Selling a Lemon under Demand Uncertainty

(Preliminary and incomplete)

Kyungmin Kim[*] and Sun Hyung Kim[†]

February 2015

Abstract

We study the dynamic pricing problem of the seller who has private information about the quality of the good, faces a sequence of randomly arriving buyers, but is uncertain about the arrival rate of buyers. We characterize a class of equilibria in which the high-quality seller insists on a constant price and show that the low-quality seller's equilibrium pricing strategy is to offer the high price initially and drop it to the optimal price under symmetric information. We also show that demand uncertainty can be beneficial to the seller.

JEL Classification Numbers: C73, C78, D82, D83, L15.

Keywords: Adverse selection; market for lemons; demand uncertainty; dynamic pricing

1 Introduction

Consider a seller who wishes to sell a used car. Due to her experience with the car, she is likely to be better informed about the quality of the car than buyers, which creates the lemons problem (Akerlof, 1970). This produces risks for buyers, thereby complicating their purchase decisions. On the other hand, the seller is likely to have uncertainty about demand. A seller rarely has precise information about the aggregate state of the economy or the general popularity of her car. This is arguably the reason why most used car sellers refer to price information services such as Kelly Blue Book. There is an extensive literature both on adverse selection and on demand uncertainty. To our knowledge, however, the interplay between adverse selection and demand uncertainty has not been investigated in the literature yet. The goal of this paper is to understand the interplay by studying a formal trading environment that features both adverse selection and demand uncertainty.

[*]University of Iowa. Contact: kyungmin-kim@uiowa.edu

[†]University of Iowa. Contact: sunhyung-kim@uiowa.edu

We consider a dynamic pricing problem facing a seller who wishes to sell an indivisible object. She sets a price and can adjust it at any point in time at no cost. Buyers arrive sequentially, observe the posted price, and decide whether to purchase the good or not. The seller has private information about the quality of the good, which is either high or low. We focus on the case where adverse selection is severe enough to impede market inefficiency (in particular, some trade delay is unavoidable). On the other hand, she faces uncertainty about demand. Specifically, she is uncertain whether the arrival rate of buyers is high or low.

We characterize a class of equilibria in which the seller insists on a constant price if her good is of high quality. The assumption about the high-quality seller's pricing behavior implies that the low-quality seller can choose only between two prices, one which is charged by the high-quality seller and the other which is optimal conditional on her type being revealed. This mitigates severe equilibrium multiplicity, due to the signaling nature of the model, as well as gives tractability to the analysis. Still, there is a continuum of prices that can be employed as the high price. Instead of selecting a particular equilibrium, we characterize all such equilibria.

We show that in any equilibrium the low-quality seller's optimal pricing strategy is a simple stopping rule: she offers the high price until some threshold time and then permanently switch to the low price. Severe adverse selection does not allow the high price to always be accepted by buyers. It also cannot be the case that buyers never accept the high price, as it implies that the low-quality seller would never offer the high price, which in turn implies that buyers would always accept the high price. Given buyers' random purchase decisions at the high price, the low-quality seller faces the following trade-off: if she posts the high price, then her payoff is higher conditional on sale, but a sale takes longer. The optimality of the stopping rule follows from this trade-off and the observation that the seller's belief about the demand state deceases conditional on no sale.

The optimality of such a pricing strategy is a joint consequence of adverse selection and demand uncertainty. With no (or mild) adverse selection, there exists an equilibrium in which the two seller types completely pool and play an identical pricing strategy. In the absence of demand uncertainty, as we formally show in Section 3, the low-quality seller's optimal strategy is indeterminate: equilibrium does not impose enough restrictions to discipline the low-quality seller's optimal pricing strategy. For example, the low-quality seller may randomize between the high price and the low price with a constant probability over time. She may even switch from the low price to the high price at some threshold time.

We also find that, contrary to the conventional wisdom, demand uncertainty can be beneficial to the seller. Uncertainty implies that an agent can take an action that is optimal only on average, not the action that is optimal conditional on each state. This is the reason why an agent's value function with respect to his belief is convex and learning (i.e., reducing uncertainty) is beneficial in standard decision problems under uncertainty. Our counterintuitive result stems from the strategic nature of our model. If buyers' purchase strategies were independent of demand uncertainty, the seller would necessarily prefer learning about demand. However, buyers' purchase behavior does depend on the presence of demand uncertainty. Therefore, the seller could benefit from uncertainty.

