Repeated Bargaining with Persistent Private Information

John Kennan[1]
University of Wisconsin-Madison and NBER

April 1997

[1]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706; jkennan@macc.wisc.edu. This paper grew out of a joint research project with Robert Wilson, to whom I am grateful for many stimulating discussions. I also thank Larry Ausubel, Andreas Blume, V. V. Chari, Ray Deneckere, Pat Kehoe, Rody Manuelli, Larry Samuelson and seminar participants at the Federal Reserve Bank of Minneapolis, the NBER Summer Institute, the University of Essex, the Stanford Institute for Theoretical Economics, the University of North Carolina, the University of Michigan and the University of Wisconsin for valuable comments. The National Science Foundation provided research support.

1. **Introduction**

Repeated bargaining relationships are important in many economic contexts. An obvious example is the continuing relationship between a union and an employer, involving periodic negotiation of contracts determining the price and quantity of labor services for a period of a few years. More generally, repeated contracts arise in international trading relationships, and in various intermediate product industries: examples include vineyards selling grapes to wineries under contract, and mining companies selling coal to electric utilities.

The typical situation in these negotiations is surely that each side knows more about its own reservation price than its opponent does. If the reservation prices are correlated across contracts, information that is revealed in the negotiation of one contract has strategic value in subsequent negotiations, and each side must take this into account in choosing an optimal bargaining strategy. This paper analyzes the strategic role of serially correlated private information on one side of a repeated bargaining relationship.

Recent work on labor contracts has emphasized the possibility of explaining collective bargaining outcomes in terms of the truth-telling constraints arising when either party to a bargaining game is endowed with unverifiable private information.[2] Meanwhile the empirical literature on labor contracts has shown convincingly that the outcome of the current negotiation is substantially influenced by what happened when the previous contract was negotiated.[3] At this point, we do not have any model that could be used to explain the regularities found in the labor contract data: that is one motivation for this paper.

A trader who is uncertain about the other side's valuation faces the standard monopoly tradeoff between prices and quantities. In the most basic case there are two choices: a pooling offer that ensures that trade will occur, at a relatively unfavorable price, or a screening offer that ensures a high price at the risk of a failure to trade. When the relationship is repeated one would expect the valuations to have both permanent and transitory components. This complicates the decision on whether to play soft or hard, because the hard strategy reveals more information, and information about present valuations will be valuable in future negotiations. This naturally leads to cycles: a trader who screens this time and loses

---

[2]See, for example, Sobel and Takahashi (1983), Hayes (1984), Fudenberg, Levine and Ruud (1985), Kennan (1986), Hart (1989), Kennan and Wilson (1989, 1993), Card (1990), and Cramton and Tracy (1992)

[3]See, for example, Riddell (1979, 1980), Card (1988, 1990), and Ingram, Metcalf and Wadsworth (1991).

will be pessimistic in subsequent negotiations, but the pooling offers induced by this pessimism do not reveal new information, and so the pessimism wears off.

In order to obtain concrete results it is necessary to make some strong simplifying assumptions. First, there is private information only on one side, the other side's valuation being common knowledge. Second, the bargaining rules offset the informational advantage by giving the uninformed party the right to make offers that the informed party must either accept or reject, so the equilibrium involves screening, with little room for signaling. For concreteness, say an uninformed seller makes take-it-or-leave-it offers to an informed buyer (with the understanding that if these labels are reversed nothing changes except that high prices become low prices, etc). Moreover, the analysis is restricted to the case where the seller can commit to one take-it-or-leave-it offer in each negotiation, although the seller has no commitment power regarding future negotiations. Third, the buyer's valuation is a two-state Markov chain, which is the simplest process that includes both transient and permanent components. At one extreme, no transitions occur, so that the high-valuation buyer is wary of revealing its type, because of the "ratchet effect": once the high valuation is revealed, the seller will claim the entire rent in all future negotiations. This situation was analyzed by Hart and Tirole (1988). At the other extreme, the current valuation is purely transitory, and the model reduces to a sequence of one-shot screening negotiations. Between these extremes the model generates cyclic equilibria. The paper focuses on the strategic complications arising when both parties are forward-looking, so that while the seller thinks about the value of learning the current state of demand in order to choose more profitable prices in the future, the buyer makes a similar calculation from the opposite point of view. The main result characterizes a class of Markov-Perfect equilibria that exhibit cycles.

In some respects, this paper is related to the literature on learning and experimentation in markets. For example, the seller faces a tradeoff between actions that are myopically optimal, and actions that improve future payoffs by revealing information about the state of demand. Aghion, Bolton, Harris and Jullien (1991) focus on whether an agent who faces such a tradeoff will eventually learn all there is to know about the environment. Rustichini and Wolinsky (1995) analyze a model in which a monopoly seller faces a nonstrategic buyer with a rectangular demand curve driven by a symmetric Markov chain. Here the seller decides whether it is worthwhile to learn the buyer's valuation always, by making screening offers, or never, by making pooling offers, or sometimes, by using the information revealed by a previous screening offer to determine whether a screening offer is worthwhile now. This is also the problem analyzed in this paper, but for the more complicated situation in which both the buyer and the seller both behave strategically, so that the seller is trying to learn something that the buyer may wish to conceal. Bergemann and Valimaki (1996) consider a monopsony buyer facing two competing sellers, where the buyer's valuation (i.e. the quality) of each seller's product is a stochastic process. This is a two-armed bandit problem in which the bandits are smart, but there is no private information involved.

In earlier work Blume (1990) and Vincent (1997) analyzed the effects of private information that arrives while a bargaining game is in progress. Blume considered a two-type model where the low type can temporarily assume the valuation of the high type, emphasizing that even if the informed party can only accept or reject offers made by the uninformed party, there is an important signaling aspect of the negotiations. The model in this paper differs from Blume's in two respects: both types change valuations, and the game involves repeated contract negotiations, as opposed to a final sale. In Vincent's model the buyer has a linear demand curve with an intercept driven by a Markov chain and the seller is precluded from using two-part tariffs and must instead set a price and let the buyer choose quantity. In this situation the buyer can signal a low valuation by purchasing a quantity that is less than the myopic optimum, but nevertheless positive, with the result that pooling equilibria are difficult to sustain.

2. **An Infinite Horizon Markov Model of Repeated Negotiations**

Consider an infinite sequence of contract negotiations between a buyer and a seller where the rent to be divided in each T-period contract follows a simple Markov process, with transitions that are observed privately by the buyer. Both sides maximize the present value of expected income, with a common discount factor $\Delta=\delta^T$. For example, if v is the present value of a firm's revenue during the term of a labor contract, net of all nonlabor costs, and if w_0 is the highest wage available to workers during a strike, then the rent is $v-w_0$. Assume that the rent follows a two-state Markov chain with continuation probabilities ρ and σ over the length of a contract, so that if the rent is low now ($v=v_L$), it will again be low T periods hence with probability ρ, and if the rent is high now ($v=v_H$) it will be high again next time with probability σ.[4] It is convenient to use w_0 as the origin and v_H-v_L as the unit, so relabel $\theta = (v_L-w_0)/(v_H-v_L)$ as the low rent, with $1+\theta$ as the high rent and zero as the seller's opportunity cost. Then the model is summarized by the four parameters $(\rho,\sigma,\theta,\Delta)$.

[4]This notation may be taken as a summary of multiple transitions in the Markov chain during the term of the contract: for instance if the chain makes one transition per period, and the basic transition matrix is A_0, then the transition matrix from one contract to the next is $A = A_0^T$, and ρ and σ may be interpreted as the diagonal elements of this matrix. Under this interpretation, v_L and v_H denote expected present values over the life of the contract, given the state in the initial period.

The rules of bargaining are simple: the uninformed party makes an offer, and if this is rejected there is no trade until this contract period expires. Thus the seller has full commitment power within the current contract (but no commitment power across contracts). A more interesting, but more complicated, possibility is to allow one offer per period, so that there might be as many as T offers per negotiation.

In the absence of any historical information, the probability of the low valuation is that implied by the stationary distribution of the Markov chain, i.e.

$$\mu = \frac{1-\sigma}{1-\sigma+1-\rho} = \frac{1-\sigma}{1-\varphi}$$

where $\varphi = \rho+\sigma-1$ measures the degree of persistence in the Markov chain. Let $\rho(s)$ denote the probability that the Markov chain is in the low state after s transitions, given that the current state is low, and let $\sigma(s)$ be the probability of being in the high state after s transitions, given that the current state is high. Then

$$\begin{aligned} \rho(s) &= \mu + \varphi^s[1-\mu] \\ \sigma(s) &= 1-\mu+\varphi^s\mu \end{aligned}$$

The parameter φ governs the extent to which successive contract negotiations are linked. Assume $\varphi \geq 0$, so that the probabilities $\rho(s)$ and $\sigma(s)$ do not oscillate. If $\varphi = 0$ information is completely transitory, so that any inference that the seller might draw from the current contract negotiation will be irrelevant by the time the next contract is negotiated. At the other extreme, if $\varphi = 1$ the current information is entirely permanent. Under the interpretation that φ summarizes multiple transitions during the term of a contract, with $\varphi = \varphi_0^T$, the linkage across contracts is made weaker if the contract length is increased, for a given value of φ_0.

A natural equilibrium of this game is a renewal process based on the outcome of screening offers made by the seller. If the buyer accepts an offer revealing that the rent is currently high, the continuation game is the same as it was the last time such a revelation was made, and similarly if the buyer rejects sufficiently many offers to convince the seller that the rent is currently low.

In each contract negotiation there are two possibilities from the seller's point of view. If information is sufficiently persistent (φ is high) and if the seller has inferred from a recent negotiation that the rent was low, it will be optimal to make a pooling offer. Alternatively, if the seller believes that the high-rent state is sufficiently likely, a screening offer will be worthwhile; this offer will be acceptable to the buyer if the rent is currently high, and unacceptable if the rent is low. Each offer that the seller makes must leave either the high or low buyer type indifferent between acceptance and rejection. If this were not so, the offer could be improved, from the seller's point of view, without changing the buyer's decision. A screening offer is just acceptable to the high type, and unacceptable to the low type. A pooling offer is just acceptable to the low type, and more than acceptable to the high type.

If the buyer accepts a screening offer, the seller will infer that the rent is high, and so the seller will screen again when the next contract is negotiated (unless perpetual pooling is optimal). Of course the buyer knows that acceptance of a screening offer weakens its bargaining position next time, so the offer must be sufficiently generous to compensate for this. If the buyer rejects a screening offer, on the other hand, the seller infers that the rent is currently low, and it may then be optimal to make a pooling offer next time, and perhaps again the time after that, and so on. A key feature of the equilibrium is the number of pooling offers, K-1, made by the seller in the sequence of contracts following rejection of a screening offer.

After rejection of a screening offer the seller concludes that the rent is currently low, and the seller then makes pooling offers in the next K-1 negotiations, followed by a screen in the K^{th} negotiation. If the buyer accepts an offer revealing a high rent now, the seller screens again when the next contract comes up. This is sketched in Figure 1, which represents varying degrees of pessimism for the seller. At one extreme, the seller believes the buyer has a low valuation now with probability ρ, because there was a screen in the previous negotiation and the buyer was found to be poor. In this situation the seller pools now, and pools again in K-1 successive negotiations until the probability of the low type has decayed past the screening threshold ζ^*. At the other extreme, the seller is most optimistic after the buyer was found to be rich in the previous negotiation; then the probability of the low state now is only $1\text{-}\sigma$.

Cyclic equilibria of this kind have been analyzed by Kennan (1995), and by Rustichini and Villamil (1996). A nonstrategic version in which buyers are not forward-looking was analyzed by Rustichini and Wolinsky (1995). The model considered in Kennan (1995) is more general than the model in this paper, in that the seller makes a sequence of offers within each contract negotiation. This leads to complications that precluded a full characterization of equilibrium, although numerical examples were computed. As will be seen below, reducing the problem by allowing only one offer per contract leads to much sharper results, without making things so simple as to be uninteresting. Rustichini and Villamil also assumed one offer per contract, and, in addition, they restricted the Markov chain to be symmetric ($\rho=\sigma$). Their main result was that cyclic equilibria exist if the degree of persistence is sufficiently high (i.e. if φ is close to 1). This is puzzling since Hart and Tirole (1988) had previously shown (for discount factors exceeding .5) that only pooling equilibria survive in the limiting case when the buyer's valuation is permanent (ρ=σ=φ=1). But these results are in fact consistent, because the Rustichini and Villamil equilibrium is a weak Perfect Bayesian equilibrium, in which the seller has strange beliefs off the equilibrium path (beliefs that are not consistent with the approximating sequence of completely mixed strategies and Bayesian beliefs used to claim that the equilibrium is a sequential equilibrium). This paper indicates that once the seller's beliefs are straightened out, cyclic equilibria are incompatible with complete persistence of the high valuation, because a seller who is virtually certain that the buyer's current valuation is high will not change this belief just because a single screening offer has been rejected. Instead, as in the standard static bargaining model with permanent valuations, a long sequence of rejected screening offers is needed to overcome the seller's initial optimism.

