Repeated Bargaining with Persistent Private Information

John Kennan[1]
University of Wisconsin-Madison

November, 1999

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 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. Because there is persistence in the Markov chain generating the surplus, a successful demand induces the seller to make another aggressive demand in the next negotiation, since the buyer's acceptance reveals that the current surplus is large. 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, after several such offers have been accepted, the seller is optimistic enough to again make an aggressive demand, creating an 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@ssc.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, Steve Tadelis 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. If the two excellent referees have not been thanked already, I thank them too. 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 valuation than its opponent does. If the valuations are correlated over time, information revealed in the negotiation of one contract has strategic value in subsequent negotiations, and this must be taken into account in choosing an optimal bargaining strategy. The paper analyzes the strategic role of serially correlated private information in this context.

A trader who is uncertain about the other side's valuation faces the usual monopoly tradeoff between prices and quantities. In the most basic case there are two choices: a pooling offer that ensures trade, at a low price, or a screening offer that specifies a high price at the risk of a failure to trade. In repeated bargaining with persistent (but not permanent) private information this choice is more complicated, because screening reveals information that will be valuable in future negotiations. This naturally leads to cycles: an unsuccessful screening offer induces pessimism in subsequent negotiations, but the pooling offers induced by this pessimism do not reveal information, and so the pessimism wears off.

Recent work on labor contracts has emphasized the possibility of explaining collective bargaining outcomes in terms of the incentive-compatibility constraints arising when either party has unverifiable private information. This literature deals with static bargaining, in the sense that the game ends as soon as a contract is signed.[2] The main emphasis in applications of the static model is on the relationship between wages and strike durations: if a strike is a screening device, then an employer must endure a long strike to reach a low-wage contract. 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. In particular, the probability of a strike rises to about 39% if there was a short strike in the previous contract negotiation, compared with probabilities between 10% and 20% following

[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). Hart and Tirole (1988) analyzed repeated bargaining, but with a static information structure.

either no strike or a long strike.[3] Since each negotiation is about dividing future surpluses, it is difficult to see why past outcomes should affect current negotiations.

Such history-dependent outcomes admit a natural explanation if private valuations are serially correlated. Indeed, if a short strike reflects an early concession by the employer, then the union infers that the employer's valuation is likely to be high again in the next negotiation, so the union makes a screening offer; then in the (relatively unlikely) event that the valuation has fallen, a strike will ensue. On the other hand a long strike reflects a refusal by the employer to concede to a screening offer. Then the union infers that the employer's valuation is likely to be low again in the next negotiation, so a pooling offer is optimal, implying that long strikes are followed by peaceful settlements.

The paper builds a dynamic bargaining model in which this intuition can be analyzed. To obtain results, some strong simplifying assumptions are needed. There is private information only on one side; the uninformed party is called the seller, and the informed party is the buyer. 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; at the other extreme, the current valuation is purely transitory, and the model reduces to a sequence of one-shot screening negotiations. The seller has the right to make offers that the buyer must accept or reject, so the ability to signal is severely restricted (but not eliminated, as will be seen). The seller can commit to a single offer in each negotiation, but cannot make any commitment regarding future negotiations. Although this entails a considerable sacrifice of realism (with respect to labor contracts, for example), suppressing the details of the screening process within each contract negotiation provides a much sharper focus for the analysis of linkages across contracts, which is the main point of the paper.

The analysis centers on the strategic interactions arising from forward-looking behavior: while the seller thinks about the value of learning the current valuation 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 transition probabilities and the discount factor). The equilibrium is unique in the sense that it has the same payoffs as every other equilibrium that satisfies a set of properties motivated by the heuristic description of cyclic screening and pooling given above.

[3]See, Riddell (1979, 1980), Card (1988, 1990), and Ingram, Metcalf and Wadsworth (1993). The relationship between strike incidence and previous strike duration in U.S. was discovered by Card (1988), who also showed that the relationship is not merely due to unobserved heterogeneity across bargaining pairs. Kennan (1995) confirmed the same relationship in Canadian data, and outlined the theoretical explanation mentioned in the text.

A novel feature of the equilibrium is that an informational rent accrues to both buyer types. This contrasts with the usual result that a pooling equilibrium drives the low type down to a reservation utility level, conceding an informational rent to the high type. The key to this result is the inference that the seller would draw if a pooling offer were to be rejected. No matter what the inference might be, the high-valuation buyer's continuation value is strictly less than the value of accepting the pooling offer, while the low-valuation buyer is indifferent between acceptance and rejection. A forward induction argument then implies that rejection would convince the seller that the current valuation is low. This means that the buyer has the option of resetting the seller's belief to its most pessimistic level, thereby ensuring that pooling offers will be made in the next few contracts. This option is valuable to the low buyer, because the Markov chain might make a transition to the high valuation while pooling offers are being made. Thus the pooling price must be below the low valuation, in order to buy out the value of this option. In fact, the price must be declining while the cycle of pooling offers is in progress, because the option to reset the seller's belief becomes more valuable to the buyer as the next screening offer comes closer.

A surprising implication of this result is that the potential presence of the high-valuation buyer can actually make the seller worse off. This arises because the seller cannot commit to refrain from screening in the future, when the high valuation becomes sufficiently likely. Screening wastes part of the pie while increasing the seller's share of what remains, and although this will be a good deal for the seller when screening is called for, the prospect of future screening means that the seller must concede an informational rent during the pooling phase. In some cases the seller would do better under a contract that binds unless both parties agree to renegotiate. But this would require an enforcement mechanism that can be trusted to prevent the seller from ever raising prices, and the lack of such a long-term enforcement mechanism is presumably one reason why we see repeated negotiations covering a few years at a time.

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. 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 (1999) consider the more general issue of whether a monopolist will choose to learn in a changing environment, in which demand is driven by a Markov process. Harris and Holmstrom (1987) consider a two-state model in which the optimal decision depends on the current state, and the state probabilities are determined by a Markov chain, which is costly to observe. They show that because information is most valuable when the two candidate decisions seem equally good, the optimal policy is to buy information only when the prior belief does not put heavy weight on either state, with the implication that there is a deterministic bound on the length of time between information purchases. 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. 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. That is also the problem analyzed here, but for the more complicated situation in which the buyer and the seller both act 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.

Blume (1998) 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 positive quantity that is below the myopic optimum, with the result that pooling equilibria are difficult to sustain.

The paper proceeds as follows. After the infinite-horizon game is formally defined in the next section, a two-period version is briefly 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 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 in this contingency, namely that the current valuation must have been low. 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 belief space into intervals indexed by the number of pooling offers that would be needed to reach the threshold (starting from a the most 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. Then the equilibrium strategies are defined, and there is an analysis of conditions ensuring 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.

2. 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 (where ρ_i denotes the probability of remaining in state i). 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 the surplus in period t be $n_t + \theta$, where n_t is either 0 or 1.[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 $\rho_i(s)$ be the probability that the Markov chain is in state i after s periods, starting from state i. Then

$$\rho_L(s) \; = \; \mu + \phi^s\left[1-\mu\right] \, , \quad \rho_H(s) \; = \; 1-\mu + \phi^s\mu$$

where $\phi \equiv \rho_L + \rho_H - 1$, and $\mu \equiv (1-\rho_H)/(1-\phi)$ is the invariant probability of the low state. The persistence parameter ϕ governs the link between successive negotiations: Prob $[n_t = 0 \,|\, n_{t-1}] = \phi(1-n_{t-1}) + 1-\rho_H$. It is assumed that $\phi \geq 0$; if $\phi = 0$, any inference drawn from the current negotiation will be irrelevant when the next contract is negotiated; at the other extreme, if $\phi = 1$ the current information is permanent.

For each period t, the actions $a_t = (p_t,q_t)$, payoffs $u(a_t)$ and $v(a_t)$, and history h_t are specified as follows:

Actions

first 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 payoffs in period t are $u(a_t) = q_t\, p_t$ for the seller, and $v(a_t) = q_t\,(n_t + \theta - p_t)$ for the buyer.

---

[4]For example, if v is a firm's revenue net of nonlabor costs, and if w_0 is the wage available to workers during a strike, the surplus is v-w_0. In the new units the low surplus is $\theta = (v_L-w_0)/(v_H-v_L)$, and the high surplus is $1+\theta = (v_H-w_0)/(v_H-v_L)$.

Both sides maximize expected present values, with a common discount factor δ; these are written as

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

History

The sequence of actions chosen by nature is	n^{t-1}	=	$\{n_1,n_2,...n_{t-1}\}$
The sequence of actions chosen by the seller is	p^{t-1}	=	$\{p_1,p_2,...p_{t-1}\}$
The sequence of actions chosen by the buyer is	q^{t-1}	=	$\{q_1,q_2,...q_{t-1}\}$
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	=	$\{n^t, p^t, q^{t-1}\}$

Strategies

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 $\mathbb{R}$. 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 $\sigma = (\sigma^S,\sigma^B)$ determines a stochastic process $A(\sigma)$ for a_t, taking values in $\mathbb{R} \times \{0,1\}$. This gives a complete description of what will happen if the players follow σ. The *path* of σ, denoted by $\underline{\sigma}$, is a support of this process.[6] A history is *on the path of σ* if the set of 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 is on the path of $A(\sigma\,|\,h_s)$.

