Repeated Bargaining with Persistent Private Information

John Kennan[1]
University of Wisconsin-Madison

December, 1998

Abstract

The paper analyzes repeated contract negotiations involving the same buyer and seller where the contracts are linked because the buyer has persistent (but not fully permanent) private information. (The main application is labor contracts, where the employer has private information about the value of labor services sold by the union). The size of the surplus being divided is specified as a two-state Markov chain with transitions that are synchronized with contract negotiation dates. Equilibrium involves information cycles triggered by the success or failure of aggressive demands made by the seller. A successful demand induces the seller to again make an aggressive demand in the next negotiation, because the buyer's acceptance reveals that the current surplus is large, and because there is persistence in the Markov chain generating the surplus. Rejection of an aggressive demand, on the other hand, leads the seller to be pessimistic about the size of the surplus in the next contract, so the seller makes a "soft" offer that is sure to be accepted. Then, several contracts later, the Markov chain has made enough transitions to make the seller optimistic enough to again make an aggressive demand, and the result of this demand re-starts the information cycle. An interesting feature of this cycle is that the soft price is not constant, but declines as the cycle continues, so as to offset the buyer's option value of re-starting the cycle when the current state is bad. An explicit mapping is given for the relationship between the basic parameters and the equilibrium prices and quantities; in particular, there is a closed-form solution for the threshold belief that makes the seller indifferent between hard and soft offers.

---

[1]Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706; jkennan@macc.wisc.edu. Revisions will be at http://www.ssc.wisc.edu/~jkennan/research/cyc_ab.htm . This paper grew out of a joint research project with Robert Wilson, to whom I am grateful for many stimulating discussions. I thank Ray Deneckere, Rody Manuelli, Larry Ausubel, Dirk Bergemann, Andreas Blume, V. V. Chari, Drew Fudenberg, Pat Kehoe, Peter Norman, 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, Yale University, the University of Valencia 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. Repeated contracts also arise in international trading relationships, and in 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. This literature deals with static bargaining, in the sense that the game ends as soon as a contract is signed.[2] Hart and Tirole (1988) analyzed repeated bargaining, but with a static information structure. 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] To analyze this relationship, a dynamic bargaining model is needed. Obviously, the model must allow repeated bargaining, but in addition, it seems desirable to allow for the arrival of new private information: 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 pool or screen, because screening reveals 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 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 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 informed party's ability to signal is severely

---

[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).

restricted (although not eliminated, as will be seen). For concreteness, say an uninformed seller makes take-it-or-leave-it offers to an informed buyer (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 a single 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. 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 interactions 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 is the derivation of a cyclic equilibrium in which the strategies and beliefs are simple functions of the fundamental parameters (the buyer's valuations, the Markov transition probabilities and the discount factor). The equilibrium is unique in the sense that it has the same expected payoffs as every other equilibrium that satisfies a set of properties motivated by the heuristic description of cyclic screening and pooling given above.

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. Keller and Rady (1997) consider the more general issue of whether a monopolist will choose to learn in a changing environment, in which the demand function is driven by a Markov process. Rustichini and Wolinsky (1995) analyze a model in which a seller faces a nonstrategic buyer with a rectangular demand curve driven by a 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 here, but for the more complicated situation in which 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 (1998) 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. **Screening Cycles**

A natural equilibrium of the bargaining 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 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.

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, based on the results of the previous negotiation, and the transition matrix of the Markov chain; the probability of the low state now given that the state was low last time is denoted by ρ_L, while the probability of the low state now given that the state was high last time is $1-\rho_H$, which is less than ρ_L by assumption. At one extreme, the seller's belief that the buyer has a low valuation now is ρ_L, 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, $\rho_L(K)$, 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 $1-\rho_H$.

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 led to complications that precluded a systematic equilibrium analysis, although numerical examples were computed. As will be seen below, reducing the problem by allowing only one offer per contract leads to sharper results, without making things so simple as to be uninteresting. Rustichini and Villamil also assumed one offer per contract, and their main result was that cyclic equilibria exist if the degree of persistence is sufficiently high. This is puzzling since Hart and Tirole (1988) had shown (for discount factors exceeding ½) that only pooling equilibria survive in the limiting case when the buyer's valuation is permanent ($\rho_L = \rho_H = 1$). But these results are in fact compatible, because the Rustichini and Villamil equilibrium is merely a weak Perfect Bayesian equilibrium, in which the seller's beliefs off the equilibrium path are not consistent with the buyer's strategy. 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

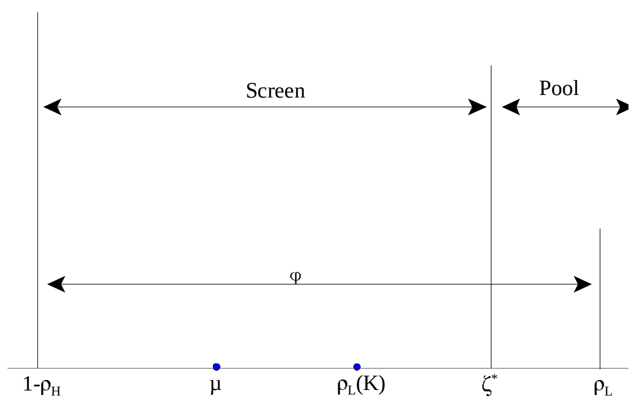

At this point a preview of the results may be useful. After the infinite-horizon game is formally defined in the next section, a two-period version is analyzed in order to illustrate the basic idea of the screening threshold, and to discuss the possibility that screening might be extended over more than one period. The two-period model also shows the variety of equilibria that can be obtained if the seller's belief following rejection of a pooling offer is unrestricted, and it shows how a forward induction argument can be used to select a specific belief, namely that the buyer's current valuation must have been low if a pooling offer was rejected. The results for the two-period model are then used to motivate a set of properties defining what is meant by cyclic screening equilibrium in the infinite-horizon game. The screening threshold naturally divides the seller's belief space into intervals that are indexed by the number of pooling offers that would be needed to reach the threshold (starting from a relatively pessimistic belief), and this index is used to construct a state variable that drives both players' strategies. After computing the value functions in terms of the state variable, the equilibrium screening threshold is derived as a function of the fundamental parameters: this is one of the main results of the paper. The strategies are then spelled out in detail, and optimality is checked, with emphasis on the conditions needed to ensure that the seller does not find it profitable to extend screening over more than one period. The other main result is that the cyclic screening equilibrium always exists in a region of the parameter space defined by a simple inequality involving the discount factors and the transition probabilities, together with a lower bound on the opportunity cost of screening. Numerical examples are then presented, and the paper closes with a brief discussion of equilibria supported by alternative beliefs off the equilibrium path.

3. An Infinite Horizon Markov Model of Repeated Negotiations

Consider an infinite sequence of contract negotiations between a buyer and a seller where the surplus to be divided in each contract is determined by a two-state Markov chain with continuation probabilities ρ_L and ρ_H. The realizations of the surplus are seen only by the buyer. It is convenient to use the seller's opportunity cost as the origin and the difference between the high and low surplus as the unit, so let $\theta > 0$ be the low surplus, and let $1+\theta$ be the high surplus.[4] Both sides maximize the present value of expected income, with a common discount factor δ. Thus the model is summarized by the four parameters $(\theta,\rho_L,\rho_H,\delta)$.

The rules of bargaining are simple: the seller 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).

Let $\rho_L(s)$ be the probability that the Markov chain is in the low state after s transitions, given that the current state is low, and let $\rho_H(s)$ be the probability of the high state after s transitions, given that the current state is high. Then

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

where $\varphi = \rho_L + \rho_H - 1$ measures the degree of persistence in the Markov chain, and μ is the invariant probability of the low valuation, i.e.

$$\mu = \frac{1-\rho_H}{1-\rho_H+1-\rho_L} = \frac{1-\rho_H}{1-\varphi}$$

The parameter φ governs the extent to which successive negotiations are linked. It is assumed that $\varphi \geq 0$, which implies that $\rho_L(s)$ and $\rho_H(s)$ do not oscillate. If $\varphi = 0$ information is purely transitory, so that

[4]For example, if v is 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, the surplus is $v-w_0$. Then in the new units the low surplus is $\theta = (v_L-w_0)/(v_H-v_L)$, and the high surplus is $1+\theta = (v_H-w_0)/(v_H-v_L)$.

any inference drawn 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 permanent.

The bargaining game is now laid out formally.

Actions

The actions $a_t = (n_t, p_t, q_t)$ taken in period t are as follows:

first, nature chooses a number n_t, and the surplus is $n_t + \theta$, where n_t is either 0 or 1;

then the seller chooses a price, p_t, from the real line;

then the buyer chooses a quantity, q_t, which is either 0 or 1.

Payoffs

The seller's payoff in period t is $u(a_t) = q_t p_t$.

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

Both the seller and the buyer maximize the expected present value of future payoffs, with a common discount factor δ. The present values of the seller's and the buyer's payoffs are written as

$$\tilde{U}(a) = \sum_{t=1}^{\infty} \delta^t p_t q_t$$

$$\tilde{V}(a) = \sum_{t=1}^{\infty} \delta^t \left[n_t + \theta - p_t \right] q_t$$

History

The history of the game at the beginning of period t is as follows:

The sequence of actions chosen by nature is	n^{t-1}	=	$\{n_1, n_2, \ldots n_{t-1}\}$
The sequence of actions chosen by the seller is	p^{t-1}	=	$\{p_1, p_2, \ldots p_{t-1}\}$
The sequence of actions chosen by the buyer is	q^{t-1}	=	$\{q_1, q_2, \ldots q_{t-1}\}$
The public history (i.e. the history available to the seller) is	h_t^0	=	$\{p^{t-1}, q^{t-1}\}$
The history (when the buyer chooses q_t) is	h_t	=	$\{n^t, p^t, q^{t-1}\}$
The buyer's private history (the history not available to the seller) is			$\{n^t\}$

Strategies

Nature's strategy, σ^N, is a Markov chain, with transition matrix $\begin{pmatrix} \rho_L & 1-\rho_L \\ 1-\rho_H & \rho_H \end{pmatrix}$. That is, $n_t = 1$ with probability $1-\rho_L+\varphi n_{t-1}$, and $n_t = 0$ with probability $\rho_L - \varphi n_{t-1}$.

Let $\mathcal{H}_t$ be the set of possible histories at t, and let $\mathcal{H}_t^0$ be the set of public histories. A behavioral strategy for the seller, σ^S, is a sequence of functions σ_t^S from $\mathcal{H}_t^0$ to the set of probability distributions on the real line. The seller's strategy set Σ^S is the set of such sequences.

A behavioral strategy for the buyer, σ^B, is a sequence of functions σ_t^B from $\mathcal{H}_t$ to $\Delta(\{0,1\})$, the set of probability distributions on $\{0,1\}$. The buyer's strategy set Σ^B is the set of such sequences.[5]

Each strategy profile $\sigma = (\sigma^N,\sigma^S,\sigma^B)$ determines a stochastic process $A(\sigma)$ for a_t, taking values in $\{0,1\}\times R\times\{0,1\}$.[6] This gives a complete description of what will happen if the players follow σ. The *path* of σ, denoted by $\underline{\sigma}$, is the support of this process – the set of sample paths having positive probability. A history is *on the path of σ* if the set of all sample paths beginning with this history has positive probability. The continuation of $A(\sigma)$ following any history h_t, written as $A(\sigma|h_t)$, is itself a stochastic process that describes future actions following h_t, and, for s < t, a history h_t is *on the continuation path of σ from h_s* if h_t has positive probability according to $A(\sigma \mid h_s)$.

Beliefs

A belief-system Π is a sequence of mappings from $R^{t-1}\times\{0,1\}^{t-1}$ to $\Delta(\{0,1\}^t)$; each term specifies a probability distribution $\Pi_t(n^t|h_t^0)$ over the (finite) set of possible realizations of n^t, for each possible realization of h_t^0.

Since n_t is Markov, the only part of the seller's belief relevant for the future is the marginal distribution over n_t. The abbreviation $\zeta_t = \Pi_t[n_t = 0|h_t^0]$ will be used to summarize the seller's degree of pessimism regarding the buyer's current valuation.[7]

Equilibrium

There is no general definition of sequential or perfect Bayesian equilibrium for games with infinite strategy sets, but a straightforward adaptation of the definitions in Kreps and Wilson (1982) and Fudenberg and Tirole (1991) yields a suitable definition of equilibrium for the game analyzed here.

---

[5]The notation $\sigma_t^B(q \mid h_t)$ will be used to mean the probability that $q_t = q$, where q is 0 or 1.

[6]Suppose the history is $h_t = \{n^t,p^t,q^{t-1}\}$. After the buyer chooses q_t, nature's next move is $n_{t+1} = 1$ w.p. $1-\rho_H + n_t\phi$ and $n_{t+1} = 0$ otherwise. The price next period is a random variable with distribution $\sigma^S(h_{t+1}^0)$. The quantity next period is a random variable with distribution $\sigma^B(h_{t+1})$. This gives the joint distribution of (n,p,q) in period t+1, and repeated application of these calculations give the joint distribution for all future periods.

[7]This an abuse of notation: $\Pi_t[n_t = 0|h_t^0]$ represents the marginal distribution over n_t, so $\Pi_t[n_t = 0 \mid h_t^0] = \sum_{\{v^t \in \{0,1\}^t \mid v_t = 0\}} \left[\Pi_t(v^t \mid h_t^0)\right]$. Similarly, $\Pi_t[n_{t-1} = 0|h_t^0]$ will be used to represent the marginal distribution over n_{t-1}, i.e., $\Pi_t[n_{t-1} = 0 \mid h_t^0] = \sum_{\{v^t \in \{0,1\}^t \mid v_{t-1} = 0\}} \left[\Pi_t(v^t \mid h_t^0)\right]$, for n $\in$ {0,1}.

Consistency

The strategy-belief pair (σ,Π) is *consistent* if, for all t, and for all public histories h^0_t

(a)

$$\Pi_t[n_t = 0 \mid h_t^0] \;=\; 1 - \rho_H + \phi\Pi_t[n_{t-1} = 0 \mid h_t^0]$$

(b) and for all prices p_t, and for all private histories n^t, and for q = 0,1

$$\left[\sum\nolimits_{\nu^t \in \{0,1\}^t} \Pi_t(\nu^t|h_t^0)\sigma_t^B(q \mid \nu^t, p^t, q^{t-1})\right]\bar{\Pi}_{t+1}(n^t|p^t,q^{t-1},q) \;=\; \Pi_t(n^t|h_t^0)\,\sigma_t^B(q \mid n^t,p^t,q^{t-1})$$

where $\bar{\Pi}_{t+1}(n^t \mid h^0_{t+1}) = \Pi_{t+1}(n^t,1 \mid h^0_{t+1}) + \Pi_{t+1}(n^t,0 \mid h^0_{t+1})$, the marginal belief regarding n^t as of period t+1.

That is, (a) the belief in period t about n_t is the belief in period t about n_{t-1} updated by the Markov transition probabilities; and (b) the revision of the seller's beliefs after seeing the buyer's action satisfies a Bayesian updating equation of the form $\Pi_t(q)\,\Pi_{t+1}(n \mid q) = \Pi_t(n \cap q)$, where q stands for the buyer's action, and n stands for the private history. Condition (b) must hold for all prices, including those that are not on the path of σ: even after a public history that is inconsistent with σ, the belief system Π specifies a distribution on n^t, and if the public history reveals no further deviations from σ, the seller's subsequent beliefs are determined by Bayes rule. For example, Rustichini and Villamil (1996) describe an equilibrium in which the buyer's strategy always rejects prices that are above the support of the seller's strategy (regardless of whether the current valuation is high or low), and yet if such a price were to be offered and rejected, the seller would infer that the buyer's valuation is low. This violates consistency: the term in square brackets in (b) above collapses to 1, and $\sigma^B_t(.)$ is also 1, so consistency requires that the posterior belief following rejection is the same as the prior belief.[8]

Sequential Optimality

For a given strategy profile σ, the buyer's expected payoff, conditional on the history h_t, is

$$V(\sigma|h_t) \;=\; E_{A(\sigma|h_t)}\tilde{V}(a)$$

where E refers to expectations over all continuations of the stochastic process $A(\sigma)$, starting from h_t. For a given strategy-belief pair (σ,Π), the seller's expected payoff, conditional on h^0_t, is

$$U(\sigma,\Pi|h_t^0) \;=\; \sum\nolimits_{\nu^t \in \{0,1\}^t} \Pi_t(\nu^t|h_t^0)\,E_{A(\sigma|\nu^t,h_t^0)}\tilde{U}(a)$$

[8]The Kreps-Wilson (1982) definition of consistency (for finite games) also requires that beliefs following probability-zero events can be rationalized as the limit of a sequence of strategy-belief pairs such that each strategy in the sequence is fully mixed, and so each belief system in the sequence is fully determined by Bayes rule. It is easy to check that this additional requirement is satisfied for all of the strategies considered below.

where E refers to expectations over all continuations of A(σ), starting from the public history h_t^0 and the private history ν^t.