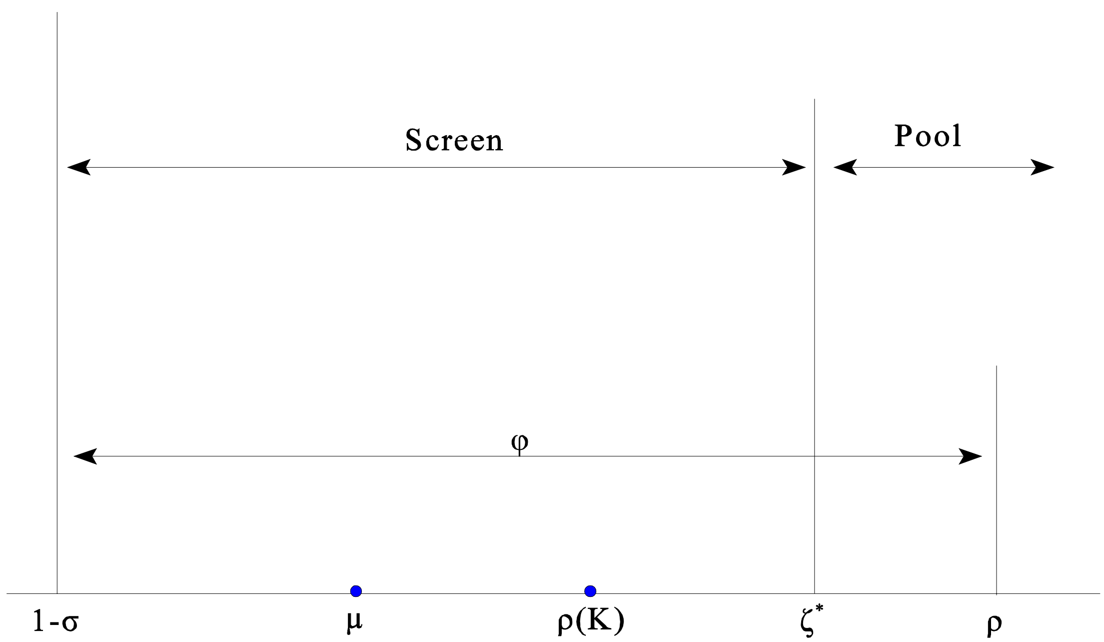

Figure 1: The Screening and Pooling Cycle

Actions and Strategies

The bargaining game is now laid out formally. The buyer's informational advantage is represented by introducing a fictional player called nature whose actions are seen by the buyer but not by the seller.

Nature's Actions

Each period, nature chooses 0 or 1.

The Seller's Actions

Each period, the seller's set of feasible actions (i.e prices) is the real line.

The Buyer's Actions

Each period, the buyer's set of feasible actions (i.e. quantities) is {0,1}.

History

The history of the game, at the beginning of period t, is in three parts:

The sequence of actions chosen by nature is	$n^t = \{n_1,n_2,...n_t\}$
The sequence of actions chosen by the seller is	$p^t = \{p_1,p_2,...p_t\}$
The sequence of actions chosen by the buyer is	$q^t = \{q_1,q_2,...q_t\}$
The seller's history (i.e. the history available to the seller) is	$h_t^0 = \{p^{t-1},q^{t-1}\}$
The buyer's history is	$h_t = \{p^t,q^{t-1},n^t\}$

Feasible Strategies

Nature's strategy is a Markov chain, with transition matrix A given by

$$A = \begin{pmatrix} \rho & 1-\rho \\ 1-\sigma & \sigma \end{pmatrix}$$

That is, $n_t = 1$ with probability $1\text{-}\rho+\varphi n_{t-1}$, where $\varphi = \sigma+\rho-1$, and $n_t = 0$ with probability $\rho\text{-}\varphi n_{t-1}$. It is assumed that φ is nonnegative.

A feasible strategy for the seller in period t is a function from h_t^0 to the set of probability distributions on the real line. The seller's strategy set is the set of sequences of such functions.

A feasible strategy for the buyer in period t is a function from h_t to the interval [0,1] (specifying the probability that $q_t = 1$). The buyer's strategy set is the set of sequences of such functions.

Payoffs

The seller's payoff in period t is $q_t\, p_t$.

The buyer's payoff in period t is $q_t\,(n_t + \theta - p_t)$. That is, the buyer's valuation is high ($v = 1+\theta$) when $n = 1$, and low ($v = \theta$) when $n = 0$.

Both the seller and the buyer maximize the expected present value of future payoffs, using the common discount factor Δ. The buyer's expectations in period t are based directly on n_t, while the seller's expectations are based on beliefs about n_t, represented by a variable ζ_t denoting the probability that $n_t = 0$.

3. Cyclic Pricing

Starting from any point in the game, the expected future payoffs for any given sequence of prices and quantities depend only on the current state of nature, n_t. Following Maskin and Tirole (1994), it is natural to consider equilibria in which the strategies depend on the history of the game only to the extent that it affects the seller's beliefs about n_t. The main results in this paper deal with equilibria that approximately satisfy this criterion. Such equilibria involve cycles in which the seller periodically makes screening offers, with a finite sequence of pooling offers after any offer is rejected. The seller's strategy is based on a belief ζ about the current state of nature, in relation to a threshold belief, labeled ζ^*, such that the seller makes screening offers whenever the probability of the low valuation falls below the threshold, and pooling offers otherwise. The central analytical task is to determine ζ^* from the basic parameters $(\theta,\rho,\sigma,\Delta)$. A cyclic equilibrium exists only if ζ^* lies above the invariant probability μ, so that although the seller might not be optimistic enough to make a screening offer immediately after inferring that the buyer's valuation was low last time, repeated pooling offers eventually lead to screening. If instead ζ^* lies below μ, then in the long run the seller makes only pooling offers in equilibrium: this is discussed in Section 5 below.

The State Variable

The obvious state variable for this game is the pair (n,ζ) consisting of nature's current action, and the seller's best estimate of this action. This yields strategies for the seller that are step functions, and it is more convenient to simply use the steps as the state variable. This is done as follows. Given that the threshold ζ^* lies above the stationary belief μ, let K be the number of transitions needed to drive ζ below ζ^*, starting from a belief that the current state is low. Define the function κ as

$$\kappa(\zeta^*) \; = \; k \;\; if \;\; \rho(k) \le \zeta^* < \rho(k-1) \;\; , \;\; k = 1,2,\ldots,\infty$$

where $\rho(\infty) = \mu$. Then $K = \kappa(\zeta^*)$. For s = 0,1,...,K-1 define

$$\zeta^*(s) \; = \; \mu + \varphi^{-s}(\zeta^* - \mu)$$

Note that $\rho(K) \le \zeta^*(0) = \zeta^* < \rho(K\text{-}1)$, and that

$$\rho(K-s) \ \leq \ \zeta^*(s) \ < \ \rho(K-s-1)$$

In particular, $\zeta^*(K\text{-}2) < \rho \leq \zeta^*(K\text{-}1)$. Define I_0 as the half-open interval $[1\text{-}\sigma,\zeta^*)$, I_s as the open interval $(\zeta^*(s\text{-}1),\zeta^*(s))$, for s = 1,2,...,K-2, and I_{K-1} as the half-open interval $(\zeta^*(K\text{-}2),\rho]$. Then the intervals $\{I_0,I_1,...,I_s,...I_{K-2},I_{K-1}\}$, together with the set of endpoints $\{\zeta^*(s)$: s = 0,1,...,K-2$\}$, partition the interval $[1\text{-}\sigma,\rho]$.[5] This partition will be used to define a state variable, τ, that governs the buyer's and the seller's strategies. Heuristically, τ is the number of pooling offers left before the next screening offer. The interpretation is that τ summarizes a calculation that uses the strategies and the history of the game to keep track of the seller's beliefs. When this calculation puts ζ in the interval I_s, the state is τ = s. When $\zeta = \zeta^*(s)$, the state is τ = s+ξ, for some number ξ in the unit interval, meaning that there is an immediate randomization that sets τ = s+1 with probability ξ, and τ = s with probability 1-ξ. Thus as ζ increases from 1-σ to ρ, τ increases from 0 to K (where K is reached only if $\rho(K)$ coincides with ζ^*, and $\zeta = \rho$, with $\xi = 1$). This construction is illustrated in Figure 2.

[5]More precisely, I_{K-1} is the half-open interval unless $\rho(K)$ coincides with ζ^*; in this case ρ coincides with $\zeta^*(K\text{-}1)$, and I_{K-1} is the open interval $(\rho(2),\rho)$, and the partition includes the K endpoints $\{\rho(K),...\rho\}$.

Beliefs and States

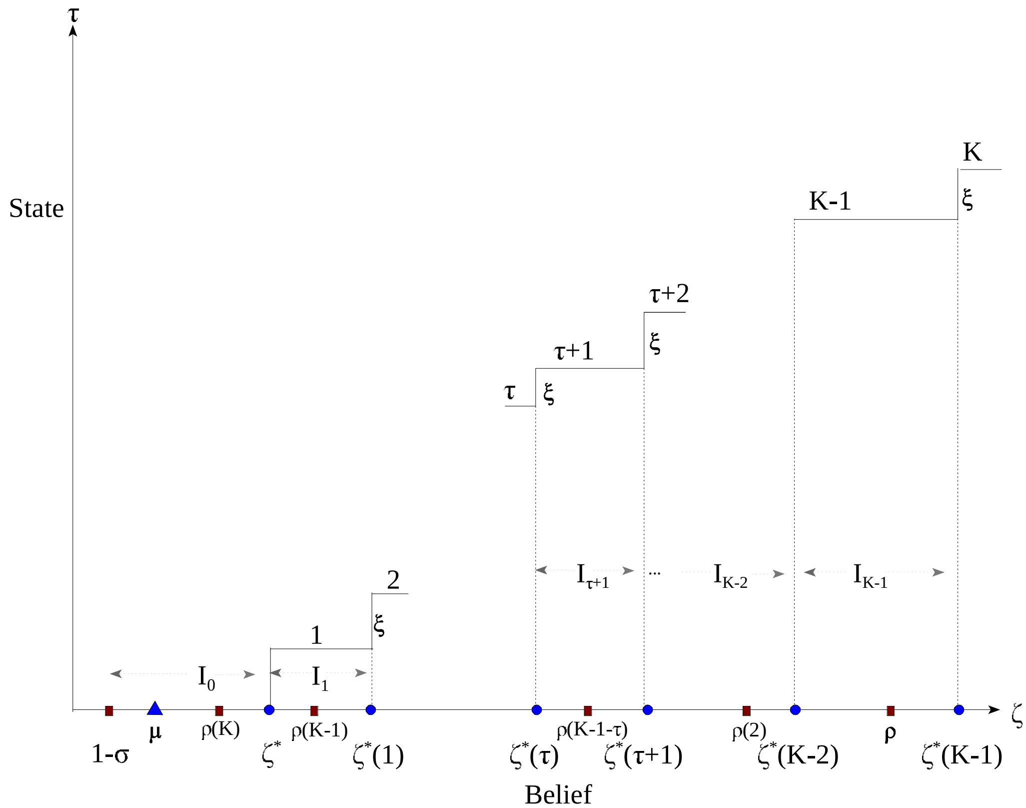

Figure 2

Cyclic Strategies

The strategies used in a cyclic equilibrium will be defined in terms of the buyer's reservation prices, where $P_n(\tau_0,\tau_1)$ denotes the reservation price for a buyer of type n, if a rejected offer implies continuation from τ_0, and an accepted offer implies continuation from τ_1. An explicit representation of the buyer's reservation price function will be derived below, but an abstract definition is enough for the moment.

Definition:

A *cyclic price system* is a pair of functions $P_0(\tau_0,\tau_1)$ and $P_1(\tau_0,\tau_1)$, defined for integer values of τ_0 and τ_1 between 0 and K, with the following properties:

(i) $P_n(\tau_0,\tau_1)$ is decreasing in τ_0, and increasing in τ_1, for n=0,1;

(ii) $0 < P_0(K,K) \le P_1(K,0)$

It will be useful to have an abbreviated notation for some of these prices. First extend the notation by defining $P_n(\tau+\lambda,\tau_1) = (1-\lambda)P_n(\tau,\tau_1) + \lambda P_n(\tau+1,\tau_1)$ for λ in [0,1]. Then for a given cyclic price system define the *screening price* as $p_H = P_1(K-1+\lambda,0)$. Also define the *pooling prices* as $p_L(\tau) = P_0(K-1+\lambda,\tau_1)$, where $\tau_1 = \max(\tau-1,0)$. The relevance of λ here is that there is a nontrivial region of the parameter space in which equilibrium requires a cycle lasting more than K-1 periods but less than K periods, which is achieved by randomization, but in the standard case λ can be set to zero and ignored.