Beliefs

A belief-system Π is a sequence of mappings from $\mathbb{R}^{t-1} \times \{0,1\}^{t-1}$ 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. This will be represented by $\zeta_t = \Pi_t[n_t = 0\,|\,h_t^0]$.[7]

---

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

[6]That is, a set of sample paths with probability 1. Thus $A(\sigma)$ may have many supports, and one is identified as the path. In practice, all strategies for the seller considered in this paper are supported on a finite set of prices, so there is no real ambiguity.

[7]This an abuse of notation: $\Pi_t[n_t = 0\,|\,h_t^0]$ is the marginal distribution, so $\Pi_t[n_t = 0\,|\,h_t^0] = \sum_{\{\nu^t \in \{0,1\}^t\,|\,\nu_t = 0\}} \Pi_t(\nu^t\,|\,h_t^0)$.

Similarly, $\Pi_t[n_{t-1} = 0\,|\,h_t^0]$ will be used to represent the marginal distribution over n_{t-1}.

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.

Consistency

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

(a)

$$\Pi_t[n_t = 0 \mid h^0_t] = 1 - \rho_H + \phi\, \Pi_t[n_{t-1} = 0 \mid h^0_t]$$

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

$$\left[\sum_{\nu^t \in \{0,1\}^t} \Pi_t(\nu^t \mid h^0_t)\, \sigma^B_t(q \mid \nu^t, p^t, q^{t-1})\right] \bar{\Pi}_{t+1}(n^t \mid p^t, q^{t-1}, q) = \Pi_t(n^t \mid h^0_t)\, \sigma^B_t(q \mid n^t, p^t, q^{t-1})$$

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 belief 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. This must hold for all prices: even after a history that is inconsistent with σ, the seller has a belief about n^t, and if there are no further deviations, subsequent beliefs are determined by Bayes rule.[8]

Sequential Optimality

For a given strategy-belief pair (σ,Π), the expected payoffs, conditional on history, are

$$V(\sigma \mid h_t) = E_{A(\sigma \mid h_t)} \tilde{V}(a), \quad U(\sigma, \Pi \mid h^0_t) = \sum_{\nu^t \in \{0,1\}^t} \Pi_t(\nu^t \mid h^0_t)\, E_{A(\sigma \mid \nu^t, h^0_t)} \tilde{U}(a)$$

where E refers to expectations with respect to the stochastic process $A(\sigma \mid h_t) \equiv A(\sigma \mid \nu^t, h^0_t)$.

Definition (Sequential Optimality)

(a) σ is *sequentially optimal for the buyer* if, for all t, and all histories h_t, and all $\hat{\sigma} \in \Sigma^B$

$$V(\sigma^S, \sigma^B \mid h_t) \geq V(\sigma^S, \hat{\sigma} \mid h_t)$$

(b) (σ,Π) is *sequentially optimal for the seller* if, for all t, and all public histories h^0_t, and all $\hat{\sigma} \in \Sigma^S$

$$U(\sigma^S, \sigma^B, \Pi \mid h^0_t) \geq U(\hat{\sigma}, \sigma^B, \Pi \mid h^0_t)$$

[8]For example, Rustichini and Villamil (1996) describe an equilibrium in which the buyer always rejects prices that are above the screening price (regardless of whether the current valuation is high or low), and yet if such a price were to be rejected, the seller would infer that the 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.

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

Definition A0: A strategy-belief pair (σ,Π) is an equilibrium if it is consistent and sequentially optimal.

3. **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). 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, and the buyer says yes or no. 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. In any sequential equilibrium of the two-period game, the continuation equilibrium in the last period is based on the threshold $\zeta_0^* = 1/(1+\theta)$. If $\zeta_0 < \zeta_0^*$, then the equilibrium involves screening: the seller chooses the price $1+\theta$, which the high-valuation buyer accepts and the low-valuation buyer rejects. If $\zeta_0 > \zeta_0^*$, the equilibrium involves pooling: the price is θ, and 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. 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 information revealed by the buyer's first-period action. Assume that $\phi\zeta_1 + 1\text{-}\rho_H < \zeta_0^*$, so that the seller is optimistic enough to screen in the last period if a pooling offer was 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 buyer might reject a pooling offer, and the seller's belief in the last period is then not determined by Bayes rule. 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 surely accepted, or screens, by offering a price that is rejected by the low-valuation buyer; or else the seller randomizes between these two alternatives. Thus equilibria can be labeled as $\tau_1 = 1$ (pooling), or $\tau_1 = 0$ (screening), or $0 < \tau_1 < 1$ (randomization).
- If there are equilibria with $\tau_1 = 0$, they all have the same equilibrium path:
  - either $p_1 = 1+\theta-\delta\rho_H$, $q_1 = n_1$, $p_0 = q_1 + \theta$, $q_0 = n_0$ (screening)

- or $p_1 = 1+\theta$, $q_1 = \hat{q}n_1$, $p_0 = 1 + \theta$, $q_0 = n_0$ (extended screening), where $\hat{q}$, the probability that the high buyer accepts p_1, is such that if this offer is rejected, then the belief in the last period is ζ_0^*.
- If there is an equilibrium with $\tau_1 = 1$, there is a continuum of such equilibria, with different first-period prices. In this case the equilibrium path satisfies $p_1 \in [\theta-\delta(1-\rho_L),\theta]$, $q_1 = 1$, $p_0 = 1 + \theta$, $q_0 = n_0$.

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. The reason for the multiplicity of pooling equilibria can seen by considering three 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 $p_1 \leq \theta$, $\qquad p_1 = \theta$

P2: $\ell(p_1) = 1$ for $p_1 \leq \theta$, $\qquad p_1 = \theta - \delta(1-\rho_L)$

P3: $\ell(p_1) = 0$ for $p_1 \leq \theta - ½\delta(1-\rho_L)$, $\ell(p_1) = 1$ for $\theta - ½\delta(1-\rho_L) < p_1 \leq \theta$, $\qquad p_1 = \theta - ½\delta(1-\rho_L)$

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 the rejection as a 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 pooling price is θ. But in P2, 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 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 example 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 $½\delta(1-\rho_L)$.[9]

A notable feature of these equilibria is that the seller's strategy is not Markov: the function $\ell(p_1)$, which determines p_0, is not constant on 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.

[9]It might seem that there is something wrong here: can't the seller do better by raising the price? No, because if the buyer rejects any higher price, 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 $\theta - ½\delta(1-\rho_L)$ would be accepted by the low buyer.

Implications of Forward Induction.

In examples P1 and P2, the high buyer gets an immediate payoff of at least 1 by accepting the first-period offer, while the payoff from rejection is no more than $\delta\rho_H < 1$. On the other hand the low buyer's equilibrium payoff is zero in P1, while the payoff from rejection of p_1 would be $\delta(1-\rho_L)$ 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 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 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 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 screening cycle in the infinite-horizon game is sketched in Figure 1, which represents varying degrees of pessimism for the seller, based on the results of the previous negotiation. At one extreme, the seller's belief that the buyer has a low valuation now is ρ_L, following rejection of a screening offer in the previous negotiation. In this situation the seller pools now, and 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 acceptance of a screening offer; then the probability of the low state now is only $1-\rho_H$.

Figure 1: The Screening and Pooling Cycle

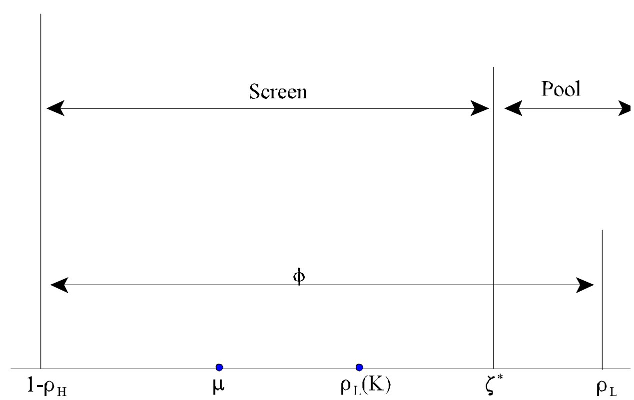

The analysis proceeds by imposing a set of "reasonable" properties that are motivated by the heuristic discussion in the Introduction, tempered by the analysis of the two-period game. The aim is to show that the equilibrium sketched in Figure 1 is not an arbitrary choice: in fact it is shown that all equilibria

[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. 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.

satisfying these properties are payoff-equivalent. The properties are roughly as follows. The strategies are cyclic, and are driven by the current belief of the seller. The forward induction argument used 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. Extended screening equilibria[11] are ruled out: this involves a substantive restriction on the fundamental parameters which is spelled out later, and in particular it rules out cases in which the buyer's valuations are very persistent.[12] 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.