Our paper contributes to the fast-growing literature on dynamic adverse selection.[1] Particularly close to our paper is Lauermann and Wolinsky (2011). They study a model that is similar to our benchmark case without demand uncertainty. Among various differences in modeling, one crucial difference is their assumption that (translating their model into our context) quality concerns only buyers' values, not the seller's cost (i.e., $c_L = c_H$). This is the reason why, in the absence of buyers' signals about the seller's type, equilibrium is complete pooling in their model, while it always features some separation in our model. More importantly, our main objective is to study the interaction between adverse selection and demand uncertainty, while they are mainly interested in the ability of prices to aggregate dispersed information in such an environment.

Mason and Välimäki (2011) consider a similar dynamic pricing problem under demand uncertainty and find the optimal solution similar to ours. The main difference is that adverse selection is absent and, therefore, buyers do not behave strategically in their model: each buyer has private value for the seller's good and accepts any price below his value. In other words, buyers' purchase behavior is exogenously given in their model, while it is endogenously derived in our model. This gives rise to, for example, the following difference: the seller in their model immediately settles on the low price if the initial probability of the high demand state is below a certain threshold, while the low-quality seller in our model begins with the high price no matter how small probability she assigns to the high demand state (because buyers' purchase behavior is adjusted so that the high price is optimal initially). Nevertheless, we significantly benefit from their analysis, because given buyers' purchase behavior, the low-quality seller's optimal pricing problem is formally identical to their problem.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 study the benchmark case without demand uncertainty. Section 4 analyzes the main model with demand uncertainty.

[1] See Evans (1989); Vincent (1989, 1990); Taylor (1999); Janssen and Roy (2002); Deneckere and Liang (2006); Hörner and Vieille (2009); Daley and Green (2012) for some seminal contributions. Most papers in this literature consider the case where uninformed players make price offers to informed players, mainly to avoid equilibrium multiplicity due to signaling. Two exceptions are Lauermann and Wolinsky (2011) and Gerardi, Hörner and Maestri (2014), both of which study the opposite case where informed players make price offers to uninformed plays. The former focuses on undefeated equilibria (Mailath, Okuno-Fujiwara and Postlewaite, 1993), while the latter characterizes the set of all equilibrium payoffs.

2 The Model

2.1 Physical Environment

A seller wishes to sell an indivisible object. Time is continuous and indexed by t. The time the seller arrives at the market is normalized to 0. At each point in time, the seller posts a price. Buyers arrive sequentially according to a Poisson process of rate λ. Upon arrival, each buyer observes the posted price and decides whether to purchase the good or not. If the buyer purchases, then trade takes place between the buyer and the seller, and the game ends. If not, the buyer leaves, while the seller continues the game. The seller's discount rate is given by $r > 0$.

The seller's good is either of low quality (L) or of high quality (H). If the good is of quality $a = H, L$, then it yields flow payoff rc_a to the seller (while she retains it) and flow payoff rv_a to a buyer (once he purchases it). In other words, the reservation value of the good to the seller is c_a, while its value to a buyer is v_a. A high-quality unit is more valuable to both the seller and buyers (i.e., $v_H > v_L$ and $c_H > c_L$). There are always gains from trade (i.e., $v_a > c_a$ for both $a = H, L$), but the quality of the good is known only to the seller. It is common knowledge that the seller's good is of high quality with probability q_0 at the beginning of the game. We normalize c_L to 0.

The seller has uncertainty about the demand state. Specifically, the arrival rate of buyers λ is either λ_h (high demand) or λ_l (low demand), where $\lambda_h > \lambda_l > 0$. The demand state is known neither to the seller nor to buyers. It is commonly known that the good is in high demand (i.e., $\lambda = \lambda_h$) with probability μ_0 at the beginning of the game, and the realization of the demand state is independent of the quality of the seller's good.

All agents are risk-neutral and maximize their expected utility. If a buyer purchases the good at price p at time t and the good is of quality a, then the buyer receives payoff $v_a - p$, while the seller obtains $(1 - e^{-rt})c_a + e^{-rt}p$.

2.2 Strategies and Equilibrium

We consider the following information structure:[2] the seller does not observe the arrival of buyers, while each buyer observes only the price posted at the time of his arrival. The former implies that the seller cannot tell whether the failure of sale is due to no arrival of buyers or due to buyers' refusal to accept the posted price. The latter implies that buyers' beliefs and strategies are independent of their arrival time and, therefore, stationary over time from the seller's viewpoint.[3] Both

[2]It is well-known that in this type of dynamic game, buyers' information about the seller's histories plays a crucial role. See, for example, Swinkels (1999), Taylor (1999), and Hörner and Vieille (2009).

[3]This is a common modeling assumption in the literature on dynamic adverse selection (see, e.g., Lauermann and Wolinsky, 2011; Zhu, 2012; Kim and Pease, 2014). For different approaches, which give rise to non-stationary dynamics, see, for example, Hörner and Vieille (2009) and Kim (2014).

give tractability to the analysis.