Definition: A *cyclic threshold strategy for the seller* maps the state to the current price in each period, using a cyclic price system, as follows:

if $\tau \ge 1$ then $p = p_L(\tau)$

if $\tau = 0$ then $p = p_H$

Definition: A *cyclic threshold strategy for the buyer* maps the state and the current price to the current quantity in each period, using a cyclic price system, as follows:

For n = 0:

if $p \le p_L(\tau)$ then q = 1 (pooled acceptance)

if $p > p_L(\tau)$ then q = 0 (low buyer rejects above the pooling price)

For n = 1:

if $p \le P_1(K-1,0)$ then q = 1

if $p > P_1(\tau-1,0)$ then q = 0 (pooled rejection if p is very high)

if $P_1(s+1,0) < p \le P_1(s,0)$, where $\tau-1 \le s \le K-2$ then q is random:

q = 1 with probability υ, and q = 0 with probability $1-\upsilon$,

where

$$\upsilon = \frac{1}{1-\hat{\zeta}}\left[1 - \frac{\hat{\zeta}}{\zeta^*(s+1)}\right]$$

and $\hat{\zeta}$ is the buyer's assessment of the seller's current belief, which is computed by applying the seller's belief system to the seller's history. The seller's belief system, which will be spelled out later, implies that $\hat{\zeta} \leq \zeta^*(s+1) \leq 1$, so the probability ν is well-defined.[6]

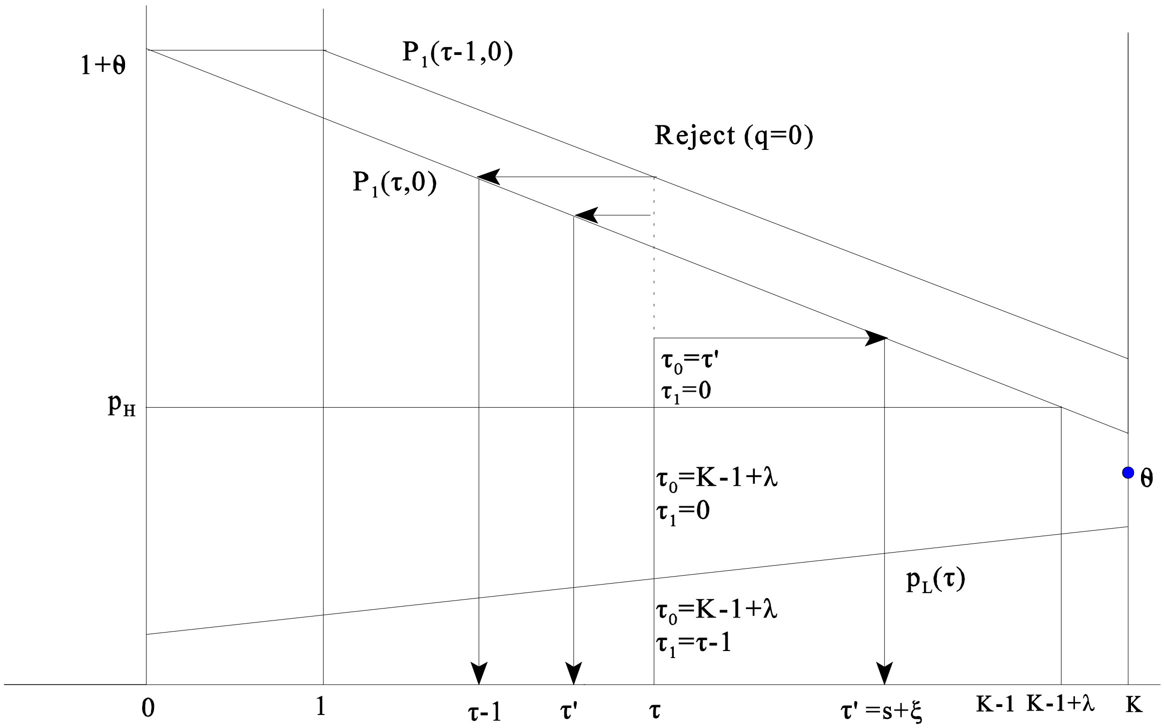

Figure 3: Randomization by the high-valuation Buyer

Figure 3 shows part of the buyer's strategy, covering the situation that arises when the current valuation is high and the seller names a price above the screening price p_H. This requires randomization by the buyer, because if all buyers would reject such a price, rejection would decrement τ by 1, in which case acceptance would be more attractive for the high buyer; but if all high buyers would accept such a

[6]The reader who is keeping track of λ might wonder why prices slightly above p_H are accepted by the high-valuation buyer, and, if they are accepted, why the seller would offer p_H instead of $P_1(K\text{-}1,0)$. The answer is in two parts. First, optimality of the seller's behavior is checked in Section 4 below, where the condition that governs this particular deviation by the seller is labeled P_K. Second, examples can be found in which this condition fails, and in that case an equilibrium can be constructed in which the seller offers $P_1(K\text{-}1,0)$ when $\zeta = 1\text{-}\sigma$.

price, then rejection would restart the pooling sequence, and this would be more attractive to the high buyer.[7]

Definition: A *cyclic updating rule* uses the buyer's actions q to determine the next state τ', as follows:

if $q = 0$, then $\tau' = K-1+\lambda$

if $q = 1$, then $\tau' = \max(\tau-1,0)$

where $0 \leq \lambda \leq 1$, and the notation $\tau' = K-1+\lambda$ means $\tau' = K$ with probability λ, and $\tau' = K-1$ with probability $1-\lambda$.

Cyclic Pricing Processes

A matched pair of cyclic threshold strategies for the seller and the buyer, together with a cyclic updating rule, determine a renewal process for the state variable τ.

Definition: A *cyclic pricing process* is a pair of cyclic pricing strategies (for the seller and the buyer), using the same cyclic price system, together with a cyclic updating rule, such that the strategies and the updating rule all use the same values of K and λ.

A cyclic pricing process gives a complete description of how the game will evolve from any given values of the buyer's valuation n and the state variable τ, if the players follow their prescribed strategies. First the seller's strategy gives p as a function of τ, then the buyer's strategy gives q as a function of (τ,n,p), then the updating rule determines the next state as a function of q, and the process continues. In addition, the updating rule determines the next state if the buyer rejects the current offer, even if the buyer's strategy prescribes acceptance. Note that the updating rule pays no attention to prices, because prices that are consistent with the seller's strategy are rejected (according to the buyer's strategy) if $n = 0$ and $\tau = 0$, and accepted otherwise. Thus the updating rule determines the players' continuation values from

[7] See Maskin and Tirole (1994) for a two-period version of this argument. Rustichini and Villamil (1996) specify that both buyer types reject prices above p_H, but the seller nevertheless infers that the buyer's valuation is low if this situation arises. This is consistent with weak Perfect Bayesian equilibrium only because the situation does not arise on the equilibrium path. Given any completely mixed strategy for the seller that approximates the equilibrium strategy, the correct Bayesian inference is that rejection of a price above p_H conveys no information, since the buyer's equilibrium strategy rejects such prices unconditionally. Thus the Rustichini-Villamil equilibrium is not a sequential equilibrium.

any state, for given cyclic pricing strategies. It also determines the continuation values if the buyer deviates, but the continuation after deviations by the seller is left open for the moment.

Equilibrium Continuation Values

Let $U_L(\tau)$ and $U_H(\tau)$ denote the seller's continuation values from state τ, where U_L applies if the buyer's current valuation is low, and U_H if it is high; and let $g(\tau)$ be the difference between these. Denote the corresponding values for the buyer as $V_L(\tau)$, $V_H(\tau)$ and $d(\tau)$, and the joint continuation values ($J=U+V$) as $J_L(\tau)$, $J_H(\tau)$ and $j(\tau)$. These can all be computed by first determining the joint values, and then the buyer's continuation values, yielding the seller's continuation values as a residual.

Consider first the joint continuation value from state τ when the buyer's current valuation is low. This is the present value of τ periods paying θ if the buyer's valuation is low, and $1+\theta$ if the valuation is high, plus the continuation value from a screening offer after τ periods. Thus

$$J_L(\tau) \;=\; \sum_{s=0}^{\tau-1} \Delta^s\big[\theta+1-\rho(s)\big] \;+\; \Delta^\tau\Big[J_L(0)+[1-\rho(\tau)]\,j(0)\Big]$$

If the state is τ and the buyer's current valuation is high, the joint continuation value is as above, but with $\sigma(s)$ in place of $1\text{-}\rho(s)$. Note that $\sigma(s)+\rho(s)-1 = \varphi^s$. This yields

$$j(\tau) \;=\; 1+\beta j(\tau-1)$$

where β is a discounted persistence parameter defined as $\beta = \Delta\varphi$. The joint continuation values from accepted and rejected screening offers are given by

$$\begin{aligned} J_L(0) \;&=\; (1-\lambda)J_L(K)+\lambda J_L(K+1)-\theta \\ J_H(0) \;&=\; J_L(0)+j(0) \;=\; 1+\theta+\Delta\Big[J_L(0)+\sigma j(0)\Big] \end{aligned}$$

Define the discounted present value sums D, B and R as

$$\begin{aligned} D(\tau) \;&=\; \sum_{s=0}^{\tau-1}\Delta^s \;=\; \frac{1-\Delta^\tau}{1-\Delta} \\ B(\tau) \;&=\; \sum_{s=0}^{\tau-1}\beta^s \;=\; \frac{1-\beta^\tau}{1-\beta} \\ R(\tau) \;&=\; \sum_{s=0}^{\tau-1}\Delta^s\rho(s) \;=\; \sum_{s=0}^{\tau-1}\Delta^s\big[\mu+(1-\mu)\varphi^s\big] \;=\; \mu D(\tau)+(1-\mu)B(\tau) \end{aligned}$$

Then

$$
\begin{aligned}
J_L(\tau) &= D(\tau)(1+\theta) - R(\tau) + \left[1 - D(\tau)(1-\Delta)\right] J_L(0) + \left[(1-\beta)R(\tau) - (1-\Delta\sigma)D(\tau)\right] j(0) \\
&= J_L(0) + R(\tau)\left[(1-\beta)j(0) - 1\right]
\end{aligned}
$$

where the second equality uses the above equation for $J_H(0)$. Then the equations for $J_L(0)$ and $J_H(0)$ can be written as

$$
\begin{aligned}
J_L(0) &= J_L(K+\lambda) - \theta = J_L(0) + R(K+\lambda)\left[(1-\beta)j(0) - 1\right] - \theta \\
(1-\Delta)J_L(0) &= 1 + \theta - (1-\Delta\sigma)j(0)
\end{aligned}
$$

and these can be solved for $J_L(0)$ and j(0):

$$
\begin{aligned}
j(0) &= \frac{1+r\theta}{1-\beta} \\
J_L(0) &= \frac{1+\theta-(1-\Delta\sigma)\frac{1+r\theta}{1-\beta}}{1-\Delta}
\end{aligned}
$$

where r = 1/R(K+λ). Finally,

$$
\begin{aligned}
j(\tau) &= B(\tau) + \beta^{\tau} j(0) \\
&= j(0) - B(\tau) r\theta
\end{aligned}
$$

This gives a complete description of how the joint continuation values are determined from the basic parameters (θ,Δ,ρ,σ), for arbitrary values of K and λ. The derivation uses the strategies only to verify that the equilibrium is a cycle of pooling and screening offers: accepted screening offers are repeated until there is a rejection, then there is a sequence of K-1+λ accepted pooling offers, followed by a sequence of screening offers and so on. The joint values are summarized in the top panel of Table 1.

Table 1: Value Functions

	Low	High	Difference
Joint	$J_L(\tau) = \dfrac{1+\theta-[1-\Delta\sigma]\,j(0)}{1-\Delta} + R(\tau)r\theta$	$J_H(\tau) = \dfrac{1+\theta-\Delta(1-\sigma)j(0)}{1-\Delta} + [R(\tau)-B(\tau)]r\theta$	$j(0) = \dfrac{1+r\theta}{1-\beta}$ $j(\tau) = j(0) - r\theta B(\tau)$
Buyer	$V_L = \dfrac{1-\rho}{(1-\Delta)\varphi}\,d(0)$	$V_H(\tau) = \dfrac{\sigma-\beta}{(1-\Delta)\varphi}\,d(0) + B(\tau)b$	$d(0) = \dfrac{1-b}{1-\beta}$ $d(\tau) = d(0) + bB(\tau)$
Seller	$U_L(\tau) = \dfrac{1+\theta-(1-\Delta\sigma)g(0) - \frac{\sigma}{\varphi}[1-b]}{1-\Delta} + R(\tau)r\theta$ $U(\zeta,\tau) = \dfrac{1+\theta}{1-\Delta} - \dfrac{\sigma(1-b)}{\varphi(1-\Delta)} - \dfrac{\Delta(1-\sigma)}{1-\Delta}g(0) - \zeta g(0) + R(\tau)r\theta - (1-\zeta)B(\tau)[b+r\theta]$	$U_H(\tau) = \dfrac{1+\theta-\Delta(1-\sigma)g(0) - \frac{\sigma}{\varphi}[1-b]}{1-\Delta} + R(\tau)r\theta - B(\tau)[b+r\theta]$	$g(0) = \dfrac{b+r\theta}{1-\beta}$ $g(\tau) = g(0) - [b+r\theta]B(\tau)$
Notation	$R(K,\lambda) = \mu D(K,\lambda) + (1-\mu)B(K,\lambda)$ $D(K,\lambda) = \dfrac{1-\Delta^K(1-\lambda+\lambda\Delta)}{1-\Delta}$ $\quad B(K,\lambda) = \dfrac{1-\beta^K(1-\lambda+\lambda\beta)}{1-\beta}$ $\mu=(1-\sigma)/(1-\varphi)$ $\quad b=1/B(K,\lambda)$, $r=1/R(K,\lambda)$, $B(\tau)=B(\tau,0)$, $R(\tau)=R(\tau,0)$		

Cyclic Prices, and Continuation Values for the Buyer

The next step is to consider prices that are *tight* in the sense that the low buyer is indifferent between accepting or rejecting each pooling offer, and the high buyer is indifferent between accepting and rejecting each screening offer.

Definition: a cyclic pricing process is *tight* if the buyer's continuation value does not depend on q, when either n = 1 and $\tau = 0$, or n = 0 and $\tau > 0$.

If the price system is tight, the differences $d(\tau)$ between the high and low buyer's continuation values from state τ can be computed without knowing the prices, by considering what happens when screening offers are rejected and pooling offers are accepted, by both buyer types. The low buyer's continuation values can be computed as the value of a strategy that sets q = n. These continuation values can then be used to determine tight prices. The results of this calculation are given in the following lemma (the proof is in the Appendix, along with other proofs).

Lemma B1:

Suppose that the strategies constitute a tight cyclic price system. The buyer's continuation values satisfy

$$V_L(\tau) = \frac{1-\rho}{\varphi}\,\frac{1-b}{(1-\Delta)(1-\beta)}$$

$$d(\tau) = \frac{1-\beta^{\tau} b}{1-\beta} = d(0) + B(\tau)\,b$$

where b = 1/B(K+λ).

The low buyer's continuation value does not depend on τ, because any rejected offer implies τ'=K-1+λ next time, and tight pricing means that the low buyer's value is the value of rejecting. In any case, the low buyer is not worth anything unless Δ is positive, and there is some chance of making a transition to the high valuation in the future.

