non-positive profits. Hence, in these equilibria, it is as if sellers with high-quality units had left the market — as under classical adverse selection.

Proposition 4.1 *Suppose $\alpha < 1$. Prices $p_L = v_L$ and $p_H > v_H$ constitute an equilibrium price pair if and only if $v_L \leq w_H$. In all such equilibria $\rho_L = 1$ and $\rho_H = 0$.*

¹¹The buyer's information sets are of the form (p, q_θ) and each such set consists of the nodes $\{(\theta, \varepsilon, p) : \varepsilon \in \mathbb{R}\}$. Perfect Bayesian equilibrium requires that the belief system μ assigns probability $F(x)$ to each subset of nodes $\{(\theta, \varepsilon, p_\theta) : \varepsilon \leq x\}$. However, this consistency requirement is only cosmetic: the error term ε is irrelevant to the buyer when $\alpha = 1$.

Proof: To prove necessity, suppose $(\mu, (v_L, p_H), b)$ is such an equilibrium. Condition (9) implies that sellers with high-quality units obtain zero profit. By condition (8), a deviation from p_H to a price $p < v_L$ sufficiently close to v_L guarantees a profit arbitrarily close to $v_L - w_H$. Hence, for $p_H > v_H$ to be an equilibrium price, it must be that $v_L - w_H \leq 0$. As for sufficiency, assume $v_L \leq w_H$. Let (μ, s, b) be such that $p_L = v_L, p_H > v_H$, let

$$b(p, \tilde{q}) = \begin{cases} 1 & \text{if } p \leq v_L, \\ 0 & \text{otherwise,} \end{cases}$$

and let μ satisfy [B1] and [B2] (again [B3] has no bite in separating equilibria). At information sets (p, \tilde{q}) of the buyer off the path of (s, b) (*i.e.*, with $p \notin \{p_L, p_H\}$), let μ assign probability one to the node (L, ε_L, p) : the buyer then believes the item is of low quality. It is easy to verify that (μ, s, b) constitutes a perfect Bayesian equilibrium. First, the belief requirements are fulfilled by construction. Second, consider the information sets of the buyer. At information sets (v_L, \tilde{q}) , the buyer is indifferent between buying and not buying. Hence, b is a best reply at all these information sets. At information sets (p_H, \tilde{q}) , the buyer can only suffer from purchasing the item: the unique best reply is not to buy. Hence, b is a best reply also at these information sets. At information sets (p, \tilde{q}) off the equilibrium path, *i.e.*, with $p \notin \{p_L, p_H\}$, the buyer believes the item is of low quality, so buying at prices $p < v_L$ and not buying at prices $p > v_L$ is optimal: b is a best reply also at these information sets. Third, consider the information sets of the seller. At information set $\{(L, \varepsilon) : \varepsilon \in \mathbb{R}\}$, a price $p \leq v_L$ yields profit $p - w_L$, so in this price range the optimal price is v_L , yielding profit $v_L - w_L > 0$. A price $p > v_L$ results in no trade and hence profit zero. Hence, the price $p_L = v_L$ is indeed a best reply at this information set. At information set $\{(H, \varepsilon) : \varepsilon \in \mathbb{R}\}$, a price $p \leq v_L$ yields profit $p - w_H \leq 0$, since $v_L \leq w_H$, and a price $p > v_L$ yields no sale and hence profit zero. Hence, any price $p > v_H$ is indeed a best reply at this information set.

To establish the last claim in the proposition, consider any equilibrium (μ, s, b) with $p_L = v_L$ and $p_H > v_H$. It follows as in the proof of the preceding proposition that $b(v_L, \tilde{q}) = 1$ for almost every $\tilde{q} \in \mathbb{R}$, and from (9) that $b(p_H, \tilde{q}) = 0$ for every $\tilde{q} \in \mathbb{R}$. **End of proof.**

In other words, if $w_H \geq v_L$, then there are equilibria where high-quality sellers withdraw from the market by outpricing themselves, regardless of information precision $\alpha \in [0, 1)$. In particular, such equilibria exist even as buyers' information is made arbitrarily precise, that is, as $\alpha \rightarrow 1$, despite the fact that no such equilibria exist under symmetric information, that is, when $\alpha = 1$ (see section 3.2). In this sense, the classical result for symmetric equilibrium is non-robust: the slightest noise in buyers' quality observations allows for equilibria where adverse selection hits at full force. At the other end of the information spectrum, when $\alpha = 0$, proposition 4.1 establishes that outpricing equilibria exists also under one-sided asymmetric information, granted $w_H \geq v_L$. Note, in particular, that this result allows for equilibria with complete adverse selection even when $w_H \leq \bar{v}$, the condition for the existence of pooling equilibria in standard models of one-sided asymmetric information. We should add that we do not find these separating

equilibria very plausible. For if a seller with a high-quality unit were certain that no buyer will buy at a price above their highest willingness to pay, v_H , then such a seller couldn't lose, but might gain, by deviating to a price p below v_H (but above w_H). More formally, none of the equilibria in proposition 4.1 survive two rounds of elimination of weakly dominated strategies. To see this, note that any buyer strategy b that assigns $b(p, \tilde{q}) > 0$ to some $p > v_H$ and $\tilde{q} \in \mathbb{R}$ is weakly dominated by the buyer strategy b' that agrees with b everywhere except at the information set (p, \tilde{q}) , where b' instead takes the value zero. After all such weakly dominated buyer strategies have been eliminated from the game all seller strategies $s = (p_L, p_H)$ with $p_H > v_H$ become weakly dominated by any seller strategy $s' = (p_L, p)$ for which $p \in (w_H, v_H]$. All separating equilibria resulting in total adverse selection can hence be discarded if this refinement is deemed appropriate to the application at hand. Alternatively: in view of the irrationality of prices $p > v_H$, as well as prices $p < w_L$, one could argue that the seller's strategy set should be $[w_L, v_H]^2$, instead of, as here, all of \mathbb{R}^2 . Any of these approaches will lead to the elimination of the whole set of equilibria with total adverse selection.

4.2 Partial adverse selection

We will construct a continuum of separating with partial adverse selection by focusing on threshold strategies for the buyer, that is, strategies according to which the buyer buys at the high price, $p_H = v_H$ if and only if the quality signal passes above a certain threshold, specific to each such equilibrium. For this purpose, we first define a set T of relevant threshold levels for the buyer:

Lemma 4.2 *Suppose $\alpha < 1$. The following two inequalities together define a non-empty interval $T \subset \mathbb{R}$ of scalars t :*

$$\Pr[\tilde{q} > t \mid \theta = L] \leq \frac{v_L - w_L}{v_H - w_L}, \quad (10)$$

$$\Pr[\tilde{q} > t \mid \theta = H] \geq \frac{v_L - w_H}{v_H - w_H}. \quad (11)$$

Proof: Recall that the c.d.f. for ε is F , and notice that, for any $t \in \mathbb{R}$,

$$\Pr[\tilde{q} > t \mid \theta] = 1 - F\left(\frac{t - \alpha q_\theta}{1 - \alpha}\right)$$

so (10) and (11) are equivalent with, respectively,

$$F\left(\frac{t - \alpha q_L}{1 - \alpha}\right) \geq \frac{v_H - v_L}{v_H - w_L} \quad (12)$$

and

$$F\left(\frac{t - \alpha q_H}{1 - \alpha}\right) \leq \frac{v_H - v_L}{v_H - w_H}. \quad (13)$$

Case 1: If $w_H \geq v_L$, then the right-hand side of (13) is at least 1, and hence (13) is met. Condition (12) is equivalent with

$$t \geq \alpha q_L + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right).$$

Case 2: If $w_H < v_L$, then $t \in \mathbb{R}$ satisfies (12) and (13) if and only if

$$\alpha q_L + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right) \leq t \leq \alpha q_H + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_H} \right).$$

By hypothesis, $q_L < q_H$, $w_L \leq w_H$, and F is strictly increasing, so the lower bound on T is a smaller number than its upper bound. **End of proof.**

As seen in the proof, the interval T is of the form $[t_0, +\infty)$ if $v_L \leq w_H$, and of the form $[t_0, t_1]$ if $v_L > w_H$, where

$$t_0 = \alpha q_L + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right) \tag{14}$$

and

$$t_1 = \alpha q_H + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_H} \right). \tag{15}$$

In the first case, $v_L \leq w_H$, sellers of high-quality units cannot make a positive profit at the price v_L , while in latter case they can. We also note that $t_0 \leq t_1 < +\infty$. We use lemma 4.2 to show that for every parameter combination (satisfying the assumptions made in section 2) there exists a set of separating equilibria where (a) the price of each quality equals buyers' willingness to pay for that quality, (b) items of low quality are sold with probability one, and (c) items of high quality are sold only with a positive probability below one. In these equilibria, the buyer's strategy is to buy from a seller who posts the low price, $p_L = v_L$, irrespective of the perceived quality \tilde{q} , but to buy from a seller who posts the high price, $p_H = v_H$, if and only if the perceived quality exceeds an threshold $t \in T$. More exactly:

Proposition 4.3 *Suppose $\alpha < 1$. All equilibria with $(p_L, p_H) = (v_L, v_H)$ have $\rho_L = 1$ and $\max \left\{ 0, \frac{v_L - w_H}{v_H - w_H} \right\} \leq \rho_H < 1$. There exists a continuum of equilibria, parametrized by $t \in T$, where $(p_L, p_H) = (v_L, v_H)$ and*

$$\rho_H = \Pr [\tilde{q} > t \mid \theta = H] = 1 - F \left(\frac{t - \alpha q_H}{1 - \alpha} \right). \tag{16}$$

Proof: In order to prove the first claim, suppose (μ, s, b) is an equilibrium with $(p_L, p_H) = (v_L, v_H)$. By [B2] and [S], the buyer buys with probability 1 at all prices $p < v_L$, irrespective of her signal \tilde{q} . Hence, by choosing $p < v_L$ arbitrarily close to v_L , a seller with a low-quality item can guarantee himself a profit arbitrarily close to $v_L - w_L$.

