Repeated Signalling Games and Dynamic Trading Relationships

by

Daniel R. Vincent[1]

Abstract

A seller of a nondurable good repeatedly faces a buyer who is privately informed about the position of his demand curve. The seller offers a price in each period. The buyer chooses a quantity given the price. The quantity demanded reveals information about the buyer. An equilibrium is characterized with the feature that buyer types separate completely in the first period. This equilibrium uniquely satisfies a modified refinement of the Cho-Kreps criterion. Despite the immediate separation, the buyer distorts his behavior throughout the game. The requirements to signal types can raise the utility of all types of informed players.

Keywords: Repeated signalling games, refinements of equilibria, bilateral monopoly, bargaining, ratchet effect.

Journal of Economic Literature Classification Numbers: C73,D43,D82,L14.

Running Head: Repeated Signalling Games.

April, 1997

1. Introduction

Playing hard to get is time-honored in markets as it is in love. The coy and clever buyer knows that betraying too much eagerness to a seller can often place him at a disadvantage as their relationship develops. However, feigned indifference comes at a cost. Delayed consumption destroys irrevocably some opportunities for satisfaction. A careful buyer must always balance his wish for immediate gratification with a caution against betraying his true desires; an interested seller must balance her desire to benefit from the current transaction with the need to extract information about the future of the relationship.

In dynamic games, this phenomenon gives rise to the so-called ratchet effect. When an uninformed agent learns information early in a game, she can be expected to exploit it subsequently to her opponent's disadvantage. One result is that the cost of inducing information revelation grows the longer the trading relationship is expected to persist. The consequence of this behavior has been found to suppress the revelation of information in dynamic games (see Freixas, Guesnerie and Tirole 1985, Hart and Tirole 1988 or Laffont and Tirole 1987). These results have generally been derived in models in which the uninformed agent is in an unusually strong strategic position either because of her contracting power or because of very simple preferences of the informed agent which offer agents of one type very little scope to separate from agents of another type. If the consequences of revealing information to an uninformed agent are severe, then we can expect little information revelation. And, in dynamic trading games, if the only way agents can separate is through the stark choice of zero purchases or single unit purchases, again we can expect little information revelation. What happens in a different contracting environment and when preferences are such that the possibilities for screening are richer?

In Section 2, a dynamic contracting game is described in which two agents desire to trade a divisible, non-durable good or a service, period by period. The buyer has private information about the position of his demand curve. Since the relationship is dynamic, the problem of an uninformed seller is twofold: to extract surplus from the current trade; and to extract information that may be exploited in future trades. Allowing an agent in a bilateral monopoly the sole right to offer non-linear contracts yields that player a substantial amount of strategic power. This power and the temptation to extract surplus in later periods reduces the ability to extract information in the current period. One way to relax this stark formulation is to restrict the seller to another commonly observed type of offers, linear contracts in which the seller offers a price and the buyer chooses a quantity.

This game possesses a perfect Bayesian equilibrium with a very stark feature. For a large subset of the parameter space, the equilibrium path is characterized by immediate information revelation by both types. Despite the revelation, though, economic behavior continues to be distorted along the equilibrium path as the low type buyer is forced to continue to convince the seller that he is indeed a low type. The low type buyer separates from the high type by selecting a quantity determined by a demand curve lower than his true demand in every period but the last even though at intermediate stages of the game the seller has acquired enough information to know with probability one whether the buyer is a high or low type. In this sense, the equilibrium illustrates that HOW an agent knows information as well as WHAT the agent knows can play an economic role.

Dynamic games in which the informed player has a large strategy space are typically plagued by a large set of equilibria. This model is no exception and, indeed, there also exist other separating equilibria as well. It is natural, therefore, to investigate the plausibility of this particular equilibrium. In Section 4, I show that if the model is extended to allow for any small but positive probability that a buyer's type might change in every period, then the equilibrium characterized in Section 3 is the unique one to survive the iterative application of a well-known refinement of perfect Bayesian equilibrium.

Section 5 examines other features of this equilibrium. As a result of the persistent concern that a seller may revise her beliefs, the equilibrium behavior confers a surprising benefit on both buyer types. In signalling games, one type of informed agent often incurs a cost due to the asymmetry of information as he attempts to separate from the other type. In this environment, though, it is shown that *both* types of buyers can benefit from the presence of asymmetric information. In a repeated context, the fear that the uninformed agent may update in an unfavorable way following some actions of the informed agent can serve as a valuable commitment device in earlier stages of the game that allows the agent to commit to strategies that would otherwise not be credible.