Definition (Sequential Optimality)

(a) The strategy profile σ is *sequentially optimal for the buyer* if, for all t, and for all histories h_t, and for all strategies $\mathfrak{s}^B \in \Sigma^B$

$$V(\sigma^N,\sigma^S,\sigma^B|h_t) \geq V(\sigma^N,\sigma^S,\mathfrak{s}^B|h_t)$$

(b) The strategy-belief pair (σ,Π) is *sequentially optimal for the seller* if, for all t, and for all public histories h_t^0, and for all strategies $\mathfrak{s}^S \in \Sigma^S$.

$$U(\sigma^N,\sigma^S,\sigma^B,\Pi|h_t^0) \geq U(\sigma^N,\mathfrak{s}^S,\sigma^B,\Pi|h_t^0)$$

The pair (σ,Π) is *sequentially optimal* if both (a) and (b) are true.

Definition A0: An equilibrium is a consistent and sequentially optimal strategy-belief pair (σ,Π).

A Two-Period Game

To motivate the equilibrium analysis for the infinite-horizon game, it is useful to consider a truncated version with just two periods. The set of sequential equilibria of the two-period game is fully characterized in Kennan (1998), and the main results are summarized here. Denote the first period as t = 1, and the last as t = 0. The last period is a one-shot game in which the seller chooses a price, the buyer says yes or no, and the game ends. Let ζ_0 be the probability that the buyer's valuation is low, as perceived by the seller at the beginning of the last period. It is easy to see that in any sequential equilibrium of the two-period game, the continuation equilibrium in the last period is based on the threshold $\zeta_0^* = \frac{1}{1+\theta}$. If $\zeta_0 < \zeta_0^*$, then the equilibrium involves screening: the seller chooses the price 1+θ, which the high-valuation buyer accepts and the low-valuation buyer rejects. If $\zeta_0 > \zeta_0^*$, then the equilibrium involves pooling: the seller chooses the price θ, which both buyer types accept. In the borderline case $\zeta_0 = \zeta_0^*$, there are two pure strategy equilibria: the seller may choose either the pooling or the screening price, being indifferent between the two. Let τ_0 be a state variable representing the probability that the seller makes a pooling offer, so that $\tau_0 = 1$ if $\zeta_0 > \zeta_0^*$, and $\tau_0 = 0$ if $\zeta_0 < \zeta_0^*$. Then the seller's strategy in the last period is fully described by τ_0, as a function of the history. This is a trivial example of the state variable that will be used to construct strategies for the infinite-horizon game.

Now consider the first period. Let ζ_1 be the prior probability that the buyer's valuation is low, at the beginning of the game. The seller's belief ζ_0 in the last period will be an update of ζ_1, based on the transition probabilities ρ_L and ρ_H, and on any information that may be inferred from the buyer's action in the first period. Assume that $\varphi\zeta_1 + 1\text{-}\rho_H < \zeta_0^*$, meaning that the update leaves the seller optimistic enough to screen in the last period if a pooling offer is made and accepted in the first period.

The equilibrium analysis is complicated by the buyer's ability to signal, even though the buyer can only accept or reject offers. The difficulty arises because the buyer might reject a pooling offer, and the seller's belief in the last period is then not determined by Bayes rule.

The full description of the set of sequential equilibria requires too much space to be included here: see Kennan (1998). The set is never empty, and it may include several types of equilibria; examples of each type will be given, with emphasis on features that carry over to the infinite-horizon game. The details of these equilibria are given in Appendix A. The set of equilibria may be characterized as follows.

- In any equilibrium, the seller either pools in the first period, by offering a price that is accepted for sure, or screens, by offering a price that is rejected for sure by the low-valuation buyer; or else the seller randomizes between these two alternatives. Thus equilibria can be labeled as $\tau_1 = 1$ (pooling in the first period), or $\tau_1 = 0$ (screening in the first period), or $0 < \tau_1 < 1$ (randomization).
- If there are equilibria with $\tau_1 = 0$, they all have the same equilibrium path:
  - either $p_1 = 1+\theta\text{-}\delta\rho_H$, $q_1 = n_1$, $p_0 = q_1 + \theta$, $q_0 = n_0$ (a screening equilibrium)
  - or $p_1 = 1+\theta$, $q_1 = \hat{q}n_1$, $p_0 = 1 + \theta$, $q_0 = n_0$ (a partial screening equilibrium), where $\hat{q}$, the probability that the high-valuation buyer accepts the first-period offer, is such that if this offer is rejected, then the seller's belief in the last period is ζ_0^*.
- If there is an equilibrium with $\tau_1 = 1$, then there is a continuum of such equilibria, with different pooling prices in the first period. In this case the equilibrium path satisfies $p_1 \; \varepsilon \; [\theta\text{-}\delta(1\text{-}\rho_L),\theta]$, $q_1 = 1$, $p_0 = 1 + \theta$, $q_0 = n_0$.

First-Period Pooling Equilibria

Given that p_0 is either θ or $1+\theta$ in any equilibrium, define $\ell(p_1)$ as the probability (specified by σ^S) that $p_0 = \theta$, if the price p_1 has been rejected. Then in any equilibrium with $\tau_1 = 1$, the seller's strategy specifies $p_1 = p_L$, where the pooling price p_L satisfies

$$p_L \; = \; \max\left\{ p : p \leq \theta - \delta(1-\rho_L)\,\ell(p) \right\}$$

Table A1 in Appendix A shows three alternative equilibria, labeled P1, P2 and P3, in which the seller pools in the first period, and screens in the last period. The strategies in these equilibria differ only in the specification of $\ell(p_1)$ for $p_1 \leq \theta$, as follows

P1: $\ell(p_1) = 0$ for $p_1 \leq \theta$, $\quad p_L = \theta$
P2: $\ell(p_1) = 1$ for $p_1 \leq \theta$, $\quad p_L = \theta - \delta(1\text{-}\rho_L)$
P3 $\ell(p_1) = 0$ for $p_1 \leq \theta - ½\delta(1\text{-}\rho_L)$, $\ell(p_1) = 1$ for $\theta - ½\delta(1\text{-}\rho_L) < p_1 \leq \theta$, $\quad p_L = \theta - ½\delta(1\text{-}\rho_L)$

These equilibria differ in the way the seller's strategy reacts when the buyer rejects an offer that should have been accepted. In P1, the seller treats a rejected pooling offer as a random mistake, and screens in the last period as if no information had been revealed in the first period. Then since the seller's action in the last period does not depend on whether the pooling offer is accepted or rejected, the low-valuation buyer accepts if and only if the price is below the current valuation θ, so the pooling price is θ.

In P2, on the other hand, rejection of a pooling offer is interpreted as a signal that the buyer's current valuation is low, and in light of this the seller pools in the last period. The option to reject the pooling offer is then valuable to the low buyer, because the buyer's valuation might switch in the last period. The pooling price p_L is below the current valuation by the amount $\delta(1-\rho_L)$, which is just enough to offset the value of the buyer's option to induce a pooling offer in the last period.

The point of the P3 equilibrium is that the low buyer strictly prefers to accept the pooling offer in the first period, because rejection would lead to a screening offer in the last period, and thus a zero payoff for the buyer in both periods, while acceptance yields a current payoff of $\frac{1}{2}\delta(1-\rho_L)$.[9]

First-Period Screening Equilibria

Table A1 shows two equilibria in which the seller screens in the first period. In the equilibrium labeled "S" the seller offers $p_1 = 1+\theta-\delta\rho_H$, and the high-valuation buyer accepts for sure, so there is a clean separation of the two buyer types. In the equilibrium labeled "PS" there is partial screening: the seller offers $p_1 = 1+\theta$, and this is rejected with positive probability by the high buyer.

Screening Thresholds

These alternative equilibria are associated with alternative specifications of the function $\ell(p_1)$, and the equilibria are valid only over specific ranges of the prior belief ζ_1. Given ℓ, the pooling price can be computed, and the seller's expected payoff can be calculated as a function of ζ_1. Then the value of pooling can be compared with the value of screening (i.e. choosing $p_1 = 1+\theta-\delta\rho_H$), and with the value of partial screening ($p_1 = 1+\theta$). These comparisons are illustrated in Figures 2a and 2b.

---

[9] It might seem that there is something wrong here: can't the seller extract this surplus from the low buyer by raising the price in the first period? No, because if the buyer rejects any higher price in the first period, the seller infers the low valuation, and makes a pooling offer in the last period; and since this is profitable for the low buyer, no price higher than $\theta - \frac{1}{2}\delta(1-\rho_L)$ would be accepted by the low buyer.

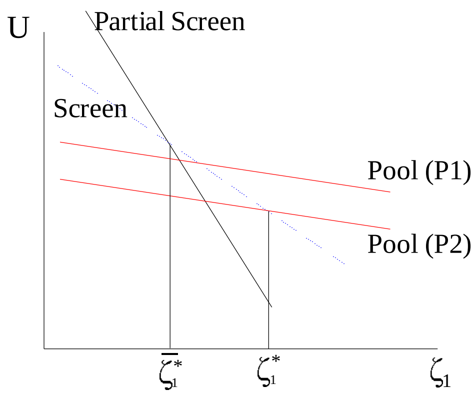

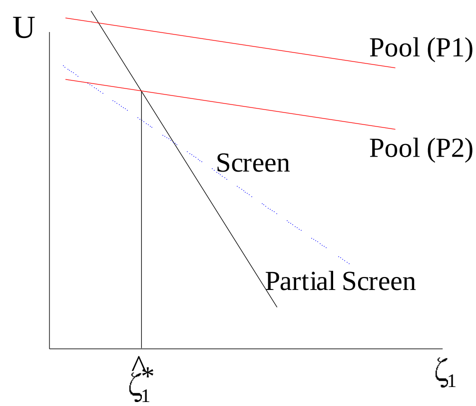

Figure 2a: expected payoffs for the seller

Figure 2b: expected payoffs for the seller

The seller's expected payoffs are linear in ζ_1, becoming progressively less sensitive to ζ_1 as the strategy shifts from partial screening to screening, and then to pooling. Three thresholds are shown: screening overtakes partial screening at $\bar{\zeta}_1^*$, pooling overtakes partial screening at $\hat{\zeta}_1^*$, and pooling overtakes screening at ζ_1^*: formulae for the thresholds in terms of the basic parameters are given in Appendix A.

The set of equilibria is determined by the upper envelope of the two screening lines and the band between the two pooling lines in Figure 2. For example, suppose the parameters $(\theta,\rho_L,\rho_H,\delta)$ are such that these lines lie in the positions shown in Figure 2a. Then if $\zeta_1 > \zeta_1^*$, there is no screening equilibrium, and there is a continuum of pooling equilibria; if $\zeta_1 < \bar{\zeta}_1^*$ there is a unique (partial screening) equilibrium path; if ζ_1 is slightly above $\bar{\zeta}_1^*$ there is a unique (screening) equilibrium path; and otherwise screening and pooling equilibria co-exist, together with equilibria in which the seller randomizes in the first period.

An important feature of the strategies in each of these equilibria is that the seller's strategy is not Markov: it uses the function $\ell(p_1)$ to determine p_0, and $\ell(p_1)$ is not constant as p_1 varies over the interval $(1+\theta-\delta\rho_H,1+\theta]$, even though p_1 is not payoff-relevant in the last period. This is unavoidable: as Maskin and Tirole (1994) point out, in bargaining games with a finite number of types there is generally no equilibrium in Markov strategies.

Implications of Forward Induction.

In examples P1 and P2, the high-valuation buyer gets an immediate payoff of at least 1 by accepting the first-period offer, whereas the maximum possible payoff from rejection is $\delta\rho_H < 1$. On the other hand the low-valuation buyer's equilibrium payoff is zero in example P1, whereas rejection of p_1 would yield an expected payoff of $\delta(1-\rho_L)$ if the seller believed that rejection signaled the low valuation. Forward induction then implies that the seller should indeed interpret an unexpected rejection as a signal that the buyer's valuation was low, since the low-valuation buyer might gain by rejection, while the high-valuation buyer necessarily loses. This means that the equilibrium in example P1 is not a forward-induction equilibrium, in the sense of Cho (1987).[10] The same argument rules out all pooling equilibria with $p_1 > \theta - \delta(1-\rho_L)$, so that P2 is the only pooling equilibrium that survives forward induction. The upper envelope of the lines in Figure 2 is then a kinked line, so that there is a unique equilibrium path for (almost) any value of ζ_1.

The most surprising feature of the forward induction equilibrium in the two-period game is that, if there are informational rents, they accrue to *both* types of the buyer. This contrasts with the usual result that a pooling equilibrium drives the low type down to the reservation utility level, conceding an informational rent to the high type. This result is explored further in Kennan (1998), and it plays a central role in the analysis of the infinite-horizon game below.

4. Cyclic Screening Equilibria

The analysis of the two-period game shows that the class of sequential equilibria includes various possibilities that seem artificial, and unlikely to arise in real bargaining situations. The analysis of the infinite-horizon game therefore proceeds by imposing a set of "reasonable" properties that are motivated by the heuristic discussion in Section 2, tempered by the analysis of the two-period game. These properties are roughly as follows. The strategies are cyclic, and are driven by the current belief of the seller. The forward induction argument used to select an equilibrium in the two-period game motivates a similar restriction on the seller's belief following rejection of a pooling offer: such a rejection would convince the seller that the buyer's current valuation is low. Partial screening equilibria[11] are ruled out: this involves a complicated restriction on the fundamental parameters which will be spelled out later, and in particular it rules out cases in which the buyer's valuations are very persistent. Prices are assumed to be tight, in the sense that any offer made by the seller leaves either the high buyer or the low buyer indifferent between acceptance and rejection; moreover, the buyer's continuation value following rejection is stationary.

The main result is the existence of a "cyclic screening equilibrium" satisfying these properties. The proof is constructive, and yields an explicit mapping from the basic parameters to the equilibrium path.

[10]Cho (1987) defines forward induction equilibrium as a refinement of sequential equilibrium in general (finite) games; this generalizes the Cho-Kreps (1987) Intuitive Criterion, which is formally defined only for signaling games. The consistency condition is strengthened to include "introspective" consistency, meaning that after a deviation by an informed player, the belief system puts no weight on types of that player which would receive less than the equilibrium payoff, if everyone behaves optimally, unless all types would lose. This definition can be used here, even though the seller's action space is infinite; Cho's existence results are then not available, but existence is proved directly in this paper.

[11]That is, equilibria in which the seller makes a screening offer not just once, but several times in succession, before inferring that the buyer's current valuation is low.

The Screening Threshold

Given a belief system Π, let $\zeta_t = \Pi_t[n_t = 0 \mid h_t^0]$ be the (marginal) probability that the buyer's current valuation is low after some history. The basic property of a cyclic screening equilibrium is that the seller's strategy compares the current belief ζ_t with a threshold belief, labeled ζ^*, and the strategy makes screening offers whenever ζ_t falls below ζ^*, and pooling offers otherwise. The main analytical task is then to determine ζ^* from the basic parameters $(\theta,\rho_L,\rho_H,\delta)$. A cyclic equilibrium exists only if ζ^* lies above the invariant probability μ, so that although the seller might not be optimistic enough to screen after inferring that the buyer's valuation was low last time, repeated pooling offers eventually lead to screening.[12]

When defining the screening threshold it is necessary to pay special attention to prices above the screening price, as was seen in the analysis of the two-period game. Recall that $\underline{\sigma}^S$ denotes the support of the seller's strategy. For any strategy σ^S, define the price ceiling $\overline{\sigma}^S$ as the highest price in $\underline{\sigma}^S$:

$$\overline{\sigma}^S = \sup\left\{ p : p \in \underline{\sigma}^S(h_t^0), h_t^0 \in \mathcal{H}_t^0, t > 0 \right\}$$

Definition A1: The strategy-belief pair (σ,Π) has the *screening threshold property* if it is consistent, and if there is a "threshold" $\zeta^* > \mu$, and a probability λ^*, such that, for all t and for all histories h_t such that $p_t \in \underline{\sigma}_t^S(h_t^0)$:

(a) if $\zeta_t < \zeta^*$, then $\underline{\sigma}_t^B(h_t) = \{n_t\}$
(b) if $\zeta_t > \zeta^*$, then $\underline{\sigma}_t^B(h_t) = \{1\}$.
(c) if $\zeta_t = \zeta^*$, and if $p_s \leq \overline{\sigma}^S$ for all $s < t$, then $\underline{\sigma}_t^B(h_t) = \{1\}$ with probability λ^*, and $\underline{\sigma}_t^B(h_t) = \{n_t\}$ with probability $1-\lambda^*$.