Hence, for $p_L = v_L$ to be optimal for the seller, it must be that $b(v_L, \tilde{q}) = 1$ for almost every $\tilde{q} \in \mathbb{R}$. To prove the second claim, let $(p_L, p_H) = (v_L, v_H)$, choose any $t \in T$, and let

$$b(p, \tilde{q}) = \begin{cases} 1 & \text{if } p \leq v_L, \text{ or if } p = v_H \text{ and } \tilde{q} > t, \\ 0 & \text{otherwise.} \end{cases}$$

This defines a pure-strategy pair, (s, b) , with

$$\rho_H = \Pr[\tilde{q} > t \mid \theta = H] = 1 - F\left(\frac{t - \alpha q_H}{1 - \alpha}\right) < 1.$$

Let μ be any belief system satisfying [B1] and [B2] ([B3] has no bite in separating equilibria). At information sets (p, \tilde{q}) of the buyer off the path of (s, b) , that is, where $p \notin \{v_L, v_H\}$, let μ assign probability 1 to the node (L, ε_L, p) . In other words, the buyer then believes the item is of low quality. We verify that (μ, s, b) is an equilibrium in three steps. First, consider the buyer's information sets. At information sets (v_L, \tilde{q}) and (v_H, \tilde{q}) , the buyer's expected payoff is zero regardless of her choice, for any $\tilde{q} \in \mathbb{R}$. Hence, b is a best reply at all information sets on the path of (s, b) . At information sets (p, \tilde{q}) off the path, the buyer believes the item is of low quality, so buying at a price $p \leq v_L$ and not buying at a price $p > v_L$ is a best reply also at such information sets. Secondly, consider the information set $\{(L, \varepsilon) : \varepsilon \in \mathbb{R}\}$ of the seller. A price $p \leq v_L$ would yield profit $p - w_L$, so in this price range the optimal price is v_L , yielding profit $v_L - w_L > 0$. A price $p > v_L$, differing from v_H , would result in no trade, and hence profit zero. The price $p = v_H$, by contrast, would result in trade with probability $\Pr[\tilde{q} > t \mid \theta = L]$. However, using (10), the resulting expected profit would not exceed

$$\frac{v_L - w_L}{v_H - w_L}(v_H - w_L) = v_L - w_L.$$

Hence charging the price v_L is indeed a best reply at this information set. Thirdly, consider the information set $\{(H, \varepsilon) : \varepsilon \in \mathbb{R}\}$ of the seller. A price $p \leq v_L$ would yield the profit $p - w_H$, so in this price range the optimal price is v_L , yielding profit $v_L - w_H$. A price $p > v_L$, differing from v_H , would result in no trade, and hence profit zero. The price $p = v_H$, by contrast, results in trade with probability $\Pr[\tilde{q} > t \mid \theta = H]$. Using (11), the resulting expected profit does not fall short of

$$\frac{v_L - w_H}{v_H - w_H}(v_H - w_H) = v_L - w_H.$$

Hence charging price v_H is indeed a best reply at this information set. It remains to prove the claim that $\rho_H \geq \max\left\{0, \frac{v_L - w_H}{v_H - w_H}\right\}$ in all separating equilibria. The inequality is trivially met if $w_H \geq v_L$, so suppose $w_H < v_L$. If $\rho_H < \frac{v_L - w_H}{v_H - w_H}$, then a seller with a high-quality unit would earn a higher profit by deviating to any price $p < v_L$ sufficiently close to v_L , since the buyer would surely buy at such a price and this the deviation profit would be arbitrarily close to $v_L - w_H$. **End of proof.**

We would like to make some comments to this proposition. First, the continuum of parametrized equilibria in the statement of the proposition are Pareto ranked: the lower the threshold $t \in T$, the higher is the probability ρ_H for trade in high-quality items, and thus the higher is the expected profit to sellers — the expected utility to buyers being zero in all these equilibria (see equation (16) and note that the right-hand side is decreasing in t). Secondly, there are other separating equilibria with the same prices as those in the proposition: buyers may use other (mixed or pure) strategies than step functions. Since buyers are indifferent between buying and not buying at the equilibrium prices, the only role that buyers' strategies play is to deter sellers from deviating, and this can be done in more than one way. Thirdly, step-functions of the kind used by the buyers in the above equilibria are optimal decision rules in a variety of decision problems related to the one faced by the buyer in the present model. Suppose, thus, that a decision-maker observes a signal of the form (1), where the true quality q and the noise term ε are statistically independent random variables. Then the signal distribution from a unit stochastically dominates those from units with lower quality. Moreover, if the probability distribution of the noise term has (a version of) the monotone likelihood ratio property (shared by for instance the normal, exponential and Gumbel distributions), then the conditionally expected quality, $\mathbb{E}(q | \tilde{q})$, is an increasing function of the signal \tilde{q} (see e.g. Mattsson *et al.* (2004) for an investigation of a general class of such decision problems). In such cases it is optimal to use decision rules of the threshold form. We finally examine the comparative statics properties of these separating equilibria with respect to the precision α of the buyer's signal. By (14), the conditional probability ρ_H for a high-quality unit to be sold in equilibrium is bounded from above by

$$\bar{\rho} = 1 - F \left[F^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right) - \frac{\alpha (q_H - q_L)}{1 - \alpha} \right]. \quad (17)$$

This upper bound $\bar{\rho}$ deters low-quality sellers from deviating to the high price: the probability of selling low quality at the high price is kept sufficiently low to make such a deviation non-profitable. Note that $\bar{\rho}$ increases with α , from a positive value, $(v_L - w_L) / (v_H - w_L)$, when $\alpha = 0$ towards 1 as $\alpha \rightarrow 1$. In other words, the more precise the signal is, the higher is the upper bound on the equilibrium trading probability for high-quality units, and hence also on the realized gains of trade. As the signal precision approaches zero, the upper bound falls to a positive value, allowing for some trade in high-quality units even when buyers have no information at all about the quality of the unit at hand. By contrast, as the signal precision approaches one, this upper bound tends to 1, the equilibrium trading probability under symmetric information (see section 3.2). By contrast, the lower bound for the conditional equilibrium probability of trade, for a high-quality unit, is independent of buyer information $\alpha < 1$:

$$\underline{\rho} = \max \left\{ 0, \frac{v_L - w_H}{v_H - w_H} \right\}. \quad (18)$$

This lower bound keeps sellers of high-quality units from deviating to the low price, at which they can make a sure profit of $v_L - w_H$. We note that $\underline{\rho}$ is positive if and only if

$w_H < v_L$. In particular, even as we approach the case of symmetric information, that is, as $\alpha \rightarrow 1$, the lower bound stays put at its constant value below one, while the upper bound was seen to approach one. In this sense, the classical result of full realization of gains of trade under symmetric information is non-robust: the slightest buyer uncertainty about quality allows for a whole set of equilibria with partial adverse selection, with accompanying welfare losses. At the other end of the information spectrum, that is when $\alpha = 0$, we note yet another difference from the standard analysis. For while under one-sided asymmetric information high-quality units have until recently been supposed not to be traded when $w_H < \bar{v}$, they are here traded with positive probability, bounded from below by a positive number when $w_H < v_L$. Ellingsen (1997) establishes the existence of separating equilibria with positive probability of trade in high-quality items in a related model (see section 7.1 for a more exact comparison). As for equilibrium gains of trade, they are all going to the seller: $\mathbb{E}[u] = 0$ in all separating equilibria. As for equilibrium gains of trade and profits, we note that

$$v_L - \lambda w_L - (1 - \lambda) w_H \leq W^* = \mathbb{E}[\pi] = \lambda(v_L - w_L) + (1 - \lambda)(v_H - w_H) \bar{\rho} \quad (19)$$

The dependence on the upper bound $\bar{\rho}$ on buyer information α is illustrated in Figure 1 below, drawn for $w_L = 0$, $q_L = 1$, $q_H = 3$, $\lambda = 0.25$ and $\varepsilon \sim N(2.5, 1)$.¹² We see that $\bar{\rho}$ rises from $(1 - 0) / (3 - 0) = 1/3$ towards 1, and that $\bar{\rho}$ is close to 1 for a wide range of precision levels α . By contrast, $\underline{\rho}$ is a constant number between zero (if $w_H \geq 1$) and $1/3$ (if $w_H = w_L = 0$). In particular, the lower bound is zero when $w_H = 2$, that is, when the gains of trade are equal for low and high quality units, which is one of the assumptions of Ellingsen (1997).

