

Article

A Scent of Lemon—Seller Meets Buyer with a Noisy Quality Observation

Mark Voorneveld¹ and Jörgen W. Weibull^{1,2,*}

¹ Department of Economics, Stockholm School of Economics, Box 6501, 113 83 Stockholm, Sweden

² Department of Economics, École Polytechnique, 91 128 Palaiseau Cedex, France

* Author to whom correspondence should be addressed; E-Mail: nejw@hhs.se; Tel.: +46-8-736 92 04; Fax: +46-8-31 32 07.

Received: 17 November 2010; in revised form: 8 February 2011 / Accepted: 14 March 2011 /

Published: 18 March 2011

Abstract: We consider a market for lemons in which the seller is a monopolistic price setter and the buyer receives a private noisy signal of the product's quality. We model this as a game and analyze perfect Bayesian equilibrium prices, trading probabilities and gains of trade. In particular, we vary the buyer's signal precision, from being completely uninformative, as in standard models of lemons markets, to being perfectly informative. We show that high quality units are sold with positive probability even in the limit of uninformative signals, and we identify some discontinuities in the equilibrium predictions at the boundaries of completely uninformative and completely informative signals, respectively.

Keywords: lemons; adverse selection; noisy quality signals; two-sided incomplete information

1. Introduction

In many markets, prospective buyers are imperfectly informed about the quality of the items for sale. Sellers, on the other hand, often know much about the quality of their products. In his classical “market for lemons” paper, George Akerlof showed that such informational asymmetry, when taken to the extreme, can lead to adverse selection. Although there is room for gains of trade in all qualities if only buyers were able to distinguish these from each other, the complete lack of such discriminatory power—the usual assumption—may imply crowding-out of high quality items, leaving

only the “lemons” in the market; see Akerlof [1]. That paper initiated a surge of economic analyses of markets with asymmetric information, where typically the seller has better information than the buyer. For example, a used-car owner usually knows more about his car than the prospective buyer and a job applicant usually knows more about her ability than the employer.¹

In this paper, we introduce a fairly rich form of asymmetric information in such situations of bilateral trade. In addition to observing the price posted by the seller, the buyer receives a noisy private signal about the product’s quality. The noisy signal has a given precision, known also to the seller, and a continuum of different values. The actual value is known only to the buyer. In the metaphorical language of our title: we endow buyers with some capability to detect the scent of lemon. We believe this is a realistic and relevant extension. Consumers typically have the opportunity to study the object for sale, thus getting an impression of its quality, and employers usually interview job candidates, thus obtaining an impression of the candidates’ abilities, *etc.* The impression may be more or less strong, and hence the consumer or employer may be more or less convinced of the quality of the object or the ability of the job candidate. So even if the signal precision is fixed and there are only two levels of quality or ability, as we here assume, the continuum signal assumption allows for a continuum range of degrees of confidence, represented as a continuum range of posterior probabilities that the object is of high quality (or the job candidate is of the high-ability type).

By varying the precision of the signal, the present model spans the whole continuum of consumer capabilities to recognize quality, from one extreme to the other. At the one end, when the signal precision is maximal, consumers recognize the true quality without mistake. At the other end, when the signal precision is minimal, it is as if all consumers have no sense of smell and are completely unable to recognize quality. Instead, they have to base their purchasing decision on their prior belief and the price chosen by the seller. This is the case treated in Akerlof’s classical model. Central questions in the present study are what happens to the set of equilibria when one varies the signal precision, and, in particular, how robust the standard predictions are for the two extreme cases of “symmetric” information (maximal signal precision) and “one-sided” asymmetric information (minimal signal precision—as in Akerlof’s approach).

Our starting point is a standard lemons model: an indivisible unit is available in two qualities, low and high. Buyers are interested in buying at most one unit. The seller knows the quality of the unit for sale. For each quality, the seller’s valuation is lower than buyers’ willingness to pay, so there are potential gains of trade in both qualities. When the seller and buyer meet, the seller sets a take-it-or-leave-it price for his unit. The buyer observes this price and receives a noisy quality signal. The seller knows the conditional probability distribution for signals, conditional upon the true quality, but does not know the buyer’s signal. For the sake of clarity and transparency we focus on an “elementary” situation—only one seller facing one buyer. As the reader will see, even an analysis of such simple situations requires substantial work. We analyze the corresponding non-cooperative game, using pure-strategy (weak) perfect Bayesian equilibrium as solution concept. We characterize equilibrium prices and trading probabilities for all degrees of signal precision.

¹We cannot possibly do justice to this rich literature here. The interested reader is referred to Riley [2] for a survey and to Mas-Colell *et al.* [3] for a textbook treatment.

Our main results are the following. The equilibria are either separating or pooling. Separating equilibria come in two varieties. In partial adverse-selection equilibria, high-quality items are traded with positive probability below one and low quality units are sold with probability one. There is a continuum of such equilibria, even in the limit as the buyer's signal precision is taken down to zero. Hence, high-quality units are sold with positive probability also when buyers have no information at all about the quality of the unit at hand. By contrast, in total-adverse equilibria, high quality units are not sold at all. However, those equilibria are not robust to the elimination of weakly dominated strategies. In pooling equilibria, buyers purchase at the going price only if their quality signal exceeds a certain threshold (determined in equilibrium). Sellers with low-quality units thus have a lower probability of selling than sellers with high-quality goods. In our metaphorical language: a scent of lemon is more likely to come from a low-quality item than from a high-quality item. The equilibrium selling probability for low quality is sufficiently high to deter such sellers from lowering their price. In general, there is a continuum of pooling equilibrium prices. Also here, some discontinuities arise. Firstly, price-pooling at buyers' average valuation—an equilibrium in the classical adverse-selection model (no sense of smell)—need not be an equilibrium outcome if the buyer has even the slightest sense of smell. Secondly, even when this is an equilibrium outcome, there is a discontinuity in trading probabilities. In the classical adverse-selection model, both low and high quality are traded with probability one, while here even a very noisy negative quality signal will deter a buyer from purchase.

Although there certainly is a related literature, to the best of our knowledge the present model is new. In Milgrom and Roberts [4], sellers signal their quality by way of a posted price and observed expenditures on advertisement. The seller knows the combined price-advertisement signal received by buyers, so also that is a model of incomplete information on only one side of the market. Bagwell and Riordan [5] 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, again the seller knows the buyer's signal. Schlee [6] considers a model in which there is quality uncertainty on both sides of the market, but where buyers and sellers are symmetrically informed; quality information is public. Hence, none of these models has two-sided incomplete information as we have in the present model: the seller does not know the quality signal received by the buyer and the buyer does not know the quality of the item for sale. The studies that are closest to our analysis appear to be Ellingsen [7] and Bester and Ritzberger [8]. A discussion of how our model and results are related to theirs is given in Section 3.5.