Under the information structure, the seller's (pure) offer strategy is a function $p : \{L, H\} \times \mathcal{R}_+ \to \mathcal{R}_+$, where $p(a, t)$ represents the price the type-a seller posts at time t. Buyers' beliefs about the seller's type are a function $q : \mathcal{R}_+ \to [0, 1]$, where $q(p)$ is the probability that buyers assign to the high type when the posted price is p. Similarly, their (mixed) purchase strategies are a function $\sigma_B : \mathcal{R}_+ \to [0, 1]$, where $\sigma_B(p)$ is the probability that buyers accept p.

A tuple (p^*, q^*, σ_B^*) is a (perfect Bayesian) equilibrium of the trading game if the following conditions hold:

- Seller optimality: for each $a = H, L$,

$$p^*(a, \cdot) \in argmax_{p(a,\cdot)} E[(1 - e^{-r\tau})c_a + e^{-r\tau} p(a, \tau)],$$

  where τ is the time at which a buyer purchases the good.

- Buyer optimality: for each $p \in \mathcal{R}_+$, $\sigma_B^*(p) > 0$ only when

$$q^*(p)v_H + (1 - q^*(p))v_L - p \geq 0.$$

- Belief consistency: for each $p \in \mathcal{R}_+$, $q^*(p)$ is obtained from p^* and σ_B^* by Bayes' rule whenever possible.

2.3 Assumptions

We focus on the case where adverse selection is so severe that market inefficiencies are unavoidable. Specifically, we maintain the following assumption, which is common in the literature.

Assumption 1

$$q_0 v_H + (1 - q_0)v_L < c_H. \tag{1}$$

If this assumption is violated, then there is an equilibrium in which both seller types always offer a price in $[c_H, q_0 v_H + (1 - q_0)v_L]$ and all buyers always accept the price, which is an efficient market outcome. Note that, since $v_H > c_H$, the assumption implies that $v_L < c_H$ and q_0 is sufficiently small.

Because of its signaling nature, the game suffers from severe equilibrium multiplicity. To avoid multiplicity, as well as for tractability, we restrict attention to the equilibria of the following structure:

Assumption 2 *The high-type seller always offers a price $p_H \in [c_H, v_H)$.*

In order to support this strategy of the high-type seller, it suffices to assume that buyers believe that all other prices are offered only by the low-type seller. This assumption immediately excludes the trivial equilibria in which the high-type seller always makes a losing offer (above v_H). In addition, it does not allow the high-type seller to adjust her price over time. As explained in the next section, in the absence of demand uncertainty, this incurs no loss of generality in characterizing the set of equilibrium outcomes. For the case of demand uncertainty in Section 4, it plays a crucial role in simplifying the analysis.

Finally, for equilibrium existence, we make the following assumption:

Assumption 3 *Buyers accept v_L with probability* 1.

It is a strictly dominant strategy for a buyer to accept $p < v_L$, because his expected value is bounded below by v_L. Therefore, the seller can ensure trade with the next arriving buyer by posting a price arbitrarily close to v_L. For equilibrium existence, it is necessary that buyers accept v_L even if they assign probability 1 to the low type and, therefore, are indifferent between accepting and rejecting v_L.

Assumptions 2 and 3 imply that there are effectively only two prices, v_L and p_H: the high type always offers p_H, while the low type chooses between v_L and p_H at each point in time. To save notation, in what follows, we describe the low-type seller's offer strategy by a function $\sigma_S : \mathcal{R} \to [0, 1]$, where $\sigma_S(t)$ denotes the probability that the low-type seller offers p_H at time t. In addition, we use σ_B to denote only the probability that each buyer accepts p_H (i.e., $\sigma_B \equiv \sigma_B(p_H)$ from now on).

3 Equilibrium under No Demand Uncertainty

We first characterize the case where there is no demand uncertainty (i.e., $\lambda_h = \lambda_l$). This provides us with a benchmark, which allows us to identify the effects due to demand uncertainty in the next section. In addition, it helps us explain some basic concepts and tools also used in the next section.

3.1 Buyers' Consistent Beliefs

In our model, a buyer's belief about the seller's type departs from his prior belief q_0 for two reasons. First, the very fact that he met the seller provides information about the seller's type. In particular, ceteris paribus, the low-type seller trades relatively faster than the high-type seller, because the former may offer v_L, while the latter insists on p_H. This means that the high type stays relatively longer than the low type, and thus the seller who is still available on the market is more likely to be the high type. We denote by q^I buyers' beliefs at this stage and refer to them as their *interim* beliefs.

Second, the posted price also conveys information about the seller's type. Since buyers' beliefs (and optimal purchase decisions) following v_L are trivial, we focus on their beliefs conditional on p_H. For notational simplicity, we use q^* to denote them and refer to them as buyers' *ex post* beliefs.