The Screening Threshold

Given a belief system Π, let $\zeta_t = \Pi_t[n_t = 0 \,|\, h_t^0]$ be the 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.[13] 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 repeated pooling offers eventually lead to screening.[14]

When defining the screening threshold it is necessary to pay particular 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 $\bar{\sigma}^S$ as the highest price in $\underline{\sigma}^S$:

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

Definition A1: The 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 $p_t \in \underline{\sigma}_t^S(h_t^0)$:

[11]That is, equilibria in which the seller makes a screening offer several times in succession, before inferring that the current valuation is low. Such equilibria are familiar from the static bargaining literature, so this paper focuses on equilibrium features that arise only in the context of repeated negotiations, such as cyclic pricing and informational rents accruing to the low buyer.

[12]It would be desirable to extend the analysis to cover the part of the parameter space excluded by this restriction, but it would also be very difficult, as can be seen from the attempt made in Kennan (1995).

[13]More precisely, the seller's strategy specifies a screening offer after any history h_t^0 such that $\Pi_t[n_t = 0 \,|\, h_t^0] < \zeta^*$.

[14]If instead ζ^* lies below μ, then in the long run the seller makes only pooling offers: see Section 10 below.

(a) if $\zeta_t < \zeta^*$, then $q_t = n_t$

(b) if $\zeta_t > \zeta^*$, then $q_t = 1$

(c) if $\zeta_t = \zeta^*$, and if $p_s \leq \bar{\sigma}^S$ for all s < t, then $q_t = 1$ with probability λ^*, and $q_t = n_t$ with probability $1-\lambda^*$.

That is (a) if the seller is optimistic, there is a screening offer, and (b) if the seller is pessimistic, there is a pooling offer. Also, (c) on the borderline between these, the seller randomizes, provided there has been no previous deviation by the seller that exceeded the price ceiling.[15]

Property A1 is enough to determine the probability of trade in all future periods after any history such that the price has never exceeded $\bar{\sigma}^S$. 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.

The State Variable

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

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

Then 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 $\rho_L(K^*) \leq \zeta^* < \rho_L(K^*-1)$, 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 $\tau = \Psi_\lambda(\zeta)$ from [0,1] to $\mathbb{R}$ defined as follows. If $\zeta < \zeta^*$, then $\tau = 0$; if $\zeta^*(s-1) < \zeta < \zeta^*(s)$, for some positive integer s, then $\tau = s$; otherwise $\zeta = \zeta^*(s)$ for some nonnegative integer s, and then $\tau = s+\lambda$. The mapping $\tau = \Psi^*(\zeta)$, is defined as $\Psi_{\lambda^*}(\zeta)$.

[15]The reason for (c) is that it may be necessary to allow the buyer's strategy to respond to an excessive price p_s by randomizing if $n_t = 1$, and to support this the seller's strategy must randomize later in a way that depends on p_s.

Figure 2: Beliefs and States

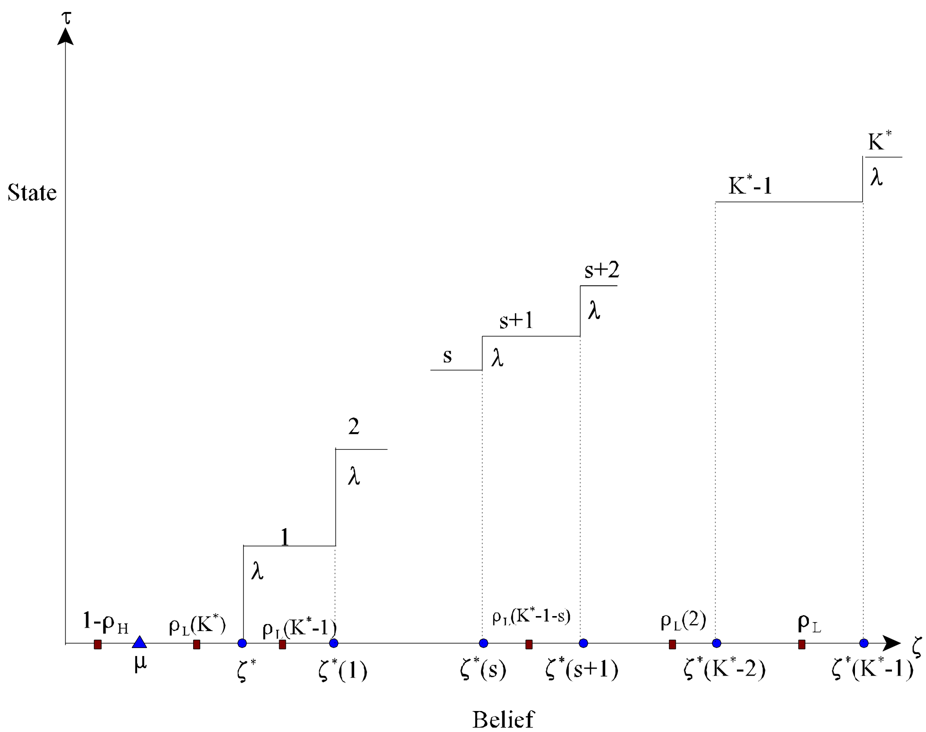

The function Ψ_λ is illustrated in Figure 2. The point of the definition is that if (σ,Π) satisfies A1, with screening threshold (ζ^*,λ^*), then after any history on the path of σ, the state $\tau = \Psi^*(\zeta)$ is sufficient to determine the continuation path: the seller will make pooling offers for the next τ periods, and the belief will then cross the screening threshold.[16]

The equilibrium cycle sketched in Figure 1 can now be described in detail. Consider first the case in which ζ^* is not on the path of σ.[17] Then A1 implies the following cycle. After any history such that $\zeta < \zeta^*$, the seller's offer is accepted if the current valuation is high, and rejected otherwise: the state is $\tau = \Psi^*(\zeta) = 0$ for all such histories. Consistency implies that the next value of ζ is $\zeta' = 1\text{-}\rho_H$ if a screening offer is accepted, and $\zeta' = \rho_L$ if it is rejected. Since $1\text{-}\rho_H < \mu < \zeta^*$, the seller again makes a screening offer in the next period after a screening offer is accepted, and this continues until an offer is rejected. Then

---

[16]Only Ψ^* is relevant for the moment, because $\tau = \Psi^*(\zeta)$ on the equilibrium path. The mapping Ψ_λ is used later when behavior off the equilibrium path is specified: 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 $\Psi^*(\zeta)$.

[17]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.

$\zeta' = \rho_L$, and (if $\rho_L > \zeta^*$) A1 implies a pooling offer. After this, consistency implies $\zeta'' = \rho_L(2)$, 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, and so on until τ reaches 0.

Now consider the case in which $\zeta^* = \rho_L(K^*)$ for some integer $K^* \geq 1$.[18] 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+λ^* pooling offers (meaning K^*-1 for sure, and one more with probability λ^*), followed by a screening offer. The state is $\tau = \Psi^*(\rho_L) = K^*-1+\lambda^*$ after a screening offer is rejected, and the countdown to the next screening offer continues as above; when τ reaches λ^* there is one more pooling offer with probability λ^*.[19]

5. Equilibrium Continuation Values

The equilibrium continuation values will now be determined as functions of the state variable τ, for any pair (σ,Π) that satisfies A1 and the other properties mentioned above. First, it is shown that if A1 holds, 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 used to compute the buyer's continuation values, yielding the seller's values as a residual. Finally, the threshold ζ^* and the randomization probability λ^* are obtained as functions of the basic parameters.

The first result shows that A1 determines a Markov chain for (n,τ). Proofs of the main results are in the Appendix (proofs of the other results are available from the author's web page).

Lemma M:

For any public history h_t^0, let $\zeta_t = \Pi_t[n_t = 0 \,|\, h_t^0]$ and $\tau_t = \Psi^*(\zeta_t)$. If (σ,Π) satisfies A1, then after any history h_t such that $p_s \leq \bar{\sigma}^S$ for all s < t, the values of n and τ on the continuation path of σ from h_t are

---

[18]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 pooling cycle, with $\zeta = \zeta^*$ at this point.

[19]The timing of this randomization is arbitrary. A public randomization could be used to set $\tau = K^*$ with probability λ^* and $\tau = K^*$-1 with probability 1-λ^*, following rejection of a screening offer. This does not affect the probability of trade, and the effect on prices is merely a change in the timing of payments, with no effect on present values.

determined by a Markov chain with state space $\{0,1\} \times \{0,\lambda^*,1+\lambda^*,...K^*-1+\lambda^*\}$ 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 $\tau \geq 1$, 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, $q_t = n_t$ if $\tau_t = 0$, and $q_t = 1$ if $\tau_t \geq 1$, while if $\tau_t = \lambda^*$ then $q_t = 1$ with probability λ^* and $q_t = n_t$ with probability $1-\lambda^*$. Thus (since joint payoffs depend only on n and q) 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 current valuation is low or high, with $j(\tau) = J^H(\tau) - J^L(\tau)$. Define the discounted persistence parameter β as $\beta = \delta\phi$, and for nonnegative integers s define the following discounted sums:

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

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 A1. The joint values are shown 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 $p_s \leq \bar{\sigma}^S$ for all $s < t$, let $\zeta_t = \Pi_t[n_t = 0 \,|\, h_t^0]$ and $\tau = \Psi^*(\zeta_t)$. Then

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

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 derived in Lemma M. Let $V^L(\tau)$ and $V^H(\tau)$ denote the buyer's continuation values from state τ, depending on whether the current valuation is low or high, and let $d(\tau) = V^H(\tau) - V^L(\tau)$.