¹²We then have $\mathbb{E}[\varepsilon] = \bar{q} = 2.5$ and

$$F^{-1}\left(\frac{v_H - v_L}{v_H - w_L}\right) = F^{-1}(2/3) \approx 2.93.$$




Figure 1: The range of trading probabilities for high quality.

5 Pooling equilibria

We saw in section 3.2 that pooling equilibria do not exist in the boundary case $\alpha = 1$ of symmetric information. We here focus on the case $\alpha < 1$. By Remark 3.2, it suffices to consider prices p in the half-open interval $[v_L, v_H)$. First, consider the decision problem faced by a buyer who receives a quality signal $\tilde{q} \in \mathbb{R}$. The conditionally expected value to the buyer of the unit at hand, given the signal, is

$$\mathbb{E}[v \mid \tilde{q}] = \mu_L(\tilde{q})v_L + [1 - \mu_L(\tilde{q})]v_H, \quad (20)$$

where $\mu_L(\tilde{q})$ is the conditional probability that the item is of low quality, defined in equation (5).¹³ For any price $p \in \mathbb{R}$, let $B(p) \subset \mathbb{R}$ denote the subset of signals \tilde{q} at which the conditionally expected quality is not less than p :

$$B(p) = \{\tilde{q} \in \mathbb{R} : \mathbb{E}[v \mid \tilde{q}] \geq p\}. \quad (21)$$

If $p_L = p_H = p$, then consistency [B3] and sequential rationality [S] together dictate it is optimal for the buyer to make a purchase if and only if his signal \tilde{q} belongs to the set $B(p)$.¹⁴ The following proposition characterizes the set of prices $p \in [v_L, v_H)$ that can be supported in pooling equilibria.

¹³Recall that if $\alpha = 0$, then $\mu_L(\tilde{q}) = \lambda$ for all \tilde{q} .

¹⁴Indeed, it is the buyer's unique best reply to buy if $\mathbb{E}[v \mid \tilde{q}] > p$.

Proposition 5.1 *Suppose $\alpha < 1$. A price $p \in [v_L, v_H)$ is a pooling equilibrium price if and only if (22) and (23) hold.*

$$(p - w_L) \Pr [\tilde{q} \in B(p) \mid \theta = L] \geq v_L - w_L \quad (22)$$

$$(p - w_H) \Pr [\tilde{q} \in B(p) \mid \theta = H] \geq \max \{0, v_L - w_H\} \quad (23)$$

Proof: To prove necessity, let (μ, s, b) be a pooling equilibrium with price $p \in [v_L, v_H)$. By sequential rationality, [S], $b(p, \tilde{q}) = 0$ if $\tilde{q} \notin B(p)$. Hence, a low-quality seller's expected profit, π_L , satisfies

$$\pi_L \leq (p - w_L) \Pr [\tilde{q} \in B(p) \mid \theta = L]. \quad (24)$$

Moreover, such a seller can guarantee himself a profit arbitrarily close to $v_L - w_L$ by choosing prices $p' < v_L$. Hence,

$$v_L - w_L \leq \pi_L \leq (p - w_L) \Pr [\tilde{q} \in B(p) \mid \theta = L]$$

proving (22). Likewise, a high-quality seller's expected profit, π_H , satisfies

$$\pi_H \leq (p - w_H) \Pr [\tilde{q} \in B(p) \mid \theta = H]. \quad (25)$$

Moreover, such a seller can guarantee himself zero profit by setting $p' > v_H$ and a profit arbitrarily close to $v_L - w_H$ by choosing price $p' < v_L$. Hence,

$$\max \{0, v_L - w_H\} \leq \pi_H \leq (p - w_H) \Pr [\tilde{q} \in B(p) \mid \theta = H]$$

proving (23). To prove sufficiency, define (μ, s, b) by setting $p_L = p_H = p \in [v_L, v_H)$,

$$b(p', \tilde{q}) = \begin{cases} 1 & \text{if } p' < v_L \text{ or if } p' = p \text{ and } \tilde{q} \in B(p), \\ 0 & \text{otherwise,} \end{cases}$$

and let μ satisfy [B1] and [B3] ([B2] has no bite in pooling equilibria). At information sets (p', \tilde{q}) of the buyer off the path of (s, b) , that is, where $p' \neq p$, let μ assign probability 1 to the node (L, ε_L, p') : the buyer then believes the item is of low quality. This strategy profile is a perfect Bayesian equilibrium. First, the belief requirements are fulfilled by construction. Second, when the price is p , the buyer plays a best reply, since by definition of $B(p)$ she buys whenever this gives nonnegative expected utility. At prices $p' \neq p$, she believes the item is of low quality, and her strategy, to buy iff the price p' is below v_L , is optimal. Conditions (22) and (23) imply that the seller has no profitable unilateral deviations: a price $p' \geq v_L$, $p' \neq p$, results in zero trade, and inequalities (24) and (25) are met with equality under (μ, s, b) . **End of proof.**

As is shown in the proof, the first (second) condition in the proposition guarantees that sellers with low (high) quality units do not have an incentive to deviate. We note that condition (23) fails if $p < w_H$. Hence, not surprisingly, such prices can not be supported in pooling equilibria — sellers with high-quality units are then better off keeping their unit. By contrast, the same condition, (23), is trivially met if $v_L \leq w_H \leq p$. In such cases,

sellers with high-quality units make a nonnegative profit at the going price p , and cannot make positive profits by deviating to prices $p \leq v_L$ — prices at which it is rational for buyers to buy irrespective of their signal. An immediate consequence of this proposition is that if $w_H \leq v_L$, then $p = v_L$ is a pooling equilibrium price. For in this case sellers are willing to part with high-quality units, and it is rational for buyers to buy irrespective of the signal (the quality cannot be lower than the going price). Formally: $B(v_L) = \mathbb{R}$, and the probabilities on the left-hand sides of (22) and (23) are both one, and thus both conditions are met for $p = v_L$.¹⁵ The reason why this low price is a pooling equilibrium price is simple: downwards price deviations do not pay off, and buyers do not trust sellers who post prices above the going price.

For the boundary case $\alpha = 0$ proposition 5.1 implies that pooling equilibria exist if and only if $w_H \leq \bar{v}$. To see this, first note that, by equation (5), we then have $\mu_L(\tilde{q}) = \lambda$ for all signals \tilde{q} . Consequently, $\mathbb{E}[v | \tilde{q}] = \bar{q}$ for all \tilde{q} . In other words, the signal being completely uninformative, the buyer expects the good to be of average value. The buyer's acceptance set thus is the whole real line, $B(p) = \mathbb{R}$ for all prices $p \leq \bar{v}$, while for higher prices it is the empty set. It is easily verified that the conditions in the proposition are met if and only if

$$\max\{v_L, w_H\} \leq p \leq \bar{v}. \quad (26)$$

Hence, the set of pooling equilibrium prices is non-empty if and only if $w_H \leq \bar{v}$. Moreover, when $\alpha = 0$, then all pooling equilibria are efficient — they all realize the full potential gains of trade: $\rho_L = \rho_H = 1$ and $W^* = \hat{W}$. The only difference between the alternative pooling equilibria in this special case ($\alpha = 0$) is the share of the gains of trade that befalls the buyer: consumer surplus being higher the lower the pooling price. Is the above result for pooling equilibrium prices p at $\alpha = 0$ robust with respect to α ? Assume, first that $\alpha > 0$, $\max\{v_L, w_H\} < \bar{v}$, and consider any price $p < \bar{v}$ satisfying (26). As $\alpha \rightarrow 0$, $\mu_L(\tilde{q}) \rightarrow \lambda$ for all quality signals \tilde{q} , by (4) and (5). Hence, $\mathbb{E}[v | \tilde{q}] \rightarrow \bar{q}$ as $\alpha \rightarrow 0$. Consequently, p is a pooling equilibrium price for sufficiently small α , since $\mathbb{E}[v | \tilde{q}] - p > 0$ for α close to 0, by continuity. Moreover, ρ_L and ρ_H converge to 1. Hence, the above result for prices p satisfying $\max\{v_L, w_H\} \leq p < \bar{v}$ are robust with respect to α . As will be shown below, however, the result is discontinuous with respect to α when $p = \bar{v}$. In order to study intermediate degrees of signal precision, $\alpha \in (0, 1)$, we make the simplifying assumption that the error term is normally distributed.