Given $\sigma_S(\cdot)$ and σ_B^*, the trading (exit) rate of the low-type seller is equal to $\lambda(\sigma_S(t)\sigma_B^* + 1 - \sigma_S(t))$, while that of the high-type seller is equal to $\lambda\sigma_B^*$. This means that the probability that the low-type seller stays on the market until time t is equal to $e^{-\int_0^t \lambda(\sigma_S(t)\sigma_B^*+1-\sigma_S(t))dt}$, while that of the high-type seller is equal to $e^{-\int_0^t \lambda\sigma_B^* dt}$. Since a seller can be interpreted to be randomly drawn from the space $\{L, H\} \times \mathcal{R}_+$, buyers' interim beliefs can be calculated as follows:

$$q^I = \frac{q_0 \int_0^\infty e^{-\int_0^x \lambda\sigma_B^* dx} dt}{q_0 \int_0^\infty e^{-\int_0^x \lambda\sigma_B^* dx} dt + (1-q_0)\int_0^\infty e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx}},$$

which is equivalent to

$$\frac{q^I}{1-q^I} = \frac{q_0}{1-q_0} \frac{\int_0^\infty e^{-\int_0^t \lambda\sigma_B^* dx} dt}{\int_0^\infty e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}. \tag{2}$$

Conditional on no trade, the low-type seller offers p_H with probability $\sigma_S(t)$ at time t. Combining this with the probability of no trade by time t, derived above, the unconditional probability that the low-type seller offers p_H to a buyer is equal to

$$\frac{\int_0^\infty \sigma_S(t) e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}{\int_0^\infty e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}.$$

Since the high-type seller always offers p_H, it follows that

$$\frac{q^*}{1-q^*} = \frac{q^I}{1-q^I} \frac{1}{\frac{\int_0^\infty \sigma_S(t) e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}{\int_0^\infty e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}}. \tag{3}$$

Combining (2) and (3) yields

$$\frac{q^*}{1-q^*} = \frac{q_0}{1-q_0} \frac{\int_0^\infty e^{-\int_0^t \lambda\sigma_B^* dx} dt}{\int_0^\infty \sigma_S(t) e^{-\int_0^t \lambda(\sigma_S(x)\sigma_B^*+1-\sigma_S(x))dx} dt}. \tag{4}$$

3.2 Equilibrium Characterization

In equilibrium, buyers must randomize between accepting and rejecting p_H (i.e., $\sigma_B^* \in (0, 1)$), while the low-type seller must offer both v_L and p_H with positive probabilities. If buyers always accept p_H, then the low-type seller strictly prefers offering p_H to v_L. Since both seller types play an identical strategy, $q^* = q^I = q_0$. But then, Assumption 1 implies that buyers' expected payoffs are strictly negative. To the contrary, if buyers always reject p_H, then the low-type seller strictly prefers offering v_L to p_H. This implies that buyers believe that the seller who offers p_H is the high type with probability 1 (i.e., $q^* = 1$) and, therefore, accept $p_H (< v_H)$ with probability 1, which is a contradiction. By the same argument, it is also clear that it cannot be that the low-type seller offers only v_L or p_H.

Formally, the following two conditions must be satisfied:

$$q^*(v_H - p_H) + (1 - q^*)(v_L - p_H) = 0 \Leftrightarrow \frac{q^*}{1 - q^*} = \frac{p_H - v_L}{v_H - p_H},$$

and

$$\frac{\lambda \sigma_B^*}{r + \lambda \sigma_B^*} p_H = \frac{\lambda}{r + \lambda} v_L \Leftrightarrow \sigma_B^* = \frac{r v_L}{(r + \lambda) p_H - \lambda v_L}.$$

The first equation corresponds to buyers' indifference between accepting and rejecting p_H, while the second corresponds to the low-type seller's indifference between p_H (left-hand side) and v_L (right-hand side).

Combining these two conditions with (4), the following result is straightforward.

Proposition 1 *For each $p_H \in [c_H, v_H)$, there exists a continuum of payoff-equivalent equilibria. In any equilibrium, $q^* = \frac{p_H - v_L}{v_H - v_L}$ and $\sigma_B^* = \frac{r v_L}{(r+\lambda) p_H - \lambda v_L}$. The low-type seller's strategy $\sigma_S(t)$ supports an equilibrium if and only if it satisfies* (4). *The low-type seller's expected payoff is equal to $\frac{\lambda v_L}{r+\lambda}$, while the high-type seller's expected payoff is equal to $\frac{r c_H + \lambda \sigma_B^* p_H}{r + \lambda \sigma_B^*}$.*

The high-type seller's expected payoff increases in p_H. This is due to the single crossing property: the low-type seller is indifferent among different p_H's. Since the high-type seller has a higher reservation value, she is more willing to trade off a lower probability of acceptance for a higher price.