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

That is, rejection of any price below $\bar{\sigma}^S$ convinces the seller that the current valuation is low. This captures two features of the cyclic equilibrium described in the Introduction. First, the high buyer accepts screening offers for sure, ruling out equilibria in which screening extends over several periods. In Section 8, this will be translated into a restriction on the basic parameters $(\theta,\rho_L,\rho_H,\delta)$. Second, rejection of a pooling offer would convince the seller that the current valuation is low. This can be rationalized by forward induction, as in the two-period game. In fact, 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, and this is not true for the low-valuation buyer.

Under A2, prices and quantities in earlier periods have no relevance for future payoffs. This motivates

Definition A3: The 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 such that $p_t \leq \bar{\sigma}^S$, $E_{A(\sigma|h_t,q_t=0)}\tilde{V}(a) = n_t V^H(0) + (1-n_t) V^L$.

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

Definition A4: The profile σ has the *tight pricing property* if for all t and all 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).

That is, for any pooling offer on the continuation path of σ, the low buyer is indifferent between acceptance and rejection, and for any screening offer on the path, the high buyer is indifferent.[20]

Properties A1, A3 and A4 together imply that for any history that has not breached the price ceiling, the low buyer's continuation value is a constant. To see this, note from A1 that the seller makes either a pooling or a screening offer after any such history. Then A3 immediately implies that the low buyer's continuation value is V^L for any history such that a screening offer is made; and A4 implies that the low buyer's continuation value when a pooling offer is made is the value of rejecting, which is again V^L.

[20]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. On the other hand, it is possible that A4 is implied by A0-A3.

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

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^0_t, let $\zeta_t = \Pi_t[n_t = 0 \mid h^0_t]$ 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

$$\underline{\sigma}^S_t(h^0_t) = \{p_H\}, \qquad \text{if } \tau_t = 0$$
$$\underline{\sigma}^S_t(h^0_t) = \{p_L(\tau_t)\}, \quad \text{if } \tau_t \geq 1$$

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

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

Two implications of this result are worth noting. First, 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. Second, changes in θ are absorbed entirely by the seller, with no effect on the buyer's 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 equilibrium 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 pair (σ,Π), the seller infers that the 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.[21] 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 $g(\tau) = U^H(\tau) - U^L(\tau)$. The seller's expected continuation value at the point of making a screening offer is the weighted average $U(\zeta,\tau)$ of the state-contingent values.

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 (in Section 7) ζ^* is used to construct a strategy-belief pair (σ^*,Π^*) that satisfies properties A1-A4. Finally (in Section 8) it is shown that this pair satisfies the sequential optimality conditions in A0, subject to some restrictions on the parameters; this completes the proof that cyclic screening equilibria exist.

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

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

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 determines 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 (σ,Π) satisfies A0-A4. Then $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$, where Z is defined as

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

where

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

---

[21]Alternatively, the seller's value function can be computed directly from the equilibrium 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, because in this case the joint values are assessed differently by the two players.

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 lies in [0,1]. 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 $\theta \leq \theta_\infty$, which ensures that $\zeta^* \geq \mu$. If $\theta > \theta_\infty$ there is no cyclic equilibrium, because screening is too expensive: this is discussed in Section 10.

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

7. Cyclic Screening Strategies

Suppose (σ,Π) satisfies A1-A4, with a screening threshold (ζ^*,λ^*) determined by the mapping Z. Then the equilibrium prices and quantities are fully determined: (n,τ) follows the Markov chain specified in Lemma M, the seller offers $p_L(\tau)$ when τ is positive, and 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. Then it will be verified that the construct does indeed satisfy properties A1-A4. The last step is to show that the strategies are sequentially optimal: this finally demonstrates the existence of cyclic screening equilibria.

The strategies will be written in terms of the state variable $\tau = \Psi_\lambda(\zeta)$, so it is necessary to say how λ is determined for every realization of the public history. It is also necessary to specify the belief system associated with (ζ^*,λ^*), and these are closely related. First, for $\tau \in [0, K^*-1+\lambda]$ define the prices $\bar{p}_H(\tau)$ as

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

so that $p_H = \bar{p}_H(K^*-1+\lambda^*)$. These prices will be used to determine the continuation in case the seller offers a price above p_H. The equilibrium beliefs are then defined as follows.

Definition: Given a screening threshold (ζ^*,λ^*), the *threshold belief system* $\Pi^*(\zeta^*,\lambda^*)$ is a sequence of mappings $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$ determining (ζ_t,λ_t) as functions of h^0_t, 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 rule, which uses $(\zeta_t,\lambda_t) = \Pi^*_t(h^0_t)$ 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 $p_t = \bar{p}_H(s+\lambda_0)$, $\tau_t-1 \le s+\lambda_0 < K-1+\lambda^*$, 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 $p_t > \bar{p}_H(\tau_t-1)$ then $\zeta_{t+1} = \phi\zeta+(1-\phi)\mu$, and $\lambda_{t+1}=\lambda^*$

(b) When $q_t = 1$:

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

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

In addition to specifying the belief as a function of the public history, this definition also specifies the probability that the seller will make a pooling offer if the belief coincides with the threshold.[22] 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, 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 A1.

The Seller's Strategy

Definition: The *cyclic pricing rule* associated with (ζ^*,λ^*) is the mapping $p = P^*(\tau)$ defined 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, and $p = p_H$ with probability $1-\tau$

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

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 p_H and $p_L(\tau)$, and these determine P^*. Then K^*, λ^*, p_H, $p_L(\tau)$ and $\bar{p}_H(\tau)$ and the numbers $\zeta^*(s) = \mu+\phi^{-s}(\zeta^*-\mu)$ 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.

---

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

The Buyer's Strategy

Definition: The *cyclic trading rule* associated with (ζ^*,λ^*) is the mapping $q=Q^*(\tau,p,n,\zeta)$ defined 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 $q = 1$, and 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 (ζ^*,λ^*) is the sequence of mappings 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 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.[23]

[23]One might wonder why prices slightly above p_H are accepted by the high 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 there is an equilibrium in which the seller does indeed offer $\bar{p}_H(K^*-1)$ when $\zeta = 1-\rho_H$. For some parameter values this equilibrium involves incomplete screening: when $\zeta = 1-\rho_H$ the high buyer rejects the screening offer with positive probability, and $\tau' = K^*-1$, so that (if $K^* > 1$) the seller makes a pooling offer next period even though this period's valuation was not fully revealed. There are also equilibria in which $\zeta^* = \rho_L(K^*)$, with $\tau' = K^*-1$ if a screening offer is rejected when $\zeta = 1-\rho_H$, and $\tau' = K^*$ if a screening offer is rejected when $\zeta = \zeta^*$. Note that such equilibria violate A1.

Figure 3: Randomization by the high-valuation Buyer (the next state is τ_0 if q=0, and τ_1 if q=1)

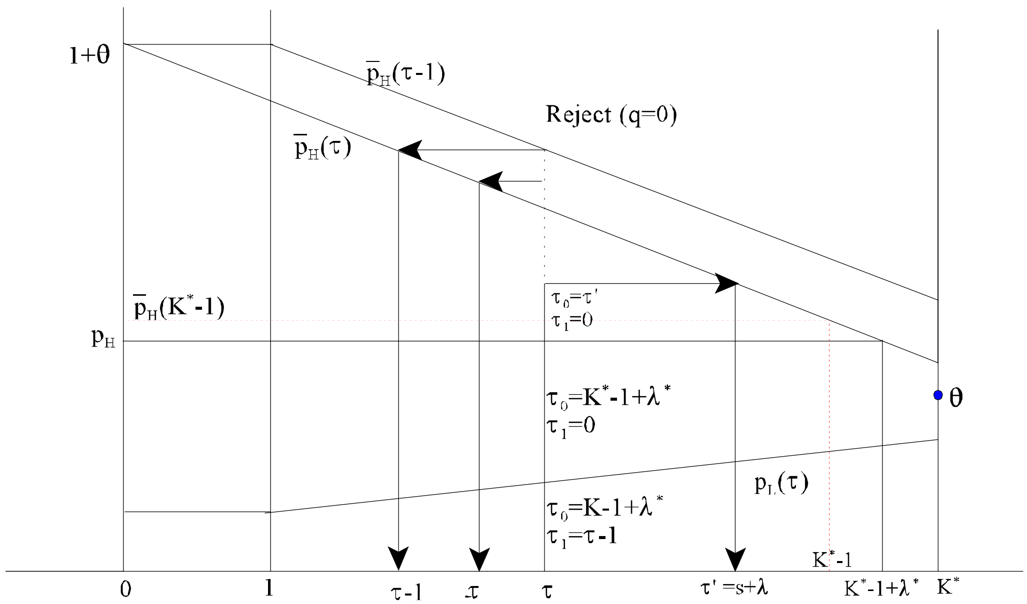

Figure 3 shows how the buyer's strategy responds to a price p above p_H, given that n = 1. The response must be random, because if all buyers reject, the end of the pooling cycle comes one period closer, in which case acceptance would be more attractive for the high buyer; but if all high buyers accept, then rejection would restart the pooling sequence, and this would be more attractive to the high buyer. The arrows in the diagram show how the next state τ_q depends on p and q.