Another noteworthy feature of the buyer's value function is that changes in θ have no effect, for given values of K and λ: changes in θ are absorbed entirely by the seller. This is a local result, however, since although K is fixed with respect to marginal changes in θ, the screening threshold is not, as will be shown below.

Prices

The buyer's value function immediately yields an explicit version of the reservation price function used to construct the equilibrium strategies. Recall that $P_n(\tau_0,\tau_1)$ represents the highest price that a buyer of type n would accept, if rejection implies continuation from τ_0, and acceptance implies continuation from τ_1. Thus

$$\Delta[V_L(\tau_0) + (1-\rho)d(\tau_0)] + n\beta d(\tau_0) = n + \theta - P_n(\tau_0,\tau_1) + \Delta[V_L(\tau_1) + (1-\rho)d(\tau_1)] + n\beta d(\tau_1)$$

where the left side is the value of rejecting, and the right side is the value of accepting. So

$$P_n(\tau_0,\tau_1) = n + \theta - [\Delta(1-\rho) + n\beta]\, b\,[B(\tau_0) - B(\tau_1)]$$

This formula satisfies the monotonicity requirements used in defining a cyclic price system.[8] The screening price is

$$p_H = 1 + \theta - \Delta\sigma \frac{B(K-1+\lambda)}{B(K+\lambda)}$$

Thus the screening price is the high valuation if K=1 and λ is zero, but in general the screening price lies below this level.

The partial screening prices used in the buyer's strategy are given by

$$P_1(\tau,0) = 1 + \theta - \Delta\sigma \frac{B(\tau)}{B(K+\lambda)}$$

The pooling prices are determined by the difference, as seen by the low buyer type, between continuation from τ_0 following rejection, or from τ-1 following acceptance. Thus

$$p_L(\tau) = \theta - \Delta(1-\rho) \frac{B(K-1+\lambda) - B(\tau-1)}{B(K+\lambda)}$$

---

[8]Note that $P_0(K,K) = \theta$, and $P_1(K,0) = \theta + 1 - \Delta\sigma\frac{B(K)}{B(K+\lambda)} \geq \theta$.

with the convention that B(-1) = 0.

An interesting aspect of the price system is that the pooling prices are below the low valuation θ. This is because the buyer has the option of rejecting any offer, and rejection means that the pooling sequence will be re-started. This option is worth something because even if the buyer's current valuation is low there is the prospect of making transitions to the high valuation while the pooling sequence is in progress. The result is that the pooling price has to be smaller than θ in order to cover the option value associated with restarting the pooling sequence. Moreover, for $\tau > 2$, the surplus yielded to the low buyer type increases as the pooling sequence comes closer to the end, because the buyer is more tempted to push the restart button when screening is imminent.

The Seller's Beliefs

The seller's belief system is a sequence of functions, $\{\Psi_t\}$, that determine the probability of the low state as a function of the history of the game. The state variable τ, which governs the buyer's and the seller's strategies, implicitly uses Ψ_t to locate the seller's current belief inside an interval I_S, or at an endpoint $\zeta^*(S)$. The belief system and the state are defined sequentially. In the first period there is no history, ζ_0 is the seller's prior belief that the buyer's initial valuation is low, and $\tau = S$, where $\zeta_0 \,\varepsilon\, I_S$ or $\zeta_0 = \zeta^*(S)$. Thereafter if $\zeta = \Psi_{t-1}(h_{t-1})$ summarizes the history of the game as of the beginning of period t, and τ is the state, then the corresponding values at the beginning of period t+1 will be determined from ζ, τ and the price and quantity in period t, according to the following definition:

Definition:

A *cyclic belief system* uses a cyclic price system and the current actions (p,q) to map the current values of ζ and τ to new values ζ' and τ', as follows:

(a) When q= 0:

if $p \le p_H$	then $\zeta' = \rho$,	and $\tau' = K-1+\lambda$
if $p = P_1(s+\xi,0)$, $\tau-1 \le s+\xi < K-1+\lambda$,	then $\zeta' = \zeta^*(s)$,	and $\tau' = s+\xi$[9]
if $p > P_1(\tau-1,0)$	then $\zeta' = \varphi\zeta+(1-\varphi)\mu$,	and $\tau' = \max(\tau-1,-1)$

[9]Here s is an integer, and $0 \le \xi \le 1$.

(b) When q = 1:

if $p \leq p_L(\tau)$ then $\zeta' = \varphi\zeta+(1-\varphi)\mu$ and $\tau' = \max(\tau-1,0)$

if $p > p_L(\tau)$ then $\zeta' = 1-\sigma$, and $\tau' = 0$

The update when the buyer rejects a price above p_H is illustrated in Figure 3 above.

The Seller's Continuation Values

A seller using a cyclic belief system infers that the buyer's current valuation is high if a screening offer p_H is accepted, and infers the low valuation if p_H is rejected. Thus the seller's and the buyer's expectations are identical after the buyer has responded to a screening offer, and the seller's continuation values can be obtained as the difference between the joint values and the buyer's values.[10] The results of these calculations are given in Table 1. Before the buyer responds to an offer the seller does not know the current valuation, so the seller's expected continuation value at the point of making a screening offer is a weighted average of the state-contingent values in Table 1, with weights determined by ζ_{t-1}.

4. Equilibrium

So far, everything rests on the presumption that there is a threshold belief ζ^* governing the seller's choice between screening and pooling offers. This threshold determines K, the length of the pooling cycle, which in turn determines the continuation values for the buyer and the seller. In particular, the seller's payoffs from screening and pooling are ultimately determined by the value of ζ^* used in the buyer's and seller's strategies. So there must be a fixed point: using ζ^* to determine the strategies, and computing the seller's payoffs from screening and pooling as the seller's belief ζ varies, it must be that screening and pooling yield the same continuation value for the seller when $\zeta = \zeta^*$.

Let $u[p,n,\tau]$ be the seller's expected continuation value if the price p is offered when the state is τ and the buyer's valuation is n. For any τ, the value of a screening offer is

$$u[p_H,n,\tau] = U_L(0) + n g(0)$$

For $\tau \geq 1$, the value of a pooling offer is

[10] Alternatively, the seller's value function can be computed directly from the equilibrium pooling and screening prices, using the cyclic updating rule, following the method used above to obtain the joint values. This alternative procedure must be used in the more general setting where the buyer and the seller have different discount factors (e.g. the extreme case where one side of the market is nonstrategic), because in this case the joint values are assessed differently by the two players.

$$u[p_L(\tau), n, \tau] \;=\; U_L(\tau) + n\,g(\tau)$$

Thus for $\tau \geq 1$, the value of n that would make the seller indifferent between pooling and screening is given by

$$n_\tau^* \;=\; \frac{U_L(\tau) - U_L(0)}{g(0) - g(\tau)}$$
$$\;=\; \frac{R(\tau)\,r\theta}{B(\tau)\,[\,b + r\theta\,]}$$

Recall that a cyclic pricing process governed by (ζ^*,λ) includes an updating rule that determines the continuation from any given values of τ and q. This rule implies that if the seller were to make a pooling offer when $\tau = 0$, the continuation would be identical to the continuation from a pooling offer at $\tau = 1$, namely $\tau' = 0$ if the offer is accepted, and $\tau' = K-1+\lambda$ if it is rejected. Thus the seller would be indifferent between pooling and screening if $n = n_1^*$, so the screening threshold is $1-n_1^*$, which can be written as

$$\zeta^* \;=\; \frac{1}{1 + \frac{\theta}{G(K+\lambda)}}$$

where the increasing function G(k) is defined as R(k)/B(k). These computations are summarized in Figure 4. Note that $n_\tau^* \geq n_1^*$, since $R(\tau) \geq B(\tau)$, implying that when $\zeta \geq \zeta^*$, the seller's equilibrium continuation value in state τ exceeds the value of deviating to a screening offer.

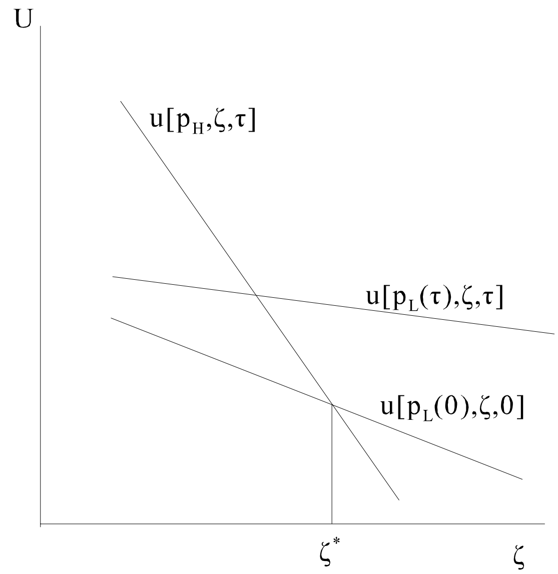

Figure 4: The Screening Threshold

Mapping the Parameters to the Equilibrium Strategies

The above formula does not give a closed-form expression for ζ^*, because K depends on ζ^*, and λ is unspecified. Moreover the formula is valid only if $\zeta^* \geq \mu$. A closed-form expression can be constructed, however. Define the function Z mapping the basic parameters $(\theta,\rho,\sigma,\Delta)$ to the (possibly randomized) screening threshold (ζ^*,λ) as follows

$$(\zeta^*,\lambda) = \begin{cases} \left(\dfrac{1}{1+\frac{\theta}{G(k)}},0\right) & \text{if } G(k)\bar{\rho}(k-1) \leq \theta \leq G(k)\bar{\rho}(k) \\ \left(\rho(k),\lambda_k\right) & \text{if } \quad G(k)\bar{\rho}(k) \leq \theta \leq G(k+1)\bar{\rho}(k) < \theta_\infty \end{cases}, \quad k = 1,2,3, \ldots$$

where

$$\bar{\rho}(k) = \frac{1-\rho(k)}{\rho(k)}$$

$$G(k) = \frac{R(k)}{B(k)} = 1 + \mu\left[\frac{D(k)}{B(k)} - 1\right]$$

$$\theta_\infty = \frac{1-\Delta\sigma}{1-\Delta}\,\frac{1-\mu}{\mu}$$

and λ_k is defined as the solution of the equation $G(k+\lambda_k)\,\bar{\rho}(k) = \theta$, i.e.

$$\lambda_k = \frac{B(k)\left[\theta - G(k)\bar{\rho}(k)\right]}{\Delta^k\left[1-\rho(k)\right]-\beta^k\theta}$$

Each component of the function Z takes values in the unit interval. The ζ^* component is continuous and monotonically decreasing in θ (although it generally has an infinite number of flat segments). The λ component increases from 0 to 1 as θ increases from $G(k)\bar{\rho}(k)$ to $G(k+1)\bar{\rho}(k)$. For $\Delta = 0$ the function collapses to $\zeta^* = 1/(1+\theta)$, with $\lambda=0$. Finally, the domain of the function is restricted by the condition $\theta \le \theta_\infty$, which ensures that $\zeta^* \ge \mu$. If θ exceeds this bound there is no cyclic equilibrium, because screening is too expensive: this is discussed in the next section.

Definition: An *optimal screening cycle* is a pair of cyclic threshold strategies, governed by the same values of ζ^* and λ, such that

(i)
$$(\zeta^*,\lambda) = Z(\theta,\rho,\sigma,\Delta)$$

and

(ii)
$$\frac{1+\theta}{\zeta^*(s)} - \frac{\theta[R(s)-B(s)]}{R(K+\lambda)\zeta^*(s)} - \frac{B(s)}{B(K+\lambda)\zeta^*} - \frac{\sigma}{(1-\sigma)\varphi}\left[1-\frac{B(s)}{B(K+\lambda)}\right] \ge 0 \quad (\mathbf{P}_s)$$

for $1 \le s \le K$, where $K = \kappa(\zeta^*)$ and $\zeta^*(K) = 1$.[11]

Condition (i) immediately implies that if an optimal screening cycle exists, it is unique. The role of the inequalities in (P_s) is to ensure that the seller cannot gain by setting a price above the screening price, that is, by deviating to a partial screening offer. This is discussed further below.

[11]If $\zeta^* = \rho(K)$, then $\zeta^*(K) = 1$, and otherwise $\zeta^*(K) = 1$ just a convention.

Examples

A basic example of an optimal screening cycle of length 2 is laid out in Table 2. When K = 2 and $\lambda = 0$, the only condition that needs to be checked is P_1 (since P_K holds with equality when λ=0). This condition indeed holds in the example (although it fails for slightly smaller values of θ, such as 4/9). The equilibrium price-quantity pairs are generated by a Markov chain with three states: (255,1), (768,0) and (768,1). The low price is always followed by the high price, and the high price is repeated if q = 1, and otherwise the high price is followed by the low price.

Table 2: An Optimal Screening Cycle

Parameters									Equilibrium	
$\theta = ½$, $\Delta = ⅛$, $\rho = ¾$, $\sigma = ¾$, $\mu = ½$, $\varphi = ½$									$\zeta^* = \frac{35}{52}$, λ=0, K = 2; $\mathbf{P}_1$: $\frac{281}{13090} > 0$	
	Continuation Values[a]									
	Joint		Buyer		Seller				Beliefs	Prices
State[b]	$J_L(\tau)$	$J_H(\tau)$	V_L	$V_H(\tau)$	$U_L(\tau)$	$U_H(\tau)$	$U(\tau)$		$\zeta(\tau)$	p_H, $p_L(1)$
τ=0	6,528	103,632	2,240	6,160	4,288	97,472	74,176		$\frac{2}{8}$	$\frac{768}{544}$
τ=0	6,528	103,632	2,240	6,160	4,288	97,472	39,232		$\frac{5}{8}$	$\frac{768}{544}$
τ=1	35,088	103,632	2,240	64,960	32,848	38,672	34,304		$\frac{6}{8}$	$\frac{255}{544}$

[a]Continuation values are scaled up by the factor 62,475 (in order to provide exact results that can easily be compared across states).
[b]The two rows with $\tau = 0$ differ only in the seller's beliefs: the first is reached after a successful screening offer, and the second is reached after a pooling offer.