It should also be mentioned that the model is different from Spence's [9] signalling model, again a model of "one-sided" incomplete 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. By contrast, in our model the seller cannot, by assumption, affect the signal, and the signal is positively correlated with the value of the unit to the buyer. Of related interest are two recent papers on signalling games à la Spence [9]. Feltovich *et al.* [10] consider a model with three types and costly signalling. Interestingly, if the receiver, in addition to observing the signal of the sender, receives an exogenous quality signal, there arise equilibria where medium types make an effort to signal that they are not of low quality, whereas high types make no such effort, relying on the exogenous signal to reveal their exclusivity. Alós-Ferrer and Prat [11] consider

a dynamic signalling model where education plays a role as quality signal and worker productivity is revealed on the job as the result of firms' updating. This gives rise to a large set of equilibria and the Riley equilibrium is no longer uniquely selected by the Intuitive Criterion.

In our model, the seller is a monopolist who offers an indivisible item for sale at a take-it-or-leave-it (TOL) price. TOL pricing is known to be an expected-profit maximizing mechanism in situations of quotation marks "one-sided" incomplete information when there is only one buyer, granted the item is of a given, and to the seller known, quality; see Riley and Samuelson [12]. Therefore, we find a study of TOL pricing relevant also in present context. To the best of our knowledge, the mechanism literature has not yet provided an answer to the question whether TOL pricing is optimal also in settings with two-sided incomplete information, such as here (see Maskin and Tirole [13]).

The paper is organized as follows. Section 2 formalizes the model and provides certain necessary conditions on prices to be consistent with pure-strategy perfect Bayesian equilibrium. Section 3 contains the equilibrium analysis including a study of the boundary cases of completely uninformative and fully informative signals, respectively. Section 4 illustrates some points, and provides some additional results, for the special case of normally distributed error terms. Section 5 concludes. All proofs are in the appendix.

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* (or *willingness to pay* for) the two qualities to be $v_L = q_L$ and $v_H = q_H$, respectively. The seller has one unit of the good for sale, and knows the quality of his or her unit. The seller's *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: $\bar{q} = \lambda q_L + (1 - \lambda) q_H$. Thus, this is also the buyer's average valuation, $\bar{v} = \bar{q}$.

The seller sets a take-it-or-leave-it price p for his unit, and he may condition the price on its quality. The buyer is informed of the price p , but cannot observe the quality q of the item at hand. Instead, the buyer observes a signal \tilde{q} that is positively correlated with quality. We parameterize this relationship in such a way that the traditional cases of symmetric information and the classical adverse-selection model become special cases. More precisely, the buyer observes the noisy signal

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

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 $f : \mathbb{R} \rightarrow \mathbb{R}_+$.² 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 buyers' 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

²To be more precise, we assume F to be absolutely continuous: $F(z) = \int_{-\infty}^z f(z) dz$ for all $z \in \mathbb{R}$. Note also that the game would remain infinite even if one were to restrict the signal to be binary.

quality without error. In the opposite extreme case, $\alpha = 0$, buyers observe pure noise, $\tilde{q} = \varepsilon$, and thus have no information at all about the quality of the unit at hand. This corresponds to the standard model of adverse selection, pioneered by Akerlof [1], in which the uninformed party receives no signal at all.³ In all cases between these two extremes, buyers make imperfect observations of quality, and the seller knows this. However, the seller does not know the quality signal \tilde{q} received by the buyer.⁴

The seller is an expected-profit maximizer, where $\pi = p - w$ is the profit to a seller with valuation $w \in \{w_L, w_H\}$ who sells his 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.

Under 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

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

One of the main concerns from a welfare viewpoint is to what extent these potential gains are realized in equilibrium, and how the realized gains of trade depend 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 $\tilde{q} = \alpha q_\theta + (1 - \alpha)\varepsilon$, and decides whether to buy the item or not. A pure strategy for 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} .⁵
4. The game ends and the players receive their payoffs,

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

A decision node for the seller in the infinite game tree 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

³See section 3.5 for a discussion of the subtle difference between pure noise and no signal.

⁴The alternative and more conventional representation $\tilde{q} = q + \varepsilon$ of a noisy signal does not permit “pure noise” as a well-defined boundary case.

⁵More generally, the buyer may randomize, in which case the range of the function b is expanded to $[0, 1]$. However, we here focus on pure strategies.

(infinite) 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, for $\alpha = 1$ the quality signal \tilde{q} always coincides with the true quality: $\tilde{q} \in \{q_L, q_H\}$. Hence, in this case the buyer information set (p, q_θ) contains infinitely many nodes of the game tree, namely (θ, ε, p) for all $\varepsilon \in \mathbb{R}$. However, in this classical boundary case of perfect precision, the random variable ε is irrelevant to the buyer’s decision making.

As noted above, this game differs from the standard “lemons model” (as, for example, presented in chapter 13 in Mas-Colell *et al.* [3]), in two respects: (a) the seller is a price setter rather than price taker, and (b) 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 the buyer’s quality signal. Note also that seller and buyer “types” are correlated: the buyer’s “type” is the random variable $\tilde{q} = \alpha q_\theta + (1 - \alpha) \varepsilon$, where $\theta \in \{L, H\}$ is the seller’s “type.”

2.2. Equilibrium

We solve this game by adapting the usual notion of (weak) perfect Bayesian equilibrium to the present setting (see, *e.g.*, Fudenberg and Tirole [14] or Mas-Colell *et al.* [3]). Such an equilibrium specifies a strategy for each player and beliefs for each player over the nodes in the player’s 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). As it requires best replies also off the induced path of play, this concept is more stringent than (Bayesian) Nash equilibrium, but is weaker than sequential equilibrium since it imposes no constraints on beliefs at information sets off its path.⁶ We will focus on pure strategies.

In the case of an imprecise signal, $\alpha < 1$, this equilibrium requirement can be broken down to four conditions (the case $\alpha = 1$ is analyzed in Section 3.1). Following Kreps and Wilson [15], let a *belief system* be a function μ that assigns to each information set a probability distribution over the nodes in that information set. A *pure strategy* for the seller is a pair $s = (p_L, p_H) \in \mathbb{R}^2$, and a *pure strategy* for the buyer is a (Borel measurable) function $b : \mathbb{R}^2 \rightarrow \{0, 1\}$. We will call such a triplet (μ, s, b) an *equilibrium* if it meets conditions [B1]-[B3] and [SR] 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 $\{(\theta, \varepsilon) : \varepsilon \leq x\}$.