That is (a) if the seller is optimistic, then there will be a screening offer, and (b) if the seller is pessimistic, there will be a pooling offer. Also, (c) if the seller's belief is on the borderline between these two, there will be a pooling offer with probability λ^*, and a screening offer with probability $1-\lambda^*$, provided there has been no previous deviation by the seller that exceeded the price ceiling.[13]

Property A1 is enough to determine the probability of trade in all future periods after any history such that the price has never exceeded $\overline{\sigma}^S$. This will be shown by constructing a state variable τ that summarizes the continuation possibilities associated with the seller's current belief; then it will be shown that A1 determines a Markov chain for (n,τ), and this Markov chain determines the probability of trade.

---

[12]If instead ζ^* lies below μ, then in the long run the seller makes only pooling offers: see Section 5 below.

[13]The reason for part (c) is that, as in the two-period game, it may be necessary to allow the strategies to respond to an excessive price offer p_s by having the high-valuation buyer randomize, and to support this it may be necessary to have the seller randomize in some later period in a way that depends on p_s.

The State Variable

Suppose (σ,Π) satisfies A1. Given ζ^*, define the points $\zeta^*(s)$, for nonnegative integers s, as

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

Then if $n_t = 0$ with probability $\zeta^*(s)$, the transition probabilities for n imply $n_{t+s} = 0$ with probability ζ^*. So if the current belief is $\zeta_t = \zeta^*(s)$, consistency requires that the belief will be ζ^* after s periods, if no new information is revealed in the meantime. Let $K^* \equiv \kappa(\zeta^*)$ be the number of transitions needed to drive ζ below ζ^*, starting from $\zeta = 1$. Then K^* is defined by the inequalities $\rho_L(K^*) \leq \zeta^* < \rho_L(K^*\text{-}1)$, and

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

The seller's belief can now be partitioned as follows

Definition T: Given probabilities ζ^* and λ, the *threshold state variable* Ψ_λ is a mapping $\boldsymbol{\tau} = \Psi_\lambda(\zeta)$ from [0,1] to $\mathbb{R}$ defined as follows. If $\zeta < \zeta^*$, then $\boldsymbol{\tau} = 0$; if $\zeta^*(s\text{-}1) < \zeta < \zeta^*(s)$, for some positive integer s, then $\boldsymbol{\tau} = s$; otherwise $\zeta = \zeta^*(s)$ for some nonnegative integer s, and then $\boldsymbol{\tau} = s+\lambda$.

The mapping $\boldsymbol{\tau} = \Psi^*(\zeta)$, is defined as $\Psi_{\lambda*}(\zeta)$.

Figure 3: Beliefs and States

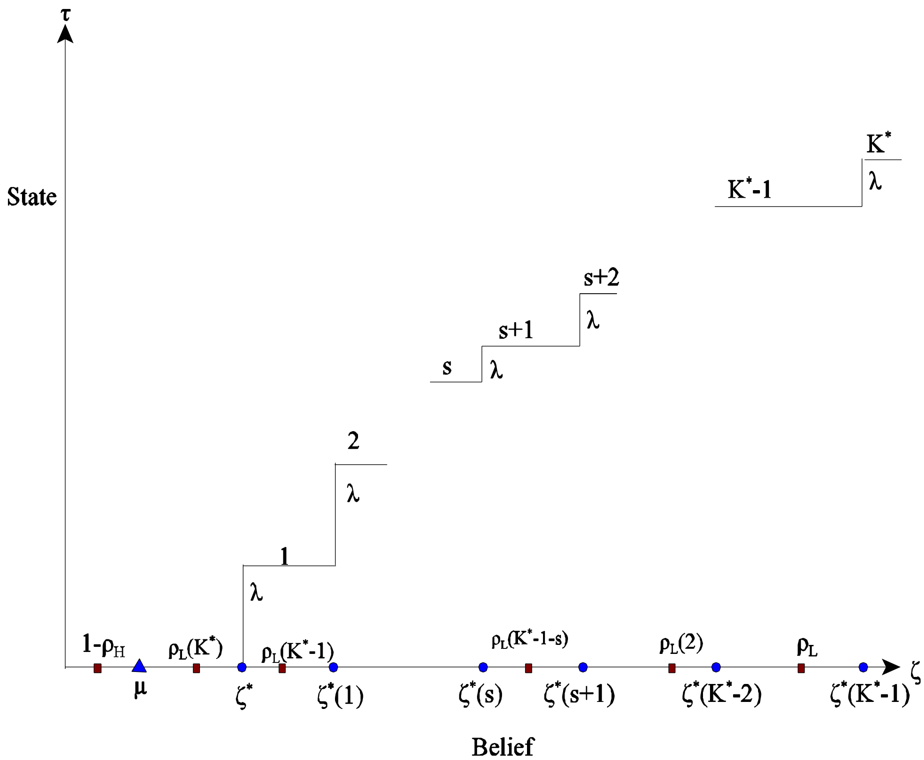

The function Ψ_λ is illustrated in Figure 3. The point of the definition is that if (σ,Π) satisfies A1, with screening threshold (ζ^*,λ^*), then after any history on the path of σ, the state $\tau = \Psi^*(\zeta)$ is sufficient to determine the continuation path: the seller will make pooling offers for the next τ periods, and (because pooling offers reveal no information) the belief will then cross the screening threshold.[14]

The equilibrium cycle sketched in Figure 1 above can now be described in detail. Consider first the case in which the threshold ζ^* is not on the path of σ.[15] Then A1 implies the following cycle. After any history such that $\zeta < \zeta^*$, the seller makes an offer that is accepted if the buyer's current valuation is high, and rejected otherwise: the state is $\tau = \Psi^*(\zeta) = 0$ for all such histories. Consistency implies that the next

[14]Only Ψ^* is relevant for the moment, because $\tau = \Psi^*(\zeta)$ on the equilibrium path. The definition of Ψ_λ will be used later when behavior off the equilibrium path is specified. In particular, if the seller deviates to a price above $\bar{\sigma}^S$, then $\tau = \Psi_\lambda(\zeta)$, for some probability λ that depends on the deviant price, until the next screening offer is made; after this, τ reverts to $\Psi^*(\zeta)$.

[15]That is, $\Pi_t[n_t = 0 \,|\, h_t^0] \neq \zeta^*$ for all histories h_t^0 that have positive probability under the strategy profile σ.

value of ζ must be $\zeta' = 1-\rho_H$ if a screening offer is accepted, and $\zeta' = \rho_L$ if it is rejected. Since $1-\rho_H = \mu(1-\varphi) < \mu < \zeta^*$, it follows that the seller again makes a screening offer in the next period after a screening offer is accepted, and this continues until a screening offer is rejected. Rejection of a screening offer implies $\zeta' = \rho_L$, and if $\rho_L > \zeta^*$ A1 then implies that the seller makes a pooling offer. Following this, consistency implies that the seller's belief moves to $\zeta'' = \rho_L(2)$, and so on until $\zeta = \rho_L(K^*) < \zeta^*$ after K^*-1 pooling offers (and $K^* = 1$ if $\rho_L < \zeta^*$, meaning that the seller is never pessimistic enough to make a pooling offer). At this point the cycle repeats itself, and so *ad infinitum*.

After any history such that $\zeta > \zeta^*$, the seller makes a sequence of pooling offers, and ζ evolves toward the invariant point μ, until it reaches or crosses ζ^*. The state variable counts the number of pooling offers remaining in the current cycle: thus $\tau = K^*-1$ if a screening offer was rejected in the previous period, with $\tau = K^*-2$ in the next period (after acceptance of a pooling offer), and so on until τ reaches 0.

Now consider the alternative case in which ζ^* is on the path of σ, meaning that $\zeta^* = \rho_L(K^*)$ for some integer $K^* \geq 1$.[16] Again, after any history such that $\zeta < \zeta^*$, screening offers are made until one is rejected, with $\tau = 0$ while this is going on. Rejection of a screening offer starts a sequence of $K^*-1+\lambda^*$ pooling offers (meaning K^*-1 for sure, and one more with probability λ^*), followed by a screening offer. The state is $\tau = \Psi^*(\rho_L) = K^*-1+\lambda^*$ following rejection of a screening offer, and the countdown to the next screening offer continues as above; when τ reaches λ^* the seller's strategy randomizes so that there is one more pooling offer with probability λ^*, and an immediate screening offer with probability $1-\lambda^*$.[17]

5. Equilibrium Continuation Values

The equilibrium continuation values will now be determined as functions of the state variable τ, for any strategy-belief pair (σ,Π) that satisfies A1 and the other "reasonable" properties mentioned above. First, it is shown that if A1 holds, then the joint continuation value after any history that has not violated the price ceiling depends only on the current value of n, and on the value of τ derived from the current belief. Simple formulae for the joint continuation values are derived, for given values of (ζ^*,λ^*). Then the remaining properties of (σ,Π) are introduced, and used to compute the buyer's continuation values. This yields the seller's values as a residual. Finally, the screening threshold ζ^* and the randomization probability λ^* can be obtained as functions of the basic parameters.

The first result shows that A1 determines a Markov chain for (n,τ). (Proofs are in Appendix B).

---

[16]It might seem that this is a knife-edge that is not worth worrying about. Unfortunately, it is not: it will be shown that there is a nontrivial region of the parameter space in which the only way to construct a cyclic screening equilibrium is to let the seller randomize between pooling and screening at the end of each cycle of pooling offers, with $\zeta = \zeta^*$ at this point.

[17]The timing of the randomization here is arbitrary. For example, one could use a public randomization to assign $\tau = K^*$ with probability λ^* and $\tau = K^*-1$ with probability $1-\lambda^*$, following rejection of a screening offer. This has no effect on the probability of trade, and although it does affect the pooling prices, this is merely a change in the timing of payments, with no effect on present values.

Lemma M:

For any public history h_t^0, let $\zeta_t = \Pi_t[n_t = 0 \mid h_t^0]$ and $\tau_t = \Psi^*(\zeta_t)$. If (σ,Π) satisfies A1, then after any history h_t such that $p_s \le \bar{\sigma}^S$ for all s < t, the values of n and τ on the continuation path of σ from h_t are determined by a Markov chain with state space $\{0,1\} \times \{0,\lambda^*,1+\lambda^*,\ldots K^*-1+\lambda^*\}$ and the following transition probabilities from (n,τ) to (n',τ'):

if n = 1,	then	$n' = 0$ w.p. $1-\rho_H$,	$n' = 1$ w.p. ρ_H,
and		$\tau' = \max(\tau-1,0)$;	
if n = 0,	then	$n' = 0$ w.p. ρ_L,	$n' = 1$ w.p. $1-\rho_L$,
and if $\tau = 0$,		then $\tau' = K^*-1+\lambda^*$	
if $\tau = \lambda^*$,		then $\tau' = K^*-1+\lambda^*$ w.p. $1-\lambda^*$	$\tau' = 0$ w.p. λ^*
if $\tau \ge 1$,		then $\tau' = \tau-1$	

Under the conditions of Lemma M, the Markov chain for (n,τ) determines the probability of trade in all future periods. In fact, if $\tau_t = 0$ then $q_t = n_t$, if $\tau_t \ge 1$ then $q_t = 1$, and if $\tau_t = \lambda^*$ then $q_t = 1$ with probability λ^* and $q_t = n_t$ with probability $1-\lambda^*$. Since the joint payoffs in each period depend only on n and q, this immediately implies that the joint continuation values are determined by n and τ.

Let $J^L(\tau)$ and $J^H(\tau)$ be the joint continuation values from state τ, depending on whether the buyer's current valuation is low or high; and let $j(\tau) = J^H(\tau) - J^L(\tau)$. Define the discounted persistence parameter β as $\beta = \delta\varphi$, and for nonnegative integers s define the discounted sums D, B and R as

$$D(s) = \sum_{i=0}^{s-1} \delta^i = \frac{1-\delta^s}{1-\delta} \;, \quad B(s) = \sum_{i=0}^{s-1} \beta^i = \frac{1-\beta^s}{1-\beta}$$

$$R(s) = \sum_{s=0}^{s-1} \delta^i \rho_L(i) = \sum_{i=0}^{s-1} \delta^i\left[\mu + (1-\mu)\varphi^i\right] = \mu D(s) + (1-\mu)B(s)$$

Also, for $\tau = s+\lambda$ where $\lambda \in [0,1]$, define $D(\tau) = D(s) + \lambda\delta^s$, and similarly for $B(\tau)$ and $R(\tau)$.

The following result determines the joint continuation values from the basic parameters $(\theta,\delta,\rho_L,\rho_H)$, for any strategy-belief pair (σ,Π) that satisfies the screening threshold property The joint values are summarized in the top panel of Table 1.

Proposition J1:

Suppose that the pair (σ,Π) satisfies the screening threshold property. For any public history h_t^0 such that $p_s \le \bar{\sigma}^S$ for all s < t, let $\zeta_t = \Pi_t[n_t = 0 \mid h_t^0]$ and $\tau = \Psi^*(\zeta_t)$. Then

$$j(\tau) = \frac{1}{1-\beta} + \frac{\beta^\tau}{1-\beta}\,\frac{\theta}{R(K^*+\lambda^*)} \quad \text{and} \quad J^L(\tau) = \frac{1+\theta-[1-\delta\rho_H]\,j(0)}{1-\delta} + \frac{R(\tau)\theta}{R(K^*+\lambda^*)}$$

where $\tau = \Psi^*(\zeta_t)$ and $\zeta_t = \Pi_t[n_t = 0 \mid h_t^0]$.

Table 1: Value Functions

	Low	High	Difference
Joint	$J^L(\tau) = \dfrac{1+\theta-[1-\delta\rho_H]\,j(0)}{1-\delta} + R(\tau)\,r^*\theta$	$J^H(\tau) = \dfrac{1+\theta-\delta(1-\rho_H)\,j(0)}{1-\delta} + [\,R(\tau)-B(\tau)\,]\,r^*\theta$	$j(0) = \dfrac{1+r^*\theta}{1-\beta}$ $j(\tau) = j(0) - r^*\theta B(\tau)$
Buyer	$V^L = \dfrac{1-\rho_L}{(1-\delta)\varphi}\,d(0)$	$V^H(\tau) = \dfrac{\rho_H-\beta}{(1-\delta)\varphi}\,d(0) + B(\tau)\,b^*$	$d(0) = \dfrac{1-b^*}{1-\beta}$ $d(\tau) = d(0) + b^*B(\tau)$
Seller	$U^L(\tau) = \dfrac{1+\theta-(1-\delta\rho_H)g(0) - \frac{\rho_H}{\varphi}[1-b^*]}{1-\delta} + R(\tau)\,r^*\theta$	$U^H(\tau) = \dfrac{1+\theta-\delta(1-\rho_H)g(0) - \frac{\rho_H}{\varphi}[1-b^*]}{1-\delta} + R(\tau)\,r^*\theta - B(\tau)[b^*+r^*\theta]$	$g(0) = \dfrac{b^*+r^*\theta}{1-\beta}$ $g(\tau) = g(0) - [\,b^*+r^*\theta\,]\,B(\tau)$
	$U(\zeta,\tau) = \dfrac{1+\theta}{1-\delta} - \dfrac{\rho_H(1-b^*)}{\varphi(1-\delta)} - \dfrac{\delta(1-\rho_H)}{1-\delta}g(0) - \zeta g(0) + R(\tau)\,r^*\theta - (1-\zeta)B(\tau)[b^*+r^*\theta]$		
Notation	$\mu = (1-\rho_H)/(1-\varphi)$ $R(K,\lambda) = \mu D(K,\lambda) + (1-\mu)B(K,\lambda)$ $b^* = 1/B(K^*,\lambda^*)$, $r^* = 1/R(K^*,\lambda^*)$ $B(\tau) = B(\tau,0)$, $R(\tau) = R(\tau,0)$	$D(K,\lambda) = \dfrac{1-\delta^K(1-\lambda+\lambda\delta)}{1-\delta}$	$B(K,\lambda) = \dfrac{1-\beta^K(1-\lambda+\lambda\beta)}{1-\beta}$

The Buyer's Continuation Values

Three additional properties are now introduced, and it is shown that for any pair (σ,Π) that satisfies these properties, the buyer's continuation values are determined by the Markov chain for (n,τ) derived in Lemma M. Let $V^L(\tau)$ and $V^H(\tau)$ denote the seller's continuation values from state τ, depending on whether the buyer's current valuation is low or high, and let $d(\tau) = V^H(\tau) - V^L(\tau)$.

Definition A2: The pair (σ,Π) has the *monotonic beliefs* property if for all t and for all histories h_t, if $p_t \le \bar{\sigma}^S$ and $q_t = 0$, then $\Pi_{t+1}[n_t = 0 \,|\, h^0_{t+1}] = 1$.

This means that rejection of any price below the ceiling $\bar{\sigma}^S$ convinces the seller that the current valuation is low. This can be rationalized by an appeal to forward induction, as in the two-period game. In particular, no matter what inference the seller might draw, the continuation value following rejection of a pooling offer in a cyclic screening equilibrium is strictly lower for the high-valuation buyer than the equilibrium value. This is not true for the low-valuation buyer, and so forward induction would convince the seller that the valuation is low if a pooling offer were to be rejected.