Random Screening

An equilibrium with random screening is shown in Figure 5, which plots the function Z over the relevant range, with parameter values $(\theta,\rho,\sigma,\Delta) = (5/34,41/44,7/22,1/2)$. The plot shows the piecewise-linear function

$\frac{1}{\theta} = \frac{1}{G(K)}\frac{\zeta^*}{1-\zeta^*}$. This function jumps at an inconvenient spot, which is why randomization by the seller is needed in equilibrium. The equilibrium screening threshold in this example is $\zeta^* = \rho(3) = 641/704$, with $\lambda = .10889$. That is, any rejected offer is followed by two pooling offers, and then another pooling offer with probability λ or a screening offer with probability $1-\lambda$. The example satisfies the conditions of Proposition E1 below (P_1 evaluates to .0656, and P_{K-1} evaluates to 79/1320 < 120/1320). The screening price is $p_H = .988214$, and the first partial screening price is $P_1(K-1,0) = .988244$.

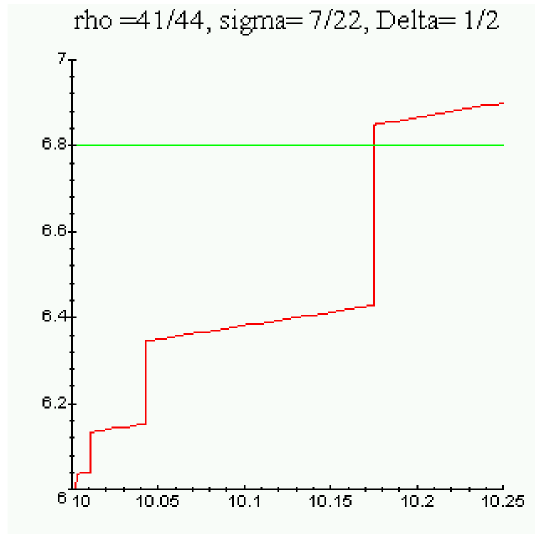

Figure 5: Random Screening Equilibrium

An Example Showing the Effects of Limited Commitment

A surprising feature of the equilibrium in some examples is that the presence of the high buyer type makes the seller worse off *ex ante*. A seller who could commit to pooling offers in every period would be worth $\theta/(1-\Delta)$, while in an optimal screening cycle the seller may be worth much less. An extreme example is shown in Table 3, using a discount factor close to 1. This may be interpreted as a limiting case in which both the seller's offers and the Markov chain transitions are made in rapid succession, with the result that the degree of persistence is negligible in real time.

Table 3: An Example of the Effect of Limited Commitment									
Parameters							Equilibrium		
$\theta = 8/5$, $\Delta = .9999$, $\rho = 3/4$, $\sigma = 3/4$, $\mu = 1/2$, $\varphi = 1/2$							$\zeta^* = .528$, $\lambda=0$, $K=5$ $\mathbf{P}_1$: $.795 > 0$ $\mathbf{P}_{K\text{-}1}$: $\frac{11}{32} < \frac{16}{32}$		
	Continuation Values								
	Joint		Buyer		Seller			Beliefs[b]	Prices
State	$J_L(\tau)$	$J_H(\tau)$	V_L	$V_H(\tau)$	$U_L(\tau)$	$U_H(\tau)$	$U(\tau)$	$\zeta(\tau)$	p_H, $p_L(\tau)$
$\tau=0$	18,691.85	18,694.77	4,837.79	4,838.76	13,854.06	13,856.01	13,855.52	16	1.87
$\tau=0$	18,691.85	18,694.77	4,837.79	4,838.76	13,854.06	13,856.01	13,855.00	33	1.87
$\tau=1$	18,692.31	18,694.77	4,837.79	4,839.28	13,854.52	13,855.49	13,854.97	34	1.13
$\tau=2$	18,692.65	18,694.88	4,837.79	4,839.53	13,854.86	13,855.35	13,855.07	36	1.38
$\tau=3$	18,692.94	18,695.06	4,837.79	4,839.66	13,855.15	13,855.40	13,855.24	40	1.51
$\tau=4$	18,693.20	18,695.26	4,837.79	4,839.73	13,855.41	13,855.53	13,855.44	48	1.57
[b]Probabilities are scaled up by the factor 64									

Coase (1972) conjectured that a seller who could not commit to hold the line on prices for any appreciable length of time would be forced to sell at the lowest demand price in the market. Here the seller does even worse than in the Coase conjecture: the highest value achieved by the seller is 13,856.01, as compared with 16,000 for a seller who faces the low-valuation buyer for sure.[12]

Existence of Equilibrium

Existence of an optimal screening cycle is established by the following two propositions, for a non-trivial region of the parameter space. The first result provides a simplified test of whether an optimal screening cycle exists for given parameter values.

[12]Similar results can be shown for moderate values of the discount factor (e.g. $\Delta = ½$, with $\theta = 1$, and $\rho = \sigma = 5/8$).

Proposition E1

Suppose $(\zeta^*,\lambda) = Z(\theta,\rho,\sigma,\Delta)$ and the following two conditions hold

$$\frac{1+\theta}{\zeta^*(1)} - \frac{b}{\zeta^*} - \frac{\sigma(1-b)}{(1-\sigma)\varphi} \geq 0 \qquad (\mathbf{P_1})$$

and

$$\varphi^{K-1}\left[\frac{\sigma}{(1-\sigma)\varphi} - \mu\right] \leq 1-\mu \qquad (\mathbf{P_{K-1}})$$

Then the cyclic screening strategies governed by (ζ^*,λ) constitute an optimal screening cycle. Moreover, if $\theta' \varepsilon (\theta,\theta_\infty)$, an optimal screening cycle also exists at $Z(\theta',\rho,\sigma,\Delta)$.

Proposition E2

Suppose the parameters ρ, σ and Δ satisfy the inequality

$$\frac{1-\rho}{\Delta(1-\sigma)} + \frac{1-\rho}{1-\Delta} > \frac{\sigma}{1-\varphi} - \varphi \qquad (\mathbf{P_\infty})$$

Then there is a number θ_0 such that for any $\theta \ \varepsilon \ (\theta_0,\theta_\infty)$ an optimal screening cycle exists.

Propositions E1 and E2 are proved in the appendix.[13] Note that the inequality P_{K-1} must hold if K is large, but it cannot hold unless $K \geq 3$. In fact, if K = 2 the inequality can be written as

$$\frac{\sigma}{1-\sigma} \leq 1 - \mu(1-\varphi) = \sigma$$

which is impossible if σ is positive.

The main result of the paper is the following:

Theorem 1:

---

[13]The labels on the inequalities in Propositions E1 and E2 do not quite match the label P_s used in the definition of an optimal screening cycle: although P_1 is just a rearrangement of P_s with s=1, P_{K-1} is stronger than P_s with s=K-1, and P_∞ is a rearrangement of P_s with s=1 and K=∞.

An optimal screening cycle is a sequential equilibrium.

The proof is given in a series of Lemmas below. The function Z defines the screening threshold ζ^* and the randomization probability λ for any given parameter vector $(\theta,\rho,\sigma,\Delta)$. The buyer's and seller's cyclic screening strategies are fully determined by $(\theta,\rho,\sigma,\Delta)$ and (ζ^*,λ), so there is a well-defined mapping from the basic parameters to the equilibrium strategies. The updating rule is pinned down by the assumption that the continuation after a rejected pooling offer is the same as the continuation after a rejected screening offer. The main difficulty in establishing equilibrium is that the seller's strategy (constructed in the manner just described) is generally not optimal if the inequality P_∞ fails, or if θ is small.

Although an optimal screening cycle is not a Markov equilibrium with respect to payoff-relevant information, it is nearly so. The only defect is that if the seller deviates to a price above the screening price, the continuation depends on what this deviant price was, until the slate is wiped clean at the end of the current pooling cycle. This kind of defect seems inevitable in bargaining models with a finite number of types.[14]

Optimality of the Buyer's Strategy

Given the current state τ and the current price p, the buyer's strategy uses the cyclic belief system to predict the next state τ_q, meaning τ_0 following rejection and τ_1 following acceptance of the seller's offer (where τ_q may be random). Then p is compared with the reservation price $P_n(\tau_0,\tau_1)$. It is easy to show, using the buyer's equilibrium continuation values, that no profitable (one-period) deviation is available.

Lemma B2:

Given the cyclic belief system, the buyer's strategy is optimal.

Proof:

Let $V_{nq}(\tau)$ be the continuation value of a type-n buyer choosing q=0 or q=1. Then

$$V_{nq}(\tau) \;=\; q[n+\theta-p] + \Delta E_\tau\Big[V_L(\tau_q) + (1-\rho)\,d(\tau_q) + n\varphi\, d(\tau_q)\Big]$$

[14]See the discussion in Maskin and Tirole (1994), for example.

Given τ_0 and τ_1, $V_{n0}(\tau)$ does not depend on p, while $V_{n1}(\tau)$ is decreasing in p, so a reservation price strategy is optimal. All that needs to be shown is that the buyer's strategy uses the correct prediction of τ_0 and τ_1, for any given values of τ and p. This is true by construction.

Optimality of the Seller's Strategy

According to the buyer's strategy, all prices below the current pooling price $p_L(\tau)$ are acceptable, and according to the seller's belief system, all such acceptances are uninformative, and thus imply the same stochastic process for future payoffs. Therefore all prices below $p_L(\tau)$ are dominated by $p_L(\tau)$ from the seller's point of view. At the other extreme, all prices above $P_1(\tau\text{-}1,0)$ are rejected, with no effect on the seller's beliefs, so these prices are dominated by the pooling price, which yields the same future payoffs, plus some current profit. Any price in the interval $(p_L(\tau),p_H]$ is rejected by the low buyer and accepted by the high buyer. Again, these prices imply the same future payoffs for the seller, so the prices in this interval are all dominated by p_H.

Thus the only relevant one-period deviations for the seller involve exchanging the pooling price $p_L(\tau)$ and the screening price p_H, or else charging a price above p_H and below $P_1(\tau\text{-}1,0)$. If any profitable deviation existed it would be dominated by one of these, so if none of these is profitable the seller's strategy is optimal.

In equilibrium, the seller's belief ζ falls below ζ^* when $\tau = 0$, and it follows from the definition of ζ^* that the seller cannot gain by deviating to $p_L(0)$ in this situation. Conversely, $\zeta \geq \zeta^*$ when $\tau > 0$, and as was pointed out above, the seller cannot gain by deviating to p_H in this situation, because $n^*_\tau \geq n^*_1$.

Partial Screening

It remains to be shown that the seller cannot gain by deviating to a price above the screening price p_H, with the implication that screening is incomplete. Recall that the definition of an optimal screening cycle includes the inequalities (P_s) which are supposed to ensure that partial screening offers are not profitable. Lemma S1 shows how this works.

Lemma S1:

If the inequalities (P_s) hold, and if $P_1(\tau,0) > p > p_H$, then $u[p,\zeta,\tau] \leq u[p_H,\zeta,\tau]$, for all τ.

The following lemma verifies that the seller's beliefs are consistent with Bayesian updating in an optimal screening cycle. Here, following Kreps and Wilson (1982), consistency includes the requirement that the seller's belief when a pooling offer is rejected can be rationalized as a Bayesian inference for a neighborhood of fully mixed buyer strategies around the equilibrium strategy.

Lemma S2:

Given a cyclic pricing process determined by ζ^*, the cyclic belief system determined by ζ^* is consistent.

5. **Alternative Equilibria**

A key assumption in the above analysis of optimal screening cycles is that rejection of a pooling offer would lead the seller to believe that the buyer's current valuation is low. This is reasonable, since the buyer would be indifferent in this case, whereas a high-valuation buyer who rejects a pooling offer leaves money on the table.[15] On the other hand, another equally reasonable seller might take the view that rejection of a pooling offer invalidates all previously held beliefs, since these beliefs implied that the offer would certainly not be rejected, and yet it was. Then the seller might reboot using the invariant distribution of the buyer's Markov chain, with $\zeta = \mu$ and $\tau = 0$ (assuming $\zeta^* > \mu$). In general, the seller is free to believe anything after rejection of a pooling offer, but an arbitrary choice of beliefs would be of little interest. A leading alternative to the specification used in this paper is to simply treat rejected pooling offers as if they had been accepted. Rustichini and Villamil (1996) adopted this alternative, arguing that the buyer might tremble, so that there is a small chance that the buyer would reject an offer by mistake. Then, if it is also assumed that the chance of a mistake is unrelated to its cost, it makes sense for the seller to draw no inference from rejection of a pooling offer.

In general, the continuation after a rejected pooling offer can be represented by a transition matrix T which maps the current state τ and the current actions (p,q) to the next state τ'. Then the cyclic updating rule is replaced by the following rule:

if $p=p_H$	and $q=0$,	then $\tau' = K-1+\lambda$
if $p=p_L(\tau)$	and $q=0$,	then $\tau' = s$ with probability $T_{\tau s}$, $s=0,1,2...K$
if $q=1$,		then $\tau' = \max(\tau-1,0)$

[15] Thus the equilibrium would pass a test analogous to the Intuitive Criterion for two-period signaling games proposed by Cho and Kreps (1987).

The matrix T is only used to reboot if the equilibrium crashes because a pooling offer is rejected. The method used above to compute optimal screening cycles can be applied in this more general context as well. The joint values are unchanged, because the equilibrium path is unchanged, but the buyer is generally worth less. Moreover, the differential between pooling and screening continuation values for the seller is affected by T, so the equilibrium screening threshold must be recalculated.