⁶The present game has infinitely many information sets and thus departs from the usual setting for sequential equilibrium. For exact definitions, see Fudenberg and Tirole [14] and Mas-Colell *et al.* [3].

[B2] If $p_L \neq p_H$: 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 $p_L = p_H$: at each buyer information set (p, \tilde{q}) where $p = p_L$, μ assigns probability

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

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)).⁷

[SR] 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 with positive probability under the strategy profile (s, b) , the buyer’s beliefs are consistent with (s, b) in the sense 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.⁸ The posterior probability $\mu_H(\tilde{q})$ can be thought of as the buyer’s degree of confidence, after inspection, that the item is of high quality. Finally, condition [SR] requires sequential rationality, as defined in Kreps and Wilson [15].

In any equilibrium $(\mu, (p_L, p_H), b)$, let ρ_L and ρ_H denote the *trading probabilities* for low and high quality units, that is,

$$\rho_\theta = \mathbb{E}[b(p_\theta, \tilde{q}) \mid \theta] \tag{6}$$

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

$$W^* = \lambda (v_L - w_L) \rho_L + (1 - \lambda) (v_H - w_H) \rho_H \tag{7}$$

An equilibrium $(\mu, (p_L, p_H), b)$ is called *separating* if $p_L \neq p_H$ and *pooling* if $p_L = p_H$.

2.3. Necessary Conditions for Equilibrium

The following lemma provides necessary conditions for equilibrium that follow from first principles. For instance, it is sequentially rational for the buyer to buy if the seller asks a price below the buyer’s willingness to pay for a low-quality unit, and it is sequentially rational not to buy if the seller asks a price above the buyer’s willingness to pay for a high-quality unit:

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

⁷In 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).

⁸Expressed 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} .

and

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

These simple observations rule out a wide range of price pairs from equilibrium:

Lemma 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 $p_L > v_L$
- (c) $p_L = v_L$ and $p_H \in (v_L, v_H)$
- (d) $p_L = p_H \geq v_H$

It follows from this lemma that, in separating equilibria, the price of low-quality units equals buyers' willingness to pay for such units (by (a) and (b)), and the price of high-quality units is never below the buyers' willingness to pay for such units (by (c)). Hence, when trade occurs in separating equilibria, the seller reaps all the gains of trade. This is not surprising, since by assumption the seller is a monopolist who can commit to take-it or leave-it offers. Moreover, 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 (by (a) and (d)). In sum:

Remark 1 The only remaining candidates for pure-strategy equilibria, when $\alpha < 1$, are separating equilibria with $p_L = v_L$ and $p_H \geq v_H$, and pooling equilibria with $p_L = p_H \in [v_L, v_H)$.

3. Analysis

In this section we analyze the game's equilibria and come to our main purpose: an investigation of the robustness of the equilibria—in particular the equilibrium prices and trading probabilities—in the classical boundary cases of uninformative signals ($\alpha = 0$) and fully revealing signals ($\alpha = 1$) via a systematic study of the equilibria near these boundaries. We proceed as follows: section 3.1 describes the equilibria when $\alpha = 1$. Moving on to the case where $\alpha < 1$, we establish conditions for the existence of separating equilibria and investigate their nature. In particular, we will show that low-quality items are sold with probability one in all separating (pure-strategy) equilibria, and we will identify two varieties of such equilibria. In one variety, high-quality units are not sold at all, while in the other variety they are sold with positive probability below one. We will refer to the first variety, studied in Subsection 3.2, 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, studied in subsection 3.3, will be called *partial adverse selection*. Subsection 3.4 characterizes the game's pooling equilibria, and Subsection 3.5, finally, discusses related models.

3.1. Symmetric Information

Under symmetric information, *i.e.*, when $\alpha = 1$, buyers know exactly the quality of the goods they buy. Pure-strategy perfect Bayesian equilibrium then requires the buyer to make a purchase if and only

if the posted price does not exceed the quality signal $\tilde{q} = q_\theta$, where q_θ is the quality of the unit at hand.⁹ At price $p = \tilde{q}$, the seller is indifferent between buying and not buying. The seller's best reply to the buyer's strategy is thus to set the price equal to the buyer's willingness to pay if the buyer strategy is to buy, and otherwise to set any price that is not lower than that. However, it is inconsistent with pure-strategy equilibrium that the buyer does not buy, since in this case the seller could marginally reduce the price and make a positive profit. Hence, the buyer always buys at the price $p = \tilde{q}$ and the seller sets prices $(p_L, p_H) = (v_L, v_H)$. Consequently, both qualities are traded with probability one in equilibrium, $\rho_L = \rho_H = 1$, and the seller absorbs all gains of trade.

3.2. Separating Equilibria with Total Adverse Selection

If $v_L \leq w_H$, then a seller with a high-quality unit cannot make a profit by selling at buyers' willingness to pay for low quality. Indeed, in this case, and only then, do there exist separating equilibria in which the price of high-quality items exceeds buyers' willingness to pay for high quality, v_H . Since it is suboptimal to buy at such prices, there is no trade in high-quality goods in such equilibria. 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 the seller sets the price as low as v_L , resulting in non-positive profits. In these equilibria, it is as if sellers with high-quality units left the market—as in Akerlof's [1] model of adverse selection.

Proposition 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 $\rho_H = 0$.*

This type of outpricing can occur for any degree of buyer signal precision $\alpha < 1$. Consequently, such equilibria exist even as buyers' information is made arbitrarily precise, that is, in the limit as $\alpha \rightarrow 1$, despite the fact that no such equilibria exist under symmetric information, that is, when $\alpha = 1$ (see Section 3.1). 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 1 establishes that outpricing equilibria exist whenever $w_H \geq v_L$. In particular, this result allows for equilibria with complete adverse selection even when $v_L \leq w_H < \bar{v}$, that is, when a pooling equilibrium would result in positive profits to sellers with high quality units. Such pooling equilibria do exist in standard adverse-selection models; see Section 3.4 for an analysis of pooling equilibria in the present model.

We should add that we do not find these separating equilibria plausible. For a seller with a high-quality couldn't lose, but might gain, by deviating to a price $p \in (w_H, v_H)$. More formally, the equilibria in Proposition 1 do not 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, all seller

⁹The buyer's information sets then 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 $\alpha = 1$.

strategies $s = (p_L, p_H)$ with $p_H > v_H$ become weakly dominated by any seller strategy $s' = (p_L, p)$ with $p \in (w_H, v_H]$.

3.3. Separating Equilibria with Partial adverse Selection

We here consider pure-strategy separating equilibria in which the seller’s strategy is $s = (v_L, v_H)$. We first outline the main arguments and then make the conclusions precise in a lemma and a proposition.