Property A2 implies that the history of prices and quantities in earlier periods has no relevance for future payoffs. It then makes sense that the continuation values from this point should not depend on history. This motivates the following

Definition A3: The strategy profile σ has the *stationary values property* if there are two numbers, $V^H(0)$ and V^L such that, for all t and for all histories h_t:

if $p_t \le \bar{\sigma}^S$, then $E_{A(\sigma|h_t,q_t=0)}\tilde{V}(a) = \begin{cases} V^H(0) & \text{if } n_t = 1 \\ V^L & \text{if } n_t = 0 \end{cases}$

That is, for every offer on the continuation path of σ from any history, the buyer's continuation value following rejection depends only on the current valuation (and not on previous prices and quantities, or on calendar time).

Definition A4: The strategy profile σ has the *tight pricing property* if for all t and for all histories h_t: if $p_t \in \underline{\sigma}^S_t(h^0_t)$, and either (i) $\underline{\sigma}^B_t(h_t) = \{1\}$ with $n_t = 0$, or (ii) $\underline{\sigma}^B_t(h_t) = \{n_t\}$ with $n_t = 1$, then

$$E_{A(\sigma\,|\,h_t,q_t=0)}\tilde{V}(a) \;=\; E_{A(\sigma\,|\,h_t,q_t=1)}\tilde{V}(a)$$

where the expectations are taken over all continuations of the stochastic process $A(\sigma)$ starting from h_t followed by $q_t = 0$ (on the left) or $q_t = 1$ (on the right).

Tight pricing means that, for any pooling offer on the continuation path of σ, the buyer is indifferent between acceptance and rejection if the current valuation is low; and for any screening offer on the path, the buyer is indifferent if the current valuation is high. It might seem intuitively obvious that this is implied by sequential optimality for the seller, but the P3 equilibrium in the two-period model shows that the intuition is wrong.[18]

Properties A1, A3 and A4 together imply that for any history that has not breached the price ceiling $\bar{\sigma}^S$ the low-valuation buyer's continuation value is a constant, V^L. To see this, note that A1 implies that the seller makes either a pooling or a screening offer after any such history. The stationary values property immediately implies that the low buyer's continuation value is V^L for any history such that a screening offer is made; and the tight pricing property implies that the low buyer's continuation value when a pooling offer is made is the value of rejecting the offer, which is again V^L.

The next result shows that any strategy-belief pair satisfying A1-A4 implies the same screening and pooling prices, and the same value function for the buyer.

Lemma B1:

Suppose that (σ,Π) satisfies A1-A4. For any public history h_t^0, let $\zeta_t = \Pi_t[n_t = 0 \mid h_t^0]$ and $\tau_t = \Psi^*(\zeta_t)$. Then for any history h_t such that $p_s \leq \bar{\sigma}^S$ for all $s < t$,

(a) the price specified by σ^S depends only on τ_t, and

$$\underline{\sigma}_t^S(h_t^0) = \{p_H\}, \qquad \text{if } \tau_t = 0$$

$$\underline{\sigma}_L^S(h_t^0) = \{p_L(\tau_t)\}, \quad \text{if } \tau_t \geq 1$$

where $K^* = \kappa(\zeta^*)$ and the screening price p_H and the pooling prices $p_L(\tau_t)$ are defined as

$$p_H = 1 + \theta - \delta\rho_H \frac{B(K^* - 1 + \lambda^*)}{B(K^* + \lambda^*)}$$

$$p_L(\tau_t) = \theta - \delta(1 - \rho_L) \frac{B(K^* - 1 + \lambda^*) - B(\tau_t - 1)}{B(K^* + \lambda^*)}$$

with the understanding that B(-1) is defined as zero, so $p_L(\tau) = p_L(0)$ for $\tau \leq 1$.

(b) the buyer's continuation values depend only on n_t and τ_t, and

[18]On the other hand, it might be possible to show that A4 is implied by A0-A3.

$$V^L = \frac{\delta(1-\rho_L)}{(1-\delta)(1-\beta)} \frac{B(K^*-1+\lambda^*)}{B(K^*+\lambda^*)}$$

$$d(\tau_t) = \frac{\beta}{1-\beta} \frac{B(K^*-1+\lambda^*)}{B(K^*+\lambda^*)} + \frac{B(\tau_t)}{B(K^*+\lambda^*)}$$

Remarks

I. 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.

II. Changes in θ are absorbed entirely by the seller, with no effect on the buyer's continuation value, for given values of K^* and λ^*. 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.

Informational Rents

A novel feature 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, with the result that the pooling sequence is re-started. This is worth something to the low buyer, because of the prospect of making transitions to the high valuation while the pooling sequence is in progress, so the pooling price has to be below θ in order to cover the option value associated with restarting the pooling sequence. Moreover, for $\tau > 2$, this informational rent 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 Continuation Values

For any consistent strategy-belief pair (σ,Π), the seller infers that the buyer's current valuation is high if a screening offer is accepted, and infers the low valuation if it 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.[19] The results of these calculations are given in Table 1, where $U^L(\tau)$ and $U^H(\tau)$ denote the seller's continuation values from state τ, and $g(\tau) = U^H(\tau) - U^L(\tau)$. 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, with weights determined by ζ_t.

[19]Alternatively, the seller's value function can be computed directly from the equilibrium pooling and screening prices, following the method used above to obtain the joint values. This 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 is nonstrategic), because in this case the joint values are assessed differently by the two players.

6. The Equilibrium Screening Threshold

This section shows how the screening threshold ζ^* and the randomization probability λ^* are derived from the basic parameters $(\theta,\rho_L,\rho_H,\delta)$. Then ζ^* is used to construct a strategy-belief pair (σ^*,Π^*) that satisfies properties A1-A4. Finally it is shown that this pair satisfies the sequential optimality conditions in A0, which completes the proof that cyclic screening equilibria exist.

Proposition T0: Suppose that the pair (σ,Π) satisfies properties A0-A4. Then

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

where ζ^* and λ^* are given by A1, $K^* = \kappa(\zeta^*)$, and the increasing function $G(\tau)$ is defined as $R(\tau)/B(\tau)$.

Proposition T0 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 $\zeta^* \geq \mu$. On the other hand, given the parameters (ρ_L,ρ_H,δ), the formula can be used to define θ as a function of ζ^* and λ^*. The next result shows how this function can be inverted to obtain a closed-form expression for the pair (ζ^*,λ^*).

Proposition T1: Suppose that the pair (σ,Π) satisfies properties A0-A4. Then $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$, where the function Z is defined as follows

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

where

$$\bar{\rho}_L(k) = \frac{1-\rho_L(k)}{\rho_L(k)} , \quad G(k) = \frac{R(k)}{B(k)} = 1 + \mu\left[\frac{D(k)}{B(k)} - 1\right]$$

$$\theta_\infty = \frac{1-\delta\rho_H}{1-\delta}\,\frac{1-\mu}{\mu}$$

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

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

Each component of the function Z takes values in the unit interval. The ζ^* component is continuous and 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 $G(k)\bar{\rho}_L(k)$ to $G(k+1)\bar{\rho}_L(k)$. 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 Section 9.

Proposition T1 implies that the pooling prices $p_L(\tau)$ (defined in Lemma B1 above) are always positive. Since $p_L(\tau)$ is increasing in τ it suffices to show this for $\tau = 0$. If $\theta \le \bar{\rho}_L(1)$ then $\zeta^* = 1/(1+\theta)$ and $K^* = 1$, so $p_L(0) = \theta$. Otherwise $\theta > \bar{\rho}_L(1) > \delta(1-\rho_L)$, and $p_L(0) > \theta - \delta(1-\rho_L) > 0$.

7. Cyclic Screening Strategies

Suppose (σ,Π) is a strategy-belief pair satisfying properties A1-A4, with an associated screening threshold (ζ^*,λ^*) determined from the fundamental parameters by the mapping Z. Then the equilibrium path of σ is fully determined: (n,τ) follows the Markov chain specified in Lemma M, the seller offers the pooling prices $p_L(\tau)$ when τ is positive, and the screening price p_H when τ is zero, and the buyer accepts unless p_H is offered when n is zero. Thus all equilibria satisfying A1-A4 are payoff-equivalent.

This section shows how to complete the construction of cyclic screening strategies by specifying behavior and beliefs off the equilibrium path. Once this has been done, it will be verified that the construct does indeed satisfy properties A1-A4. The last step is then to show that the strategies are sequentially optimal: this finally demonstrates the existence of cyclic screening equilibria.

The screening price p_H and the pooling prices $p_L(\tau)$ are determined by Lemma B1. For $\tau \in [0,K^*-1+\lambda]$ define the partial screening prices $\bar{p}_H(\tau)$ as

$$\bar{p}_H(\tau) = 1 + \theta - \delta\rho_H \frac{B(\tau)}{B(K^* + \lambda^*)}$$

$$p_H = \bar{p}_H(K^* - 1 + \lambda^*)$$

The strategies will be written in terms of the state variable $\tau = \Psi_\lambda(\zeta)$, so it is necessary to specify how λ is determined for every possible realization of the public history. It is also necessary to specify the belief system associated with the screening threshold (ζ^*,λ^*), and these are closely related.

Definition: Given a screening threshold (ζ^*,λ^*), the *threshold belief system* $\Pi^*(\zeta^*,\lambda^*)$ is a sequence of mappings $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$ determining the seller's belief ζ_t and the randomization probability λ_t as functions of the public history h_t^0, as follows. For t = 1, ζ_1 is the seller's prior belief and $\lambda_1 = \lambda^*$. For t > 1, (ζ_t,λ_t) is determined by repeated application of the following transition rule, which uses $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$ and $\tau_t = \Psi_{\lambda_t}(\zeta_t)$, together with p_t and q_t, to specify $(\zeta_{t+1},\lambda_{t+1})$:

(a) When $q_t = 0$:

if $p_t \le p_H$ then $\zeta_{t+1} = \rho_L$, and $\lambda_{t+1} = \lambda^*$

if $p_t = \bar{p}_H(s+\lambda_0)$, $\tau_t\text{-}1 \le s+\lambda_0 < K\text{-}1+\lambda^*$, then $\zeta_{t+1} = \zeta^*(s)$, and $\lambda_{t+1} = \lambda_0$, where λ_0 solves $p_t = \bar{p}_H(s+\lambda_0)$

if $p_t > \bar{p}_H(\tau_t\text{-}1)$ then $\zeta_{t+1} = \varphi\zeta+(1\text{-}\varphi)\mu$, and $\lambda_{t+1} = \lambda^*$

(b) When $q_t = 1$:

if $p_t \le p_L(\tau_t)$ then $\zeta_{t+1} = \varphi\zeta+(1\text{-}\varphi)\mu$ and $\lambda_{t+1} = \lambda_t$

if $p_t > p_L(\tau_t)$ then $\zeta_{t+1} = 1\text{-}\rho_H$, and $\lambda_{t+1} = \lambda^*$

This definition is non-standard in two respects. First, in addition to specifying the belief as a function of the public history, it also specifies the probability that the seller will make a pooling offer if the belief coincides with the threshold.[20] Second, the definition is incomplete in that it does not say how the buyer's current action changes the seller's beliefs about the entire history of the buyer's valuations, but the missing data can easily be filled in using Bayes' rule, and in any case beliefs about the past are irrelevant for equilibria that satisfy the screening threshold property.

The update when the buyer rejects a price above p_H is illustrated in Figure 4 below.

The Seller's Strategy

Definition: The *cyclic pricing rule* associated with the screening threshold (ζ^*,λ^*) is the mapping $p = P^*(\tau)$ that determines the price p as a function of the state τ, as follows:

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

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

if $0 \le \tau \le 1$ then $p = p_L(1)$ with probability τ

$p = p_H$ with probability $1\text{-}\tau$

Definition: The *cyclic pricing strategy* $\sigma^S(\zeta^*,\lambda^*)$ associated with the screening threshold (ζ^*,λ^*) is the sequence of mappings $p_t = P^*\circ\Psi^*\circ\Pi_t^*(h_t^0)$ from public histories to prices, defined by the composition of the cyclic pricing rule P^*, the threshold state variable Ψ^*, and the threshold belief system Π^*.

[20]This can be rationalized as a secondary belief system used by the seller as a tie-breaking rule when $\zeta = \zeta^*$

This defines a strategy for the seller that is uniquely determined by the basic parameters $(\theta,\rho_L,\rho_H,\delta)$. First, the function Z determines ζ^* and λ^*, with $K^* = \kappa(\zeta^*)$. Next, K^* and λ^* are used to obtain the screening price p_H and the pooling prices $p_L(\tau)$, and these determine P^*. Then K^*, λ^*, p_H, $p_L(\tau)$ and $\bar{p}_H(\tau)$ and the numbers $\zeta^*(s) = \mu + \varphi^{-s}(\zeta^*-\mu)$ are used to construct the mapping Π^*_t, for every t. Finally, for every t, and for every public history h^0_t, the state τ_t is found by applying Ψ_λ to $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$, and the price is set by applying P^* to τ_t.

The Buyer's Strategy

Definition: The *cyclic trading rule* associated with the screening threshold (ζ^*,λ^*) is the mapping $q = Q^*(\tau,p,n,\zeta)$ that determines the quantity q as a function of the state τ, the price p, the private state n and the belief ζ, as follows:

For n = 0:

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

if $p > p_L(\tau)$ then q = 0

For n = 1:

if $p \le \bar{p}_H(K^*-1)$ then q = 1

if $p > \bar{p}_H(\tau-1)$ then q = 0

if $\bar{p}_H(s+1) < p \le \bar{p}_H(s)$, where $\tau-1 \le s \le K^*-2$ then q = 1 w. p. υ, and q = 0 w. p. $1-\upsilon$, where

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

Definition: The *cyclic trading strategy* $\sigma^B(\zeta^*,\lambda^*)$ associated with the screening threshold (ζ^*,λ^*) is the sequence of mappings from histories to quantities, defined by the composition of the cyclic trading rule Q^* with the cyclic belief system Π^* and the threshold state variable Ψ, given by $q_t = Q^*(\tau_t,p_t,n_t,\zeta_t)$, where $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$ and $\tau_t = \Psi(\zeta_t,\lambda_t)$

The strategy $\sigma^B(\zeta^*,\lambda^*)$ is uniquely determined by the basic parameters $(\theta,\rho_L,\rho_H,\delta)$. The components described above for the seller's strategy determine the trading rule Q^*, and the quantity is chosen by applying Q^* to (τ_t,p_t,n_t,ζ_t). Since the belief system implies $\zeta_t \le \zeta^*(s+1) \le 1$, the probability ν is well-defined.[21]

[21]One 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 $\bar{p}_H(K^*-1)$. The answer is in two parts. First, optimality of the seller's behavior is checked in Proposition S3 below, where the condition that deters a deviation to $\bar{p}_H(K^*-1)$ is labeled X_K. Second, examples can be found in which X_K fails, and in that case an equilibrium can be constructed

Figure 4 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 bring the end of the pooling cycle one period closer, in which case acceptance would be more attractive for the high buyer; but if all high buyers would accept such a price, then rejection would restart the pooling sequence, and this would be more attractive to the high buyer.

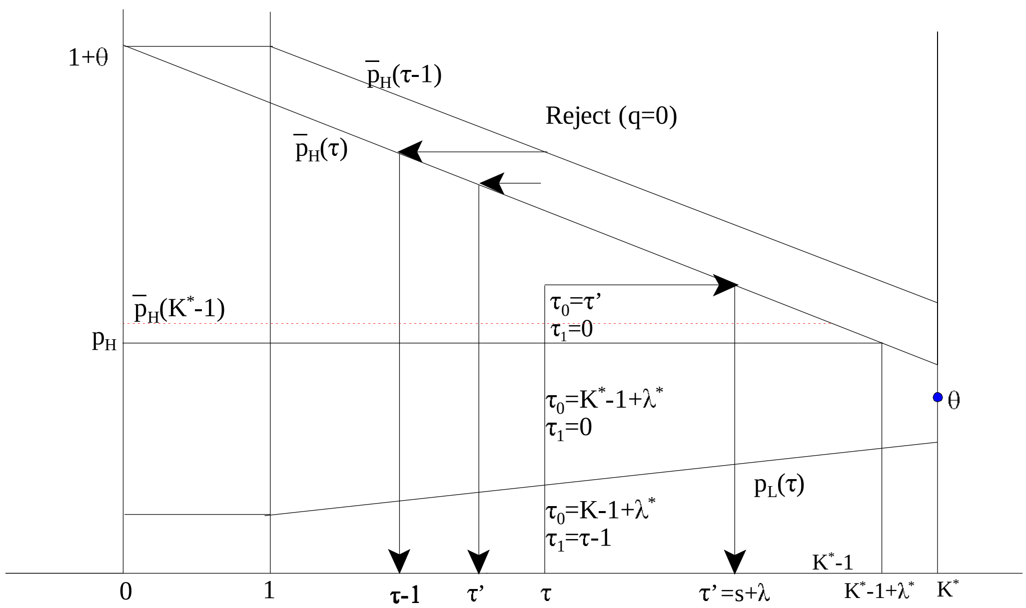

Figure 4: Randomization by the high-valuation Buyer

The complete equilibrium construct can now be formally defined.