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 $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$.

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:

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

Corollary:

The strategy-belief pair (σ^*,Π^*) satisfies the immediate signaling property A2.

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 $J^L(\tau_t) + (1\text{-}n_t)j(\tau_t)$.

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 profile σ^* is sequentially optimal for the buyer. After any history, the buyer's continuation value is $V^L + n\,d(\tau)$, where n and τ are current values, and V^L and $d(\tau)$ are as defined in Lemma B1

Corollary:

The profile σ^* satisfies properties A3 and A4.

Lemma 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 pair (σ^*,Π^*) implies that the seller's continuation value is $U^L(\tau_t) + (1\text{-}\zeta_t)g(\tau_t)$, where the functions $g(\tau)$ and $U^L(\tau)$ are defined by

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

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

To complete the proof that (σ^*,Π^*) is an equilibrium, it remains only to show sequential optimality for the seller. First, Lemma S2 shows that there is no profitable one-period deviation below p_H, and in particular that the seller cannot gain by deviating from a pooling to 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 $p \leq p_H$

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

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, provided that the parameters satisfy the following:

Definition A5:

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

$$\frac{1+\theta - r^*\theta\,[R(s)-B(s)]}{\zeta^*(s)} - B(s)(b^*+r^*\theta) - \frac{\rho_H}{(1-\rho_H)\phi}\left[1-b^*B(s)\right] \geq 0 \qquad (\mathbf{X_s})$$

for $1 \leq s \leq K^*$, where $(\zeta^*,\lambda^*) = Z(\theta,\rho_L,\rho_H,\delta)$, $K^* = \kappa(\zeta^*)$, $b^* = 1/B(K^*+\lambda^*)$, $r^* = 1/R(K^*+\lambda^*)$ and $\zeta^*(K^*) = 1$.[24]

Proposition S3:

If $(\theta,\rho_L,\rho_H,\delta)$ satisfies A5, then (σ^*,Π^*) is sequentially optimal for the seller.

The proof (in the Appendix) shows that the left side of A5 is the difference between the equilibrium continuation value and the value of deviating to a price p such that $\bar{p}_H(s) < p \leq \bar{p}_H(s-1)$, when $\zeta = 1-\rho_H$.

8. Existence of Equilibrium

To prove existence of the cyclic screening equilibrium (σ^*,Π^*), it remains only to show that A5 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 A5.

Proposition E1

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

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

and

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

Then $(\theta',\rho_L,\rho_H,\delta)$ satisfies A5 for any $\theta' \in [\theta,\theta_\infty)$.

Note that X_1 is a rearrangement of the inequality X_s in A5 with $s=1$, while $\bar{X}_{K-1}$ is stronger than X_s with $s=K^*-1$. Also, $\bar{X}_{K-1}$ 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 $\rho_H \leq \rho_H(1-\rho_H)$, which is impossible for $\rho_H > 0$. But this does not mean that X_1 fails when $K^*=2$, as is illustrated in Table 2 below.

The second result gives a condition on the parameters (ρ_L,ρ_H,δ) guaranteeing that the complete screening conditions hold when θ is large (i.e. when the cost of screening is high).

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

Proposition E2

Suppose the parameters ρ_L, ρ_H and δ satisfy the inequality[25]

$$\frac{1-\rho_L}{\delta\,(1-\rho_H)} + \frac{1-\rho_L}{1-\delta} > \frac{\rho_H}{1-\phi} - \phi \qquad (X_1^\infty)$$

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 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 valuations are permanent, an increase in δ implies more screening. The following proposition shows that this is valid more generally, even when the buyer is forward-looking.

Proposition L:

Suppose that $(\theta,\rho_L,\rho_H,\delta)$ satisfies A5 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.

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.

Thus 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.

---

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

Table 2: An Optimal Screening Cycle

Parameters		**Equilibrium**
$\theta = ½$, $\delta = ⅛$, $\rho_L = ¾$, $\rho_H = ¾$, $\mu = ½$, $\phi = ½$		$\zeta^* = \frac{35}{52}$, $\lambda^*=0$, $K^* = 2$ 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, $p_L(1)$
$\tau=0$	6,528	103,632	2,240	6,160	4,288	97,472	74,176	$\frac{2}{8}$	$\frac{24}{17}$
$\tau=0$	6,528	103,632	2,240	6,160	4,288	97,472	39,232	$\frac{5}{8}$	$\frac{24}{17}$
$\tau=1$	35,088	103,632	2,240	64,960	32,848	38,672	34,304	$\frac{6}{8}$	$\frac{8}{17}$

[a]The values are scaled up by 62,475, to give exact results that are easy to compare across states.

[b]The two rows with $\tau = 0$ differ only in the seller's beliefs: the first is reached after a successful screening offer, and the second after a pooling offer.

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 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).

Random Screening

An equilibrium with random screening is shown in Figure 4, which plots the function Z over the relevant range, with parameters $(\theta,\rho_L,\rho_H,\delta)$ = (5/34,41/44,7/22,1/2). The plot shows the piecewise-linear function $\Theta = \frac{1}{G(K^* + \lambda^*)}\frac{\zeta}{1-\zeta}$. This function jumps at an inconvenient spot, which is why randomization by the seller is needed in equilibrium. The screening threshold is $\zeta^* = \rho(3)$ =641/704, with λ^* = .10889. Thus a rejected offer is followed by two pooling offers, and then another pooling offer with probability λ^* or a screening offer with probability $1-\lambda^*$. The example satisfies the E1 conditions 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 extended screening price is $\bar{p}_H(K-1) = .988244$.

Figure 4: Random Screening Equilibrium
(ρ_L=41/44,ρ_H=7/22, δ=½)

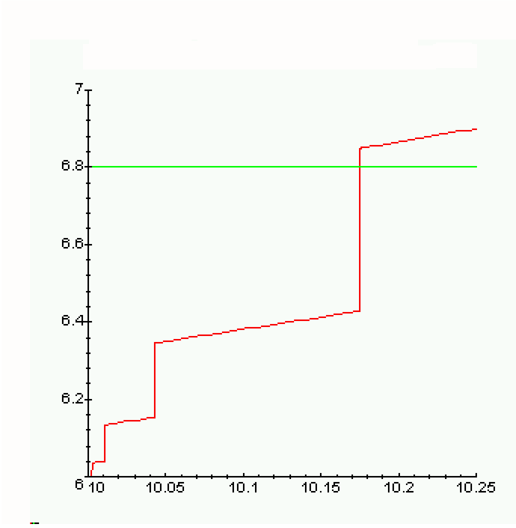

An Example Showing the Effects of Limited Commitment

A surprising feature of the equilibrium is that the presence of the high buyer type can make the seller worse off *ex ante*. An extreme case is shown in Table 3, using a discount factor close to 1. Coase (1972) conjectured that a monopoly seller of a durable good who could not commit to hold the line on prices for any length of time would be unable to obtain any monopoly profit. In the case of the rectangular demand curve used in this paper, the implication is that the seller can do no better than the pooling price, with a value of 16,000. But in the cyclic screening equilibrium the seller does even worse than this: the highest value achieved is 13,856, so that the seller would prefer to face the low-valuation buyer for sure. 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 large informational rent. The seller would be happy to guarantee that future prices will never exceed θ (as if the high buyer type did not exist). But this guarantee is credible only if there is a mechanism that will enforce even the most ancient contracts, because there will come a time after which the seller will always wish to renege.

Table 3: An Example of the Effect of Limited Commitment									
Parameters						Equilibrium			
θ = 8/5, δ = .9999, ρ_L = 3/4, ρ_H = 3/4, μ = 1/2, ϕ = 1/2						ζ^* = .528, λ^*=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, $p_L(\tau)$
τ=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
τ=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
τ=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
τ=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
τ=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
τ=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									

A Cyclic Screening Equilibrium with High Persistence

For any given value of θ, condition X_1 generally fails as ρ_H approaches 1, meaning that the seller would deviate from the cyclic screening strategy so as to extend screening over more than one period.[26] On the other hand, Proposition E2 indicates that cyclic screening equilibria exist even if ρ_H is arbitrarily close to 1, provided that θ is large. This is illustrated in Table 4.