5.1 Normally distributed noise

In the special case when the perturbation ε of the consumer's quality signal is normally distributed, the necessary and sufficient conditions (22) and (23) can be written explicitly in terms of the primitives of the model. For this purpose, suppose that $\alpha \in (0, 1)$ and let

$$\varphi(p) = \frac{\alpha(q_H - q_L)}{2(1 - \alpha)\sigma} + \frac{(1 - \alpha)\sigma}{\alpha(q_H - q_L)} \ln \left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)} \right]. \quad (27)$$

¹⁵This is true even when $\alpha = 0$.

Clearly $\varphi(p) \in \mathbb{R}$ for all $p \in (v_L, v_H)$, and φ is increasing in p on the interval (v_L, v_H) , with $\lim_{p \downarrow v_L} \varphi(p) = -\infty$, $\lim_{p \uparrow v_H} \varphi(p) = +\infty$, and

$$\varphi(\bar{v}) = \frac{\alpha(q_H - q_L)}{2(1 - \alpha)\sigma}. \quad (28)$$

Proposition 5.2 *Suppose that $\alpha \in (0, 1)$ and $\varepsilon \sim N(\bar{q}, \sigma^2)$. A price $p \in \mathbb{R}$ is a pooling equilibrium price if and only if $\max\{v_L, w_H\} \leq p < v_H$ and*

$$\int_{\varphi(p)}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \geq \frac{v_L - w_L}{p - w_L} \quad (29)$$

Proof: Condition (23) is violated if $p < w_H$, so no such price is a pooling equilibrium price. We observed earlier that $p = v_L$ is a pooling equilibrium price for all $\alpha \in (0, 1)$ if $v_L \geq w_H$. To see that $p = v_L$ satisfies inequality (29), note that its left-hand side then equals 1 and so does the right-hand side.

It thus only remains to study prices $p \in (v_L, v_H) \cap [w_H, +\infty)$. Using (5), we note that $\tilde{q} \in B(p)$ iff

$$\lambda(q_L - p)f(\varepsilon_L) + (1 - \lambda)(q_H - p)f(\varepsilon_H) \geq 0,$$

where ε_θ is defined in (4) and f is the density of $N(\bar{q}, \sigma^2)$. Hence, $\tilde{q} \in B(p)$ if and only if

$$f\left(\frac{\tilde{q} - \alpha q_H}{1 - \alpha}\right) \geq \frac{\lambda(p - q_L)}{(1 - \lambda)(q_H - p)} f\left(\frac{\tilde{q} - \alpha q_L}{1 - \alpha}\right),$$

or, equivalently,

$$\exp\left[-\left(\frac{\tilde{q} - \alpha q_H}{1 - \alpha} - \bar{q}\right)^2 / (2\sigma^2)\right] \geq \exp[\beta(p)] \exp\left[-\left(\frac{\tilde{q} - \alpha q_L}{1 - \alpha} - \bar{q}\right)^2 / (2\sigma^2)\right],$$

where

$$\beta(p) = \ln\left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)}\right].$$

The exponential function being strictly increasing, the inequality above holds iff

$$\left(\frac{\tilde{q} - \alpha q_H}{1 - \alpha} - \bar{q}\right)^2 \leq \left(\frac{\tilde{q} - \alpha q_L}{1 - \alpha} - \bar{q}\right)^2 - 2\sigma^2\beta(p).$$

Using (1), we obtain

$$\begin{aligned} \rho_L &= \Pr[\tilde{q} \in B(p) \mid \theta = L] \\ &= \Pr\left[\left(\varepsilon - \frac{\alpha(q_H - q_L)}{1 - \alpha} - \bar{q}\right)^2 \leq (\varepsilon - \bar{q})^2 - 2\sigma^2\beta(p)\right] \\ &= \Pr\left[\varepsilon - \bar{q} \geq \frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2\beta(p)(1 - \alpha)}{\alpha(q_H - q_L)}\right] \end{aligned}$$

and

$$\begin{aligned}
\rho_H &= \Pr[\tilde{q} \in B(p) \mid \theta = H] \\
&= \Pr\left[(\varepsilon - \bar{q})^2 \leq \left(\varepsilon + \frac{\alpha(q_H - q_L)}{1 - \alpha} - \bar{q}\right)^2 - 2\sigma^2\beta(p)\right] \\
&= \Pr\left[\varepsilon - \bar{q} \geq -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)\beta(p)}{\alpha(q_H - q_L)}\right]. \tag{30}
\end{aligned}$$

As expected, $\rho_H > \rho_L$. For any $p \in (v_L, v_H) \cap [w_H, +\infty)$, inequality (23) is met whenever inequality (22) is met. Thus, a necessary and sufficient condition for such a price to be a pooling equilibrium price is

$$\Pr\left[\varepsilon - \bar{q} \geq \frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)\beta(p)}{\alpha(q_H - q_L)}\right] \geq \frac{v_L - w_L}{p - w_L},$$

or, equivalently, (29). **End of proof.**

We first illustrate this result by way of plotting the set of pooling equilibrium prices associated with different values of $\alpha \in (0, 1)$, using the same parameters as in figure 1 ($w_L = 0$, $q_L = 1$, $q_H = 3$, $\lambda = 0.25$ and $\sigma = 1$). For $w_H \leq v_L$, the graph of this equilibrium correspondence is the area between (and on) the curve and the horizontal line $p = v_L = 1$. For $w_H > v_L$ the equilibrium graph is the part of that area where $p \geq w_H$. We note that for $\alpha > 0.5$ and $w_H \leq v_L$, only v_L and prices just above that level are pooling prices. Higher pooling equilibrium prices are hard to maintain when buyers receive relatively precise quality signals. We also note that the curve's maximum exceeds \bar{v} , the buyer's willingness to pay for average quality, and that this maximum is obtained for a positive α (approximately 0.1), while the highest pooling price at $\alpha = 0$ is \bar{v} . Hence, in our model pooling equilibrium is possible even when $w_H > \bar{v}$, but only if buyers' signal carries some (and not too much) information about the true quality (α is positive but small). The reason why such equilibria require buyers' quality signal to have some precision is that otherwise they would expect all units to be of average quality, and thus would be willing to pay only \bar{v} .




Figure 2: Combinations of price and signal precision that are consistent with pooling equilibrium.

Having thus studied a numerical example, we now turn back to the general case. First note that the set of pooling equilibrium prices, if non-empty, is an interval, and that this interval has $\max\{v_L, w_H\}$ as its left end point, while inspection of equation (27) shows that its right end point is decreasing in λ . Thus, the larger the share of “lemons,” the lower is the highest pooling price. Secondly, the upper bound on the pooling price also depends in a systematic way on α , the precision of buyers’ quality perception. The following observation follows from (27) and (29): The set of pooling equilibrium prices approaches the set

$$P_0 = \begin{cases} [\max\{v_L, w_H\}, \bar{v}] & \text{if } \frac{v_L - w_L}{\bar{v} - w_L} < \frac{1}{2} \\ [\max\{v_L, w_H\}, \bar{v}) & \text{if } \frac{v_L - w_L}{\bar{v} - w_L} \geq \frac{1}{2} \end{cases} \quad (31)$$

as $\alpha \rightarrow 0$. To see this, note that for $p = \bar{v}$, the left-hand side of the inequality in (29) approaches the value $1/2$ from above. The intuition for this is that consumers are willing to buy at prices $p < \bar{v}$ when the precision of their signal is close to zero, since in the limit their posterior beliefs are identical with their prior, that is, they expect average quality. However, at $p = \bar{v}$, they are indifferent between buying and not buying when $\alpha = 0$, while even an imprecise signal carries some information, and thus, due to the symmetry of the normal distribution, they will optimally buy only with probability $1/2$ in the limit as $\alpha \rightarrow 0$.

Conversely, let buyers’ signal precision go to one. First, for all $p \in (v_L, v_H)$, $\varphi(p)$ is continuous in α , and tends to $+\infty$ as α approaches 1. Hence, for α sufficiently close to 1, any such price p ceases to be a pooling price: with P_1 denoting the set of pooling equilibrium prices in the limit as $\alpha \rightarrow 1$, we have

$$P_1 \cap (v_L, v_H) = \emptyset. \quad (32)$$

Hence, if $w_H > v_L$, then $P_1 = \emptyset$. We have already shown that $p = v_L$ is a pooling equilibrium price for all $\alpha \in (0, 1)$ if $v_L \geq w_L$, and hence we also have $v_L \in P_1$ in this case. How do trading probabilities and expected gains of trade depend on buyers' signal precision α in pooling equilibria? The following equations follow immediately from the above proof:

$$\rho_L = \int_{\varphi(p)}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \quad (33)$$

and

$$\rho_H = \int_{\psi(p)}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx, \quad (34)$$

where

$$\psi(p) = -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)\sigma} + \frac{(1 - \alpha)\sigma}{\alpha(q_H - q_L)} \ln \left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)} \right]. \quad (35)$$

As expected, $\psi(p) < \varphi(p)$, and hence $\rho_H > \rho_L$. Moreover, the logarithmic factor that appears in the defining equations for both $\varphi(p)$ and $\psi(p)$ is negative for prices below \bar{v} , positive for prices above \bar{v} , and equal to zero for $p = \bar{v}$. This allows us to make some qualitative statements about the effect of changes in α on the integration bounds $\varphi(p)$ and $\psi(p)$ and hence on the conditional trading probabilities ρ_L and ρ_H . Suppose, thus, that $p \in (v_L, v_H)$ is a pooling equilibrium price.