4 Equilibrium under Demand Uncertainty

In this section, we analyze the main model with demand uncertainty. We characterize the set of equilibria in which the high-type seller offers a single price $p_H \in [c_H, v_H)$ and compare it to the corresponding set under no demand uncertainty.

4.1 Optimal Pricing

In the presence of demand uncertainty, the seller's problem is no longer stationary. Since all buyers play an identical purchase strategy, conditional on the state, she faces a stationary decision problem. However, she learns about the state over time. In particular, sine she does not observe the arrival of buyers, conditional on no trade, she becomes more pessimistic about the demand state and assigns lower probabilities to the high demand state as time goes by. This provides an incentive for the low-type seller to adjust her price over time. In particular, she begins by offering p_H, which yields higher profits to her if accepted, but has a lower acceptance probability. Once she becomes sufficiently pessimistic about the demand state, she switches to v_L, which yields lower profits but ensures fast trade.

Let $\mu(t)$ denote the low-type seller's belief about the demand state (i.e., the probability that she assigns to the high demand state). The evolution of $\mu(t)$ depends on the low-type seller's pricing strategy. Suppose buyers accept p_H with probability σ_B. If the low-type seller offers p_H, then $\mu(t)$ evolves according to

$$\mu(t+dt) = \frac{\mu(t)e^{-\lambda_h \sigma_B dt}}{\mu(t)e^{-\lambda_h \sigma_B dt} + (1-\mu(t))e^{-\lambda_l \sigma_B dt}},$$

which is equivalent to

$$\dot{\mu}(t) = -\mu(t)(1-\mu(t))(\lambda_h - \lambda_l)\sigma_B. \tag{5}$$

Similarly, if the offer is v_L, then

$$\dot{\mu}(t) = -\mu(t)(1-\mu(t))(\lambda_h - \lambda_l).$$

Clearly, the low-type seller's belief decreases faster in the latter case. Intuitively, when the offer is v_L, sale does not occur only when no buyer has arrived, while with price p_H, it may also be because of a buyer's refusal to accept the price. Therefore, the failure of sale is a stronger signal about the low demand state when the seller offers v_L.

The seller's optimal pricing strategy is clear. There is a threshold belief level $\overline{\mu}$ such that the low-type seller offers p_H if and only if her belief exceeds $\overline{\mu}$. Since $\mu(t)$ necessarily decreases over time, this means that there exists a cutoff time at which the low-type seller switches from p_H to v_L, and the low-type seller never reverts back to p_H. Of course, the cutoff time is the point at which the seller's belief is equal to $\mu(t) = \overline{\mu}$.

To explicitly solve for the solution to the low-type seller's pricing problem, let $V(\mu)$ denote her expected payoff when her belief is equal to μ. If $\mu > \overline{\mu}$, then her optimal price is p_H. Therefore, the continuous-time Bellman equation is given as follows:

$$rV(\mu) = (\mu\lambda_h + (1-\mu)\lambda_l)(p_H - V(\mu)) + \dot{V}(\mu).$$

Applying (5),

$$rV(\mu) = (\mu\lambda_h + (1-\mu)\lambda_l)(p_H - V(\mu)) - V'(\mu)\mu(1-\mu)(\lambda_h - \lambda_l)\sigma_B.$$

If $\mu \leq \overline{\mu}$, then the low-type seller offers v_L. Since her belief constantly decreases, she never switches to p_H, as if she commits to offering only v_L. It is then clear that

$$V(\mu) = \frac{(\mu\lambda_h + (1-\mu)\lambda_l)v_L}{r + (\mu\lambda_h + (1-\mu)\lambda_l)}.$$

The solution to this optimal stopping problem is known in the literature. In particular, given buyers' purchase strategies σ_B, the problem is identical to the two-point case in Mason and Välimäki (2011), who study a similar problem in which buyers have private values (implying no inference problem on the seller's type) and the distribution of their values is exogenously given (i.e., in the context of our model, σ_B is exogenously given). We translate their solution into our context and report all the results in the following proposition.