[26]Evaluating Z at ρ_H=1 shows that ζ^*=1/(1+θ); also μ = 0, so $\zeta^* > \mu$. Fix θ so that 1/(1+θ) < ρ_L (if θ is not in this region there is an unconditional screening equilibrium). Then K^* > 1, and so b^* < 1, and then X_1 fails as ρ_H approaches 1. Thus (σ^*,Π^*) is not an equilibrium. Given any initial belief the seller eventually becomes optimistic enough to screen, and the high buyer accepts, and this implies ζ'= 0. Then if the buyer follows the cyclic screening strategy, the response to p = 1+θ is q = 1 (for the high type, which is the only type entertained by the seller at this point), so the seller gains by deviating from p_H to 1+θ.

Table 4: An Example with Almost Permanent Valuations									
Parameters					Equilibrium				
θ = 100, δ = 3/4, ρ_L = .99 , ρ_H = .9999					ζ^* = .00990393, λ^*=0, K^* = 1254				
μ = 1/101 = .00990099, ϕ = .9899					x_1: 2672.682659 > 0 $\bar{x}_{K^*-1}$: .9598790345 > 0				
	Continuation Values								
	Joint		Buyer		Seller			Beliefs	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, $p_L(\tau)$
τ=0	300.110	403.9688	0.1165	2.9988	299.99	400.97	400.959911	0.00010000	100.250075
τ=0	300.110	403.9688	0.1165	2.9988	299.99	400.97	399.969946	0.00990392	100.250075
τ=1	325.860	403.9688	0.1165	3.2564	325.74	400.71	399.969946	0.00990395	99.9925
τ=2	344.980	403.9708	0.1165	3.4476	344.86	400.52	399.971884	0.00990398	99.99443
τ=1252	400.110	403.9988	0.1165	3.9988	399.99	400.00	399.993658	0.98010100	100.0
τ=1253	400.110	403.9988	0.1165	3.9988	399.99	400.00	399.993594	0.99000000	100.0

The main point of this example is that the ratchet effect does not preclude screening, despite the above-mentioned Hart-Tirole (1988) result that there are no equilibria with screening in early periods when the valuations are completely permanent, and the horizon is long. What happens in the example is that the buyer is willing to reveal the high valuation because the seller never becomes quite certain that the current valuation is high, and even a small doubt leaves the seller willing to pool, because the cost of screening is high. Thus the buyer knows that even if a screening offer is accepted now, the seller will restart the pooling cycle if a screening offer is rejected next time.

10. Alternative 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 "transient screening equilibria" with $\zeta^* < \mu$, so called because a seller who infers that the current valuation is low will never again be optimistic enough to screen. Such equilibria resemble 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. But in a transient screening equilibrium, a pooling offer at ζ^* implies pooling 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 equilibrium with $K^* = \infty$, because an accepted screening offer is followed by another screening offer. The value of a pooling offer is $\theta/(1-\delta)$, and the screening threshold is obtained by finding the value of $n = 1-\zeta$ where this matches the value of screening. This yields

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

If ζ^* is not in either the cyclic screening interval $[\mu,1]$ or the transient screening interval $[1-\rho_H,\mu]$, it must be in the "unconditional pooling" interval $[0,1-\rho_H]$, so called because even the most optimistic seller chooses pooling. At the other extreme, there is an "unconditional screening" interval $[\rho_L,1]$ where even the most pessimistic seller chooses screening. In these unconditional cases, 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, so ζ^* is the solution of the static problem, namely $\zeta^* = 1/(1+\theta)$.

For a nontrivial cyclic equilibrium, ζ^* must lie between μ (with $K^* = \infty$) and ρ_L (with $K^* = 1$). Thus

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

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

$$\frac{1-\rho_H}{\rho_H(1-\delta\rho_H)} \leq \Theta \leq \frac{(1-\rho_H)(1-\beta)}{(1-\rho_L)(1-\delta\rho_H)}$$

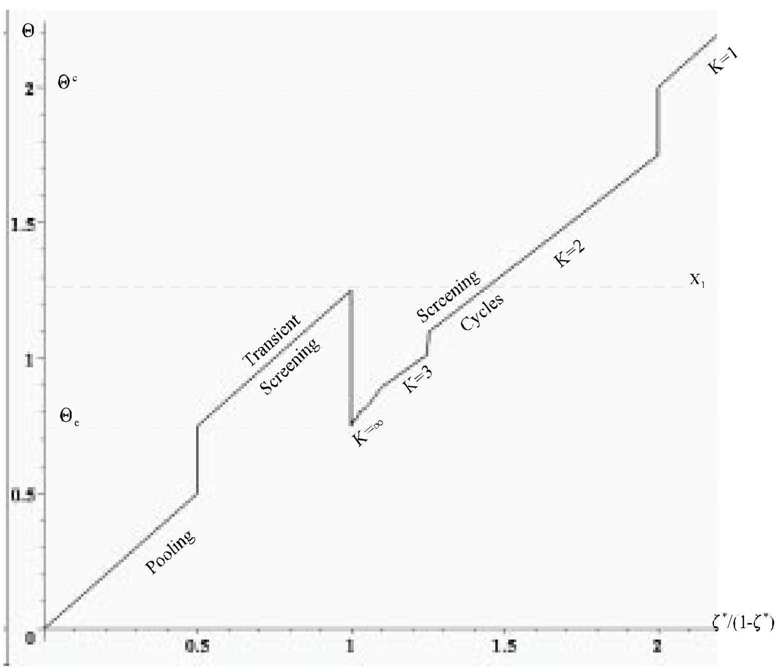

Figure 5: Alternative equilibria ($\rho_L = ⅔$, $\rho_H = ⅔$, $\delta = ½$)

Finally, the condition for an unconditional pooling equilibrium is $0 \leq \Theta < 1/\rho_H - 1$. These results are illustrated in Figure 5, which shows how ζ^* varies with θ, for a particular ordering of the critical values of Θ defined above. The diagram plots Θ against $\zeta^*/(1-\zeta^*)$, 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 complete screening conditions are satisfied by all values of Θ below the dashed line labeled X_1 (which is just above the transient screening region).

Alternative Beliefs[27]

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 (A2). A simple alternative that might seem appealing is to 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 not, the pooling price is just the low valuation, θ. But it can be shown that the derivation of the screening threshold becomes more difficult. In fact, there is no simple way to determine ζ^* from the basic parameters, so 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 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.

Cyclic equilibria of this kind have previously been analyzed by Kennan (1995), and by Rustichini and Villamil (1996). The model in Kennan (1995) allows a sequence of offers within each contract negotiation. This led to complications that precluded a systematic equilibrium analysis; in this paper, these complications are suppressed so as to focus on linkages across contracts. Rustichini and Villamil also assumed one offer per contract, and their main result was that cyclic equilibria exist if persistence is high. This is puzzling since Hart and Tirole (1988) showed (for discount factors exceeding ½) that only pooling equilibria survive in the limiting case when the buyer's valuation is permanent. But these results are in fact compatible, because the Rustichini and Villamil equilibrium is merely a weak Perfect Bayesian equilibrium,

[27] The details of these results are in an unpublished appendix, at http://www.ssc.wisc.edu/~jkennan/research/App_C_cyc.pdf.

in which the seller's beliefs off the equilibrium path are not consistent with the buyer's strategy. This paper indicates that once consistency is imposed, cyclic equilibria disappear as ρ_H approaches 1 (with θ fixed), because a seller who is virtually certain that the current valuation is high will not abandon this belief just because a single screening offer has been rejected. Instead, as in the standard static bargaining model, a long sequence of rejected screening offers may be needed to overcome the seller's initial optimism.

The main limitation of the model developed here is that such extended screening equilibria are ruled out, by excluding a portion of the parameter space in which the cost of an unsuccessful screening offer is low. A more general analysis would expand the state variable to count down the number of rejected offers needed to convince the seller to restart the pooling sequence. From the point of view of application, a more important question is whether the results obtained here can be generalized to cover the case in which the seller makes a sequence of offers within each contract negotiation, and does not infer that the current state is low until all of these offers have been rejected.

The paper does not address the larger question of why the parties cannot commit to long-term contracts: it is assumed that contract duration is given, and that the Markov chain makes one transition per contract. The literature contains no good theory of contract duration. This paper can shed some light on this question, by linking the contract duration to the degree of persistence in the buyer's private information. In particular, if the use of short-term contracts is interpreted as being due to uncertainty about whether long-term contracts will be enforced, the model shows how improvements in the enforcement technology affect the equilibrium of the bargaining game, by weakening the persistence of private information across contracts.

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, "Contract Renegotiation with Time-Varying Valuations," *Journal of Economics and Management Strategy*, Fall 1998, 435-460.

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, David Levine, and Jean Tirole, "Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information," in Roth, Alvin (ed.), *Game-Theoretic Models of Bargaining*. Cambridge: Cambridge University Press, 1985, 73-98.

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.

Harris, Milton and Bengt Holmstrom "On the Duration of Agreements," *International Economic Review*, 28(2), June 1987, 389-406.