1. If $p < \bar{v}$, the logarithm is negative, so $\varphi(p)$ is increasing in α , implying that ρ_L decreases as α increases. The probability that a buyer will accept a low-quality unit falls as the precision of his quality signal increases.
2. If $p = \bar{v}$, the logarithm is zero, so $\varphi(p)$ increases in α and $\psi(p)$ decreases in α , implying that ρ_L decreases and ρ_H increases as α increases. The probability that a buyer will accept a low-quality (high-quality) unit falls (increases) as the precision of his quality signal increases.
3. If $p > \bar{v}$, the logarithm is positive, so $\psi(p)$ is a decreasing function of α , implying that ρ_H increases as α increases. The probability that a buyer will accept a high-quality unit increases as the precision of his quality signal increases.

Equilibrium gains of trade are

$$W^* = \lambda(v_L - w_L) \int_{\varphi(p)}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx + (1 - \lambda)(v_H - w_H) \int_{\psi(p)}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx. \quad (36)$$

We consider the question whether the conditional trading probability for high quality units can be higher in a pooling equilibrium than in all separating equilibria. Formally, we ask whether $\rho_H > \bar{\rho}$ is possible. Note that

$$\begin{aligned} \rho_H &= \Pr \left[\varepsilon - \bar{q} \geq -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)}{\alpha(q_H - q_L)} \beta(p) \right] \\ &= 1 - F \left[\bar{q} - \frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)}{\alpha(q_H - q_L)} \beta(p) \right] \end{aligned} \quad (37)$$

Hence, $\rho_H \geq \bar{p}$ if and only if

$$\frac{v_H - v_L}{v_H - w_L} \leq \int_{-\infty}^{\varphi(p)} \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \leq \frac{p - v_L}{p - w_L}, \quad (38)$$

where the last inequality is necessary and sufficient for p to be a pooling equilibrium price. But the left-most term is greater than the right-most term for all $p < v_H$, so this is not possible. Conclude that the conditional trading probability for high quality units cannot be higher in a pooling equilibrium than in all separating equilibria. Since the conditional trading probability in the low quality is one in all separating equilibria, and less than one in all pooling equilibria, this implies the following

Remark 5.3 *The expected profit in every pooling equilibrium falls short of the maximal profit in separating equilibria.*

6 A behavioral refinement: naïve buyers

The (weak) Bayesian perfect equilibrium concept imposes no restrictions on beliefs off the equilibrium path. It seems natural to impose such restrictions, but not necessarily the rationalistic restrictions that have been used in the refinement literature in the past, but rather restrictions that seem behaviorally plausible. We here consider one such “behavioral” restriction, namely, that buyers form their posteriors for the quality of the item at hand on the basis of the prior λ , also at non-equilibrium prices. Let us thus call a buyer *naïve* if she also off the equilibrium path uses Bayesian updating, given the quality signal, to determine the expected quality of the item for sale — with no consideration of the possible motives a seller might have to make such a deviation. Formally, this amounts to adding the following additional requirement on the belief system:

[B4] At each buyer information set (p, \tilde{q}) with $p \notin \{p_L, p_H\}$, μ assigns the probability $\mu_L(\tilde{q})$, defined in equation (5), to the node (L, ε_L, p) .

We call a tuple $(\mu, (p_L, p_H), b)$ satisfying [B1]-[B4] and [S] a *naïve perfect Bayesian equilibrium*. Let us investigate the existence of such naïve equilibria. First, note that conditions [B1]-[B4] completely determine the belief system μ : beliefs at information sets of the seller are determined by [B1], beliefs at information sets of the buyer on the equilibrium path are determined by [B2] (if $p_L \neq p_H$) and [B3] (if $p_L = p_H$), and beliefs off the equilibrium path are determined by [B4]. Secondly, since the conditionally expected value (upon seeing the quality signal) is a convex combination of v_L and v_H , conditions (8) - (9) and consequently Lemma 3.1 remain true: the same range of price pairs as before can be ruled out *a priori*.

Thirdly, by sequential rationality [S], a buyer necessarily buys the item at price p if $\mathbb{E}[v \mid \tilde{q}] > p$ and is indifferent between buying and not buying if $\mathbb{E}[v \mid \tilde{q}] = p$. Hence,

the expected payoff to a seller with a unit of quality $\theta \in \{L, H\}$ who deviates to a price $p \notin \{p_L, p_H\}$ is at least

$$(p - w_\theta) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta]. \quad (39)$$

A necessary condition for $(\mu, (p_L, p_H), b)$ to be a naïve equilibrium is thus that the inequality

$$\rho_\theta(p_\theta - w_\theta) \geq (p - w_\theta) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta] \quad (40)$$

holds for all prices $p \in \mathbb{R}$, signals $\tilde{q} \in \mathbb{R}$, and types $\theta \in \{L, H\}$, where ρ_θ is the equilibrium trading probability for type θ . Likewise, a sufficient condition for (40) to hold is that

$$\rho_\theta(p_\theta - w_\theta) \geq (p - w_\theta) \Pr[\mathbb{E}[v \mid \tilde{q}] \geq p \mid \theta] \quad (41)$$

for all prices $p \in \mathbb{R}$, signals $\tilde{q} \in \mathbb{R}$, and types $\theta \in \{L, H\}$, where ρ_θ is the equilibrium trading probability for type θ .

Under reasonable regularity conditions, these two conditions are identical. To see this, first recall from (20) and (5) that

$$\begin{aligned} \mathbb{E}[v \mid \tilde{q}] &= v_H - (v_H - v_L) \mu_L(\tilde{q}) \\ &= v_H - (v_H - v_L) \frac{\lambda}{\lambda + (1 - \lambda) f(\varepsilon_H) / f(\varepsilon_L)}, \end{aligned} \quad (42)$$

where, by (4),

$$f(\varepsilon_H) / f(\varepsilon_L) = f\left(\frac{\tilde{q} - \alpha q_H}{1 - \alpha}\right) / f\left(\frac{\tilde{q} - \alpha q_L}{1 - \alpha}\right). \quad (43)$$

In other words, the buyer's conditional expectation for the quality of the unit at hand, given the buyer's quality signal, is a strictly increasing function of the associated likelihood ratio for high quality. Hence, if the error distribution has the monotone likelihood ratio property (MLRP), then this conditional expectation is a non-decreasing function of the signal \tilde{q} .

Formally, let

[MLRP] The ratio $f(x - a) / f(x)$ is strictly increasing in x , for all $a > 0$.¹⁶

Under this condition, $\mathbb{E}[v \mid \tilde{q}]$ is a strictly increasing function of \tilde{q} , mapping the real line to the interval (q_L, q_H) . Hence, since by hypothesis \tilde{q} has a continuous probability distribution (no atoms), also the random variable $\mathbb{E}[v \mid \tilde{q}]$ has a continuous probability distribution. In particular, the probability that the buyer will be indifferent then has zero probability: $\Pr[\mathbb{E}[v \mid \tilde{q}] = p] = 0$ for all prices $p \in \mathbb{R}$. Consequently, conditions (40) and (41) are identical.

By the monotone convergence theorem, the limit of $\mathbb{E}[v \mid \tilde{q}]$ as $\tilde{q} \rightarrow +\infty$ exists. Let

$$\hat{v} = \lim_{\tilde{q} \rightarrow +\infty} \mathbb{E}[v \mid \tilde{q}] \quad (44)$$

¹⁶See Mattsson et al. (2004) for a systematic investigation of this version of the MLRP in the context of discrete choice.

A critical property of the error distribution for the existence of naive perfect Bayesian equilibria, is whether or not $\hat{v} = v_H$. For in this case the probability is positive that the buyer will buy at any given price $p < v_H$, namely, if the quality signal is sufficiently strong. If instead $\hat{v} < v_H$, then the buyer will never buy at a price $p \in [\hat{v}, v_H]$; no signal is strong enough to make the conditionally expected value as high as such a price.