It follows from first principles that in equilibrium buyers always buy at the low price, $p = v_L$. For had they not done so, then a seller with a low quality unit could make a marginal price cut and sell for sure (since this would be the buyer’s unique optimal reaction), thus increasing his profit. Moreover, in equilibrium, buyers buy at the high price, $p = v_H$, with a probability that is less than one. For had they bought for sure at the high price from a seller with a high-quality unit, then they would also have bought for sure from a seller with a low-quality unit who posted that price.¹⁰ If this were the case, low-quality sellers would profitably deviate to the high price. Hence, the realization probability ρ_H is necessarily strictly less than one.

What can be said about the upper bound on the realization probability ρ_H ? Note first that $\rho_H = E [b(p_H, \tilde{q}) | \theta = H]$. Suppose that the seller has a low-quality unit but posts the high price. The probability that the buyer will purchase the unit is then $\mathbb{E} [b(p_H, \tilde{q}) | \theta = L]$. A necessary condition for equilibrium is thus that this probability is sufficiently low to deter low-quality sellers from deviating to the high price, or, more exactly, that

$$\mathbb{E} [b(p_H, \tilde{q}) | \theta = L] \leq \frac{v_L - w_L}{v_H - w_L} \tag{10}$$

Hence, the upper bound on the equilibrium realization probability ρ_H is

$$\hat{\rho}_H = \max_{b \in B} \mathbb{E} [b(p_H, \tilde{q}) | \theta = H] \tag{11}$$

where B is the subset of pure strategies b that satisfy inequality (10).

If $w_H < v_L$, so that a seller with a high-quality unit would make a (positive) profit by deviating to the low price, then another necessary condition for equilibrium is that

$$\rho_H \geq \frac{v_L - w_H}{v_H - w_H} \tag{12}$$

For had the realization probability ρ_H been lower, then a seller with a high-quality unit would earn $v_L - w_H$ by deviating to the low price, thus exceeding his expected equilibrium profit, $\rho_H (v_H - w_H)$.

It remains to establish existence of equilibria and to find an operational representation of $\hat{\rho}_H$. Let

$$t_0 = \alpha q_L + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right) \tag{13}$$

and, if $w_H < v_L$,

$$t_1 = \alpha q_H + (1 - \alpha) F^{-1} \left(\frac{v_H - v_L}{v_H - w_H} \right) \tag{14}$$

¹⁰In order to buy from a high-quality seller with probability one the buyer has to buy irrespective of the quality signal, since the error term has full support.

We note that $t_1 > t_0$, since, by hypothesis, $q_L < q_H$ and $w_L \leq w_H$. Let

$$T = \begin{cases} [t_0, t_1] & \text{if } w_H < v_L \\ [t_0, +\infty) & \text{otherwise} \end{cases} \tag{15}$$

Lemma 2 *If $\alpha < 1$, then $t \in T$ if and only if the following two conditions hold:*

$$\Pr [\tilde{q} > t \mid \theta = L] \leq \frac{v_L - w_L}{v_H - w_L} \tag{16}$$

$$\Pr [\tilde{q} > t \mid \theta = H] \geq \frac{v_L - w_H}{v_H - w_H} \tag{17}$$

We use this lemma to establish the existence of partially separating equilibria. 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} , and to buy from a seller who posts the high price, $p_H = v_H$, if and only if the perceived quality exceeds a certain threshold. More exactly:

Proposition 2 *Suppose $\alpha < 1$. All equilibria with $s = (v_L, v_H)$ have $\rho_L = 1$ and $\frac{v_L - w_H}{v_H - w_H} \leq \rho_H < 1$. For each $t \in T$ there exists such an equilibrium, with $b(p_H, \tilde{q}) = 1$ if and only if $\tilde{q} > t$, and with*

$$\rho_H = 1 - F\left(\frac{t - \alpha q_H}{1 - \alpha}\right) \tag{18}$$

We would like to make some comments to this proposition. First, the continuum of parameterized equilibria is Pareto ranked: the lower the threshold $t \in T$, the higher is the equilibrium probability ρ_H for trade in high-quality items, and thus the higher is the expected equilibrium profit to the seller. The expected utility to buyers being zero in all equilibria, this means that the equilibrium gains of trade, W^* , decrease with t and the equilibrium at t_0 Pareto dominates the others.

Secondly, 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, for instance, that a decision-maker in a non-strategic environment 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 high-quality 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 distribution), then the conditionally expected quality, $\mathbb{E}[q \mid \tilde{q}]$, is an increasing function of the signal \tilde{q} .¹¹ In such decision problems, it is optimal to use decision rules of the threshold form.

Thirdly, we examine some comparative statics properties of these separating equilibria with respect to the precision α of the buyer’s signal. By (13), the realization probability ρ_H is, for each $\alpha < 1$, bounded from above by

$$\bar{\rho}_H = 1 - F\left[F^{-1}\left(\frac{v_H - v_L}{v_H - w_L}\right) - \frac{\alpha(q_H - q_L)}{1 - \alpha}\right] \tag{19}$$

The upper bound $\bar{\rho}_H$ increases continuously with α , from a positive value when α is zero, towards one as $\alpha \rightarrow 1$. In other words, the more precise the signal is, the higher can be the equilibrium trading

¹¹See Mattsson, Voorneveld and Weibull [16] for an investigation of a class of such decision problems.

probability for high-quality units, and in the limit, as buyer information tends to perfection, the maximal trading probability for high-quality units tends to one.

What may appear as a surprise is that the present model allows for some trade in high-quality units even when buyers have no information at all about the quality of the unit at hand:

$$\alpha = 0 \quad \Rightarrow \quad \bar{\rho}_H = \frac{v_L - w_L}{v_H - w_L} \in (0, 1) \tag{20}$$

Although buyers receive only a non-informative signal, low-quality sellers do not have an incentive to deviate to the high price, since the trading probability at that price is too low. Until relatively recently, it was commonly presumed that, in classical adverse-selection models, high-quality units are not traded when $w_H > \bar{v}$. However, here they are traded with positive probability, and this is indeed even possible in the absence of a signal, a fact established by Ellingsen [7].¹²

In sum, the set of equilibrium trading probabilities for high-quality units contains the interval

$$I_0 = \left(\max \left\{ 0, \frac{v_L - w_H}{v_H - w_H} \right\}, \frac{v_L - w_L}{v_H - w_L} \right) \tag{21}$$

when $\alpha = 0$ (and in the limit as $\alpha \rightarrow 0$) and it contains the interval

$$I_1 = \left(\max \left\{ 0, \frac{v_L - w_H}{v_H - w_H} \right\}, 1 \right) \tag{22}$$

in the limit as $\alpha \rightarrow 1$. 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. In equilibrium, buyers are indifferent between buying and not buying, and may hence virtually abstain from buying, and a seller of a high-quality unit cannot gain from a marginal price cut in these equilibria, since buyers, who have less than perfect information, then believe that the unit at hand is likely to be of low quality.