Proposition 2 *Given buyers' purchase strategies σ_B, the low-type seller's optimal pricing strategy is to offer p_H if he assigns a lower probability than $\overline{\mu}$ to the high type and offer v_L otherwise, where*

$$\overline{\mu} = \begin{cases} 1, \ \textit{if} \ \frac{\lambda_h \sigma_B p_H}{r+\lambda_h\sigma_B} \leq \frac{\lambda_h v_L}{r+\lambda_h}, \\ 0, \ \textit{if} \ \frac{\lambda_l \sigma_B p_H}{r+\lambda_l\sigma_B} \geq \frac{\lambda_l v_L}{r+\lambda_l}, \\ \frac{-\lambda_l(r+\lambda_h)(\sigma_B\lambda_l(p_H-v_L)+r(\sigma_B p_H - v_L))}{(\lambda_h-\lambda_l)(r(\sigma_B\lambda_l(p_H-v_L)+r(\sigma_B p_H-v_L))+\sigma_B\lambda_h(p_H-v_L)(r+\lambda_l))}, \ \textit{otherwise}. \end{cases}$$

The low-type seller's expected payoff as a function of her belief is equal to

$$V(\mu) = \begin{cases} \left(\mu\frac{\lambda_h\sigma_B}{r+\lambda_h\sigma_B} + (1-\mu)\frac{\lambda_l\sigma_B}{r+\lambda_l\sigma_B}\right) + C(1-\mu)\left(\frac{1-\mu}{\mu}\right)^{\frac{r+\lambda_L\sigma_B}{(\lambda_H-\lambda_L)\sigma_B}}, & \textit{if} \ \mu > \overline{\mu}, \\ \frac{(\mu\lambda_h+(1-\mu)\lambda_l)}{r+(\mu\lambda_h+(1-\mu)\lambda_l)} v_L, & \textit{if} \ \mu \leq \overline{\mu}, \end{cases}$$

where

$$C = \frac{\overline{\mu}(\lambda_h - \lambda_l)\sigma_B}{\overline{\mu}(\lambda_h - \lambda_l)\sigma_B + \lambda_l\sigma_B}\left(\frac{r(\lambda_h-\lambda_l)\sigma_B p_H}{(r+\lambda_h\sigma_B)(r+\lambda_l\sigma_B)} - \frac{r v_L(\lambda_h - \lambda_l)}{(r+\lambda_h)(r+\lambda_l)}\right)\left(\frac{1-\overline{\mu}}{\overline{\mu}}\right)^{-\frac{r+\lambda_L\sigma_B}{(\lambda_H-\lambda_L)\sigma_B}}.$$

Clearly, $\overline{\mu}$ is decreasing in both σ_B and p_H, while $V(\mu)$ is increasing in them whenever $\mu > \overline{\mu}$: an increase in σ_B or p_H is necessarily beneficial to the low-type seller and increases the low-type seller's incentive to offer p_H, thereby pushing down the cutoff $\overline{\mu}$.

Let $T(\sigma_B)$ denote the length of time it takes for the low-type seller's belief to reach $\overline{\mu}$, assuming that she follows the optimal pricing strategy in Proposition 2. The following result is straightforward from Proposition 2.

Corollary 1 *The low-type seller offers p_H until time $T(\sigma_B)$ and v_L thereafter, where*

$$T(\sigma_B) = \frac{1}{(\lambda_h - \lambda_l)\sigma_B} ln\left(\frac{\mu_0}{1-\mu_0}\frac{1-\overline{\mu}}{\overline{\mu}}\right).$$

$\sigma_B T(\sigma_B)$ continuously increases from 0 to ∞ as σ_B increases from σ_B such that $\overline{\mu}(\sigma_B) = \mu_0$ to $\frac{rv_L}{rp_H + \lambda_l(p_H - v_L)}$ (i.e., the minimal σ_B such that $\overline{\mu}(\sigma_B) = 0$).

4.2 Buyers' Equilibrium Beliefs

We now derive buyers' equilibrium beliefs q^* (the probability that each buyer assigns to the high type conditional on observing p_H). As shown in the previous section, the trading rates of each type affect q^*. Since the trading rates also depend on the demand state, it is necessary to determine buyers' equilibrium beliefs about the demand state, although those beliefs do not directly influence their purchase decisions. We denote by μ^* buyers' (ex post) beliefs about the demand state (conditional on observing p_H). We first take μ^* as given and determine q^*.

Given μ^* and σ_B, buyers' beliefs about the seller's type q^* can be derived as in the previous section. As shown above, given σ_B, the low-type seller's optimal pricing strategy is uniquely determined. Therefore, it suffices to derive q^* that corresponds to the low-type seller's pricing strategies in Proposition 2: the low-type seller offers p_H until time $T(\sigma_B)$ and switches to v_L.