Definition: The *cyclic screening equilibrium* for the parameter vector $(\theta,\rho_L,\rho_H,\delta)$, is the pair $(\sigma^*,\Pi^*) = \left(\sigma^N,\sigma^S(\zeta^*,\lambda^*),\sigma^B(\zeta^*,\lambda^*),\Pi(\zeta^*,\lambda^*)\right)$, where $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$.

---

in which the seller does indeed offer $\bar{p}_H(K^*\text{-}1)$ when $\zeta = 1\text{-}\rho_H$. For some parameter values this equilibrium involves incomplete screening: when $\zeta = 1\text{-}\rho_H$ the high buyer rejects the screening offer with positive probability, and $\tau' = K^*\text{-}1$, so that (if $K^* > 1$) the seller makes a pooling offer next period even though this period's valuation was not fully revealed. It is also possible to construct equilibria in which $\zeta^* = \rho_L(K^*)$, with $\tau' = K^*\text{-}1$ if a screening offer is rejected when when $\zeta = 1\text{-}\rho_H$, and $\tau' = K^*$ if a screening offer is rejected when when $\zeta = \zeta^*$.

The following results show that if the parameters $(\theta,\rho_L,\rho_H,\delta)$ lie in a region defined by a set of inequalities derived below, then (σ^*,Π^*) is indeed an equilibrium. Lemma C shows consistency, Proposition B2 shows sequential optimality for the buyer, and Proposition S3 shows sequential optimality for the seller.

Lemma C:

The strategy-belief pair (σ^*,Π^*) is consistent

Corollary:

The strategy-belief pair (σ^*,Π^*) satisfies the monotonic beliefs property.

Lemma T2:

The strategy-belief pair (σ^*,Π^*) satisfies the screening threshold property A1.

Lemma J2:

For any t, and for any public history h_t^0, let $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$, and $\tau_t = \Psi(\zeta_t,\lambda_t)$. The strategy profile σ^* implies that the joint continuation value from period t is $J^L(\tau_t) + (1-n_t)j(\tau_t)$.

Lemma B1 shows that the buyer's value function is uniquely determined by consistency and properties A1-A4. But the tight pricing property is not easy to check. The following result shows directly that σ^* implies the value function in Lemma B1, and that σ^* is sequentially optimal for the buyer.

Proposition B2:

The strategy profile σ^* is sequentially optimal for the buyer. After any history, the buyer's continuation value depends only on the current values of n and τ. When n = 0, the value is V^L, for all τ, and when n = 1, the value is $V^L + d(\tau)$, where V^L and $d(\tau)$ are as defined in Lemma B1

Corollary:

The strategy profile σ^* satisfies properties A3 and A4.

Lemmas J2 and Proposition B2 immediately imply the following result for the seller's values.

Lemma S1:

For any t, and for any public history h_t^0, let $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$, and $\tau_t = \Psi(\zeta_t,\lambda_t)$. The strategy-belief pair (σ^*,Π^*) implies that the seller's continuation value from period t is $U^L(\tau_t) + (1-\zeta_t)g(\tau_t)$, where the functions $g(\tau)$ and $U^L(\tau)$ are defined by

$$g(\tau) = \left[\frac{1}{B(K^* + \lambda^*)} + \frac{\theta}{R(K^* + \lambda^*)} \right] \frac{\beta^\tau}{1-\beta}$$

$$U^L(\tau) = \frac{1 + \theta - (1 - \delta\rho_H) g(0)}{1-\delta} - \frac{\delta\rho_H B(K^* - 1 + \lambda^*)}{(1-\delta) B(K^* + \lambda^*)} + \frac{R(\tau)\theta}{R(K^* + \lambda^*)}$$

The proof that σ^* is sequentially optimal for the seller is in two parts. First, Lemma S2 shows that there is no profitable one-period deviation below the screening price p_H, and in particular that the seller cannot gain by making a pooling offer when $\sigma^S(\zeta^*,\lambda^*)$ specifies a screening offer, or vice versa.

Lemma S2:

Let $\hat{\sigma}^S(p)$ be a strategy for the seller that selects the price p in period t, and follows the cyclic pricing strategy $\sigma^S(\zeta^*,\lambda^*)$ in all subsequent periods. Then for any public history h_t^0, and for any $p \le p_H$

$$U\left(\sigma^S(\zeta^*,\lambda^*), \sigma^B(\zeta^*,\lambda^*), \Pi^* \mid h_t^0\right) \ge U\left(\hat{\sigma}^S(p), \sigma^B(\zeta^*,\lambda^*), \Pi^* \mid h_t^0\right)$$

The second part of the optimality proof, in Proposition S3, shows that there is no profitable one-period deviation that sets a price above p_H, given that the parameters satisfy the following:

Definition A5:

The parameter vector $(\theta,\rho_L,\rho_H,\delta)$ satisfies the *partial screening conditions* if

$$\frac{1 + \theta - r^* \theta [R(s) - B(s)]}{\zeta^*(s)} - B(s)(b^* + r^* \theta) - \frac{\rho_H}{(1-\rho_H)\varphi}\left[1 - b^* B(s)\right] \ge 0 \qquad (\mathbf{X}_s)$$

for $1 \le s \le K^*$, where $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$, $K^* = \kappa(\zeta^*)$, $b^* = 1/B(K^*+\lambda^*)$, $r^* = 1/R(K^*+\lambda^*)$ and $\zeta^*(K^*) = 1$.[22]

Proposition S3:

If $(\theta,\rho_L,\rho_H,\delta)$ satisfies the partial screening conditions, then the strategy-belief pair (σ^*,Π^*) is sequentially optimal for the seller.

This completes the proof that (σ^*,Π^*) is an equilibrium.

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

8. Existence of Equilibrium

To prove existence of the cyclic screening equilibrium (σ^*,Π^*), it remains only to show that the partial screening conditions can be satisfied. This is established by the following two results, for a non-trivial region of the parameter space. The first gives a simplified test that is sufficient for the partial screening conditions. The second gives a condition on the parameters (ρ_L,ρ_H,δ) guaranteeing that the partial screening conditions hold when θ is large (i.e. when the cost of screening is high).

Proposition E1

Suppose $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$ and the following two conditions hold

$$\frac{1+\theta}{\zeta_1^*} - \frac{b^*}{\zeta^*} - \frac{\rho_H(1-b^*)}{(1-\rho_H)\varphi} \geq 0 \qquad (\mathbf{X_1})$$

and

$$\varphi^{K^*-1}\left[\frac{\rho_H}{(1-\rho_H)\varphi} - \mu\right] \leq 1-\mu \qquad (\mathbf{\overline{X}_{K-1}})$$

Then the parameter vector $(\theta,\rho_L,\rho_H,\delta)$ satisfies the partial screening conditions, and for any $\theta' \in (\theta,\theta_\infty)$, $(\theta',\rho_L,\rho_H,\delta)$ also satisfies the partial screening conditions.

Note that X_1 is a rearrangement of the inequality X_s used in the definition of the partial screening conditions with $s = 1$, while $\overline{X}_{K-1}$ is stronger than X_s with $s = K^*-1$. Also, $\overline{X}_{K-1}$ 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

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

which is impossible if ρ_H is positive. But this does not mean that X_1 fails when $K^* = 2$, as is illustrated by the example given in Table 2 below.

Proposition E2

Suppose the parameters ρ_L, ρ_H and δ satisfy the inequality[23]

---

[23]The inequality X_1^∞ is a rearrangement of X_s with $s = 1$ and $K = \infty$.

$$\frac{1-\rho_L}{\delta(1-\rho_H)} + \frac{1-\rho_L}{1-\delta} > \frac{\rho_H}{1-\varphi} - \varphi \qquad (\mathbf{X}_1^{\infty})$$

Then there is a number θ_0 such that for any $\theta \in (\theta_0,\theta_\infty)$ a cyclic screening equilibrium exists.

9. Applications and Examples

Patience and Learning

If the seller is myopic then clearly there is no point in screening to gain information, although screening may still be profitable if the high valuation is sufficiently likely. If the seller is forward-looking the value of information provides an additional motive for screening, but a forward-looking buyer will resist the seller's attempts to learn the current valuation. Rustichini and Wolinsky (1995) showed that if the buyer is not forward-looking, and the Markov chain is symmetric and highly persistent ($\rho_L = \rho_H \approx 1$) then an increase in δ implies more screening. The following proposition shows that this result is valid more generally, and that it survives even when the buyer is forward-looking.

Proposition L:

Suppose the parameter vector $(\theta,\rho_L,\rho_H,\delta)$ satisfies the partial screening conditions for all values of δ in some interval I. Then for $\delta \in$ I the screening threshold ζ^* in the cyclic screening equilibrium at $(\theta,\rho_L,\rho_H,\delta)$ is increasing in δ, with (θ,ρ_L,ρ_H) fixed.

Proof:

The function G(k) is increasing in δ, for each k. This implies that the value of θ associated with a given value of ζ^* increases with δ, and since ζ^* is a decreasing function of θ, the result follows.

The implication of this proposition is that the screening region expands as δ increases. By the same argument, when λ^* is positive, an increase in δ implies a decrease in λ^*, which again means more screening.

A Basic Example

A cyclic screening equilibrium with a cycle of length 2 is shown in Table 2. When $K^* = 2$ and $\lambda^* = 0$, the only condition that needs to be checked is X_1 (since X_K holds with equality when $\lambda^* = 0$). 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: (8/17,1), (24/17,0) and (24/17,1). 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
$\theta = ½$, $\delta = ⅛$, $\rho_L = ¾$, $\rho_H = ¾$, $\mu = ½$, $\varphi = ½$	$\zeta^* = \frac{35}{52}$, $\lambda^*=0$, $K^* = 2$; x_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	$\frac{2}{8}$	$\frac{24}{17}$
τ=0	6,528	103,632	2,240	6,160	4,288	97,472	39,232	$\frac{5}{8}$	$\frac{24}{17}$
τ=1	35,088	103,632	2,240	64,960	32,848	38,672	34,304	$\frac{6}{8}$	$\frac{8}{17}$

[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 τ = 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 $(\theta,\rho_L,\rho_H,\delta) = (5/34,41/44,7/22,1/2)$. The plot shows the piecewise-linear function

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

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

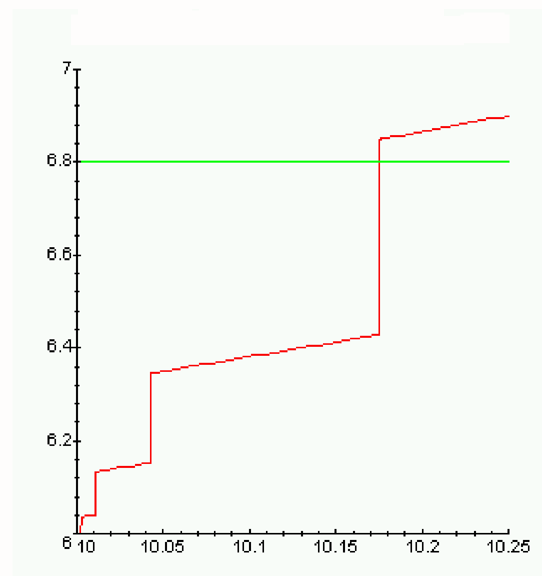

Figure 5: Random Screening Equilibrium (ρ_L=41/44,ρ_H=7/22, δ=½)

pooling offer with probability λ^* or a screening offer with probability $1-\lambda^*$. The example satisfies the conditions of Proposition E1 above (x_1 evaluates to .0608, and x_{K-1} evaluates to 79/1320 < 120/1320). The screening price is $p_H = .988214$, and the first partial screening price is $\bar{p}_H(K-1) = .988244$.

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 a cyclic screening equilibrium 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 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	
$\theta = 8/5$, $\delta = .9999$, $\rho_L = 3/4$, $\rho_H = 3/4$, $\mu = 1/2$, $\varphi = 1/2$								$\zeta^* = .528$, $\lambda^*=0$, $K^* = 5$; x_1: .795 > 0; x_{K^*-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	1.36
$\tau=2$	18,692.65	18,694.88	4,837.79	4,839.53	13,854.86	13,855.35	13,855.07	36	1.49
$\tau=3$	18,692.94	18,695.06	4,837.79	4,839.66	13,855.15	13,855.40	13,855.24	40	1.55
$\tau=4$	18,693.20	18,695.26	4,837.79	4,839.73	13,855.41	13,855.53	13,855.44	48	1.58
[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 competitive price. 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.[24] One reason for this is that a substantial piece of the pie is wasted due to unsuccessful screening offers; the other is that the low-valuation buyer gets a substantial informational rent.

10. Alternative Equilibria

This section first modifies A1 to cover equilibria in which the screening threshold ζ^* is below the invariant probability μ, so that pooling now implies pooling in all future periods. Then modifications of A2 are considered, with emphasis on what happens if the seller interprets a rejected pooling offer as an uninformative mistake.

Transient Screening Equilibria

Cyclic screening equilibria exist only if $\zeta^* \geq \mu$, which reduces to the condition $\theta \leq \theta_\infty$. If θ exceeds this bound, there are equilibria with $\zeta^* < \mu$. These may be called "transient screening equilibria," since 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 a cyclic screening equilibrium with $K^* = \infty$, with one important difference. In a cyclic equilibrium, ζ^* is the belief that leaves the seller indifferent between screening and pooling, given that pooling now implies screening next time, because the updated belief 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 $\zeta^* \geq 1-\rho_H$, the value of a screening offer in a transient screening equilibrium is exactly as it would be in a cyclic screening equilibrium with $K^* = \infty$, 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

$$U^L(0) + ng(0) \;=\; \frac{\delta\theta}{1-\delta} + n\left[1 + \frac{(1-\delta)\theta}{1-\delta\rho_H}\right]$$

The value of a pooling offer is $\theta/(1-\delta)$, and the screening threshold is obtained by finding the value of n = 1-ζ where this matches the value of screening. This yields

$$\zeta^* \;=\; \frac{1 - \dfrac{\delta(1-\rho_H)\theta}{1-\delta\rho_H}}{1 + \dfrac{(1-\delta)\theta}{1-\delta\rho_H}}$$

[24]Similar results obtain for moderate values of the discount factor (e.g. $\delta = ½$, with $\theta = 1$, and $\rho_L = \rho_H = 5/8$).

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. The details of this construction are omitted.[25]

Repeated Static Equilibria

If ζ^* is not in either the interval $[\mu,1]$ associated with cyclic screening, or the interval $[1-\rho_H,\mu]$ associated with transient screening, it must be in $[0,1-\rho_H]$. Equilibria governed by a screening threshold in this interval may be called "unconditional pooling equilibria," because even the most optimistic seller chooses pooling. In this case ζ^* equates 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 now. At the other extreme, screening cycles with $K^* = 1$ may be called "unconditional screening equilibria," since even the most pessimistic seller chooses screening, 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_L can be written as

$$\Theta_c \equiv \frac{1-\delta}{1-\delta\rho_H}\frac{1-\rho_H}{1-\rho_L} \leq \Theta \leq \frac{\rho_L}{1-\rho_L} \equiv \Theta^c$$

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

$$\Theta_t \equiv \frac{1-\rho_H}{\rho_H(1-\delta\rho_H)} \leq \Theta \leq \frac{(1-\rho_H)(1-\beta)}{(1-\rho_L)(1-\delta\rho_H)} \equiv \Theta^t$$

Finally, the condition for an unconditional pooling equilibrium is

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

---

[25]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 $1-\rho_H$ 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 $\zeta = 1-\rho_H$, it might be that $\zeta' = (1+\phi)(1-\rho_H)$ 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 $\tau=0$ or $\tau=1$ with probabilities that depend on how high the price was. Finally, transient screening equilibria must be checked against the partial screening test.