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 in British Manufacturing in the 1980s," *Industrial & Labor Relations Review*, 1993, 46(4), 704-717.

Keller, Godfrey and Sven Rady, "Optimal Experimentation in a Changing Environment," *Review of Economic Studies*, **66**(3), 1999, 475-507.

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: Proofs

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 $J^L(\tau)$ is the value of s periods paying $n+\theta$, plus the continuation value from $\tau = \lambda^*$ after s periods:

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

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 $\rho_H(i) + \rho_L(i) - 1 = \phi^i$. 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

$$J^L(0) = J^L(K^*+\lambda^*) - \theta \;;\; J^H(0) = J^L(0) + j(0) = 1 + \theta + \delta\left[J^L(0) + \rho_H j(0)\right]$$

So

$$(\tau) = D(\tau)(1+\theta) - R(\tau) + \left[1 - D(\tau)(1-\delta)\right]J^L(0) + \left[(1-\beta)R(\tau) - (1-\delta\rho_H)D(\tau)\right];\; j(0) = J^L(0) + R(\tau)\left[(1-\beta)j(0) -\right.$$

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

$$J^L(0) = J^L(0) + R(K^*+\lambda^*)\left[(1-\beta)j(0) - 1\right] - \theta \;;\; (1-\delta)J^L(0) = 1 + \theta - (1-\delta\rho_H)j(0)$$

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

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

Proposition T0:

By A1, the seller randomizes between pooling and screening after any history such that $\zeta_t = \zeta^*$ and $p_s \leq \bar{\sigma}^S$ for all s < t, so (by A0) the seller must be indifferent between these alternatives. 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 under (σ,Π) for ζ_t just above ζ^*). 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 under (σ,Π) for ζ_t just below ζ^*). Thus the seller's expected payoffs 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 equate these, i.e.

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

Proposition T1:

Fix 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 is in the first interval, then $\zeta^* = 1/(1+\theta) > \rho_L$ and $\lambda^* = 0$, meaning that the seller always makes screening offers. In the second interval, $\zeta^* = \rho_L$ and $\lambda^* = \frac{\theta - \bar{\rho}_L}{\delta\, 1 - \rho_L - \beta\theta}$, 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.

From now on, for any t, and for any public history h_t^0, let $(\zeta_t,\lambda_t) = \Pi_t^*(h_t^0)$ and $\tau_t = \Psi(\zeta_t,\lambda_t)$.

Lemma C:

For any price that is surely accepted by $\sigma^B(\zeta^*,\lambda^*)$, Bayes' rule implies that the posterior after acceptance must be the same as the prior, so the belief next time is $\zeta' = \phi\zeta_t + (1-\phi)\mu$. This covers all prices below $p_L(\tau_t)$, with $q_t = 1$. If any such price is rejected, any belief is consistent with Bayes' rule. Similarly, the seller's belief is consistent for any price that is rejected by $\sigma^B(\zeta^*,\lambda^*)$; this covers all prices above $\bar{p}_H(\tau_t\text{-}1)$. If p is accepted by $\sigma^B(\zeta^*,\lambda^*)$ when n_t=1 and rejected when n_t=0, 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 $q_t = 1$, so $\zeta' = 1\text{-}\rho_H$. This covers all p in $(p_L(\tau_t),\bar{p}_H(K^*\text{-}1)]$. The leaves only prices in $(\bar{p}_H(K^*\text{-}1),\bar{p}_H(\tau_t\text{-}1)]$. These prices are accepted with positive probability by $\sigma^B(\zeta^*,\lambda^*)$ iff $n_t = 1$, so the belief next time is $\zeta' = 1\text{-}\rho_H$ following acceptance. The belief following rejection depends on p. If $n_t = 1$, $\sigma^B(\zeta^*,\lambda^*)$ locates p in a subinterval $(\bar{p}_H(s+1),\bar{p}_H(s)]$, where $\tau\text{-}1 \le s \le K^*\text{-}2$, and uses s and $\zeta_t = \Pi^*_t(h^0_t)$ to determine the acceptance probability ν_t. The posterior belief if p is rejected 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 $(\bar{p}_H(K^*\text{-}1),\bar{p}_H(\tau_t\text{-}1)]$, which completes the proof.[28]

Proof of Corollary:

If $q_t = 0$ and $p_t \le p_H$ then $\zeta_{t+1} = \rho_L$. But consistency requires $\zeta_{t+1} = 1\text{-}\rho_H + \phi\Pi_{t+1}[n_t = 0\,|\,h^0_{t+1}]$, 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 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 $j(\tau_{t+i}) = \lambda_t j(1) + (1\text{-}\lambda_t)j(0)$, and similarly for J^L.

Proposition B2:

Given that the seller follows $\sigma^S(\zeta^*,\lambda)$, the buyer's value is the value of a dynamic programming problem with control q and state τ. For $\tau \ge 1$, the seller offers $p_L(\tau)$, and if q = 1 then $\tau' = \tau\text{-}1$, and if q = 0 then $\tau' = K\text{-}1+\lambda$. If $\tau = 0$ the seller offers p_H and if q = 1 then $\tau' = 0$, and if q = 0 then $\tau' = K\text{-}1+\lambda$. For $0 < \tau < 1$, the seller behaves as if $\tau = 1$ with probability τ, and as if $\tau = 0$ with probability $1\text{-}\tau$. 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 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 follows $\sigma^S(\zeta^*,\lambda)$. The principle of optimality yields the following functional equation for $\hat{v}^L(\tau)$ and $\hat{v}^H(\tau) \equiv \hat{v}^L(\tau) + \hat{d}(\tau)$:

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

[28]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 q =1, the buyer instead accepts with probability $1\text{-}\varepsilon^{1+n}$, while if q = 0, the buyer accepts with probability ε^{2-n}. In particular, ε^2 is the probability that a pooling offer is rejected by the high buyer and ε is the 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 ζ=1/(1+ϵ). Taking ε = 1/m yields a sequence ζ_m converging to 1, supporting the belief that if a pooling offer is rejected, the buyer's valuation is low.

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 (i.e. it is true for next period, then it is also true for this period). 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 $\tau \geq 1$, the definition of $p_L(\tau)$ implies that $\hat{V}^L(\tau)$ is the maximum of two equal numbers, and

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

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

$$\hat{V}^L(\tau) = \delta V^L + \delta\left[1-\rho_L\right]d(K-1+\lambda^*) = V^L;\ \hat{d}(\tau) = 1 + \beta\, d(\tau-1) = d(\tau)$$

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

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

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

$$\hat{V}^L(0) = \delta V^L + \delta\left[1-\rho_L\right]d(K-1+\lambda^*) = V^L;\ \hat{d}(0) = \beta\, d(K-1+\lambda^*) = d(0)$$

For $0<\tau<1$, $\hat{V}^L(\tau)=\tau\hat{V}^L(1)+(1-\tau)\hat{V}^L(0)=V^L$ and $\hat{V}^H(\tau)=\tau\,\hat{V}^H(1)+(1-\tau)\hat{V}^H(0) = V^L + d(\tau)$. Thus $\hat{V}^L(\tau) = V^L$ 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 $\tau \geq 1$, and similarly the definition of $p_L(0)$ is such that the low buyer's value at $\tau = 0$ satisfies the equation

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

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 $\bar{p}_H(s+1) < p \leq \bar{p}_H(s)$, then $q = 1$ implies $\tau' = 0$, and $q = 0$ implies $\tau' = s+\lambda$, where λ is defined by $p = \bar{p}_H(s+\lambda)$. Then the high buyer's value is

$$\delta V^L + \delta\rho_H d(s+\xi) = 1+\theta - \bar{p}_H(s+\xi) + \delta V^L + \delta\rho_H d(0)$$

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

$$\bar{p}_H(s+\xi) = 1+\theta - \delta\rho_H\left[d(s+\xi) - d(0)\right]$$

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 shown in the proof of B2, the definitions of $p_L(\tau)$ and p_H imply the tight pricing property (A3). Property A4 for $n_t = 0$ is included in B2. For $n_t = 1$, the continuation following $q_t = 0$ is the same for any value of τ_t, and A3 implies that this has value $V^H(0)$.

Lemma S2:

According to $\sigma^B(\zeta^*,\lambda^*)$, all prices below $p_L(\tau_t)$ are accepted, and such acceptances are uninformative according to Π^*, and thus imply the same stochastic process for future payoffs. Thus prices below $p_L(\tau_t)$ are dominated by $p_L(\tau_t)$. Any price in the interval $(p_L(\tau_t),p_H]$ is rejected 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 involve exchanging $p_L(\tau)$ and p_H, or else charging a price above p_H.

For any history, the continuation from $p = p_H$ 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 $\tau_{t+1} = K-1+\lambda^*$, and the value of this is $U^L(0)$; and the buyer accepts if $n_t = 1$, with $\zeta_{t+1} = 1-\rho_H$ and $\tau_{t+1} = 0$, and the value of this is $U^L(0) + g(0)$. Thus after any history the value of offering p_H is $U^L(0) + (1-\zeta_t)g(0)$, which is linear in ζ_t, and decreasing.