Following the same steps as in the previous section (and skipping the derivation of q^I),

$$\begin{aligned}\frac{q^*}{1-q^*} &= \frac{q_0}{1-q_0}\frac{\mu^*\int_0^\infty e^{-\lambda_h\sigma_B t}dt + (1-\mu^*)\int_0^\infty e^{-\lambda_l\sigma_B t}dt}{\mu^*\int_0^{T(\sigma_B)} e^{-\lambda_h\sigma_B t}dt + (1-\mu^*)\int_0^{T(\sigma_B)} e^{-\lambda_l\sigma_B t}dt} \\ &= \frac{q_0}{1-q_0}\frac{\frac{\mu^*}{\lambda_h\sigma_B} + \frac{1-\mu^*}{\lambda_l\sigma_B}}{\frac{\mu^*(1-e^{-\lambda_h\sigma_B T(\sigma_B)})}{\lambda_h\sigma_B} + \frac{(1-\mu^*)(1-e^{-\lambda_l\sigma_B T(\sigma_B)}))}{\lambda_l\sigma_B}}.\end{aligned} \tag{6}$$

q^* departs from q_0 for three reasons. First, the high-type seller stays on the market relatively longer than the low-type seller (which pushes up q^I above q_0). Second, the high-type seller is more likely to offer p_H than the low-type seller (which pushes up q^* above q^I). Finally, buyers' beliefs about the demand state also influence q^*. This last effect is unclear at this stage, because μ^* is also an endogenous variable.

We now determine μ^*. Similarly to the relationship between q^* and q_0, there are two reasons why μ^* departs from μ_0. First, trade occurs faster in the high demand state than in the low demand state. Therefore, the very fact that the seller is still available (and charges p_H) makes buyers assign a relatively lower probability to λ_h. Second, for a given length of time, the seller meets relatively more buyers in the high demand state than in the low demand state. Therefore, a buyer's arrival is more likely to occur in the high demand state than in the low demand state. This increases the probability that the demand state is high.

Applying similar arguments to those used to derive q^* in the previous section,

$$\begin{aligned} \frac{\mu^*}{1-\mu^*} &= \frac{\mu_0}{1-\mu_0}\frac{\lambda_h}{\lambda_l}\frac{q_0\int_0^\infty e^{-\lambda_h\sigma_B t}dt + (1-q_0)\int_0^{T(\sigma_B)} e^{-\lambda_h\sigma_B t}dt}{q_0\int_0^\infty e^{-\lambda_l\sigma_B t}dt + (1-q_0)\int_0^{T(\sigma_B)} e^{-\lambda_l\sigma_B t}dt} \\ &= \frac{\mu_0}{1-\mu_0}\frac{\lambda_h}{\lambda_l}\frac{\frac{q_0}{\lambda_h\sigma_B} + \frac{(1-q_0)(1-e^{-\lambda_h\sigma_B T(\sigma_B)})}{\lambda_h\sigma_B}}{\frac{q_0}{\lambda_l\sigma_B} + \frac{(1-q_0)(1-e^{-\lambda_l\sigma_B T(\sigma_B)})}{\lambda_l\sigma_B}}. \end{aligned} \tag{7}$$

The second term in the right-hand side captures the second effect in the previous paragraph, while the last term represents the first effect.

Combining (6) and (7) yields the following result.

Lemma 1 *Given σ_B, q^* is uniquely determined by*

$$\frac{q^*}{1-q^*} = \frac{q_0}{1-q_0}\frac{\frac{\mu_0}{\lambda_h}(q_0+(1-q_0)\kappa_h(\sigma_B)) + \frac{1-\mu_0}{\lambda_l}(q_0+(1-q_0)\kappa_l(\sigma_B))}{\frac{\mu_0}{\lambda_h}\kappa_h(\sigma_B)(q_0+(1-q_0)\kappa_h(\sigma_B)) + \frac{1-\mu_0}{\lambda_l}\kappa_l(\sigma_B)(q_0+(1-q_0)\kappa_l(\sigma_B))},$$

where

$$\kappa_a(\sigma_B) = 1 - e^{-\lambda_a\sigma_B T(\sigma_B)}, \textit{ for each } a = l, h.$$

4.3 Equilibrium Characterization

We close the model by characterizing buyers' equilibrium purchase strategies. In equilibrium, as in the model without demand uncertainty, buyers must be indifferent between accepting and rejecting p_H. If they always accept p_H, then the low-type seller prefers offering p_H to v_L, independently of her belief about the demand state. In this case, $q^* = q_0$, but then buyers' expected payoffs become strictly negative, due to Assumption 1. To the contrary, if buyers always reject p_H, then the low-type seller always prefers v_L to p_H. In this case, $q^* = 1$, and thus buyers strictly prefer accepting p_H to rejecting it, which is a contradiction. This gives us the last equilibrium condition:

$$q^*(v_H - p_H) + (1-q^*)(v_L - p_H) = 0 \Leftrightarrow \frac{q^*}{1-q^*} = \frac{p_H - v_L}{v_H - p_H}. \tag{8}$$