Let $\mathbf{V}_L$ be a vector containing the low buyer's continuation values $[V_L(0),...V_L(\tau)...V_L(K)]$, and let $\mathbf{d}$ be the corresponding vector containing the differences $d(\tau)$ between the values for the low and the high buyer. It is easy to see from the proof of Lemma B1 above that d is unaffected by changes in T. Define T^0 as the transition matrix obtained when the first row of T is replaced by the vector $[0,0,...0,1\text{-}\lambda,\lambda]$. Then the general formula for the low buyer's continuation values can be written as

$$V_L = \Delta(1-\rho)\left(I - \Delta \mathrm{T}^0\right)^{-1} \mathrm{T}^0 d$$

The expression for the screening threshold can be generalized as follows:

$$\zeta^* = \frac{b - v_0}{b + r\theta}$$

where v_0 is the amount the low buyer would lose if the seller were to make a pooling offer at $\tau = 0$:

$$v_0 = V_L(0) - \sum_{s=0}^{K} \tau_{0s}\left[\Delta V_L(s) + \Delta(1-\rho)\, d(s)\right]$$

which is zero in the basic specification.

Transient Screening Equilibria

Even in the basic case with $v_0 = 0$, optimal screening cycles do not provide a complete set of equilibria covering all regions of the parameter space. Optimal screening cycles exist only if $\zeta^* \geq \mu$, which reduces to the condition $\theta \leq \theta_\infty$. If θ exceeds this bound, there are equilibria with $\zeta^* < \mu$. Such equilibria will be called "transient screening equilibria," because they have the property that a seller who infers that the buyer's current valuation is low will never again be optimistic enough to make a screening offer. A transient screening equilibrium is like an optimal screening cycle with $K = \infty$, with one important difference. In an optimal screening cycle, ζ^* is defined as the belief that leaves the seller indifferent between screening and pooling, given that a pooling offer now implies a screening offer next time, because ζ' is a weighted average of ζ^* and μ, so $\zeta' < \mu$. In a transient screening equilibrium, on the other hand, the same logic implies that if a pooling offer is made at ζ^*, the seller will make pooling offers in all future periods.

If $\zeta^* \geq 1\text{-}\sigma$, the value of a screening offer in a transient screening equilibrium is exactly as it would be in an optimal screening cycle with K = ∞, because an accepted screening offer is followed by another screening offer. Thus for any value of the state variable τ, if the screening price p_H is set when the buyer's valuation is θ+n, the seller's continuation value is

$$\begin{aligned} u[p_H, n, \tau] &= U_L(0) + n g(0) \\ &= \frac{\Delta\theta}{1-\Delta} + n\left[1 + \frac{(1-\Delta)\theta}{1-\Delta\sigma}\right] \end{aligned}$$

The value of a pooling offer is

$$u[p_L(\tau), n, \tau] = \frac{\theta}{1-\Delta}$$

The screening threshold is obtained by finding the value of n = 1-ζ that equates the values of screening and pooling. This yields

$$\zeta^* = \frac{1 - \dfrac{\Delta(1-\sigma)\theta}{1-\Delta\sigma}}{1 + \dfrac{(1-\Delta)\theta}{1-\Delta\sigma}}$$

It is not difficult to construct strategies for the buyer and the seller, and a belief system for the seller, that constitute a transient screening equilibrium governed by ζ^*, using the above formula to determine ζ^* from the basic parameters (note that ρ does not directly affect ζ^*, although a decrease in ρ may cause μ to fall below ζ^*, voiding the equilibrium). The details of this construction are omitted.[16]

---

[16]The state variable construction has to be redone, because the long run state is in the pooling region. So $\zeta^*(1)$ is below ζ^*, and $\zeta^*(2)$ is below $\zeta^*(1)$, and 1-σ is in the interval $\zeta^*(k+1), \zeta^*(k)$, for some k. If the seller makes a pooling offer, the state advances toward 1, meaning the long run pooling state. But if the seller were to make a pooling offer at ζ = 1-σ, it might be that $\zeta' = (1+\varphi)(1-\sigma)$ is still in the screening region, so that a pooling offer now would be followed by a screen next time. This deviation is not profitable.

Also, if ζ is between $\zeta^*(1)$ and ζ^*, then any price above the screening price would be rejected for sure, and the continuation would be pooling from then on. But if ζ is below $\zeta^*(1)$ then a price above p_H requires randomization by the buyer, followed by $\zeta' = \zeta^*$, and τ=0 or τ=1 with probabilities that depend on how high the price was. Finally, transient screening equilibria must be checked against the partial screening test.

Repeated Static Equilibria

Finally, if ζ^* does not fall in either the interval $[\mu,1]$ associated with optimal screening cycles, or the interval $[1-\sigma,\mu]$ associated with transient screening equilibria, it can only be in the interval $[0,1-\sigma]$. Equilibria governed by a screening threshold in this interval may be called "unconditional pooling equilibria," because even the most optimistic seller chooses the pooling price. In this case ζ^* solves the static problem, equating the value of pooling and screening offers in the current period without regard to the future, because the future is not affected by what happens in the current period. At the other extreme, optimal screening cycles with K = 1 may be called "unconditional screening equilibria," since even the most pessimistic seller chooses the screening price, and again ζ^* is the solution of the static problem, namely $\zeta^* = 1/(1+\theta)$.

Classification of Equilibria

The inequalities placing ζ^* between μ and ρ can be written as

$$\Theta_c \;\equiv\; \frac{1-\Delta}{1-\Delta\sigma}\,\frac{1-\sigma}{1-\rho} \;\leq\; \Theta \;\leq\; \frac{\rho}{1-\rho} \;\equiv\; \Theta^c$$

where $\Theta = 1/\theta$. Also, the conditions for a transient screening equilibrium with $1-\sigma \leq \zeta^* \leq \mu$ can be written as

$$\Theta_t \;\equiv\; \frac{1-\sigma}{\sigma(1-\Delta\sigma)} \;\leq\; \Theta \;\leq\; \frac{(1-\sigma)(1-\beta)}{(1-\rho)(1-\Delta\sigma)} \;\equiv\; \Theta^t$$

Finally, the condition for an unconditional pooling equilibrium is

$$0 \;\leq\; \Theta \;\leq\; \frac{1-\sigma}{\sigma} \;\equiv\; \Theta^p$$

These results are illustrated in Figure 6, which shows how the screening threshold varies with θ, for a particular ordering of the critical values of Θ defined above. The most effective diagram plots the reciprocal of θ against the odds ratio $\frac{\zeta^*}{1-\zeta^*}$, since these are equal in the static case, and otherwise the relationship between them is either piecewise-linear (in the case of optimal screening cycles) or affine (in the transient screening case).

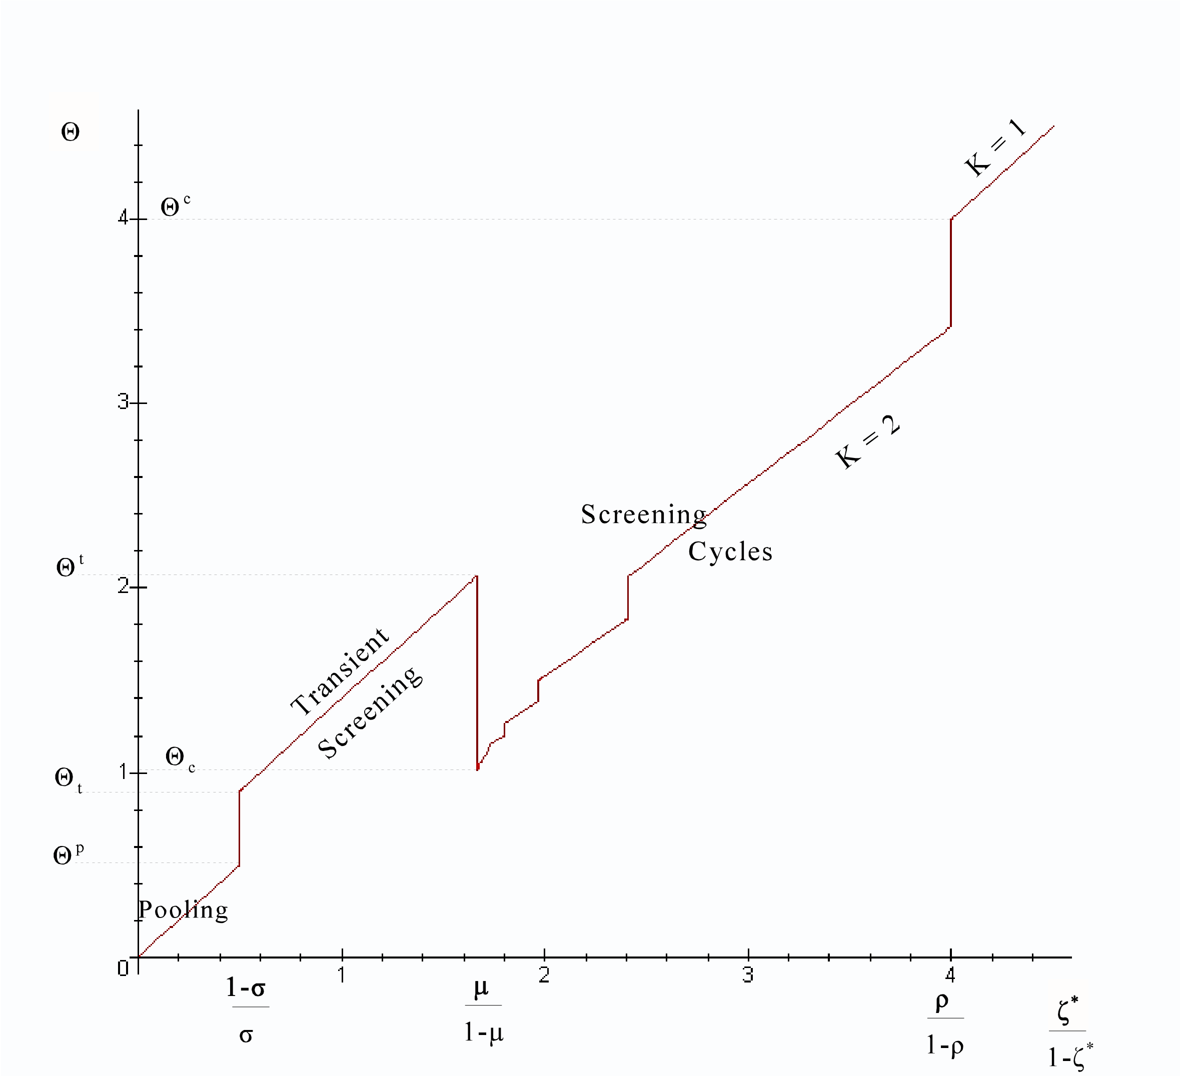

Figure 6 Alternative equilibria (ρ = 4/5, σ = ⅔, Δ = ⅔)

Figure 6 indicates that there will generally be either a unique equilibrium, or else three possible equilibria (at least one of which involves randomization by the seller). These possibilities may be classified as follows. First note that

$$0 \leq \Theta^p \leq \Theta_t \leq \Theta^t < \infty$$

This partitions the parameter space into four regions. Each region is further divided into three subregions according to where Θ lies in relation to Θ_c and Θ^c, but two of these regions are empty, because

$$\frac{\Theta_c}{\Theta^t} = \frac{1-\Delta}{1-\beta} \leq 1$$

and

$$\Theta^p = \frac{1-\sigma}{\sigma} \leq \frac{\rho}{1-\rho} = \Theta^c$$

The 10 parameter regions are illustrated in Figure 7, for fixed values of θ and Δ. In this diagram, any choice of (ρ,σ) that falls in one of the five outside regions (those containing a segment with $\rho = 1$ or $\sigma = 1$) gives a unique equilibrium, while a choice in one of the five inside regions yields three equilibria. For example, if (ρ,σ) is in the small triangular region just above the diagonal, there is an unconditional screening equilibrium, a randomized equilibrium with $\zeta^* = \mu$, and a randomized equilibrium with $\zeta^* = 1-\sigma$.