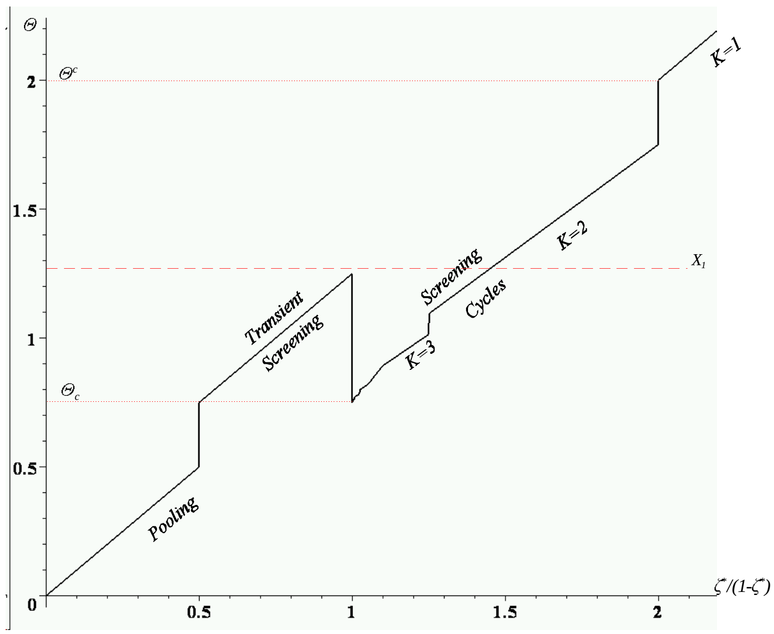

Figure 6: Alternative equilibria ($\rho_L = ⅔$, $\rho_H = ⅔$, $\delta = ½$)

These results are illustrated in Figure 6, which shows how ζ^* varies with θ, for a particular ordering of the critical values of $\Theta = 1/\theta$ defined above. The most effective diagram plots 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 screening cycles) or affine (in the transient screening case). In this example, the partial screening conditions are satisfied by all values of Θ below the dashed line labeled X_1.

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 the parameter space is divided into four regions by the inequalities $0 \leq \Theta^p \leq \Theta_t \leq \Theta^t < \infty$. 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

$$\frac{\Theta_c}{\Theta^t} = \frac{1-\delta}{1-\beta} \leq 1 \quad \text{and} \quad \Theta^p = \frac{1-\rho_H}{\rho_H} \leq \frac{\rho_L}{1-\rho_L} = \Theta^c$$

Figure 7: Alternative Equilibrium Regions (θ = ⅔, Δ = 0.7)

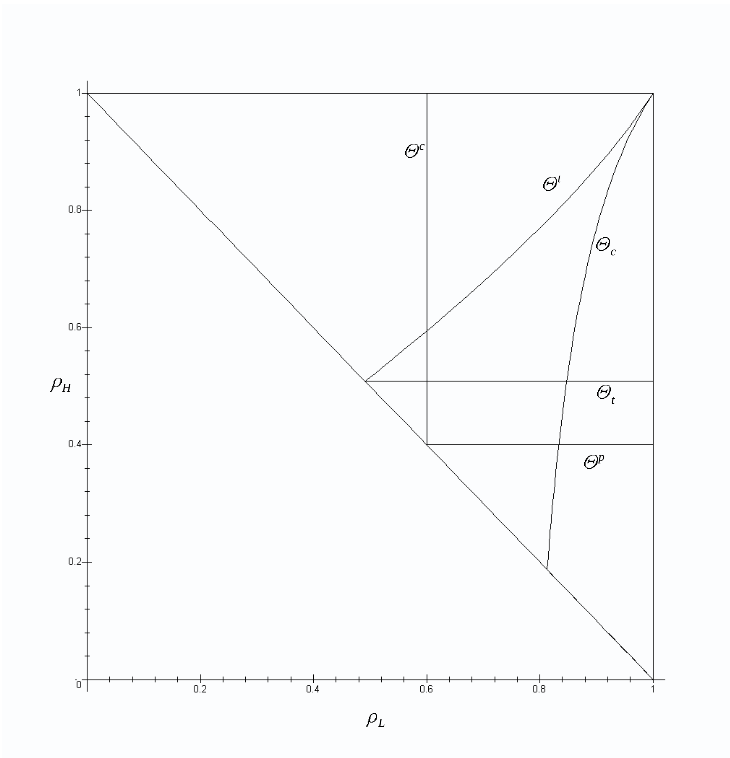

The 10 parameter regions are illustrated in Figure 7, for fixed values of θ and δ. In this diagram, any choice of (ρ_L,ρ_H) that falls in one of the five outside regions (those containing a segment with $\rho_L = 1$ or $\rho_H = 1$) gives a unique equilibrium, while a choice in one of the five inside regions yields three equilibria. For example, if (ρ_L,ρ_H) 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 $\zeta^* = 1-\rho_H$.

Alternative Beliefs

A key assumption in the above analysis is that rejection of a pooling offer would convince the seller that the buyer's current valuation is low (property A2). In the two-period model it was shown that alternative beliefs following rejection of a pooling offer can be ruled out by a forward induction argument. A leading alternative to forward induction is to simply treat rejected pooling offers as uninformative mistakes, and proceed as if they had been accepted. Then since the continuation will be the same whether the price is accepted or rejected, the pooling price is simply the low valuation, θ. But it will be shown that the derivation of the screening threshold becomes more difficult.

The following sketch shows how to derive the value functions and the screening threshold for equilibria that do not satisfy A2 (properties A1, A3 and A4 are assumed to hold, and a suitably modified form of A5 is needed to ensure that partial screening is not profitable for the seller). The continuation after a rejected pooling offer can be represented by a transition matrix α which maps the current state τ and the current actions (p,q) to the next state τ'. Then the law of motion for τ has the following structure:[26]

[26]For simplicity, it is assumed that if the equilibrium requires the seller to randomize between pooling and screening, then the state following rejected screening offers is determined by a public random variable (cf. footnote 17 above).

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 $\alpha_{\tau s}$, $s=0,1,2...K^*$
if $q = 1$, then $\tau' = \max(\tau-1,0)$

The matrix α is only used to reboot if the equilibrium crashes because a pooling offer is rejected. The method used above to compute the equilibrium prices and the value functions can be applied in this more general context as well. The joint values are unchanged, because the equilibrium quantities are unchanged, but the buyer is generally worth less. Moreover, the differential between pooling and screening continuation values for the seller is affected by α, so the 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 that d is unaffected by changes in α. Define α^0 as the transition matrix obtained when the first row of α is replaced by the vector $[0,0,...0,1-\lambda^*,\lambda^*]$. Then the general formula for the low buyer's continuation values can be written as[27]

$$V_L = \delta(1-\rho_L)\left(I-\delta\alpha^0\right)^{-1}\alpha^0 d$$

The generalized expression for the screening threshold is:

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

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

$$v_0 = V^L(0) - \sum_{s=0}^{K} \alpha_{0s}\left[\delta V^L(s) + \delta(1-\rho_L)d(s)\right]$$

which is zero in the basic specification analyzed in Section 6.

If the seller draws no inference from rejected pooling offers then $\alpha_{00} = 1$, and $\alpha_{0s} = 0$ for $s > 0$. Then the above formula for ζ^* reduces to:

$$\zeta^* = \frac{\dfrac{1-\delta\rho_H}{1-\beta} + \dfrac{(1-\rho_L)\delta^{K^*}(1-\varphi^{K^*})}{(1-\beta)(1-\varphi)D(K^*)}}{1+\dfrac{\theta}{G(K^*)}}$$

[27]Details of the following calculations are available in an unpublished appendix, at http://www.ssc.wisc.edu/~jkennan/research/App_C_cyc.pdf.

but this is not a closed-form solution, because K^* appears on the right side, and K^* depends on ζ^*. As in the basic model, the formula can be used to write θ as a function of ζ^*, with the important difference that this function is no longer monotonic, in general. Since there is no simple way to determine ζ^* from the basic parameters, the analysis of equilibria supported by these alternative beliefs is problematic.

11. 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.

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 is be suitable for empirical application: in particular, explicit solutions are obtained for equilibrium prices and quantities.

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.

Cho, In-Koo, “A Refinement of Sequential Equilibrium,” *Econometrica* 55(6), November 1987, 1367-89.

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.

Fudenberg, Drew, and Jean Tirole (1991), "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, **53**, 236-260.

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.

Keller, Godfrey and Sven Rady, “Optimal Experimentation in a Changing Environment,” GSB Research Paper No. 1443, Stanford University, July 1997.

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, “Informational Rents in Bargaining with Serially Correlated Valuations,” unpublished, University of Wisconsin-Madison, November 1998; http://www.ssc.wisc.edu/~jkennan/research.

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*, 39 (2), May 1998, 275-294.

Appendix A: The Two-Period Game

For $p_1 > \theta$, define

$$\ell(p_1) \;=\; \min\left[\max\left(\frac{1+\theta-p_1}{\delta\rho_H}, 0\right), 1\right]$$

That is, ℓ is a linear function of p_1, declining from 1 to 0, truncated above at 1 and truncated below at 0. The alternative equilibria differ only in the specification of $\ell(p_1)$ for $p_1 \le \theta$, and in the seller's choice of p_1. The strategies and beliefs are laid out in Table A1.[28]

Table A1: Equilibrium Strategies

	Last Period	First Period
Buyer	$q_0 = 1$ if $p_0 \le n_0 + \theta$ $q_0 = 0$ otherwise	If $n_1 = 0$, $q_1 = 1$ if $p_1 \le \theta-\delta(1-\rho_L)\ell(p_1)$ $q_1 = 0$ otherwise If $n_1 = 1$: $q_1 = 1$ if $p_1 \le 1+\theta-\delta\rho_H$, $q_1 = 0$ if $p_1 > 1+\theta$, $q_1 = \dfrac{1-\dfrac{\varphi\zeta_1}{\zeta_0^*-(1-\rho_H)}}{1-\zeta_1}$ if $1+\theta-\delta\rho_H < p_1 \le 1+\theta$
Seller	if $q_1 = 0$: $p_0 = \theta$ w. p. $\ell(p_1)$ $p_0 = 1+\theta$ w. p. $1-\ell(p_1)$ if $q_1 = 1$: $p_0 = 1+\theta$	***P1***: $p_1 = \theta$; ***P2***: $p_1 = \theta-\delta(1-\rho_L)$; ***P3***: $p_1 = \theta-½\delta(1-\rho_L)$ ***S***: $p_1 = 1+\theta-\delta\rho_H$ ***PS***: $p_1 = 1+\theta$

Beliefs

	$p_1 \le \theta-\delta(1-\rho)\ell(p_1)$	$p_1 > \theta-\delta(1-\rho_L)\ell(p_1)$ $p_1 \le 1+\theta-\delta\rho_H$	$p_1 > 1+\theta-\delta\rho_H$ $p_1 \le 1+\theta$	$p_1 > 1+\theta$
$q_1 = 0$	$\zeta_0 \le \zeta_0^*$ if $\ell(p_1) < 1$ $\zeta_0 \ge \zeta_0^*$ if $\ell(p_1) > 0$	$\zeta_0 = \rho_L$	$\zeta_0 = \zeta_0^*$	$\zeta_0 = \varphi\zeta_1 + 1-\rho_H$
$q_1 = 1$	$\zeta_0 = \varphi\zeta_1 + 1-\rho_H$	$\zeta_0 = 1-\rho_H$		$\zeta_0 \le \zeta_0^*$

Screening Thresholds

[28]The best pooling equilibrium for the seller is P1, while P2 is the worst, and there is a continuum of possibilities in between, with P3 in the middle of these. Figure 2a suggests there is a range of beliefs that support equilibria in which the seller randomizes between pooling and screening in the first period. Such equilibria can indeed be constructed by modifying P3 so that the pooling price is a weighted average of the prices in P1 and P2, and choosing the weight to equate the expected payoffs from pooling and screening.

The screening thresholds are defined as follows

$$\bar{\zeta}_1^* \equiv \frac{1}{1+\frac{\rho_L(1+\theta)-1}{\rho_H}\left[1+\frac{1}{\delta\left[\zeta_0^*-(1-\rho_H)\right]}\right]}$$

$$\hat{\zeta}_1^* \equiv \frac{(1+\theta-p_L)[\zeta_0^*-(1-\rho_H)]}{\varphi(1+\theta)}$$

$$\zeta_1^* \equiv \frac{1+\theta-\delta\rho_H-p_L}{1+\theta-\delta\varphi-\delta\rho_L\theta}$$

Numerical examples satisfying all of the relevant equilibrium conditions are shown in Table A2.

Table A2: Numerical Examples for the Two-Period Game

		ρ_L ¾ ρ_H ¾	θ ½	δ ½	
ζ_1		$\mathbf{q_0}$	$\mathbf{p_0}$	$\mathbf{q_1}$	$\mathbf{p_1}$
5/7	P1	n_0	3/2	1	½
	P2	n_0	3/2	1	⅜
	P3	n_0	3/2	1	7/16
⅔	S	$n_0 + (1-n_0)\chi(p_0=½)$	$p_0 = ½ + \chi(q_1=1)$	n_1	9/8
½	PS	n_0	3/2	2/5	3/2

Notes: χ is the indicator function. The pooling equilibria P1-P3 involve different choices of $\ell(p_1)$ for $p_1 \leq ½$: $\ell(p_1) = 0$ for P1, $\ell(p_1) = 1$ for P2, and $\ell(p_1) = \chi(p_1 > 7/16)$ for P3.

Appendix B: Proofs

Lemma M:

The transition probabilities from n to n' are given by σ^N. If $(n_t,\tau_t) = (1,0)$ then A1 implies a screening offer that is accepted, and consistency implies $\zeta_{t+1} = 1-\rho_H$, with $\tau_{t+1} = \Psi^*(1-\rho_H) = 0$. If $(n_t,\tau_t) = (1,\lambda^*)$ then A1 implies either a screening offer, with the same effect, or a pooling offer that is accepted, and consistency implies $\zeta_{t+1} = \varphi\zeta_t + 1-\rho_H < \zeta^*$, with $\tau_{t+1} = 0$.

If $(n_t,\tau_t) = (0,0)$ then A1 implies a screening offer that is rejected, and consistency implies $\zeta_{t+1} = \rho_L$, with $\tau_{t+1} = \Psi^*(\rho_L) = K^*-1+\lambda^*$. If $(n_t,\tau_t) = (0,\lambda^*)$ then A1 implies either a screening offer, with the same effect, or a pooling offer that is accepted, and consistency then implies $\zeta_{t+1} = \varphi\zeta_t + 1-\rho_H < \zeta^*$, with $\tau_{t+1} = 0$. Finally, if $\tau_t \geq 1$ then a pooling offer is made and accepted, and consistency implies $\zeta_{t+1} = \varphi\zeta_t + 1-\rho_H$, with $\tau_{t+1} = \tau_t - 1$.

Proposition J1:

Lemma M implies that $\tau = s+\lambda^*$ for some nonnegative integer s, and that the joint continuation value is determined by the Markov chain for (n,τ). Thus the joint continuation value $J^L(\tau)$ is the present value of s periods paying θ if n=0 and $1+\theta$ if n=1, plus the value of the continuation from $\tau = \lambda^*$ after s periods. This implies

$$J^L(\tau) = \sum_{i=0}^{s-1} \delta^i\Big[\theta+1-\rho_L(i)\Big] + \delta^s\Big[J^L(\lambda^*) + [1-\rho_L(s)]\,j(\lambda^*)\Big]$$

If $n_t = 1$, the joint continuation value $J^H(\tau)$ is as above, but with $\rho_H(i)$ in place of $1-\rho_L(i)$. Note that $\rho_H(i) + \rho_L(i) - 1 = \varphi^i$. This yields $j(\tau) = 1+\beta j(\tau-1)$. When $\tau = \lambda^*$, the continuation is as if $\tau = 0$ or $\tau = 1$, with probabilities τ and $1-\tau$, so $J^L(\lambda^*) = \lambda^* J^L(1) + (1-\lambda^*) J^L(0)$ and $j(\lambda^*) = \lambda^* j(1) + (1-\lambda^*) j(0)$.

The joint continuation values from accepted and rejected screening offers are given by

$$\begin{aligned} J^L(0) &= J^L(K^*+\lambda^*) - \theta \\ J^H(0) &= J^L(0) + j(0) = 1+\theta+\delta\Big[J^L(0) + \rho_H j(0)\Big] \end{aligned}$$

So

$$\begin{aligned} J^L(\tau) &= D(\tau)(1+\theta) - R(\tau) + \big[1 - D(\tau)(1-\delta)\big]J^L(0) + \Big[(1-\beta)R(\tau) - (1-\delta\rho_H)D(\tau)\Big]j(0) \\ &= J^L(0) + R(\tau)\big[(1-\beta)j(0) - 1\big] \end{aligned}$$

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

$$\begin{aligned} J^L(0) &= J^L(0) + R(K^*+\lambda^*)\big[(1-\beta)j(0) - 1\big] - \theta \\ (1-\delta)J^L(0) &= 1+\theta - (1-\delta\rho_H)j(0) \end{aligned}$$

and these can be solved for $J^L(0)$ and $j(0)$:

$$j(0) = \frac{1+r^*\theta}{1-\beta} \;, \quad J^L(0) = \frac{1+\theta-(1-\delta\rho_H)j(0)}{1-\delta}$$

Finally, $j(\tau) = j(0) - B(\tau)r^*\theta$, which completes the proof.