Given Proposition 2, Lemma 1, and (8), equilibrium characterization reduces to finding σ_B that satisfies

$$\frac{q_0}{1-q_0}\frac{\frac{\mu_0}{\lambda_h}(q_0+(1-q_0)\kappa_h(\sigma_B))+\frac{1-\mu_0}{\lambda_l}(q_0+(1-q_0)\kappa_l(\sigma_B))}{\frac{\mu_0}{\lambda_h}\kappa_h(\sigma_B)(q_0+(1-q_0)\kappa_h(\sigma_B))+\frac{1-\mu_0}{\lambda_l}\kappa_l(\sigma_B)(q_0+(1-q_0)\kappa_l(\sigma_B))}=\frac{p_H-v_L}{v_H-p_H}. \quad (9)$$

Proposition 3 *For each $p_H \in [c_H, v_H)$, there exists a unique equilibrium in the model with demand uncertainty. In the equilibrium, σ_B^* is such that $T(\sigma_B^*) > 0$ (i.e., $\overline{\mu}(\sigma_B^*) < \mu_0$).*

Proof. Let $\underline{\sigma}_B$ be the value such that $\overline{\mu}(\underline{\sigma}_B) = \mu_0$ (i.e., the maximal value of σ_B such that $T(\sigma_B) = 0$). In addition, let $\overline{\sigma}_B$ be the value such that $\overline{\mu}(\overline{\sigma}_B) = 0$ (i.e., the minimal value of σ_B such that $T(\sigma_B) = \infty$).

We show that equilibrium σ_B^* necessarily lies in $(\underline{\sigma}_B, \overline{\sigma}_B)$. Whenever $\sigma_B \leq \underline{\sigma}_B$, $k_a(\sigma_B) = 0$ for both $a = h, l$. Therefore, the left-hand side is necessarily larger than the right-hand side in (9). To the contrary, if $\sigma_B \geq \overline{\sigma}_B$, then $\kappa_a(\sigma_B) = 1$ for both $a = h, l$. This implies that the left-hand side in (9) reduces to $q_0/(1-q_0)$, which is strictly less than the right-hand side by Assumption 1. Finally, the left-hand side is continuous and strictly decreasing on $(\underline{\sigma}_B, \overline{\sigma}_B)$. Therefore, σ_B that satisfies (9) uniquely exists on $(\underline{\sigma}_B, \overline{\sigma}_B)$. ■

References

Akerlof, George, "The market for "Lemons": quality uncertainty and the market mechanism," *Quarterly Journal of Economics*, 1970, *83* (3), 488–500.

Daley, Brendan and Brett Green, "Waiting for news in the dynamic market for lemons," *Econometrica*, 2012, *80* (4), 1433–1504.

Deneckere, Raymond and Meng-Yu Liang, "Bargaining with interdependent values," *Econometrica*, 2006, *74* (5), 1309–1364.

Evans, Robert, "Sequential bargaining with correlated values," *Review of Economic Studies*, 1989, *56* (4), 499–510.

Gerardi, Dino, Johannes Hörner, and Lucas Maestri, "The role of commitment in bilateral trade," *Journal of Economic Theory*, 2014, *154*, 578–603.

Hörner, Johannes and Nicolas Vieille, "Public vs. private offers in the market for lemons," *Econometrica*, 2009, *77* (1), 29–69.

Janssen, Maarten C. W. and Santanu Roy, "Dynamic trading in a durable good market with asymmetric information," *International Economic Review*, 2002, *43* (1), 257–282.

Kim, Kyungmin, "Information about sellers' past behavior in the market for lemons," *mimeo*, 2014.

— **and Marilyn Pease**, "Costly Search with Adverse Selection: Solicitation Curse vs. Accelerating Blessing," *mimeo*, 2014.

Lauermann, Stephan and Asher Wolinsky, "Search with adverse selection," *mimeo*, 2011.

Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite, "Belief-based refinements in signalling games," *Journal of Economic Theory*, 1993, *60*, 241–276.

Mason, Robin and Juuso Välimäki, "Learning about the arrival of sales," *Journal of Economic Theory*, 2011, *146* (4), 1699–1711.

Swinkels, Jeroen, "Education signalling with preemptive offers," *Review of Economic Studies*, 1999, *66* (4), 949–970.

Taylor, Curtis R., "Time-on-the-market as a sign of quality," *Review of Economic Studies*, 1999, *66* (3), 555–578.

Vincent, Daniel R., "Bargaining with common values," *Journal of Economic Theory*, 1989, *48* (1), 47–62.

— , "Dynamic auctions," *Review of Economic Studies*, 1990, *57* (1), 49–61.

Zhu, Haoxiang, "Finding a good price in opaque over-the-counter markets," *Review of Financial Studies*, 2012, *25* (4), 1255–1285.