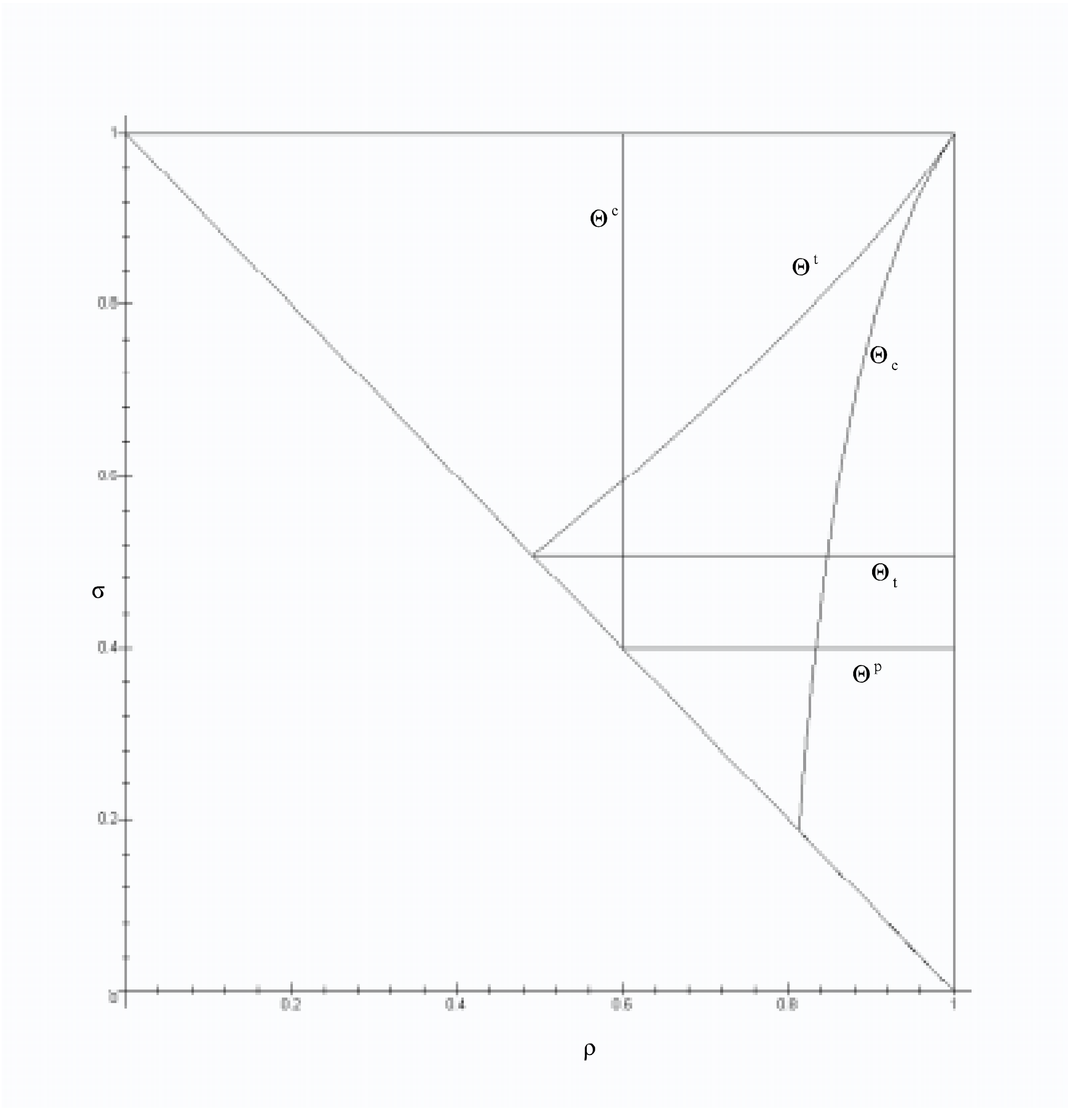

Figure 7 Alternative Equilibrium Regions ($\theta = ⅔$, $\Delta = 0.7$)

6. **Conclusion**

This paper analyzes repeated bilateral monopoly with a private stochastic process for the buyer's valuation. The main results are concerned with cyclic movements of equilibrium prices and quantities generated by a two-state Markov chain for the buyer's valuation. A novel feature of the model is that pooling offers give the buyer a surplus even in the bad state, because the buyer has the option of refusing. The sequence of pooling prices driven by the value of this option involves a gradual decline while the seller is in the pooling phase of the equilibrium cycle, and a sudden jump at the end of this phase.

The analysis is based on the idea of Markov-Perfect equilibrium. At any point in the game, the only information that is relevant for future payoffs is the buyer's current valuation. The seller's strategy is driven by a belief about this valuation, using everything that can be inferred from the buyer's recent actions in the context of the buyer's equilibrium strategy. This belief is summarized by a state variable that counts the number of pooling offers remaining before the seller will be optimistic enough to make the next screening offer. The buyer's strategy then uses this state variable together with the actual current valuation. This is a tractable structure that should be suitable for empirical application: in particular, explicit solutions are obtained for equilibrium prices and quantities, and for the value functions.

The main limitation of the model is that equilibria in which screening is extended over more than one contract are ruled out, by excluding a portion of the parameter space in which the cost of an unsuccessful screening offer is relatively low. A more general analysis would extend the state variable to count down the number of rejected offers needed to convince the seller to restart the pooling sequence, as well as the number of offers remaining in the pooling sequence. From the point of view of application, a more useful alternative is to relax the assumption that each negotiation involves just one take-it-or-leave-it offer, instead allowing a sequence of offers that ends when the seller becomes convinced that the current state is low.

References

Aghion, Philippe, Patrick Bolton, Christopher Harris and Bruno Jullien, "Optimal Learning by Experimentation," *Review of Economic Studies*, 58 (4) No. 196, July 1991, 621-654.

Bergemann, Dirk and Juuso Valimaki, "Learning and Strategic Pricing," *Econometrica*, 64 (5), September 1996, 1125-1149.

Blume, Andreas (1990), "Bargaining with Randomly Changing Valuations," Working Paper 90-22, University of Iowa.

Card, David (1988), "Longitudinal Analysis of Strike Activity," *Journal of Labor Economics*, **6**, 147-176.
Card, David (1990), "Strikes and Wages: a Test of an Asymmetric Information Model", *Quarterly Journal of Economics*, 105 (August), 625-659.

Cho, In-Koo and David M. Kreps (1987), "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics,* 102, 179-221.

Coase, Ronald H., "Durability and Monopoly," *Journal of Law and Economics,* 1972, 15, 143-149.

Cramton, Peter C. and Tracy, Joseph S. (1992), "Strikes and Holdouts in Wage Bargaining: Theory and Data," *American Economic Review*, March 1992, 100-121.

Fudenberg, Drew, David Levine, and Paul Ruud, "Strike Activity and Wage Settlements", UCLA Working Paper # 249, revised September 1985.

Fudenberg, Drew, and Jean Tirole (1983), "Sequential Bargaining with Incomplete Information about Preferences", *Review of Economic Studies*, 50: 221-47.

Hart, Oliver (1989), "Bargaining and Strikes", *Quarterly Journal of Economics*, 104: 25-44.

Hart, Oliver and Jean Tirole (1988), "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55: 509-540.

Hayes, Beth (1984), "Unions and Strikes with Asymmetric Information", *Journal of Labor Economics*, 2: 57-83.

Ingram, Peter, David Metcalf and Jonathan Wadsworth, "Strike Incidence and Duration in British Manufacturing Industry in the 1980s," Working Paper No. 88, Centre for Economic Performance, LSE, April 1991.

Kennan, John (1986), "The Economics of Strikes", *Handbook of Labor Economics*, Volume II, O. Ashenfelter and R. Layard (eds.). Amsterdam: Elsevier Science Publishers BV.

Kennan, John, "Repeated Contract Negotiations with Private Information," *Japan and the World Economy*, 7 (1995), 447-472.

Kennan, John and Robert Wilson (1989), "Strategic Bargaining Models and Interpretation of Strike Data," *Journal of Applied Econometrics,* December 1989 (Supplement), Vol. 4, S87-S130.

Kennan, John and Robert Wilson (1993), "Bargaining with Private Information," *Journal of Economic Literature*, March 1993, 45-104

Kreps, David M., and Robert Wilson, “Sequential Equilibrium”, *Econometrica*, 1982, 50, 863-894.

Maskin, Eric and Jean Tirole, “Markov Perfect Equilibria,” working paper, December 1994.

Riddell, William Craig (1979), "The Empirical Foundations of the Phillips Curve: Evidence from Canadian Wage Contract Data," *Econometrica*, 47, 1-24.

Riddell, William Craig (1980), "The Effects of Strikes and Strike Length on Negotiated Wage Settlements," *Relations Industrielles*, January 1980, volume 35, number 1, pp. 115-120.

Rustichini, Aldo and Anne P. Villamil, “Intertemporal Pricing in Markets with Differential Information,” *Economic Theory*, 8, 1996, 211-227.

Rustichini, Aldo and Asher Wolinsky, “Learning about Variable Demand in the Long Run,” *Journal of Economic Dynamics and Control*, 19 (1995), 1283-1292.

Sobel, Joel, and I. Takahashi (1983), "A Multi-stage Model of Bargaining", *Review of Economic Studies*, 50: 411-26.

Vincent, Daniel, "Repeated Signalling Games and Dynamic Trading Relationships," *International Economic Review*, 1997, forthcoming.

Appendix: Proofs

Lemma B1:

For $\tau \geq 1$, if both buyer types accept the pooling price $p_L(\tau)$, the continuation values for the high and low buyer, and the difference between them, are given by

$$\begin{aligned} V_H(\tau) &= 1 + \theta - p_L(\tau) + \Delta V_L(\tau-1) + \Delta\sigma\, d(\tau-1) \\ V_L(\tau) &= \quad\ \theta - p_L(\tau) + \Delta V_L(\tau-1) + \Delta(1-\rho)\, d(\tau-1) \\ d(\tau) &= 1 + \beta\, d(\tau-1) \end{aligned}$$

The continuation after rejection of any pooling offer is the same as the continuation after rejection of a screening offer. Given that prices are tight, this implies $V_L(\tau) = V_L(0) \equiv V_L$, for $\tau \geq 0$, and

$$\begin{aligned} V_H(0) &= \Delta V_L + \Delta\,\sigma \left[(1-\lambda)\, d(K-1) + \lambda\, d(K)\right] \\ V_L &= \Delta V_L + \Delta(1-\rho)\left[(1-\lambda)\, d(K-1) + \lambda\, d(K)\right] \\ d(0) &= \beta \left[(1-\lambda)\, d(K-1) + \lambda\, d(K)\right] \\ V_L &= \frac{(1-\rho)\, d(0)}{\varphi(1-\Delta)} \end{aligned}$$

Thus the differences $d(\tau)$ are obtained by solving the set of K+1 linear equations defined by

$$\begin{aligned} d(\tau) &= 1 + \beta\, d(\tau-1)\ ,\ \tau = 1,2,\ldots,K \\ d(0) &= \beta\left[(1-\lambda)\, d(K-1) + \lambda\, d(K)\right] \end{aligned}$$

This yields

$$\begin{aligned} d(\tau) &= \frac{1}{1-\beta} - \frac{\beta^\tau}{1-\beta^K(1-\lambda+\lambda\beta)} \\ &= \frac{1-\beta^\tau b}{1-\beta} \end{aligned}$$

This completes the proof.

Lemma S1:

Suppose $P_1(s+1,0) < p \leq P_1(s,0)$, for some integer s, with $\tau - 1 \leq s \leq K-1$. Then the buyer accepts with probability q(s), with continuation from $\zeta' = 1-\sigma$ and $\tau' = 0$, and rejects with probability 1-q(s), with continuation from $\zeta' = \zeta^*(s)$ and $\tau' = s+\xi$, where ξ is defined by

$$p = P_1(s+\xi,0)$$

The seller's current payoff $P_1(s+\xi,0)$ and the expected continuation value from next period on are both linear in ξ, so if a deviation to $P_1(s+\xi,0)$ is profitable, a deviation to $P_1(s,0)$ or $P_1(s+1,0)$ must also be profitable. Thus $\xi = 0$ can be assumed without loss of generality.

The probability q(s) that a partial screening offer is accepted is such that if the offer is rejected the seller's belief next time will be $\zeta' = \zeta^*(s)$. This means that after a rejection and before the transition from n to n' the seller's belief is $\zeta^*(s+1)$. Thus

$$\zeta^*(s+1) = \frac{\hat{\zeta}}{1-q(s)}$$

and in equilibrium $\hat{\zeta} = \zeta$.

A partial screening offer is either accepted by the high buyer, or rejected by the high buyer, or rejected by the low buyer. So

$$\begin{aligned} u[p,\zeta,\tau] &= q(s)\,u_H^a + [\,1-q(s)-\zeta\,]\,u_H^r + \zeta u_L^r \\ &= u_H^a - [1-q(s)\,]\left[u_H^a - u_H^r\right] - \zeta\left[u_H^r - u_L^r\right] \end{aligned}$$

Take the terms in this expression in reverse order. Rejection means continuation from $\tau' = s$, and from the seller's point of view the difference between n=1 and n=0 in this context is exactly the same as it would be at $\tau = s+1$. So

$$u_H^r - u_L^r = g(s+1)$$

Next if the current valuation is high, the difference between acceptance and rejection is just the difference in the joint continuation values, since the high buyer must be indifferent between acceptance and rejection. Acceptance means continuation from $\tau' = 0$ and rejection means $\tau' = s$, so

$$\begin{aligned} u_H^a - u_H^r &= J_H(0) - \left[J_H(s+1) - (1+\theta)\right] \\ &= 1+\theta + [\,B(s+1) - R(s+1)\,]\,r\theta \end{aligned}$$

Now compare the seller's continuation value from a screening offer with the value of a partial screen:

$$\begin{aligned} \Omega(s+1) &= U_H(0) - \zeta g(0) - u[p,\zeta,\tau] \\ &= -\left[u_H^a - U_H(0)\right] + [1-q(s)\,]\left[u_H^a - u_H^r\right] - \zeta\left[g(0) - g(s+1)\right] \end{aligned}$$

Consider the first term here. The joint continuation value is always the same when a screening offer is accepted, so the first term is just the difference between partial and full screening offers from the high buyer's point of view, and for any screening offer the high buyer's continuation value is the value of rejecting the offer. Therefore,

$$\begin{aligned} u_H^a - U_H(0) &= \Delta V_L + \Delta\sigma\, d(K-1+\lambda) - \Delta V_L - \Delta\sigma\, d(s) \\ &= \frac{\sigma}{\varphi}\left[\,1 - B(s+1)\,b\,\right] \end{aligned}$$

Note here that

$$d(0) \;=\; \beta\, d(K-1+\lambda)$$

$$d(s+1) \;=\; 1+\beta\, d(s) \;=\; d(0)+B(s+1)\, b$$

These results can be summarized as :

$$\Omega(s) \;=\; -\frac{\sigma}{\varphi}\left[\,1-B(s)\,b\,\right]+\frac{\zeta}{\zeta^{*}(s)}\left[1+\theta+[\,B(s)-R(s)\,]\,r\theta\right]-\zeta B(s)\left[b+r\theta\right]$$

Since $\zeta^*(s) \le 1$, and B(s) and R(s) are increasing in s, $\Omega(s)$ is an increasing linear function of ζ, so if it is positive when ζ takes its smallest value (which is $1\text{-}\sigma$) then it is always positive. But the inequality P_s is just $\Omega(s)/\zeta \ge 0$ with $\zeta = 1\text{-}\sigma$. This completes the proof.