Fourthly, recall from our second comment that the threshold functions in Proposition 2—where the item is purchased if the quality signal is sufficiently high—are optimal in a wide range of non-strategic circumstances. If the set B in the definition of the theoretical upper bound $\hat{\rho}_H$ in (11) is restricted to pure strategies that are weakly increasing functions of quality signal, then the upper bound $\bar{\rho}_H$ in (19) in fact equals the theoretical upper bound:

Proposition 3 *Suppose $\alpha < 1$. If buyers use weakly increasing pure strategies, then $\hat{\rho}_H = \bar{\rho}_H$.*

3.4. Pooling Equilibria

We saw in Section 3.1 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 1, it suffices to consider prices p in the half-open interval $[v_L, v_H)$.

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 her signal, is

$$\mathbb{E}[v \mid \tilde{q}] = \mu_L(\tilde{q})v_L + [1 - \mu_L(\tilde{q})]v_H \tag{23}$$

¹²Notice, however, that this requires that the parties deliberately randomize. In our model, the required randomization is obtained by pure strategies conditioned on the random signal.

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 $A(p) \subset \mathbb{R}$ denote the subset of signals \tilde{q} at which the conditionally expected value to the buyer is not less than p :

$$A(p) = \{\tilde{q} \in \mathbb{R} : \mathbb{E}[v | \tilde{q}] \geq p\} \tag{24}$$

If $p_L = p_H = p$, then consistency [B3] and sequential rationality [SR] together imply that it is optimal for the buyer to make a purchase if and only if his signal \tilde{q} belongs to $A(p)$.¹⁴ This set will be called the buyer’s *acceptance set* at price p .

The following proposition characterizes the set of prices that can be supported in pooling equilibria.

Proposition 4 *Suppose $\alpha < 1$. A price $p \in [v_L, v_H]$ is a pooling equilibrium price if and only if (25) and (26) hold.*

$$(p - w_L) \Pr[\tilde{q} \in A(p) | \theta = L] \geq v_L - w_L \tag{25}$$

$$(p - w_H) \Pr[\tilde{q} \in A(p) | \theta = H] \geq \max\{0, v_L - w_H\} \tag{26}$$

The first (second) condition in the proposition guarantees that a seller with a low (high) quality unit does not have an incentive to deviate. We note that condition (26) fails if $p < w_H$. Hence, not surprisingly, such prices can not be supported in pooling equilibria—a seller with a high-quality unit is then better off keeping his unit. By contrast, the same condition, (26), 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, irrespective of α (as long as $\alpha < 1$). For in this case the seller is willing to part with a high-quality unit, and it is rational for buyers to buy irrespective of the signal (their valuation cannot be lower than the going price). Formally: $A(v_L) = \mathbb{R}$, and the probabilities on the left-hand sides of (25) and (26) are both one, and thus both conditions are met for $p = v_L$.¹⁵

For the boundary case $\alpha = 0$, equation (5) implies that $\mu_L(\tilde{q}) = \lambda$ 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, $A(p) = \mathbb{R}$ for all prices $p \leq \bar{v}$, while for higher prices it is the empty set. It is easily verified that the two conditions in the proposition are met if and only if

$$\max\{v_L, w_H\} \leq p \leq \bar{v} \tag{27}$$

Hence, the set of pooling equilibrium prices is non-empty at $\alpha = 0$ if and only if $w_H \leq \bar{v}$, just as in standard adverse-selection models. 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 $W^* = \hat{W}$.

Is the equilibrium correspondence that maps buyers’ precision α to the corresponding set of pooling equilibrium prices p continuous at $\alpha = 0$? Assume that $\alpha > 0$ and $w_H < \bar{v}$, and consider any price $p < \bar{v}$ satisfying (27). 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 $\mathbb{E}[v | \tilde{q}] - p > 0$

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 $\mathbb{E}[v | \tilde{q}] > p$.

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

for α close to 0, by continuity. Moreover, ρ_L and ρ_H converge to 1. Hence, the above result for prices $p < \bar{v}$ is robust with respect to α . As will be shown in Section 4, however, the result is discontinuous with respect to α when $p = \bar{v}$.

3.5. Related Models

The case $\alpha = 0$, though closely related to standard models with adverse selection, is technically distinct from these. For in our model there is incomplete information on both sides of the market even when $\alpha = 0$, in the trivial sense that the seller does not know the buyer's *uninformative* signal. The signal makes the outcome of the buyer's pure strategies randomized. In the boundary case $\alpha = 0$, our model is close to those of Ellingsen [7] and Bester and Ritzberger [8]. 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 more than two qualities, distributed over a finite quality grid with fixed grid size. Bester and Ritzberger [8] 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 has incomplete information on both sides of the market. The differences are that (a) in their model buyers either have no signal or a perfect signal, (b) it combines incomplete information (hidden information) with imperfect information (hidden action).

The model in Akerlof [1] is equivalent with the non-strategic variant of our model where the seller is a price taker and the buyer has no quality signal. Akerlof asks at what market prices p , the same for all qualities, trade is possible. If each of the two qualities is offered with probability λ and $1 - \lambda$, respectively, then the buyer's willingness to pay will be \bar{v} : the buyer is willing to buy if and only if $p \leq \bar{v}$. Similarly, sellers are willing to sell both qualities if and only if $p \geq \max\{w_L, w_H\} = w_H$. Thus, a necessary and sufficient condition for trade in both low and high quality is $p \in [w_H, \bar{v}]$. At all such prices, the full gains of trade are realized. On the other hand, if $w_H > \bar{v}$, total adverse selection arises: high quality is driven out of the market and trade occurs only in low quality, at prices $p \in [w_L, v_L]$.

4. Normally Distributed Noise

In the final lines of subsection 3.4, we mentioned a discontinuity in the set of pooling equilibria: although $p = \bar{v}$ is a pooling equilibrium price when $\alpha = 0$, this may or may not be the case when $\alpha \rightarrow 0$. To make this explicit (indeed, see (30) below), this section studies the particular case of normally distributed error terms. The necessary and sufficient conditions (25) and (26) for pooling equilibrium can then be written explicitly in terms of the primitives of the model. Suppose that $\alpha \in (0, 1)$, and define the function $\varphi : (v_L, v_H) \rightarrow \mathbb{R}$ by

$$\varphi(p) = \frac{\alpha(q_H - q_L)}{2(1 - \alpha)\sigma} + \frac{(1 - \alpha)\sigma}{\alpha(q_H - q_L)} \ln \left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)} \right] \quad (28)$$

Clearly φ is increasing, with $\varphi(p) \rightarrow -\infty$ as $p \rightarrow v_L$ and $\varphi(p) \rightarrow +\infty$ as $p \rightarrow v_H$.