Lemma B1:

First, $\tau = 0$ if and only if $\zeta_t < \zeta^*$. Then the screening threshold property implies a screening offer, and the tight pricing property implies that the high buyer's continuation value is the value of rejecting, which is $V^H(0)$. So for any history such that $\tau = 0$, the buyer's value is V^L if $n = 0$, and $V^H(0) = V^L + d(0)$ if $n = 1$.

The price offered by the seller when $\tau = 0$ leaves the high buyer indifferent between acceptance and rejection. Acceptance means that the seller's belief next period will be $\zeta' = 1-\rho_H < \zeta^*$, which implies $\tau' = 0$, so the buyer's value next period will be V^L with probability $1-\rho_H$, and $V^H(0)$ with probability ρ_H. The high buyer's indifference means that the screening price p_H offered by the seller must satisfy the equation

$$V^L + d(0) \; = \; 1 + \theta - p_H + \delta [V^L + \rho_H d(0)]$$

Thus for any history such that $\tau = 0$, the price offered by the seller can be written as

$$p_H \; = \; 1 + \theta - \left(1 - \delta \rho_H\right) d(0) \; - \; (1-\delta) V^L$$

This proves that the price is constant over all histories such that $\tau = 0$.

Now consider any history such that $\tau = 1$. Then the screening threshold property implies that the seller makes a pooling offer, and the tight pricing property implies that the low buyer is indifferent between accepting and rejecting this offer. Acceptance means $\tau' = 0$ next period, and the low buyer's indifference means that the price $p_L(1)$ must satisfy the equation

$$V^L \; = \; \theta - p_L(1) + \delta V^L + \delta\left(1 - \rho_L\right) d(0)$$

Thus the price offered by the seller is constant over all histories such that $\tau = 1$, and this price satisfies

$$p_L(1) \; = \; \theta + \delta\left(1 - \rho_L\right) d(0) \; - \; (1-\delta) V^L$$

The high buyer's value when $p_L(1)$ is offered is the value of accepting, so the high buyer's value is constant over all histories such that $\tau = 1$, and

$$V^H(1) \; = \; V^L + d(1) \; = \; 1 + \theta - p_L(1) + \delta V^L + \delta \rho_H \, d(0)$$

A comparison of this expression with the corresponding expression for V^L yields

$$d(1) \; = \; 1 + \beta d(0)$$

After any history such that the price ceiling has not been exceeded, $\tau = \lambda^*$ implies that the seller's belief is $\zeta = \zeta^*$, and the screening threshold property implies that there will be a pooling offer with probability λ^*, and a screening offer with probability $1-\lambda^*$. If the outcome of this randomization is a pooling offer, then the argument just given for the case $\tau = 1$ can be applied, and if the outcome is a screening offer, the argument given for $\tau = 0$ can be applied. Thus after any history such that $\tau = \lambda^*$, the price is $p_L(1)$ with probability λ^*, and p_H with probability $1-\lambda^*$,and

$$d(\lambda^*) \; = \; \lambda^* d(1) + (1-\lambda^*) d(0)$$

At this point the prices and values after any history such that $\tau \in \{0,\lambda^*,1\}$ have been obtained from the values V^L and $V^H(0)$ given by the stationary values property.

Next, suppose $\tau > 1$, and suppose that, for every value of $\tau' \leq \tau-1$, the pooling prices $p_L(\tau')$ and the values $V^H(\tau')$ are constant over all histories such that the state is τ', and that these can be expressed in terms of V^L and $V^H(0)$. Then it will be shown that the same statement holds for τ, so that all of the prices and values can be found by induction, since the induction hypothesis has just been verified for $\tau = 2$, and for $\tau = 1+\lambda^*$.

For $\tau > 1$, the screening threshold property implies that the seller makes a pooling offer, and the tight pricing property implies that the low buyer is indifferent between accepting and rejecting this offer. Acceptance means $\tau' = \tau-1$ next period, and the low buyer's indifference means that the price $p_L(\tau)$ must satisfy the equation

$$V^L \ = \ \theta - p_L(\tau) + \delta V^L + \delta\left(1-\rho_L\right) d(\tau-1)$$

Thus for any history such that $\tau > 1$,

$$p_L(\tau) \ = \ \theta + \delta\left(1-\rho_L\right) d(\tau-1) \ - \ (1-\delta) V^L$$

The high buyer's value when $p_L(\tau)$ is offered is the value of accepting, so the high buyer's value for any history such that $\tau > 1$ depends only on τ, and

$$V^H(\tau) \ = \ V^L + d(\tau) \ = \ 1 + \theta - p_L(\tau) + \delta V^L + \delta \rho_H \, d(\tau-1)$$

A comparison of this expression with the corresponding expression for V^L yields

$$d(\tau) \ = \ 1 + \beta \, d(\tau-1)$$

This completes the induction step.

Now return to the case $\tau = 0$, and note that rejection of the screening offer in this case implies $\tau' = K-1+\lambda^*$. The tight pricing property then yields

$$\begin{aligned} V^H(0) \ &= \ \delta V^L + \delta \rho_H \, d(K^*-1+\lambda) \\ V^L \ &= \ \delta V^L + \delta\left[1-\rho_L\right] d(K^*-1+\lambda) \\ d(0) \ &= \ \beta \, d(K^*-1+\lambda) \end{aligned}$$

where $d(K-1+\lambda^*)$ is determined by the induction argument. So

$$V^L \ = \ \frac{1-\rho_L}{(1-\delta)\varphi} \, d(0)$$

Finally, the differences $d(\tau)$ can be obtained by solving the following set of K+1 linear equations

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

This yields

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

Proposition T0:

According to property A1, the seller randomizes between a pooling offer and a screening offer after any history h_t such that $p_s \leq \bar{\sigma}^S$ for all s < t, and $\zeta_t = \zeta^*$. Optimality implies that the seller must be indifferent between these two alternatives. Properties A2-A4 imply that the continuation after a pooling offer, conditional on n_t, is identical to the continuation after a pooling offer with $\tau_t = 1$ (i.e. the continuation specified by (σ,Π) for values of ζ_t slightly above ζ^*). And these properties also imply that the continuation after a screening offer, conditional on n_t, is identical to the continuation after a screening offer with $\tau_t = 0$ (i.e. the continuation specified by (σ,Π) for values of ζ_t slightly below ζ^*). It follows that the seller's expected payoffs in this situation are $U^L(1) + (1-\zeta^*)g(1)$ for a pooling offer, and $U^L(0) + (1-\zeta^*)g(0)$ for a screening offer, so ζ^* must be such that these are equal. This yields

$$1 - \zeta^* = \frac{U^L(1) - U^L(0)}{g(0) - g(1)} = \frac{r^*\theta}{b^* + r\theta}$$

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

Proposition T1:

Fix the parameters ρ_L, ρ_H and δ. Using Proposition T0, the values of θ can be partitioned into adjacent pairs of intervals, indexed by k. The first pair (with k=1) is $\left[0,\frac{1-\rho_L}{\rho_L}\right],\left[\frac{1-\rho_L}{\rho_L},\frac{(1+\delta\rho_L)(1-\rho_L)}{(1+\beta)\rho_L}\right]$. If θ lies in the first of these two intervals, then $\zeta^* = 1/(1+\theta) > \rho_L$ and $\lambda^* = 0$, meaning that the cost of screening is so low that the seller always makes screening offers. If θ lies in the second interval, then $\zeta^* = \rho_L$ and $\lambda^* = \frac{\theta - \bar{\rho}_L}{\delta\left[1-\rho_L\right] - \beta\theta}$, meaning that after each rejected screening offer, the seller makes another screening offer next period with probability $1-\lambda^*$. Progressively higher values of θ fall in the first interval labeled k=2, then in the second interval labeled k=2, and so on for larger values of the index k.

Lemma C:

For any t, and for any public history h_t^0, let $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$ and $\tau_t = \Psi(\zeta_t,\lambda_t)$. For any price that is surely accepted by the buyer's strategy, Bayes' rule implies that the seller's posterior following acceptance must be the same as the prior, so the belief next time must be $\zeta' = \varphi\zeta_t + (1-\varphi)\mu$. This covers all prices that do not exceed $p_L(\tau_t)$, with $q_t = 1$. If any such price is rejected, the seller's inference that $n_t = 0$ is consistent with Bayes' rule (as is any other inference). Similarly, the seller's belief is consistent for any price that is surely rejected by the buyer's strategy; this covers all prices above $\bar{p}_H(\tau_t-1)$.

If p is a price that the buyer's strategy rejects or accepts according to whether the current valuation is low or high, then the seller must believe that (1) $n_t = 0$ if $q_t = 0$, which implies $\zeta' = \rho_L$, and (2) $n_t = 1$ if $q_t = 1$, which implies $\zeta' = 1-\rho_H$. This covers all prices in the half-open interval $(p_L(\tau_t),\bar{p}_H(K^*-1)]$.

The only remaining prices are those in the interval $(\bar{p}_H(K^*-1),\bar{p}_H(\tau_t-1)]$. The buyer's strategy accepts prices in this interval with positive probability if and only if $n_t = 1$, so the seller's belief next time is $\zeta' = 1-\rho_H$ following acceptance. The belief following rejection depends on the price. If $n_t = 1$, the buyer's strategy locates the price p in a subinterval $(\bar{p}_H(s+1),\bar{p}_H(s)]$, where $\tau-1 \leq s \leq K^*-2$, and uses s and $\zeta_t = \Pi_t^*(h_t^0)$ to determine the acceptance probability ν_t. The seller's posterior belief that the current valuation is low, if the buyer rejects p, is

$$\Pi^*(n_t=0,|p,q_t=0) = \frac{\zeta_t}{1-\upsilon(1-\zeta_t)} = \zeta^*(s+1)$$

Thus $\zeta' = \zeta^*(s)$, so the belief is consistent for prices in the interval $(\bar{p}_H(K^*-1),\bar{p}_H(\tau_t-1)]$, which completes the proof.[29]

---

[29]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. Modify the buyer's strategy so that wherever the strategy specifies q =1, the buyer instead accepts with probability 1-ε1+n, while if the strategy specifies q = 0, the buyer accepts with probability ε2-n. In

Corollary: The strategy-belief pair (σ^*,Π^*) satisfies the monotonic beliefs property.

Proof:
If $q_t = 0$ and $p_t \le p_H$ then $\zeta_{t+1} = \rho_L$. But consistency requires $\zeta_{t+1} = 1-\rho_H+\varphi\Pi_{t+1}[n_t = 0\,|\,h^0_{t+1}]$, so $\Pi_{t+1}[n_t = 0\,|\,h^0_{t+1}] = 1$.

Lemma T2:
Consider any period t, and any public history h^0_t. Let $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$, and $\tau_t = \Psi(\zeta_t,\lambda_t)$. If $\zeta_t < \zeta^*$, then $\tau_t = 0$, and σ^* specifies $p_t = p_H$ and $q_t = n_t$. If $\zeta_t > \zeta^*$, then $\tau_t \ge 1$, and σ^* specifies $p_t = p_L(\tau_t)$ and $q_t = 1$. If $\zeta_t = \zeta^*$, then $\tau_t = \Psi(\zeta^*,\lambda_t) = \lambda_t$, so it suffices to show that if $p_s \le p_H$ for all $s < t$, then $\lambda_t = \lambda^*$. But this follows immediately from the definition of $\Pi^*(\zeta^*,\lambda^*)$, because in the first period $\lambda_1 = \lambda^*$, and subsequently $p_s \le p_H$ implies either $\lambda_{s+1} = \lambda_s$, or $\lambda_{s+1} = \lambda^*$, so $\lambda_t = \lambda^*$.

Lemma J2:
This follows immediately from Proposition J1 and Lemmas C and T2 if $\lambda_t = \lambda^*$. For any value of λ_t, the continuation path of σ^* is exactly the same as if $\lambda_t = \lambda^*$ except that the first time $\tau_{t+i} = \lambda_t$ the probability of a pooling offer is λ_t instead of λ^*, with $\lambda_s = \lambda^*$ for all $s > t+i$. Thus the result holds for $s > t+i$, so $j(\tau_{t+i}) = \lambda_t j(1) + (1-\lambda_t)j(0)$, and similarly for J^L.

Proposition B2:
Given that the seller is using the cyclic pricing strategy $\sigma^S(\zeta^*,\lambda)$, the buyer's value function is the value of a dynamic programming problem in which the control variable is q and the state variable is τ. For $\tau \ge 1$, the seller offers $p_L(\tau)$, and if $q = 1$ then $\tau' = \tau-1$, and if $q = 0$ then $\tau' = K-1+\lambda$. If $\tau = 0$ the seller offers p_H and if $q = 1$ then $\tau' = 0$, and if $q = 0$ then $\tau' = K-1+\lambda$. For $0 < \tau < 1$, the seller behaves as if $\tau = 1$ with probability τ, and as if $\tau = 0$ with probability $1-\tau$. This presents the buyer with a well-defined law of motion from (τ,q) to τ', for any value of $\tau \in [0,K]$, and for $q \in \{0,1\}$. Thus the buyer's continuation value depends only on the current values of n and τ.

Let $\hat{V}^L(\tau)$ and $\hat{V}^H(\tau)$ denote the buyer's optimal continuation values when the seller plays the strategy $\sigma^S(\zeta^*,\lambda)$. The principle of optimality yields a functional equation for $\hat{V}^L(\tau)$ and $\hat{V}^H(\tau) \equiv \hat{V}^L(\tau) + \hat{d}(\tau)$, which can be written as follows

$$\begin{aligned}
\hat{V}^L(\tau) &= \max\left\{\delta\hat{V}^L(K^*-1+\lambda^*)+\delta\left[1-\rho_L\right]\hat{d}(K^*-1+\lambda^*),\ \theta-p_L(\tau)+\delta\hat{V}^L(\tau-1)+\delta\left[1-\rho_L\right]\hat{d}(\tau-1)\right\},\ \ \tau\ge 1\\
\hat{V}^L(0) &= \max\left\{\delta\hat{V}^L(K^*-1+\lambda^*)+\delta\left[1-\rho_L\right]\hat{d}(K^*-1+\lambda^*),\ \theta-p_H+\delta\hat{V}^L(0)+\delta\left[1-\rho_L\right]\hat{d}(0)\right\}\\
\hat{V}^L(\tau) &= \tau\hat{V}^L(1)+(1-\tau)\hat{V}^L(0),\quad 0<\tau<1\\
\hat{V}^H(\tau) &= \max\left\{\delta\hat{V}^L(K^*-1+\lambda^*)+\delta\rho_H\hat{d}(K^*-1+\lambda^*),\ 1+\theta-p_L(\tau)+\delta\hat{V}^L(\tau-1)+\delta\rho_H\hat{d}(\tau-1)\right\},\ \ \tau\ge 1\\
\hat{V}^H(0) &= \max\left\{\delta\hat{V}^L(K^*-1+\lambda^*)+\delta\rho_H\hat{d}(K^*-1+\lambda^*),\ 1+\theta-p_H+\delta\hat{V}^L(0)+\delta\rho_H\hat{d}(0)\right\}\\
\hat{V}^H(\tau) &= \tau\hat{V}^H(1)+(1-\tau)\hat{V}^H(0),\quad 0<\tau<1
\end{aligned}$$

The claim is that $\hat{V}^L(\tau) = V^L$, and $\hat{V}^H(\tau) = V^L + d(\tau)$. It is enough to show that if this is true on the right sides of the above equations, then it is also true for the left sides. That is, if the claim is true for the buyer's values next period, then it is also true for this period's values. Then, since the payoffs are bounded, the claim can be proved by backward induction.