Propositions E1 and E2.

Given the basic parameters $(\theta,\rho,\sigma,\Delta)$, the function Z determines ζ^* and λ, and the function κ determines K, so K and $\Omega(s)$ can be considered as functions of $(\theta,\rho,\sigma,\Delta)$: each point in the basic parameter space implies a unique value of K, and unique values of $\Omega(s)$, for $1 \le s \le K$. In this appendix, the parameters ρ, σ and Δ are fixed, and K and $\omega(s)$ are considered as functions of θ, using the notation $K(\theta)$ and $\omega(s;\theta)$.

The proof of E1 uses the following result.

Lemma A: If c is positive, and Δ and φ are in the interval (0,1), then the function $f:\Re\rightarrow\Re$ defined by

$$f(s) \;=\; a\,\Delta^{s}-\frac{c}{\varphi^{s}}-(\Delta\varphi)^{s}$$

is quasiconcave.

Proof:

It will be shown that f(s) is increasing for $s < s_0$, and decreasing for $s > s_0$, where s_0 is the unique solution of the equation

$$\varphi^{s} \;=\; \chi a+\frac{c\,(1-\chi)}{(\Delta\,\varphi)^{s}} \;\;,\quad \chi \;\equiv\; \frac{\log(\Delta)}{\log(\Delta)+\log(\varphi)}$$

The left side of this equation decreases from ∞ to 0, while the right side increases from χa to ∞, as s increases from $-\infty$ to $+\infty$, so the equation does indeed have a unique solution. Also,

$$\begin{aligned} f'(s) \;&=\; a\log(\Delta)\,\Delta^{s}+c\log(\varphi)\,\varphi^{-s}-\log(\Delta\varphi)\,(\Delta\varphi)^{s} \\ &=\; -\log(\Delta\varphi)\,\Delta^{s}\left[\varphi^{s}-\chi a-\frac{c\,(1-\chi)}{(\Delta\varphi)^{s}}\right] \end{aligned}$$

The expression in brackets is decreasing in s, and it is zero at $s = s_0$. Thus, since $\log(\Delta\varphi)$ is negative, f(s) is increasing for $s < s_0$, and decreasing for $s > s_0$.

Proof of Proposition E1:

According to Lemma S1, the first part of the proposition can be stated as $\Omega(s) \geq 0$, for $1 \leq s \leq K$. This is equivalent to $\omega(s) \geq 0$, where

$$\omega(s) = \frac{\zeta^*(s)\, B(K+\lambda)\, \Omega(s)}{1-\sigma}$$

Write $\zeta^*(s)$ as $\mu + \varphi^{-s} z_0$, where $z_0 = \zeta^* - \mu$. Substituting this in $\omega(s)$ and rearranging terms yields

$$\begin{aligned}\omega(s) = &\left[\frac{1}{\zeta^*} - 1\right][R(K+\lambda) - \mu D(s)] + [1 - \mu(\Xi+1)]B(K+\lambda) + \mu\,\Xi\, B(s) \\ &+ \varphi^{-s} z_0 \left[\Xi + 1 - \frac{1}{\zeta^*}\right] B(s) - \varphi^{-s} z_0 (\Xi+1) B(K+\lambda)\end{aligned}$$

where

$$\Xi \equiv \frac{\sigma}{\varphi(1-\sigma)} - 1$$

After substituting for D(s) and B(s), $\omega(s)$ can be written in the form

$$\omega(s) = a_0 + a_\Delta \Delta^s - a_\varphi \varphi^{-s} - a_\beta \beta^s$$

where a_0 and a_Δ are irrelevant constants, and

$$\begin{aligned} a_\varphi &= \frac{z_0}{1-\beta}\left[y_0 - \beta^K(\Xi+1)\left(1 - \lambda[1-\beta]\right)\right] \\ a_\beta &= \frac{\alpha\Xi}{1-\beta} \end{aligned}$$

The coefficient a_φ is nonnegative if

$$\beta^K[\Xi+1] \leq y_0$$

This is implied by (P_{K-1}):

$$\begin{aligned} \varphi^{K-1}\left[1 + \frac{\Xi}{1-\mu}\right] &\leq 1 \\ \varphi^{K-1}[\Xi+1] &\leq 1 - \mu\left[1 - \varphi^{K-1}\right] \leq 1 \\ \beta^K[\Xi+1] &\leq 1 \leq y_0 \end{aligned}$$

Since a_β is positive, the function $\omega(s)/a_\beta$ satisfies the conditions of Lemma A, so $\omega(s)$ is quasiconcave. The next step is to show that P_{K-1} implies $\omega(K) \geq 0$ and $\omega(K-1) \geq 0$. Then (given P_1) quasiconcavity implies $\omega(s) \geq \min[\omega(1),\omega(K-1)] \geq 0$, for $1 \leq s \leq K-1$.

To show that $\omega(K) \geq 0$, write $\omega(K)$ as

$$\omega(K) \;=\; \lambda\Delta^K\left[\; 1-\mu - \left(\frac{\sigma}{\varphi(1-\sigma)} - \mu\right)\varphi^K\right] \;\geq\; 0$$

The bracketed term is clearly nonnegative if P_{K-1} holds. If $\zeta^* > \rho(K)$ with $\lambda=0$ then the relevant condition is $\omega(K-1) \geq 0$. This condition can be written as

$$\begin{aligned}\omega(K-1) \;=\; &\left[\frac{B(K-1)}{\zeta^*} + \frac{\sigma}{\varphi(1-\sigma)}\beta^{K-1}\right]\left[1-\zeta^*(K-1)\right]\\ &- \beta^{K-1}\left(\frac{\sigma}{\varphi(1-\sigma)} - 1\right) + \Delta^{K-1}\rho(K-1)\left[\frac{1}{\zeta^*} - 1\right]\\ &+ \lambda\Delta^K\left[\; 1-\mu - \left(\frac{\sigma}{\varphi(1-\sigma)}\rho - \mu\right)\varphi^K\right]\end{aligned}$$

The first term is nonnegative, and the term involving λ is nonnegative if P_{K-1} holds, so it is enough to show that the remainder is nonnegative. But $\rho(K-1) \geq \zeta^*$, so this is also implied by P_{K-1}:

$$\begin{aligned}\rho(K-1)\left[\frac{1}{\zeta^*} - 1\right] - \varphi^{K-1}\left(\frac{\sigma}{\varphi(1-\sigma)} - 1\right) \;&\geq\; 1-\rho(K-1) - \varphi^{K-1}\left(\frac{\sigma}{\varphi(1-\sigma)} - 1\right)\\ &=\; 1-\mu - \varphi^{K-1}\left(\frac{\sigma}{\varphi(1-\sigma)} - \mu\right)\end{aligned}$$

This proves the first part of E1.

To prove the second part of E1 it is enough to show that an increase in θ relaxes P_1 and P_{K-1}. It is obvious that P_{K-1} is relaxed when θ increases, because $K(\theta)$ is increasing. The definition of the function Z shows that there are two possibilities when θ increases. First, if $\bar\rho(K-1) \leq \theta/G(K) < \bar\rho(K)$, then a (small) increase in θ reduces ζ^* without disturbing K or λ. Write P_1 as

$$\frac{1}{\zeta^*(1)} + \theta\left[\frac{1}{\zeta^*(1)} - \frac{1}{R(K+\lambda)}\right] - \frac{\sigma}{(1-\sigma)\varphi} + \left[\frac{\sigma}{(1-\sigma)\varphi} - 1\right]\frac{1}{B(K+\lambda)} \;\geq\; 0 \qquad (\mathbf{P_1})$$

The first term in brackets here is positive, because $\zeta^*(1) \leq 1 \leq R(K+\lambda)$; also $1/\zeta^*(1)$ increases when ζ^* decreases, so the left side of the inequality increases when θ increases.

The other possibility is that θ satisfies $G(K) \le \theta/\bar{\rho}(K) < G(K+1)$. In this case a (small) increase in θ increases λ while K and ζ^* remain unchanged.[17] It will be shown that this increases $\omega(1)$, implying that an increase in θ relaxes P_1. Write $\omega(1)$ as

$$\omega(1) = \left[\frac{1}{\zeta^*} + \frac{\sigma}{\varphi(1-\sigma)}\big(B(K+\lambda)-1\big)\right]\big[1-\zeta^*(1)\big]$$
$$- [B(K+\lambda)-1]\left(\frac{\sigma}{\varphi(1-\sigma)}-1\right) + [R(K+\lambda)-1]\left[\frac{1}{\zeta^*}-1\right]$$

Then, since $\zeta^* = \rho(K)$ when λ is positive,

$$\frac{\partial\omega(1)}{\partial\lambda} = \frac{\sigma\beta^K\big[1-\zeta^*(1)\big]}{\varphi(1-\sigma)} - \beta^K\left(\frac{\sigma}{\varphi(1-\sigma)}-1\right) + \Delta^K[1-\rho(K)]$$

The first term here is obviously positive, since $\zeta^*(s) \le 1$. To show that the remainder is positive, divide it by Δ^K, and note that

$$1-\rho(K) - \varphi^K\left(\frac{\sigma}{\varphi(1-\sigma)}-1\right) = 1-\mu - \varphi^K\left(\frac{\sigma}{\varphi(1-\sigma)}-\mu\right)$$

is positive when P_{K-1} holds. This completes the proof.

Proof of Proposition E2:

First consider $\omega(1,\theta)$ as a function of θ. As θ increases, $K(\theta)$ increases, with $K=\infty$ for $\theta > \theta_\infty$. Also,

$$\frac{(1-\beta)\,\omega(1)}{\Delta} = \frac{\theta}{\Delta} + \varphi - \frac{\sigma\zeta^*}{1-\sigma}$$
$$= \left[\frac{1}{\Delta} + \frac{1-\sigma}{1-\Delta}\right]\frac{1-\zeta^*}{\zeta^*} + \varphi - \frac{\sigma\zeta^*}{1-\sigma}$$

Substituting $\zeta^* = \mu$ in this equation shows that the inequality $\omega(1,\theta_\infty) > 0$ is equivalent to (P_∞). It follows that there is some open interval (θ_0,θ_∞) such that $\omega(1,\theta) > 0$ for $\theta \in (\theta_0,\theta_\infty)$. Since P_{K-1} necessarily holds for large K, θ_0 can be chosen so as to ensure that P_{K-1} holds for $\theta \in (\theta_0,\theta_\infty)$. Then the conditions of Proposition E1 are satisfied for $\theta \in (\theta_0,\theta_\infty)$, which proves the result.

[17]One might well ask how K ever changes, given the assertion that it remains fixed in both cases considered above. This is merely a matter of notation. Starting from an equilibrium with $\zeta^* = \rho(K)$, increases in θ yield equilibria with the same value of ζ^* and increasing values of λ, until λ reaches 1. At this point K is augmented by one and λ returns to zero, so that the same screening threshold now satisfies $\zeta^* = \rho(K-1)$, with $\zeta^*(K-1) = 1$. Note also that the inequality P_{K-1} is equivalent to $\omega(K-1) \ge 0$ at this point.

Proof of Lemma S2:

Consider an arbitrary price p. It will be shown that if the seller were to offer p, the inference drawn by the seller from the buyer's response would be consistent with the buyer's equilibrium strategy.

First note that if the seller's current belief is ζ, the belief next period, in the absence of any new information, is

$$\zeta' = \zeta\rho + (1-\zeta)(1-\sigma) = \varphi\zeta + (1-\varphi)\mu$$

Thus if the buyer's strategy specifies that p is accepted whether the current valuation is high or low, then the above expression gives the correct Bayesian updating rule for the seller when p is chosen; the same is true if p is a price that the buyer's strategy unconditionally rejects. If the buyer's strategy rejects or accepts p according to whether the current valuation is low or high, then the seller sets $\zeta = 0$ if p is rejected, and $\zeta = 1$ if p is accepted, implying $\zeta' = 1-\sigma$ and $\zeta' = \rho$, respectively.

The buyer's strategy specifies randomization if the current valuation is high, and the seller names a price p between p_H and $P_1(\tau-1,0)$. In equilibrium, the seller's belief ζ agrees with the estimate used by the buyer to determine the probability that p is accepted. Let ζ^0 be the seller's posterior belief that the current valuation is low, if the buyer rejects p. Then

$$\zeta^0 = \frac{\zeta}{1 - \upsilon(1-\zeta)} = \zeta^*(s+1)$$

Thus $\zeta' = \zeta^*(s)$, so the seller's belief is correct following randomization by the buyer.

Finally, it must be shown that the seller's belief following a rejected pooling offer can be supported as the limiting Bayesian inference for a sequence of fully mixed buyer strategies approaching the equilibrium strategy. This is easily done. Modify the buyer's equilibrium strategy so that wherever the strategy specifies q =1, the buyer instead accepts with probability $1-\varepsilon^{1+n}$, and rejects with probability ε^{1+n}; while if the strategy specifies q = 0, the buyer accepts with probability ε^{2-n}, and rejects with probability $1-\varepsilon^{2-n}$. In particular (setting n = 1), ε^2 is the probability that a pooling offer is rejected by the high-valuation buyer and (setting n = 0) ε is corresponding probability for the low buyer.[18] Then if a pooling offer is rejected, Bayes rule implies that the probability of the low type is

$$\zeta = \frac{\varepsilon}{\varepsilon + \varepsilon^2} = \frac{1}{1+\varepsilon}$$

Taking $\varepsilon = 1/m$ yields a sequence ζ_m converging to 1 as m increases, supporting the seller's belief that if a pooling offer is rejected, the buyer's valuation is low.

[18]The point here is that the probability of making a mistake is inversely related to its cost: the high buyer is more likely to err on the high side, and the low buyer is more likely to err on the low side.