Proposition 5 *If $\alpha \in (0, 1)$ and $\varepsilon \sim N(\bar{q}, \sigma^2)$, then $p \in \mathbb{R}$ is a pooling equilibrium price if and only if $\max\{v_L, w_H\} \leq p < v_H$ and*

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

We would like to make some comments to this proposition. Firstly, 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 the buyers would expect all units to be of average quality, and thus would be willing to pay only \bar{v} . Indeed, the reader may verify that for parameter values $w_L = 0, q_L = 1, q_H = 3, \lambda = 0.25, \sigma = 1$, and α close to 0.1, pooling is possible at a price exceeding the average valuation \bar{v} .

Secondly, the set of pooling equilibrium prices, if non-empty, is an interval. This interval has $\max\{v_L, w_H\}$ as its left endpoint, while inspection of equation (28) shows that its right endpoint is decreasing in λ . Thus, the larger the share of “lemons,” the lower is the highest pooling price.

Thirdly, by (28) and (29), the set of pooling equilibrium prices approaches the set

$$P_0 = \begin{cases} [\max\{v_L, w_H\}, \bar{v}] & \text{if } \frac{v_L - w_L}{\bar{v} - w_L} < \frac{1}{2} \\ [\max\{v_L, w_H\}, \bar{v}) & \text{if } \frac{v_L - w_L}{\bar{v} - w_L} \geq \frac{1}{2} \end{cases} \tag{30}$$

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 $\alpha \rightarrow 0$. That is, if $p = \bar{v}$ is a pooling equilibrium price as $\alpha \rightarrow 0$, the corresponding trading probabilities ρ_L, ρ_H converge to $1/2$. Recall, however, from the discussion following equation (27) that at $\alpha = 0$, the trading probability for both qualities in pooling equilibrium at price \bar{v} must be equal to one. This discontinuity in the trading probabilities persists even if mixed/behavioral strategies were allowed. Firstly, the pure-strategy equilibria remain equilibria in degenerate mixed/behavioral strategies. Secondly, randomizing between buying and not buying is rational only if the buyer is indifferent. Yet after receiving the signal, the buyer is indifferent with probability zero when $\alpha > 0$, while she is indifferent with probability one when $\alpha = 0$.

Fourthly, no price above v_L is a pooling price in the limit as buyers’ signal precision α goes to one. To see this, note that, for any $p \in (v_L, v_H)$, $\varphi(p)$ is continuous in α and tends to $+\infty$ as α approaches 1. Hence, as α tends to 1, the left-hand side in (29) tends to zero, and thus any such price p ceases to be a pooling equilibrium price.

On a somewhat different note, let us compare the equilibrium profit to the seller in pooling equilibria with the maximal expected profit in separating equilibria with partial adverse selection. The expected profit in a pooling equilibrium with price p equals

$$\mathbb{E} [\pi^{pool}(p)] = \lambda (p - w_L) (1 - \Phi[\varphi(p)]) + (1 - \lambda) (p - w_H) (1 - \Phi[\psi(p)]) \tag{31}$$

where Φ is the c.d.f. of $N(0, 1)$, φ is defined in (28), and

$$\psi(p) = -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)\sigma} + \frac{(1 - \alpha)\sigma}{\alpha(q_H - q_L)} \ln \left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)} \right] \tag{32}$$

We compare this with the expected profit in separating equilibria. Since the prices are the same in all such equilibria, v_L for low quality and v_H for high, expected profits are maximal in the separating equilibrium with the maximal selling probability for high quality. By proposition 3 and (19), the maximal expected profit in separating equilibrium (see Appendix for a derivation) is

$$\mathbb{E}[\pi^{psep}] = \lambda(v_L - w_L) + (1 - \lambda)(v_H - w_H) [1 - \Phi(\xi)] \tag{33}$$

where

$$\xi = \Phi^{-1} \left(\frac{v_H - v_L}{v_H - w_L} \right) - \frac{\alpha(q_H - q_L)}{(1 - \alpha)\sigma} \tag{34}$$

These equations are useful when comparing the expected profit in pooling and separating equilibria in the limit case when the signal precision α tends to zero. Recalling that the set P_0 of pooling equilibrium prizes in this limit case is empty if $w_H > \bar{v}$ and non-empty if $w_H < \bar{v}$, we focus on the latter case. Let $\bar{w} = \lambda w_L + (1 - \lambda) w_H$ be the seller’s average reservation price. The following result establishes that, for all $\alpha > 0$ sufficiently small, there exist pooling equilibria with higher expected profit than in all separating equilibria, and that there are pooling equilibrium profits arbitrarily close to $\bar{v} - \bar{w}$ as $\alpha \rightarrow 0$. Formally:

Proposition 6 *If $w_H < \bar{v}$, then*

$$\lim_{\alpha \rightarrow 0} \left(\sup_p \mathbb{E}[\pi^{pool}(p)] \right) = \bar{v} - \bar{w} > \lim_{\alpha \rightarrow 0} \mathbb{E}[\pi^{psep}]$$

In other words, if the buyer’s average willingness to pay exceeds the seller’s valuation of high quality, then the seller prefers pooling equilibria over separating equilibria. Moreover, in the limit as the buyer’s quality signal becomes uninformative, potential gains of trade are virtually realized. The qualifier “virtually” refers to the fact (established in the proof of the proposition) that while there exist pooling equilibria that result in profits arbitrarily close to $\bar{v} - \bar{w}$, no equilibrium actually achieves this lowest upper bound exactly. Note, however, that this ambiguity is artificial; it hinges on the assumption that a price can be any real number, while in practice prices are quoted in some smallest monetary unit. In the latter case, the supremum in the statement of the proposition should be taken over the corresponding discrete set of prices, and this supremum profit will indeed be achieved by some pooling price, namely, the largest multiple of the monetary unit that does not exceed the buyer’s willingness to pay. In sum: the welfare loss due to asymmetric information is virtually zero in the limit as the buyer’s quality signal becomes uninformative. By contrast, some gains of trade are lost in all equilibria at any intermediate level $\alpha \in (0, 1)$ of signal precision.