We illustrate these two possibilities by means of two examples.¹⁷ First, let ε be normally distributed, $N(\bar{q}, \sigma^2)$. It is easily verified that this distribution has the MLRP, and, moreover, that the likelihood ratio in equation (43) tends to plus infinity as $\tilde{q} \rightarrow +\infty$. Hence, in this case $\hat{v} = v_H$. Secondly, let ε be Gumbel (or doubly exponentially) distributed:

$$F(\varepsilon) = \exp[-e^{-\tau(\varepsilon-\eta)}] \quad (45)$$

for parameters $\eta \in \mathbb{R}$ and $\tau > 0$. Then $\mathbb{E}[\varepsilon] = \eta + \gamma/\tau$, where $\gamma \approx 0.577$ is Euler's constant (see, for instance, Ben-Akiva and Lerman, 1985, p. 104). It is not difficult to verify that also this distribution has the MLRP. Moreover, the likelihood ratio in equation (43) tends to $\exp[\frac{\tau\alpha}{1-\alpha}(v_H - v_L)]$ as $\tilde{q} \rightarrow +\infty$. Hence, in this case

$$\hat{v} = v_H - \frac{\lambda(v_H - v_L)}{\lambda + (1 - \lambda) \exp[\frac{\tau\alpha}{1-\alpha}(v_H - v_L)]} < v_H. \quad (46)$$

We finally note that $\hat{v} \rightarrow v_H$ as the signal precision α tends to 1, and $\hat{v} \rightarrow \bar{v}$ as the signal precision α tends to 0. Figure 3 below shows the graph of $\tilde{q} \mapsto \mathbb{E}[v | \tilde{q}]$, for $v_L = 1$, $v_H = 3$, $\lambda = 0.25$, $\tau = 1$, and (from the left to the right) $\alpha = 0.1, 0.25, 0.5$ and 0.75 , respectively.




Figure 3: The conditionally expected value to the buyer as a function of the buyer's quality signal, for different levels of signal precision.

¹⁷See Mattsson *et al.* (2004) for more on the MLRP and similar calculations.

6.1 Total adverse selection

In this section we derive necessary and sufficient conditions for the existence of naïve equilibria in which high-quality sellers outprice themselves. Such equilibria do not exist if noise is normally distributed. However, they do exist if noise is Gumbel distributed and the buyer’s signal precision α is low.

Proposition 6.1 *Suppose $\alpha < 1$ and that the error distribution satisfies [MLRP]. There is a naïve perfect Bayesian equilibrium with prices $p_L = v_L$ and $p_H > v_H$ if and only if the following two conditions hold:*

$$v_L \leq w_H \tag{47}$$

$$\Pr[\mathbb{E}[v \mid \tilde{q}] > w_H \mid \theta = H] = 0 \tag{48}$$

In all such equilibria $\rho_L = 1$ and $\rho_H = 0$.

Proof: Necessity of (47) follows as in Proposition 4.1. To prove necessity of (48), suppose $(\mu, (p_L, p_H), b)$ is a naïve equilibrium with $p_L = v_L$ and $p_H > v_H$. Condition (9) implies that $\rho_H = 0$. Sellers with high-quality units thus obtain zero profit. Hence, for $p_H > v_H$ to be an equilibrium price, substitution in (40) yields the necessary condition

$$0 \geq (p - w_H) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] \quad \forall p \in \mathbb{R}. \tag{49}$$

The inequality is trivially satisfied when $p \leq w_H$, in which case $p - w_H \leq 0$, and when $p > v_H$, in which case $\Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid H] = 0$. For prices $p \in (w_H, v_H]$, the inequality in (49) is met if and only if

$$\Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] = 0.$$

The probability on the left-hand side of this equation is non-increasing in p , so condition (49) is equivalent with (48).

As for sufficiency, assume (47) and (48) hold. Let $(\mu, (p_L, p_H), b)$ be such that $p_L = v_L$, $p_H > v_H$, let μ be the unique belief system satisfying [B1]-[B4] and let $b(p, \tilde{q}) = 1$ if and only if $\mathbb{E}[v \mid \tilde{q}] > p$. It is easy to verify that $(\mu, (p_L, p_H), b)$ is a naïve perfect Bayesian equilibrium.

The proof of the last claim in the proposition is identical to that in Proposition 4.1.

End of proof.

Although demanding, condition (48) does not rule out the existence of naïve equilibria with total adverse selection. This depends on how “thin” the tails of the error distribution is. More precisely, if noise is normally distributed, then naïve equilibria with total adverse selection do not exist, while they do exist if noise is Gumbel distributed.

Proposition 6.2 *There are no naïve equilibria with prices $p_L = v_L$ and $p_H > v_H$ if $\alpha \in (0, 1)$ and $\varepsilon \sim N(\bar{q}, \sigma^2)$.*

Proof: Just as in our characterization of pooling equilibria under normal distributed noise, one obtains

$$\begin{aligned} & \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] = \\ & = \Pr\left[\varepsilon - \bar{q} > -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)\beta(w_H)}{\alpha(q_H - q_L)}\right] > 0, \end{aligned}$$

violating (48). **End of proof.**

Proposition 6.3 *Suppose that ε has a Gumbel distribution with expectation \bar{q} .¹⁸ There is no naïve perfect Bayesian equilibrium with prices $p_L = v_L$ and $p_H > v_H$ if $\bar{v} \geq w_H$. If $\bar{v} < w_H$, such equilibria exist for $\alpha \in (0, 1)$ sufficiently close to 0.*

Proof: If $w_H < v_L$, no such naïve equilibria exist by condition (47). Hence, we assume in the remainder of the proof that $v_L \leq w_H$. Condition (48) holds if and only if $\mathbb{E}[v \mid \tilde{q}] \leq w_H$ for all \tilde{q} , i.e., if and only if $\hat{v} = \lim_{\tilde{q} \rightarrow +\infty} \mathbb{E}[v \mid \tilde{q}] \leq w_H$. By (46), \hat{v} is strictly increasing in $\alpha \in (0, 1)$, so this is possible if and only if

$$\bar{v} = \lim_{\alpha \rightarrow 0} \hat{v} < w_H.$$

To summarize: condition (48) will be violated if $\bar{v} \geq w_H$, excluding the existence of naïve equilibria with total adverse selection. On the other hand, if $\bar{v} < w_H$, then the fact that \hat{v} in (46) is strictly increasing in α guarantees the existence of an interval $(0, \bar{\alpha})$ for which such equilibria do exist. **End of proof.**

6.2 Pooling equilibria

By (40), a pooling equilibrium price $p^* \in [v_L, v_H)$ in a naïve equilibrium should simultaneously solve

$$\max_p (p - w_L) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = L] \quad (50)$$

and

$$\max_p (p - w_H) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H]. \quad (51)$$

Since the probabilities of trade are independent of sellers' reservation prices w_L and w_H , this generically rules out all interior prices $p^* \in (v_L, v_H)$. Indeed, assuming the probabilities in (50) and (51) to be differentiable functions of prices $p \in (v_L, v_H)$, as is the case if noise is normally distributed, the first order conditions for such an internal maximum are

$$\forall \theta \in \{L, H\} : w_\theta = p + \frac{\Pr[\tilde{q} \in B(p) \mid \theta]}{\frac{\partial}{\partial p} \Pr[\tilde{q} \in B(p) \mid \theta]}. \quad (52)$$

Since the right-hand side is independent of w_θ , this generically excludes such internal solutions.

¹⁸Since $\mathbb{E}[\varepsilon] = \eta + \gamma/\tau$, there are infinitely many such distributions.

This leaves us with only one candidate for a pooling equilibrium price: $p^* = v_L$. Of course, this requires that the high-quality seller is willing to sell his product at price v_L , i.e., $w_H \leq v_L$. Yet such equilibria are ruled out if information precision is sufficiently large: in that case, a high-quality seller can rely on the information signal to reveal the high quality, so that he can benefit from setting a slightly higher price. But also if information precision is sufficiently small, such equilibria are ruled out: in that case, buyers' expected quality is close to average market quality \bar{v} , so that they are with large probability willing to buy at all prices $p \in (v_L, \bar{v})$, so that a price slightly above v_L increases expected profits. Formally:

Proposition 6.4 *Suppose that $\varepsilon \sim N(\bar{q}, \sigma^2)$ and $w_H \leq v_L$. If $\alpha \in (0, 1)$ is sufficiently close to one or sufficiently close to zero, there is no naïve equilibrium with pooling price $p^* = v_L$.*

Proof: At price $p^* = v_L$, the profit to a high-quality seller equals $v_L - w_H$. Fix $p \in (v_L, v_H)$. Just as in our characterization of pooling equilibria under normal distributed noise, one obtains

$$\begin{aligned} \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] &= \\ &= \Pr\left[\varepsilon - \bar{q} > -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)\beta(p)}{\alpha(q_H - q_L)}\right] \rightarrow 1 \end{aligned}$$

as $\alpha \rightarrow 1$. But then

$$(p - w_H) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] > v_L - w_H$$

for α sufficiently large: the high-quality seller's expected profit from deviating to p is higher than the profit at the pooling price $p^* = v_L$. Similarly, as $\alpha \rightarrow 0$, $\mu_L(\tilde{q}) \rightarrow \lambda$ for all signals \tilde{q} , so that for every price $p \in (v_L, \bar{v})$ and α sufficiently small:

$$(p - v_H) \Pr[\mathbb{E}[v \mid \tilde{q}] > p \mid \theta = H] > v_L - w_H.$$

In other words, the high-quality seller's expected profit from deviating to a price p below average quality is higher than the profit at the pooling price $p^* = v_L$. **End of proof.**

7 Comparison with established models

In order to high-light the difference between the present model and established models of symmetric and one-sided asymmetric information, we briefly review the results for the following cases: (i) no signal and monopolistic sellers, (ii) no signal and perfect competition, (iii) perfect signal and monopolistic sellers, (iv) perfect signal and perfect competition. Case (i) is closely related to Ellingsen (1997) and Bester and Ritzberger (2001), and case (ii) is the model in Akerlof (1970), while cases (iii) and (iv) are the classical cases of symmetric information.