2. The Model

In a T-period game[2], an uninformed supplier faces a buyer who is privately informed about his preferences for a non-durable good. The seller can provide the good at constant (zero) marginal cost and seeks to maximize total discounted expected profits. A buyer of type a_t in period t obtains a per period utility from a quantity, q_t, of the good purchased at a per unit price, p_t, given by

(1) $$2(a_t - p_t)q_t - q_t^2, \quad a_t \in \{a_L, a_H\}, \; q \geq 0.$$

Observe that the marginal rates of substitution between q_t and the money good differs depending on the buyer type. The buyer wishes to maximize the expected value of the T-period discounted sum of (1) (the discount factor, δ, is the same for buyer and seller) and has private information about the true value of a_t which may be high or low. The prior probability that the buyer is a high type at the beginning of the game ($Prob[a_T = a_H]$) is $\mu_T \in (0,1)$.[3] An example generating these payoffs is a market where the seller sells to a downstream retailer or importer who incurs a quadratic cost of distributing the good. In addition, the true price the retailer receives upon reselling the object, a_t, remains private to the retailer perhaps because of unobserved taxes or rebates or other linear costs.

In a myopic or static framework, the preferences of the buyer yield a demand curve of the form $q_t = max\{0, a_t - p_t\}$ and the monopolist's optimal price for $a_L \geq a_H/2$ is just a weighted combination of her monopoly price against each of the possible demand curves where the weight is given by her prior belief about the buyer type. It is never optimal for the seller to charge a price above $a_H/2$ so there is no loss of generality in restricting attention to prices below this bound. However, if $a_L < a_H/2$, even in the static game, the seller's profit function is non-concave in prices for some beliefs. In this case, her optimal price is either the weighted combination or just $a_H/2$ depending on her beliefs. Since these issues are not the focus of this paper, I rule out this possibility by maintaining the restriction, $a_L \geq a_H/2$.

The game consists of the following moves. At the beginning of each period, the seller offers a per unit price and commits herself to providing any quantity the buyer chooses at that price. The buyer then chooses a non-negative quantity. In the dynamic game, the strategic power of the seller is determined by the space of contracts available to her. The extreme case in which the seller may choose only linear pricing contracts is restrictive but it is of interest for two reasons[4]. A truly bilateral monopoly should have the characteristic that in the case of complete information, there is a non-trivial sharing of surplus. Restrictions on the strategy space are often exploited to induce this division. One way to model bilateral bargaining power would be to incorporate explicitly offer-counteroffer negotiations in each period as in Rubinstein (1982). Solving this game would be a daunting task. The restriction to linear pricing schemes allows the construction of a game which forces the sharing of some surplus and, as it turns out, is tractable. Furthermore, we often observe such restrictive pricing schemes. For example, negotiations between unions and firms often have the characteristic that a wage is determined (in this model, set by the union) and employment is then selected by the firm.

In the next Section, a perfect Bayesian equilibrium[5] of the linear pricing game is illustrated with the feature that for a large class of parameter values the buyer types separate in every period.

3. Equilibrium Behavior in the Linear Pricing Game -- the Case of Complete Separation

Dynamic signalling games typically possess many perfect Bayesian equilibria (*pBe*). The focus of this paper is on equilibria which are fully revealing in every period. It is well-known that even in "well-behaved" two-stage signalling games, fully separating equilibria may not exist if, given preferences, the signalling space is not large enough to allow one type to profitably distinguish himself from another type (Cho and Sobel 1990). Since I am interested in examining the nature of separation in many period signalling games, that issue is side-stepped by placing restrictions on the parameter space in order to allow the possibility of full and immediate revelation. The equilibrium characterized in this section is extremely simple. In each period, buyer types separate for every price by choosing a quantity determined by a linear demand curve independent of seller beliefs and of the history of the game. The high type chooses the demand corresponding to his static demand curve, the low type, in every period, chooses a quantity off a linear demand curve that is generally lower than (but parallel to) his static demand curve. The actual position of the demand curve is such that in any period if the buyer type is high, he is just indifferent between mimicking the low type in this period and for the rest of the game and choosing his high demand. As a result, the more periods remaining to the end of the game, the lower this low type equilibrium demand curve must be. (See Figure 1.)