We finally note that when the buyer’s quality signal is virtually perfect, $\alpha \rightarrow 1$, the best partially separating equilibrium gives a higher expected profit than all pooling equilibria, and this separating equilibrium realizes the potential gains of trade:

$$\lim_{\alpha \rightarrow 1} \mathbb{E}[\pi^{pool}] = v_L - \bar{w} < \lim_{\alpha \rightarrow 1} \mathbb{E}[\pi^{psep}] = \bar{v} - \bar{w} \tag{35}$$

In other words, the welfare loss due to asymmetric information vanishes in the limit as the buyer’s quality signal becomes perfect.

5. Conclusion

The goal of this paper was to study equilibria when noisy inspection is introduced in an otherwise standard “lemons” model; more specifically, when the buyer does not know the quality of the unit for sale but receives a noisy signal about its quality while the seller knows the quality but not the buyer’s signal. We believe that this extension is relevant and provides a more realistic description of asymmetric information in markets. A prospective buyer of a used car may walk around it, take a look under the hood, check whether the bumper falls off when it receives a slight kick *etc.* and thus get at least some impression of the car’s quality, rather than, as in Akerlof’s seminal analysis, no impression at all. Although the seller knows the quality of his car, he usually can only guess at the exact impression on the prospective buyer. We chose to study an elementary situation—only one seller facing one buyer at a time—devoid of all the bells and whistles that could make the model richer and more realistic, for the sake of clarity and transparency. Moreover, we focussed on pure-strategy equilibria. In the present model, this is a rich set of strategies, offering a continuum of options for the seller and allowing the buyer to condition the purchasing decision on signal values in a continuum. Arguably, this provides a behaviorally plausible form of randomization—in the form of hesitation in face of quality uncertainty.

While the focus of this paper is admittedly narrow, we felt it natural to first establish a clear-cut and transparent analysis of the set of pure strategies, before embarking on studies of richer, and thus more complex, market settings. Evidently, there are numerous relevant directions for further research. We here briefly name two:

Equilibrium selection: The model allows for multiple equilibria with distinct outcomes. This raises the question whether there are relevant refinements or stability conditions that one could impose in order to cut down the equilibrium set. This is a very interesting avenue for future research, although non-trivial, since standard refinements are defined only for finite games, while the present game is infinite.

Generalizations and extensions: Highly relevant extensions would be to analyze cases of multiple quality levels, multiple sellers and/or buyers, and a variety of pricing mechanisms.

We are currently working on a few such directions, in joint work with Yukio Koriyama. The analysis being quite demanding, it appears that generalization in one dimension typically requires simplification in another. We hope that the present study will inspire future research on market interactions in which some or all parties have noisy signals about product quality.

Acknowledgements

We thank Carlos Alós-Ferrer, Tore Ellingsen, Yukio Koriyama, Lars-Göran Mattsson, Jianjun Miao, and Zvika Neeman, and two anonymous referees for comments. Voorneveld thanks the Wallander/Hedelius Foundation and the Netherlands Organization for Scientific Research for financial support.

References

1. Akerlof, G. The market for lemons: Quality uncertainty and the market mechanism. *Quart. J. Econ.* **1970**, *84*, 488-500.

2. Riley, J. Silver signals: Twenty-five years of screening and signaling. *J. Econ. Lit.* **2001**, *39*, 432-478.
3. Mas-Colell, A.; Whinston, M.D.; Green, J.R. *Microeconomic Theory*; Oxford University Press: New York, NY, USA, 1995.
4. Milgrom, P.; Roberts, J. Price and advertising signals of product quality. *J. Polit. Econ.* **1986**, *94*, 796-821.
5. Bagwell, K.; Riordan, M.H. High and declining prices signal product quality. *Amer. Econ. Rev.* **1991**, *81*, 224-239.
6. Schlee, E.E. The value of information about product quality. *RAND J. Econ.* **1996**, *27*, 803-815.
7. Ellingsen, T. Price signals quality: The case of perfectly inelastic demand. *Int. J. Ind. Organ.* **1997**, *16*, 43-61.
8. Bester, H.; Ritzberger, K. Strategic pricing, signalling, and costly information acquisition. *Int. J. Ind. Organ.* **2001**, *19*, 1347-1361.
9. Spence, M.A. Job market signaling. *Quart. J. Econ.* **1973**, *87*, 355-374.
10. Feltovich, N.; Harbaugh, R.; To, T. Too cool for school? Signalling and countersignalling. *RAND J. Econ.* **2002**, *33*, 630-649.
11. Alós-Ferrer, C.; Prat, J. Job market signalling and employer learning. In *IZA Discussion Paper Series*; No. 3285; Institute for the Study of Labor: Bonn, Germany, 2008.
12. Riley, J.G.; Samuelson, W.F. Optimal auctions. *Amer. Econ. Rev.* **1981**, *71*, 381-392.
13. Maskin, E.; Tirole, J. The principal-agent relationship with an informed principal, II: Common values. *Econometrica* **1992**, *60*, 1-42.
14. Fudenberg, D.; Tirole, J. *Game Theory*; MIT Press: Cambridge, MA, USA, 1991.
15. Kreps, D.M.; Wilson, R. Sequential equilibria. *Econometrica* **1982**, *50*, 863-894.
16. Mattsson, L.G.; Voorneveld, M.; Weibull, J.W. Better may be worse: Some monotonicity results and paradoxes in discrete choice under uncertainty. *Theor. Decis.* **2007**, *63*, 121-151.

Appendix

Proof of Lemma 1

(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, according to [SR], 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 [SR], 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 (p_H, \tilde{q}) of the buyer. By [B3], he believes with positive probability (5) that the item is of low quality, in which case a purchase would give negative utility ($v_L - p_H = v_L - v_H < 0$), and with the residual probability that the item is of high quality, in which case a purchase would give zero utility ($v_H - p_H = 0$). The expected payoff thus being negative, the optimal decision, by [SR], 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 $p - w_L$.

Proof of Proposition 1

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 $\tilde{q} \in \mathbb{R}$.

Proof of Lemma 2

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 (16) and (17) are equivalent with, respectively,

$$F\left(\frac{t - \alpha q_L}{1 - \alpha}\right) \geq \frac{v_H - v_L}{v_H - w_L} \tag{36}$$

and

$$F\left(\frac{t - \alpha q_H}{1 - \alpha}\right) \leq \frac{v_H - v_L}{v_H - w_H}. \tag{37}$$

If $w_H < v_L$, then $t \in \mathbb{R}$ satisfies (36) and (37) if and only if

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

If $w_H \geq v_L$, then the right-hand side of (37) is at least 1, and hence (37) is met.