7.1 Pure noise signal

Set $\alpha = 0$. In this boundary case, the signal is statistically independent of the current item's quality. This is very close to certain established models of one-sided asymmetric information. Technically, however, there is two-sided asymmetric information even when $\alpha = 0$, in the trivial sense that the seller does not know the buyer's *uninformative* signal. The signal thus provides the buyer with a private randomization device on which he or she can condition the purchasing decision. Hence, the signal effectively allows the buyer to use certain mixed strategies. In this special case, our model resembles those of Ellingsen (1997) and Bester and Ritzberger (2001). The differences are that in Ellingsen's model (a) the set of feasible prices is finite, (b) the seller may randomize over these, (c) the gains of trade are the same for both qualities (in our notation: $v_L - w_L = v_H - w_H$), and (d) he also analyzes the case of many qualities, distributed over a finite quality grid with fixed grid size. Bester and Ritzberger (2001) endow the buyer with the option to inspect the true quality at a cost, an action that the seller is assumed not to observe. Hence, like our model, also their model contains a form of two-sided asymmetric information. The differences are that in that model (a) buyers either have no signal or a perfect signal, (b) it combines one-sided incomplete information (hidden information) with one-sided imperfect information (hidden action). The analysis in sections 4 and 5 led to the following three implications in the special case $\alpha = 0$:¹⁹

1. There exist separating equilibria with total adverse selection if and only if $v_L \leq w_H$. When this inequality is satisfied, all price pairs $p_L = v_L$, $p_H > v_H$ are equilibria, and $\rho_L = 1$ and $\rho_H = 0$.
2. There exist separating equilibria with partial adverse selection, such that $p_L = v_L$, $p_H = v_H$, $\rho_L = 1$ and

$$\max \left\{ 0, \frac{v_L - w_H}{v_H - w_H} \right\} \leq \rho_H \leq \frac{v_L - w_L}{v_H - w_L}.$$

3. Pooling equilibria exist if and only if $w_H \leq \bar{v}$. The set of pooling prices is

$$[\max\{v_L, w_H\}, \bar{v}],$$

and in (pure strategy) equilibrium $\rho_L = \rho_H = 1$.

7.2 Very noisy signal

Having examined the equilibrium outcomes when $\alpha = 0$, now suppose $\alpha > 0$, and let $\alpha \rightarrow 0$. In other words, let the signal be almost statistically independent of the item's quality. We then have seen that:

¹⁹In models where there is no quality signal at all, the natural point of comparison would be to allow buyers strategies that assign to any posted price a probability of buying at that price; technically, a *behavior strategy*. The set of equilibria obtained in this way shares these three properties.

1. The limit set of separating equilibria with total adverse selection is the same as for $\alpha = 0$.
2. The limit set of separating equilibria with partial adverse selection converges to that for $\alpha = 0$.
3. All prices satisfying

$$\max\{v_L, w_H\} \leq p < \bar{v},$$

are pooling equilibrium prices both when $\alpha = 0$, and in the limit as $\alpha \rightarrow 0$. However, at $p = \bar{v}$, there are two potential discontinuities, at least if ε has a unimodal and symmetric probability distribution

Discontinuity 1: If $2v_L - w_L \geq \bar{v}$, then $p = \bar{v}$ is not a pooling equilibrium price for α close to 0.

Discontinuity 2: If $2v_L - w_L < \bar{v}$, and $w_H < \bar{v}$, then $p = \bar{v}$ is a pooling equilibrium price for α close to 0, but $\rho_L, \rho_H \rightarrow 1/2$.

7.3 Perfect signal

Let $\alpha = 1$. In this boundary case, the signal is perfectly informative of product quality. In our model of monopolistic sellers, there is a unique perfect Bayesian equilibrium (see section 3.2), and we have

$$p_L = v_L, p_H = v_H \text{ and } \rho_L = \rho_H = 1.$$

However, the analysis in sections 4 and 5 shows that the set of equilibria is non-robust at $\alpha = 1$.

7.4 Almost perfect signal

For α sufficiently close to 1:

1. There exist separating equilibria with

$$p_L = v_L, p_H > v_H$$

if and only if $v_L \leq w_H$. In these equilibria, $\rho_L = 1$ and $\rho_H = 0$. We thus have *Discontinuity 3:* No such equilibria exist if $\alpha = 1$.

2. There exist separating equilibria with

$$\begin{cases} p_L = v_L, p_H = v_H \\ \rho_L = 1 \\ \max\left\{0, \frac{v_L - w_H}{v_H - w_H}\right\} \leq \rho_H \leq \bar{\rho} \end{cases}$$

where $\bar{\rho} \rightarrow 1$ as $\alpha \rightarrow 1$. *Discontinuity 4:* Equilibria with trade probabilities below 1 are absent when $\alpha = 1$.

3. No pooling equilibrium exists in the limit as $\alpha \rightarrow 1$ if $w_H > v_L$. Hence, in this case there is continuity at $\alpha = 1$. However, if $w_H \leq v_L$, then the set P of pooling equilibria converges to the singleton $\{v_L\}$ as $\alpha \rightarrow 1$. It is easily verified that $\rho_L = \rho_H \rightarrow 1$ in the latter case. *Discontinuity 5*: No pooling equilibria exist if $\alpha = 1$.

8 Randomized buyer strategies

We did not need mixed or (randomized) behavior strategies in order to establish existence of equilibria. Therefore, we did not adopt such more complex strategies. A behavior strategy for the seller is a function that assigns to each “type” $\theta \in \{L, H\}$ a probability distribution over \mathbb{R} for the price to post. We have not investigated this route. The interested reader is advised to read Ellingsen (1997) who, in a related model of one-sided asymmetric information allows for such randomization (over a finite subset of \mathbb{R}). A behavior strategy for the buyer is a function $b : \mathbb{R}^2 \rightarrow [0, 1]$ that assigns to each price-quality signal pair (p, \tilde{q}) the probability $b(p, \tilde{q}) \in [0, 1]$ that the buyer purchases the item. Our results remain largely unaffected if we allow buyers to use such randomized strategies. Indeed, the equilibrium candidates from remark 3.2 remain unchanged, as are the separating equilibria with total adverse selection. The analysis for partial adverse selection changes as follows. Let P_θ for $\theta \in \{L, H\}$ denote the probability distribution of the quality signal \tilde{q} conditional on quality q_θ (for fixed information precision α). It follows that the probability for trade at price p_H equals

$$\sigma_\theta = \int_{\tilde{q}} b(v_H, \tilde{q}) dP_\theta.$$

Thus, in line with lemma 4.2, these probabilities in a separating equilibrium with $(p_L, p_H) = (v_L, v_H)$ must satisfy

$$\begin{cases} \sigma_L & \leq \frac{v_L - w_L}{v_H - w_L} \\ \sigma_H & \geq \frac{\max\{0, v_L - w_H\}}{v_H - w_H} \end{cases}$$

Although randomized buying might increase the equilibrium gains of trade above and beyond what we found for deterministic threshold strategies, we have not succeeded in finding such buyer strategies. Also the analysis of pooling equilibria is affected, but typically only on a set of probability zero. Recall the definition (21) of the acceptance set. If $\mathbb{E}[q | \tilde{q}] > p$, then buying with probability one is the unique best reply at the buyer’s information set (p, \tilde{q}) . Conversely, if $\mathbb{E}[q | \tilde{q}] < p$, then buying with probability zero is the unique optimal response at the buyer’s information set (p, \tilde{q}) . Only in the remaining case, where $\mathbb{E}[q | \tilde{q}] = p$, is the buyer indifferent between buying and not buying and thus may optimally randomize. However, except for degenerate cases, this equality occurs with probability zero.

9 Directions for further research

The discussion in the preceding sections concerned a certain extension of the standard paradigm for asymmetric information in markets by introducing a particular form of *two-sided* asymmetric information: one party does not know the quality of the unit for sale but makes a noisy observation of its quality, and the other party knows the quality but not its impression on the first party. We think there are several interesting directions for further research on this paradigm, and we also believe that it could potentially be developed to provide a unifying framework for models of hidden information (adverse selection) and hidden action (moral hazard).