[FIGURE 1 ABOUT HERE.]

The proof that such behavior is the outcome of a perfect Bayesian equilibrium is best seen by construction. Define a monotonic sequence, $\{a_{Lt}\}$ as follows. Let $x^* = a_H(4\text{-}3\delta)/(4\text{-}\delta)$. If $a_L \leq x^*$, set $a_{Lt} = a_L$. Otherwise define a_{Lt} iteratively by $a_{L1} = a_L$ and

(2) $$a_{Lt} = a_H - \sqrt{\delta(a_H - a_{Lt-1})(3a_H - a_{Lt-1})}/2, \quad for\ t \geq 2.$$

The intercept of the low type demand curves is given by a_{Lt} in every period. It can be shown that a_{Lt} converges to x^* as t becomes large and that for a fixed t, a_{Lt} falls with increases in δ. In order to ensure that the signalling space is large enough to allow complete separation, I maintain the following restriction on parameters:

(A1) T,δ,a_L,a_H are such that $a_{LT} \geq a_H/2$.

Notice that for $\delta \leq 4/5$, **A1** holds for all values of T as long as $a_L \geq a_H/2$ holds.

In a perfect Bayesian equilibrium, we must specify seller beliefs following any history of the game. In all of what follows, I consider equilibria in which the seller's beliefs may change only following a move by the buyer. Given Bayes' rule, then, it is sufficient to characterize seller beliefs solely by the sequence, $\{\mu_t\}_{t=1}^{T}$, which is the probability the seller places on the buyer being a high type in period t just before the seller posts a price, p_t.

Theorem 1: *Assume **A1**. The following behavior can be supported as a perfect Bayesian equilibrium outcome of the game. For any price, p_t, a buyer who is a high type in period t demands $q_H = a_H - p_t$ and a low type buyer demands $q_{Lt} = a_{Lt} - p_t$. For any history resulting in seller's beliefs in period t of μ_t, the seller offers a price $p_t = (\mu_t a_H + (1-\mu_t)a_{Lt})/2$.*

Proof: For some period τ onwards define strategies as

> **H1**$_\tau$: In every period, $i \leq \tau$, for every history, for every seller belief, and for every price offer, p_i, a high type buyer in period i demands $a_H - p_i$ and a low type buyer demands $a_{Li} - p_i$.
>
> **H2**$_\tau$: In every period, $i \leq \tau$, for every history of the game, the seller offers a price, $p_i(\mu_i) = (\mu_i a_H + (1-\mu_i)a_{Li})/2$.

Clearly, **H1** and **H2** are perfect Bayesian equilibrium strategies for $\tau = 1$.[6] Suppose that they can be supported as *pBe* strategies following some period $\tau = t$-1. I show that they can also be supported as *pBe* strategies for $\tau = t$ and the theorem follows by induction.

Observe that by assumption **A1**, $p_i(\mu_i) \leq a_{Li}$ in every period, so positive quantities will be demanded in every period. **H1** and **H2** imply that in period t-1, whatever his behavior in the past, a high type expects to separate by demanding $a_H - p_{t-1}$ in period t-1 and that the subsequent equilibrium path is stationary. Thus, his payoff following period t-1 can be represented simply by some constant, v_{Ht}.

Now let p_t be a seller price offer in period t and define q_{Lt} so that

$$2(a_H-p_t)q_{Lt}-q_{Lt}^2+\delta((a_H-p_{t-1}(0))^2+v_{Ht}) = (a_H-p_t)^2+\delta((a_H-p_{t-1}(1))^2+v_{Ht})$$

That is, q_{Lt} is the highest quantity choice such that the high type is just indifferent between revealing herself now by demanding $a_H - p_t$ or demanding q_{Lt} in this period, persuading the seller he is a low type and receiving the lowest possible price, $p_{Lt}(0)$ in period t-1. By **H1**, in either case, he will demand a_H-p_{t-1} in the following period and reveal his type. Note that $q_{Lt} = a_{Lt} - p_t$ where a_{Lt} is defined by (2).