Proof of Proposition 2

In order to prove that $\rho_L = 1$, suppose (μ, s, b) is an equilibrium with $(p_L, p_H) = (v_L, v_H)$. By [B2] and [SR], 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 $\tilde{q} \in \mathbb{R}$. The inequality $\rho_H \geq \frac{v_L - w_H}{v_H - w_H}$ 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 $v_L - w_H$. The inequality $\rho_H < 1$ was established in the text.

To prove the second claim, let $(p_L, p_H) = (v_L, v_H)$, let $t \in T$, 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 (16), 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 (17), 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.

Proof of Proposition 3

Assume the buyer’s strategy at price p_H is a threshold function: there is a $t \in \mathbb{R}$ such that $b(p_H, x) = 0$ if $x \leq t$ and $b(p_H, x) = 1$ otherwise. Then the problem (11) becomes

$$\max \Pr[\tilde{q} > t \mid \theta = H] \text{ s.t. } \Pr[\tilde{q} > t \mid \theta = L] \leq \frac{v_L - w_L}{v_H - w_L}$$

Since these probabilities are decreasing in t , the constraint will hold with equality:

$$\frac{v_L - w_L}{v_H - w_L} = \Pr[\tilde{q} > t \mid \theta = L] = 1 - F\left(\frac{t - \alpha q_L}{1 - \alpha}\right)$$

Solving this equation for t yields that $t = t_0$ and that $\hat{\rho}_H$, the value of the maximization problem (11), equals the trading probability in high quality if buyers purchase at price q_H if the signal exceeds t_0 , which is exactly $\bar{\rho}_H$.

Proof of Proposition 4

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

$$\pi_L \leq (p - w_L) \Pr[\tilde{q} \in A(p) \mid \theta = L] \tag{38}$$

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

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

proving (25). Likewise, a high-quality seller’s expected profit, π_H , satisfies

$$\pi_H \leq (p - w_H) \Pr[\tilde{q} \in A(p) \mid \theta = H] \tag{39}$$

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,

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

proving (26). To prove sufficiency, define (μ, s, b) by setting $p_L = p_H = p \in [v_L, v_H)$,

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

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 $A(p)$ 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 (25) and (26) imply that the seller has no profitable unilateral deviations: a price $p' \geq v_L, p' \neq p$, results in zero trade, and inequalities (38) and (39) are met with equality under (μ, s, b) .

Proof of Proposition 5

Condition (26) 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 $\tilde{q} \in A(p)$ 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, $\tilde{q} \in A(p)$ 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

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

and

$$\begin{aligned} \rho_H &= \Pr[\tilde{q} \in A(p) \mid \theta = H] \\ &= \Pr\left[(\varepsilon - \bar{q})^2 \leq \left(\varepsilon + \frac{\alpha(q_H - q_L)}{1 - \alpha} - \bar{q}\right)^2 - 2\sigma^2\beta(p)\right] \\ &= \Pr\left[\varepsilon - \bar{q} \geq -\frac{\alpha(q_H - q_L)}{2(1 - \alpha)} + \frac{\sigma^2(1 - \alpha)\beta(p)}{\alpha(q_H - q_L)}\right] \end{aligned} \tag{40}$$

As expected, $\rho_H > \rho_L$. For any $p \in (v_L, v_H) \cap [w_H, +\infty)$, inequality (26) is met whenever inequality (25) 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).

Derivation of Equation (34)

Notice that, with F denoting the error c.d.f.:

$$F(z) = \Pr[\varepsilon \leq z] = \Pr\left[\frac{\varepsilon - \bar{q}}{\sigma} \leq \frac{z - \bar{q}}{\sigma}\right] = \Phi\left(\frac{z - \bar{q}}{\sigma}\right)$$

and

$$F^{-1}(z) = \sigma\Phi^{-1}(z) + \bar{q}$$

Hence

$$F^{-1}\left(\frac{v_H - v_L}{v_H - w_L}\right) - \frac{\alpha(q_H - q_L)}{1 - \alpha} = \bar{q} + \sigma\Phi^{-1}\left(\frac{v_H - v_L}{v_H - w_L}\right) - \frac{\alpha(q_H - q_L)}{1 - \alpha}$$

and thus, by (19):

$$\bar{\rho}_H = 1 - \Phi\left(\Phi^{-1}\left(\frac{v_H - v_L}{v_H - w_L}\right) - \frac{\alpha(q_H - q_L)}{\sigma(1 - \alpha)}\right)$$

Proof of Proposition 6

It follows from (28) and (32) that $\varphi(p)$ and $\psi(p)$ have identical limits as $\alpha \rightarrow 0$. This limit value depends on whether the term

$$\beta(p) = \ln\left[\frac{\lambda(p - v_L)}{(1 - \lambda)(v_H - p)}\right] \tag{41}$$

is negative, positive, or zero.

1. If $p < \bar{v}$, $\beta(p) < 0$, so $\varphi(p), \psi(p) \rightarrow -\infty$ as $\alpha \rightarrow 0$. Conclude from (31) that

$$\mathbb{E} [\pi^{pool}(p)] \rightarrow \lambda(p - w_L) + (1 - \lambda)(p - w_H) = p - \bar{w}$$

2. If $p = \bar{v}$, $\beta(p) = 0$, so $\varphi(p), \psi(p) \rightarrow 0$ as $\alpha \rightarrow 0$. Conclude from (31) that

$$\mathbb{E} [\pi^{pool}(p)] \rightarrow \frac{1}{2}\lambda(p - w_L) + \frac{1}{2}(1 - \lambda)(p - w_H) = \frac{1}{2}(p - \bar{w})$$

3. If $p > \bar{v}$, $\beta(p) > 0$, so $\varphi(p), \psi(p) \rightarrow +\infty$ as $\alpha \rightarrow 0$. Conclude from (31) that

$$\mathbb{E} [\pi^{pool}(p)] \rightarrow \lambda(p - w_L) \cdot 0 + (1 - \lambda)(p - w_H) \cdot 0 = 0$$

Let $p \in (v_L, \bar{v})$. Then $\varphi(p) \rightarrow -\infty$ as $\alpha \rightarrow 0$, so condition (29) for $p \in (v_L, \bar{v})$ being a pooling equilibrium price is satisfied for sufficiently small $\alpha > 0$. Together with the first observation above, this implies that

$$\limsup_{\alpha \rightarrow 0} \mathbb{E} [\pi^{pool}(p)]$$

is well-defined and equals $\bar{v} - \bar{w}$, whereas (33) and (34) yield

$$\lim_{\alpha \rightarrow 0} \mathbb{E} [\pi^{psep}] = \lambda(v_L - w_L) + (1 - \lambda)(v_H - w_H) \frac{v_L - w_L}{v_H - w_L} < \bar{v} - \bar{w}$$

© 2011 by the authors; licensee MDPI, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).