Substitute $\hat{V}^L(\tau) = V^L$ and $\hat{V}^H(\tau) = V^L + d(\tau)$ on the right sides of these equations. Then, for $\tau \ge 1$, the definition of $p_L(\tau)$ implies that $\hat{V}^L(\tau)$ is the maximum of two equal numbers, and

---

particular (setting n = 1), ε2 is the probability that a pooling offer is rejected by the high-valuation buyer and (setting n = 0) ε is the corresponding probability for the low buyer (the point here is that the probability of 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). Then if a pooling offer is rejected, Bayes rule implies that the probability of the low type is $\zeta = \dfrac{\varepsilon}{\varepsilon+\varepsilon^2} = \dfrac{1}{1+\varepsilon}$. Taking $\varepsilon = 1/m$ yields a sequence ζ_m converging to 1 as m increases, supporting the belief that if a pooling offer is rejected, the buyer's valuation is low.

$$\hat{V}^{L}(\tau) \;=\; \delta V^{L} + \delta\big[1-\rho_{L}\big]d(K-1+\lambda^{*}) \;=\; \theta - p_{L}(\tau) + \delta V^{L} + \delta\big[1-\rho_{L}\big]d(\tau-1)$$

$$\hat{V}^{H}(\tau) \;=\; \max\Big\{\delta V^{L} + \delta\rho_{H}\,d(K-1+\lambda^{*}),\; 1+\theta-p_{L}(\tau) + \delta V^{L} + \delta\rho_{H}\,d(\tau-1)\Big\}$$

$$= \max\Big\{V^{L} + \beta\,d(K-1+\lambda^{*}),\, V^{L}(\tau) + 1 + \beta\,d(\tau-1)\Big\} \;=\; V^{L} + 1 + \beta\,d(\tau-1)$$

where the last equality follows from $\beta\,[d(K\text{-}1+\lambda) - d(\tau\text{-}1)] \le 1$ for $\tau \ge 1$. Thus, for $\tau \ge 1$,

$$\hat{V}^{L}(\tau) \;=\; \delta V^{L} + \delta\big[1-\rho_{L}\big]d(K-1+\lambda^{*}) \;=\; V^{L}$$

$$\hat{d}(\tau) \;=\; 1 + \beta\,d(\tau-1) \;=\; d(\tau)$$

For $\tau = 0$, the definition of p_H implies that $\hat{V}^H(0)$ is the maximum of two equal numbers, and

$$\hat{V}^{H}(0) \;=\; \delta V^{L} + \delta\rho_{H}d(K-1+\lambda^{*}) = 1+\theta-p_{H} + \delta V^{L} + \delta\rho_{H}d(0)$$

$$\hat{V}^{L}(0) \;=\; \max\Big\{\delta V^{L} + \delta\big[1-\rho_{L}\big]d(K-1+\lambda^{*}),\; \theta - p_{H} + \delta V^{L} + \delta\big[1-\rho_{L}\big]d(0)\Big\}$$

$$= \max\Big\{V^{H}(0) - \beta\,d(K-1+\lambda^{*}),\; V^{H}(0) - 1 - \beta\,d(0)\Big\} \;=\; V^{H}(0) - \beta\,d(K-1+\lambda^{*})$$

where the last equality follows from $\beta\,[d(K^{*}\text{-}1+\lambda^{*}) - d(0)] \le 1$. Thus

$$\hat{V}^{L}(0) \;=\; \delta V^{L} + \delta\big[1-\rho_{L}\big]d(K-1+\lambda^{*}) \;=\; V^{L}$$

$$\hat{d}(0) \;=\; \beta\,d(K-1+\lambda^{*}) \;=\; d(0)$$

For $0 < \tau < 1$, $\hat{V}^L(\tau) = \tau\,\hat{V}^L(1) + (1\text{-}\tau)\hat{V}^L(0) = V^L$ and $\hat{V}^H(\tau) = \tau\,\hat{V}^H(1) + (1\text{-}\tau)\hat{V}^H(0) = V^L + d(\tau)$. This completes the proof that $\hat{V}^L(\tau) = V^L$ and $\hat{V}^H(\tau) = V^L + d(\tau)$, for all τ.

It remains only to show that the buyer responds optimally to deviant prices. Clearly, acceptance of prices below $p_L(\tau)$ is optimal for $\tau \ge 1$, and similarly the definition of $p_L(0)$ is such that the low buyer's value at $\tau = 0$ satisfies the equation

$$V^{L} \;=\; \theta - p_{L}(0) + \delta V^{L} + \delta\big[1-\rho_{L}\big]d(0)$$

where the left side is the value of rejection, and the right side is the value of acceptance. Thus acceptance of any price below $p_L(0)$ is optimal for the low buyer when $\tau = 0$.

If the seller offers a price p such that $\bar{p}_H(s+1) < p \le \bar{p}_H(s)$, then $q = 1$ implies $\tau' = 0$, and $q = 0$ implies $\tau' = s+\lambda$, where λ is defined by the equation $p = \bar{p}_H(s+\lambda)$. Then the high buyer's value is

$$\delta V^{L} + \delta\rho_{H}d(s+\xi) = 1+\theta - \overline{p}_{H}(s+\xi) + \delta V^{L} + \delta\rho_{H}d(0)$$

where the left side is the value of rejecting, and the right side is the value of accepting. The definition of $\bar{p}_H(s)$ is such that

$$\overline{p}_{H}(s+\xi) = 1+\theta - \delta\rho_{H}\big[d(s+\xi) - d(0)\big]$$

Thus the buyer is indifferent between acceptance and rejection, and randomization is optimal. Finally, if the price is above $\bar{p}_H(\tau\text{-}1)$, then q = 0 implies $\tau' = \tau\text{-}1$, and the value of rejection exceeds the value of acceptance. This completes the proof.

Corollary: The strategy profile σ^* satisfies properties A3 and A4.

Proof:

As was just shown in the proof of B2, the definitions of $p_L(\tau)$ and p_H imply the tight pricing property. The stationary values property for the case $n_t = 0$ is included in B2. For $n_t = 1$, the continuation following $q_t = 0$ is the same for any value of τ_t, including $\tau_t = 0$, and the tight pricing property implies that this has value $V^H(0)$.

Lemma S2:

For any t, and for any public history h^0_t, let $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$, and $\tau_t = \Psi(\zeta_t,\lambda_t)$.

According to the buyer's strategy $\sigma^B(\zeta^*,\lambda^*)$, all prices below the current pooling price $p_L(\tau_t)$ are accepted, and all such acceptances are uninformative according to Π^*, and thus imply the same stochastic process for future payoffs. Therefore all prices below $p_L(\tau_t)$ are dominated by $p_L(\tau_t)$ from the seller's point of view. Any price in the interval $(p_L(\tau_t),p_H]$ is rejected by $\sigma^B(\zeta^*,\lambda^*)$ if $n_t = 0$, and accepted if $n_t = 1$. 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.

For any history, if the seller offers p_H the continuation is exactly as it would be for a history such that $\tau_t = 0$. That is, the buyer rejects if $n_t = 0$, with $\zeta_{t+1} = \rho_L$ and $\tau_{t+1} = K\text{-}1+\lambda^*$, and the value of this continuation for the seller is $U^L(0)$; and the buyer accepts if $n_t = 1$, with $\zeta_{t+1} = 1\text{-}\rho_H$ and $\tau_{t+1} = 0$, and the value of this continuation for the seller is $U^L(0) + g(0)$. Thus after any history the value of offering p_H is $U^L(0) + (1\text{-}\zeta_t)g(0)$, which is linear in ζ_t, and decreasing.

Recall that for $\tau_t \le 1$, the pooling price $p_L(\tau_t)$ does not depend on τ_t. Conditional on n_t, the continuation when $p_L(\tau_t)$ is offered with $\tau_t \le 1$ is exactly as it would be if $\tau_t = 1$: the price $p_L(\tau_t) = p_L(1)$ is accepted, with $\tau_{t+1} = 0$. The seller's expected value is then $U^L(1) + (1\text{-}\zeta_t)g(1)$, which is linear in ζ_t, and decreasing less rapidly than the value of a screening offer, because $g(1) = g(0)\text{-}(b^*+r^*\theta)$. Lemma T0 showed that the threshold ζ^* is defined by the equation $U^L(0) + (1\text{-}\zeta^*)g(0) = U^L(1) + (1\text{-}\zeta^*)g(1)$, so for $\tau_t \le 1$, p_H is optimal for the seller if $\zeta_t \le \zeta^*$, and $p_L(\tau_t)$ is optimal if $\zeta_t \ge \zeta^*$.

Lemma T2 showed that (a) if $\zeta_t < \zeta^*$, then $\tau_t = 0$, and (b) if $\zeta_t > \zeta^*$, then $\tau_t \ge 1$. It follows that if $\tau_t = 0$ then case (b) is ruled out, so $\zeta_t \le \zeta^*$, and if $\tau_t \ge 1$ then case (a) is ruled out, so $\zeta_t \ge \zeta^*$, and if $0 < \tau_t < 1$ then both cases are ruled out, so $\zeta_t = \zeta^*$. This proves that p_H is optimal for the seller if $\tau_t = 0$, and $p_L(\tau_t)$ is optimal if $\tau_t = 1$, and randomization is optimal if $0 < \tau_t < 1$, so the strategy $\sigma^S(\zeta^*,\lambda^*)$ is optimal after any history such that $\tau_t \le 1$.

It remains only to show that $p_L(\tau_t)$ is optimal for the seller if $\tau_t > 1$, meaning that the value $U^L(0) + (1\text{-}\zeta_t)g(0)$ obtained by offering p_H does not exceed $U^L(\tau_t) + (1\text{-}\zeta_t)g(\tau_t)$, which is the value of following $\sigma^S(\zeta^*,\lambda^*)$. To verify this, note that

$$\begin{aligned}
U^L(\tau) + (1-\zeta)g(\tau) - U^L(0) - (1-\zeta)g(0) &= r^*\theta R(\tau) - (1-\zeta)(b^* + r^*\theta)B(\tau) \\
&= r^*\theta\left[R(\tau) - \frac{1-\zeta}{1-\zeta^*}B(\tau)\right] \\
&\ge r^*\theta B(\tau)\left[1 - \frac{1-\zeta}{1-\zeta^*}\right] = r^*\theta B(\tau)\left[\frac{\zeta-\zeta^*}{1-\zeta^*}\right]
\end{aligned}$$

since $R(\tau) \ge B(\tau)$. As was noted above, $\tau_t > 1$ implies $\zeta_t \ge \zeta^*$, so this inequality proves that $p_L(\tau_t)$ is optimal if $\tau_t > 1$.

Proposition S3:

For any t, and for any public history h^0_t, let $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$ and $\tau_t = \Psi(\zeta_t,\lambda_t)$.

Lemma S2 shows that no price below p_H yields a higher payoff than the price specified by $\sigma^S(\zeta^*,\lambda^*)$. Prices above $\bar{p}_H(\tau_t\text{-}1)$ are rejected by $\sigma^B(\zeta^*,\lambda^*)$, and such acceptances are uninformative according to Π^*, so these prices are dominated by the pooling price $p_L(\tau_t)$, which yields the same future payoffs, plus some current profit (recall the remark in the text that all pooling prices are positive). So what must be shown is that no price between p_H and $\bar{p}_H(\tau_t\text{-}1)$ improves the expected payoff.

Suppose $\bar{p}_H(s+1) < p \le \bar{p}_H(s)$, for some integer s, with $\tau\text{-}1 \le s \le K\text{-}1$. Then the buyer accepts with probability q(s), with continuation from $\zeta' = 1\text{-}\sigma$ 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 $p = \bar{p}_H(s+\xi)$. The seller's current payoff $\bar{p}_H(s+\xi)$ and the expected continuation value from next period on are both linear in ξ, so if a deviation to $\bar{p}_H(s+\xi)$ is profitable, a deviation to either $\bar{p}_H(s)$ or $\bar{p}_H(s+1)$ 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 $\zeta^*(s+1) \;=\; \dfrac{\hat{\zeta}}{1-q(s)}$ and in equilibrium $\hat{\zeta} = \zeta$.

A partial screening offer is either accepted by the high buyer, with continuation value u_H^a for the seller, or rejected by the high buyer (u_H^r) , or rejected by the low buyer (u_L^r). So the seller's expected payoff is given by

$$q(s)\,u_H^a + [\,1-q(s)-\zeta\,]\,u_H^r + \zeta u_L^r \;=\; u_H^a - [1-q(s)]\left[u_H^a - u_H^r\right] - \zeta\left[u_H^r - u_L^r\right]$$

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,

$$\begin{aligned} u_H^a - U^H(0) \;&=\; \delta V^L + \delta\rho_H d(K-1+\lambda^*) - \delta V^L - \delta\rho_H d(s) \\ &=\; \frac{\rho_H}{\varphi}\left[\,1 - B(s+1)\,b^*\right] \end{aligned}$$

Note here that

$$\begin{aligned} d(0) \;&=\; \beta\, d(K^* - 1 + \lambda^*) \\ d(s+1) \;&=\; 1 + \beta\, d(s) \;=\; d(0) + B(s+1)\,b^* \end{aligned}$$

These results can be summarized as :

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

Since $\zeta^*(s) \leq 1$, 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 $1-\rho_H$) then it is always positive. But the inequality X_s is just $\Omega(s)/\zeta \geq 0$ with $\zeta = 1-\rho_H$. This completes the proof.

Proposition E1

The proof of E1 uses the following mathematical result.

Lemma E0: If c is positive, and δ and φ are in the interval (0,1), then the function $f:\Re \to \Re$ defined by

$$f(s) \;=\; a\delta^s - \frac{c}{\varphi^s} - (\delta\varphi)^s \qquad \text{is quasiconcave.}$$

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

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

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 $s > s_0$.

In the following proofs, the parameters ρ_L, ρ_H and δ are fixed, and K^* and $\omega(s)$ are considered as functions of θ, using the notation $K(\theta)$ and $\omega(s;\theta)$.

Proof of Proposition E1:

The proof of Proposition S3 shows that 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

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

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

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

where $\Xi \equiv \dfrac{\rho_H}{\varphi(1-\rho_H)} - 1$. After substituting for D(s) and B(s), ω(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

$$a_\varphi = \frac{z_0}{1-\beta}\left[y_0 - \beta^{K^*}(\Xi+1)\left(1-\lambda[1-\beta]\right)\right]$$
$$a_\beta = \frac{\mu \Xi}{1-\beta}$$

The coefficient a_φ is nonnegative if $\beta^{K^*}[\Xi+1] \le y_0$. This is implied by (X_{K-1}):

$$\varphi^{K^*-1}\left[1 + \frac{\Xi}{1-\mu}\right] \le 1$$
$$\varphi^{K^*-1}[\Xi+1] \le 1 - \mu\left[1 - \varphi^{K^*-1}\right] \le 1$$
$$\beta^{K^*}[\Xi+1] \le 1 \le y_0$$

Since a_β is positive, the function $\omega(s)/a_\beta$ satisfies the conditions of Lemma E0, so ω(s) is quasiconcave. The next step is to show that X_{K-1} implies $\omega(K^*) \ge 0$ and $\omega(K^*-1) \ge 0$. Then (given X_1) quasiconcavity implies $\omega(s) \ge \min[\omega(1),\omega(K^*-1)] \ge 0$, for $1 \le s \le K^*-1$.

To show that $\omega(K^*) \ge 0$, write $\omega(K^*)$ as

$$\omega(K^*) = \lambda \delta^{K^*}\left[1 - \mu - \left(\frac{\rho_H}{\varphi(1-\rho_H)} - \mu\right)\varphi^{K^*}\right] \ge 0$$

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

$$\omega(K-1) = \left[\frac{B(K^*-1)}{\zeta^*} + \frac{\rho_H}{\varphi(1-\rho_H)}\beta^{K^*-1}\right]\left[1-\zeta^*(K^*-1)\right] - \beta^{K^*-1}\left(\frac{\rho_H}{\varphi(1-\rho_H)} - 1\right)$$
$$+ \delta^{K^*-1}\rho_L(K^*-1)\left[\frac{1}{\zeta^*} - 1\right] + \lambda\delta^K\left[1 - \mu - \left(\frac{\rho_H}{\varphi(1-\rho_H)}\rho_L - \mu\right)\varphi^{K^*}\right]$$

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

$$\rho_L(K^*-1)\left[\frac{1}{\zeta^*}-1\right] - \varphi^{K^*-1}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-1\right) \geq 1-\rho_L(K^*-1) - \varphi^{K^*-1}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-1\right)$$
$$= 1-\mu - \varphi^{K^*-1}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-\mu\right)$$

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 X_1 and X_{K-1}. It is obvious that X_{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 X_1 as

$$\frac{1}{\zeta^*(1)} + \theta\left[\frac{1}{\zeta^*(1)} - \frac{1}{R(K^*+\lambda^*)}\right] - \frac{\rho_H}{(1-\rho_H)\varphi} + \left[\frac{\rho_H}{(1-\rho_H)\varphi}-1\right]\frac{1}{B(K^*+\lambda^*)} \geq 0 \qquad (\mathbf{X_1})$$

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 $G(K^*) \leq \theta/\bar{\rho}_L(K^*) < G(K^*+1)$. In this case a (small) increase in θ increases λ while K^* and ζ^* remain unchanged.[30] It will be shown that this increases $\omega(1)$, implying that an increase in θ relaxes X_1. Write $\omega(1)$ as

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

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

$$\frac{\partial\omega(1)}{\partial\lambda^*} = \frac{\rho_H\beta^{K^*}\left[1-\zeta^*(1)\right]}{\varphi(1-\rho_H)} - \beta^{K^*}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-1\right) + \delta^{K^*}\left[1-\rho_L(K^*)\right]$$

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

$$1-\rho_L(K^*) - \varphi^{K^*}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-1\right) = 1-\mu - \varphi^{K^*}\left(\frac{\rho_H}{\varphi(1-\rho_H)}-\mu\right)$$

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

---

[30] 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 $\omega(K-1) \geq 0$ at this point.

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,

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

Substituting $\zeta^* = \mu$ in this equation shows that the inequality $\omega(1,\theta_\infty) > 0$ is equivalent to (X_∞). 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 X_{K-1} necessarily holds for large K^*, θ_0 can be chosen so as to ensure that X_{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.