Recall that for $\tau_t \le 1$, $p_L(\tau_t)$ 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 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 value is then $U^L(1) + (1-\zeta_t)g(1)$, which is linear in ζ_t, and decreasing less rapidly than the value of a screening offer, because $g(1) = g(0)-(b^*+r^*\theta)$. Lemma T0 showed that the threshold ζ^* is defined by the equation $U^L(0)+(1-\zeta^*)g(0) = U^L(1)+(1-\zeta^*)g(1)$, so for $\tau_t \le 1$, p_H is optimal if $\zeta_t \le \zeta^*$, and $p_L(\tau_t)$ is optimal if $\zeta_t \ge \zeta^*$.

By Lemma T2, (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 (b) is ruled out, so $\zeta_t \le \zeta^*$, and if $\tau_t \ge 1$ then (a) is ruled out, so $\zeta_t \ge \zeta^*$, and if $0 < \tau_t < 1$ then both cases are ruled out, so $\zeta_t = \zeta^*$. Thus p_H is optimal 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 $U^L(0) + (1-\zeta_t)g(0)$ obtained by offering p_H does not exceed $U^L(\tau_t) + (1-\zeta_t)g(\tau_t)$, which is the value of following $\sigma^S(\zeta^*,\lambda^*)$. To verify this, note that

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

$$\ge r^*\theta B(\tau)\left[1 - \frac{1-\zeta}{1-\zeta^*}\right] = r^*\theta B(\tau)\left[\frac{\zeta-\zeta^*}{1-\zeta^*}\right]$$

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:

Lemma S2 shows that no price below p_H yields a higher payoff than the price specified by $\sigma^S(\zeta^*,\lambda^*)$. Prices above $\bar{p}_H(\tau_t-1)$ are rejected by $\sigma^B(\zeta^*,\lambda^*)$, and such rejections are uninformative according to Π^*, so these prices are dominated by $p_L(\tau_t)$, which yields the same future payoffs, plus some current profit (recall that $p_L(\tau_t)$ is positive). What must be shown is that no price between p_H and $\bar{p}_H(\tau_t-1)$ improves the payoff. Suppose $\bar{p}_H(s+1)<p\le \bar{p}_H(s)$, for some integer s, with $\tau-1\le s\le K-1$. Then the buyer accepts with probability q(s), with continuation from $\zeta' = 1-\sigma$ and $\tau' = 0$, and rejects with probability 1-q(s), with $\zeta'=\zeta^*(s)$ 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 continuation value from next period are both linear in ξ, so if $\bar{p}_H(s+\xi)$ is a profitable deviation, 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. Next, the probability q(s) that $\bar{p}_H(s)$ is accepted is such that if the offer is rejected the belief next time will be $\zeta' = \zeta^*(s)$, so after a rejection and before the transition from n to n' the seller's belief is $\zeta^*(s+1)$. Thus $\zeta^*(s+1) = \hat{\zeta}/[1-q(s)]$, and in equilibrium $\hat{\zeta} = \zeta$. The price $\bar{p}_H(s)$ 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 value is

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

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

$$u_H^a - u_H^r = J^H(0) - \left[J^H(s+1) - (1+\theta)\right] = 1+\theta + [B(s+1) - R(s+1)]r^*\theta$$

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

$$\Omega(s+1) = U^H(0) - \zeta g(0) - u[p,\zeta,\tau] = -\left[u_H^a - U^H(0)\right] + [1-q(s)]\left[u_H^a - u_H^r\right] - \zeta\left[g(0) - g(s+1)\right]$$

Consider the first term. Since the joint continuation value is always the same when a screening offer is accepted, this is the difference between extended and one-shot 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. Therefore,

$$u_H^a - U^H(0) \;=\; \delta V^L + \delta\rho_H d(K-1+\lambda^*) - \delta V^L - \delta\rho_H d(s) \;=\; \frac{\rho_H}{\phi}\Big[1 - B(s+1)b^*\Big]$$

Note here that

$$d(0) \;=\; \beta\, d(K^*-1+\lambda^*) \text{ , and } d(s+1) \;=\; 1+\beta\, d(s) \;=\; d(0) + B(s+1)\,b^*$$

These results can be summarized as :

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

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

Proposition E1

The proof of E1 uses the following result.

Lemma E0: If $c > 0$, $0 < \delta < 1$, and $0 < \phi < 1$, then the function $f{:}\Re \to \Re$ defined by $f(s) \;=\; a\delta^s - \dfrac{c}{\phi^s} - (\delta\phi)^s$ is quasiconcave.

Proof:

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

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

The right side increases from χa to ∞, and ϕ^s decreases from ∞ to 0, as s increases from $-\infty$ to $+\infty$, so there is a unique solution. Also,

$$f'(s) \;=\; a\log(\delta)\delta^s + c\log(\phi)\phi^{-s} - \log(\delta\phi)(\delta\phi)^s \;=\; -\log(\delta\phi)\delta^s\left[\phi^s - \chi a - \frac{c(1-\chi)}{(\delta\phi)^s}\right]$$

The bracketed expression is decreasing in s, and zero at $s = s_0$, and $\log(\delta\phi) < 0$, so f is increasing for $s < s_0$, and decreasing for $s > s_0$.

In the following proofs, 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 S3 shows that the first part of E1 can be stated as $\Omega(s) \ge 0$, for $1 \le s \le K^*$. This is equivalent to $\omega(s) \ge 0$, where

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

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

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

where $\Xi = -1 + \rho_H/[\phi(1\text{-}\rho_H)]$. After substituting for D(s) and B(s), $\omega(s)$ can be written in the form

$$\omega(s) \;=\; a_0 + a_\delta\delta^s - a_\phi\phi^{-s} - a_\beta\beta^s$$

where a_0 and a_δ are irrelevant constants, and

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

The coefficient a_ϕ is nonnegative if $\beta^{K^*}\left[\Xi+1\right] \le y_0$. This is implied by $\overline{X}_{K-1}$:

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

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

Since a_β is positive, the function $\omega(s)/a_\beta$ satisfies the conditions of Lemma E0, so $\omega(s)$ is quasiconcave. The next step is to show that $\overline{X}_{K-1}$ 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

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

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

$$\begin{aligned}\omega(K-1) = {} & \left[\frac{B(K^*-1)}{\zeta^*} + \frac{\rho_H}{\phi\,(1-\rho_H)}\,\beta^{K^*-1}\right]\left[1-\zeta^*(K^*-1)\right] - \beta^{K^*-1}\left(\frac{\rho_H}{\phi\,(1-\rho_H)} - 1\right) \\ & + \delta^{K^*-1}\rho_L(K^*-1)\left[\frac{1}{\zeta^*}-1\right] + \lambda\,\delta^{K}\left[1-\mu - \left(\frac{\rho_H}{\phi\,(1-\rho_H)}\,\rho_L - \mu\right)\phi^{K^*}\right]\end{aligned}$$

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

$$\rho_L(K^*-1)\left[\frac{1}{\zeta^*}-1\right] - \phi^{K^*-1}\left(\frac{\rho_H}{\phi\,(1-\rho_H)} - 1\right) \ge 1-\mu - \phi^{K^*-1}\left(\frac{\rho_H}{\phi\,(1-\rho_H)} - \mu\right)$$

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 $\overline{X}_{K-1}$. It is obvious that $\overline{X}_{K-1}$ 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 $\bar{\rho}(K^*-1) \le \theta/G(K^*) < \bar{\rho}(K^*)$, then a (small) increase in θ reduces ζ^* without disturbing K^* or λ^*. Write X_1 as

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

The first term in brackets is positive, because $\zeta^*(1) \le 1 \le R(K+\lambda^*)$; also $1/\zeta^*(1)$ increases when ζ^* decreases, so the left side of the inequality is increasing in θ. The other possibility is that $G(K^*) \le \theta/\bar{\rho}_L(K^*) < G(K^*+1)$. Then a (small) increase in θ increases λ^* while K^* and ζ^* remain unchanged.[29] It will be shown that this increases $\omega(1)$, implying that an increase in θ relaxes X_1. Write $\omega(1)$ as

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

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

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

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

[29] One might ask how K^* ever changes, given 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 $\overline{X}_{K-1}$ is equivalent to $\omega(K-1) \ge 0$ at this point.

and note that

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

is positive when $\overline{X}_{K-1}$ holds. This completes the proof.

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,

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

Substituting $\zeta^* = \mu$ in this equation shows that (X_1^∞) is equivalent to $\omega(1,\theta_\infty) > 0$, so there is an interval (θ_0,θ_∞) such that $\omega(1,\theta) > 0$ for $\theta \in (\theta_0,\theta_\infty)$. Since $\overline{X}_{K-1}$ necessarily holds for large K^*, θ_0 can be chosen to ensure that $\overline{X}_{K-1}$ holds for $\theta \in (\theta_0,\theta_\infty)$. Then the conditions of Proposition E1 are satisfied for $\theta \in (\theta_0,\theta_\infty)$.