9.1 Equilibrium selection

Our model allows for multiple equilibria, with distinct outcomes. This raises the question whether there are natural refinements or stability conditions that one could or should impose on these equilibria. One basic refinement is structural consistency (Kreps and Wilson (1982)), that is, to require that the belief in each information set can be derived by Bayes' law from *some* strategy profile that reaches the information set with positive probability. Indeed, the belief systems used in our proofs stand this test.²⁰ Another relevant refinement is Kreps' and Cho's (1987) intuitive criterion. Adapted to the present setting, this would essentially require that if only one seller type $\theta \in \{L, H\}$ would have an incentive to deviate to some non-equilibrium price, then the buyer should place unit probability on that seller type when observing that price. Bester and Ritzberger (2001) show, in their related model (see above), that a somewhat strengthened version of the intuitive criterion cuts down their equilibrium set to a singleton. An investigation of the effects of such refinements on our equilibrium set would be a natural next step. Related belief-based refinements have been applied to different versions of Spence's (1973) signalling model, see Hellwig's (1987) and Riley's (2001) surveys. One equilibrium which stands out in these studies is the so-called Riley equilibrium (Riley (1979)). In that equilibrium, high-productivity job applicants signal their type by choosing the least costly signal needed to deter low-productivity workers from sending it. Employers' cut-off signal level for the high wage offer then is minimal among all separating equilibria. This separating equilibrium Pareto dominates all other separating equilibria in Spence's model. A number of arguments in favor of the Riley equilibrium have been raised in this literature (see Riley (1979, 2001) and Hellwig (1987)). It would therefore be interesting to investigate the relevance and power of those arguments in the present context, in particular with respect to our set of separating equilibria with partial adverse selection. Those equilibria were seen to be Pareto ranked. One might thus ask whether such arguments would single out the Pareto

²⁰The beliefs of the seller are by [B1] induced by nature's move, so it remains to check the buyer's beliefs at information sets off the equilibrium path. In our proofs (of Propositions 4.3, 4.1, and 5.1) the buyer assigns probability one to the item being of low quality at information sets (p, \tilde{q}) off the equilibrium path. This is structurally consistent with the following strategy for the seller: set $p_L = p$ and $p_H = p'$ for some $p' \neq p$.

dominant equilibrium, that is, where buyers' signal cut-off level is minimal. We have to leave this question for future research.²¹ Nöldeke and Samuelson (1997) and Jacobsen *et al.* (2001) used tools from evolutionary game theory to identify stable long-run outcomes in game-theoretic versions of Spence's signalling model. While Nöldeke and Samuelson found some support also for pooling equilibria, Jacobsen *et al.* found strong support for the Riley equilibrium. Likewise, Ania *et al.* (2002) applied similar tools to Rothschild's and Stiglitz' (1976) screening model. It would be interesting to investigate the applicability and power of those tools in the present model. A different approach that might be used for equilibrium selection in the present model is that of Maskin and Tirole (1992). They analyze a class of incomplete-information games, where an informed principal (here the seller) offers a contract to an uninformed agent (here the buyer), a contract that the agent may accept or reject. The setting is one of common values, that is, the principal's type enters directly into the payoff function of the agent — the case in our model, where the principal's type is the quality of his unit. Most of Maskin's and Tirole's analysis concerns the case when the agent has no private information. However, they also consider the case of relevance here, namely, when the agent has some private information (here, the quality signal). Maskin and Tirole characterize the allocations that result under perfect Bayesian equilibrium and identify a necessary and sufficient condition for their game to have a unique such equilibrium. We would like to investigate the applicability and power of their methods in the present setting — yet another topic for further research.

9.2 Generalizations

Having discussed potential approaches to equilibrium selection, we finally turn to potential generalizations of the model. One natural extension would be to allow for more than two quality levels. In the equilibria analyzed in this paper, buyers condition their purchasing decisions on the posted price and the quality signal. In separating equilibria, they buy at the high price only if the quality signal is good enough, and in pooling equilibria they likewise buy at the common price only if the quality signal is good enough. Sellers know this, and, knowing the quality of their own product, take the corresponding conditional buying probabilities into account in their pricing decisions. This logic does not hinge on the assumption that there are only two quality levels. We believe that most of the analysis — most likely at the cost of increased complexity — can be generalized to an arbitrary finite number of qualities. Another relevant extension would be to allow for a heterogeneous population of buyers, for instance by letting their valuations of the qualities differ (see also footnote 3).²² Such heterogeneity would add to the uncertainty that the seller has about the buyer: on top of not knowing the buyer's quality signal, the

²¹It should be noted, however, that although we identify a whole set of separating equilibria in which high quality goods are traded with positive probability, there may exist other separating equilibria where trade in high-quality items occurs with an even higher probability, equilibria which thus would Pareto dominate our equilibria, and which may be robust.

²²Alternatively, as in Bester and Ritzberger (2001), one may let buyers have different reservation values for *not* buying (here set equal to zero for all consumers).

seller would then be uncertain also about the buyer's valuation. This could be modelled by way of introducing a probability distribution for buyers' valuations of each quality. Assuming that the seller knows these distributions, it should be straight-forward to generalize the current analysis accordingly. If these value distributions are continuous, the nature of separating equilibria would change from holding all buyers indifferent between buying and not buying (as in the present version), to making them have strict preferences: with probability one the buyer's valuation will then differ from the posted price. There are numerous other natural sources of buyer heterogeneity, such as differences in signal precision; some buyers are more knowledgeable than others about the product in question. This could be modelled by letting the buyer's signal precision α be drawn from some probability distribution (known by the seller). Other directions for generalization would be to let the seller influence the quality of his product, thereby allowing for analyses of monopolists' incentives to enhance quality. This could be done by letting the seller choose the probability distribution over qualities. For example, suppose that in the present model the seller could choose λ , the probability that the product will be of low quality, and suppose that a low λ would be more costly than a high. If this choice were observed by the buyer, then such an extension should be straight-forward — the model would remain a pure adverse-selection model. By contrast, if the buyer could not observe the seller's choice of λ , the resulting model would also contain a moral hazard element. Likewise, it could be interesting to model situations in which the seller could influence α , the precision of the buyer's signal. Clearly, a seller with a high-quality unit could have an incentive to increase α . The present simple model seems to lend itself to a wide variety of such generalizations. In a similar vein, it would seem natural to allow buyers to influence their signal precision, that is, learn about the product in question before they go to the market, thereby increasing their signal precision. The buyer would then have to trade off the cost or disutility associated with such learning against its expected benefit for the subsequent purchasing decision (see Mattsson *et al.* (2004) for an analysis of such decision problems).

References

- [1] Akerlof, G. (1970): "The market for lemons: quality uncertainty and the market mechanism", *Quarterly Journal of Economics* 84, 488-500.
- [2] Ania, A.B., Tröger, T., and Wambach, A. (2002): "An evolutionary analysis of insurance markets with adverse selection", *Games and Economic Behavior* 40, 153-184.
- [3] Bagwell, K., and Riordan, M.H. (1991): "High and declining prices signal product quality", *American Economic Review* 81, 224-239.
- [4] Bester, H., and Ritzberger, K. (2001): "Strategic pricing, signalling, and costly information acquisition", *International Journal of Industrial Organization* 19, 1347-1361.

- [5] Cho, I.-K., and Kreps, D. (1987): "Signaling games and stable equilibria", *Quarterly Journal of Economics* 102, 179-221.
- [6] Ellingsen, T. (1997): "Price signals quality: The case of perfectly inelastic demand", *International Journal of Industrial Organization* 16, 43-61.
- [7] Fudenberg, D., and Tirole, J. (1991): *Game Theory*. MIT Press, Cambridge MA.
- [8] Hellwig, M. (1987): "Some recent developments in the theory of competition in markets with adverse selection", *European Economic Review* 31, 319-325.
- [9] Jacobsen, H.-J., Jensen, M., and Sloth, B. (2001): "Evolutionary learning in signalling games", *Games and Economic Behavior* 34, 34-63.
- [10] Kreps, D.M., and Wilson, R. (1982): "Sequential equilibria", *Econometrica* 50, 863-894.
- [11] Mas-Colell, A., Whinston, M.D., and Green, J.R. (1995): *Microeconomic Theory*, Oxford University Press, New York.
- [12] Maskin E. and J. Tirole (1992): "The principal-agent relationship with an informed principal, II: Common values", *Econometrica* 60, 1-42.
- [13] Mattsson, L.G., Voorneveld, M., and Weibull, J.W. (2004): "Discrete choice under uncertainty", Stockholm School of Economics WP in Economics, Finance and Statistics, No. 558.
- [14] Milgrom, P., and Roberts, J. (1986): "Price and advertising signals of product quality", *Journal of Political Economy* 94, 796-821.
- [15] Nöldeke, G., and Samuelson, L. (1997): "A dynamic model of equilibrium selection in signalling markets", *Journal of Economic Theory* 73, 118-156.
- [16] Riley, J. (1979): "Informational equilibrium", *Econometrica* 47, 331-359.
- [17] Riley, J. (2001): "Silver signals: twenty-five years of screening and signaling", *Journal of Economic Literature* 39, 432-478.
- [18] Rothschild, M., and Stiglitz, J. (1976): "Equilibrium in competitive insurance markets: An essay on the economics of imperfect information", *Quarterly Journal of Economics* 90, 629-649.
- [19] Schlee, E.E. (1996): "The value of information about product quality", *RAND Journal of Economics* 27, 803-815.
- [20] Spence, M.A. (1973): "Job market signaling", *Quarterly Journal of Economics* 87, 355-374.