In order to show that demanding a_H-p_t is optimal for the high type we need to describe the consequences of other choices. Observe that given $\mathbf{H1}_{t\text{-}1}$, the seller's strategy defined by $\mathbf{H2}_{t\text{-}1}$ is sequentially rational and determined solely by her t-1'st period beliefs. The consequences of a deviant quantity in period t, then, are determined by the effects of this quantity on the t-1'st period beliefs of the seller, μ_{t-1}. For any history and any quantity choices, $q_t \leq q_{Lt}$, the seller believes that the type is a low type in the current period with probability one and therefore $\mu_{t-1} = 0$. If $q' \geq q_{Lt}$, define $\mu^*(q')$ implicitly by

$$2(a_H-p_t)q'-q'^2+\delta((a_H-p_{t-1}(\mu^*))^2+v_H) = (a_H-p_t)^2+\delta((a_H-p_{t-1}(1))^2+v_H).$$

For any history with $q_t > q_{Lt}$, she believes that the buyer is a high type with probability at least $\mu^*(q')$ and therefore, $\mu_{t-1} \geq \mu^*$. (Note that even for $\mu_t = 0$, so the seller initially believes she is facing a high-type buyer with probability zero, whenever $q' > q_{Lt}$ is an out-of-equilibrium event, seller updating from $\mu_t = 0$ to $\mu_{t-1} > 0$ is consistent with *pBe*.) With these beliefs, the definition of μ^* ensures that the high type at least weakly prefers to demand a_H-p_t to any lower quantity. Lemma 1 in the Appendix uses the induced preferences of the two types of buyers to show via a single-crossing property that the low type strictly prefers q_{Lt} which results in the lowest plausible price in the next period to any higher quantity and the higher next period price, $p_{t-1}(\mu^*)$. Thus the behavior of both buyer types satisfy **H1** for period t, and since buyer types separate for any price, there is no dynamic role of pricing for the seller. For any seller belief in [0,1], given the separating behavior of the buyer for any p_t, the monopolist's optimization problem is exactly the same as the static problem confronting a monopolist with a linear demand curve with intercept a_H or a_{Lt}. Given **A1**, this problem is concave, $p_t(\mu_t) \leq a_H/2$ for all μ_t and the behavior described in **H2** is optimal for period t as well. Since **H1** and **H2** is satisfied for $\tau = 1$, induction implies that the behavior satisfies the conditions of a perfect Bayesian equilibrium for all periods T such that $a_{LT} \geq a_H/2$. ■

Notice that in the extreme case of $\delta = 0$, buyers discount the future completely. In this case, $x^* = a_H$ and, by definition, $a_L = a_{Lt}$ for all t. Separation implies no distortion since there is no incentive for the high-type to underdemand. Similarly with low δ's, separation will always occur and the equilibrium corresponds to the repeated static solution.

More interesting, however, is the case where the future matters because of a higher discount factor. In this case, where $x^* < a_L$, the succession of demand curves, a_{Lt}, is monotonically decreasing in t. As the length of the game increases, the low type buyer must underdemand more in order to dissuade imitation by the high type.

Observe that,

$$\lim_{T\to\infty,\delta\to 1} a_{LT} = a_H/3$$

If δ is high and T is large, it is possible that assumption **A1** is violated. This possibility is also a potential source of non-concavity in the seller's optimization problem and her pure strategy best response correspondence may not be convex-valued. In such cases, complete separation cannot be supported by the perfect Bayesian equilibrium described in Theorem 1. Some partial pooling will typically occur in early stages of the game such that $a_{LT} < a_H/2$. Initially, then, there may be some gradual learning. In these cases, the equilibrium behavior is not as simple as that of Theorem 1 since they will typically involve mixed strategies in equilibrium for the high type buyer and mixed strategies (out of equilibrium) for the seller. These complications are accounted for in an earlier version of this paper (Vincent, 1994) which provides a characterization of equilibrium strategies over the full parameter space. For values of a_H, a_L, δ, and t such that $a_{Lt} \geq a_H/2$, the equilibrium paths coincide. Otherwise, if the seller believes relatively strongly that the buyer is a high type, it may be in her best interest to offer prices which induce only gradual revelation by the high-type buyer. In these periods, behavior very similar to the original ratchet effect of Laffont and Tirole emerges. The high-type buyer reveals himself with some probability, β_t. Only as the game approaches the later periods does complete separation emerge.

Theorem 1 characterizes only one of many possible *pBe*. Even in two stage signalling games, pooling equilibria often can coexist with separating equilibria. This is true here for $T = 2$ and, therefore, for $T > 2$ as well. In two-period games, many such pooling equilibria fail to survive common belief-based refinements of perfect Bayesian equilibria. I show next that a similar approach may be applied in the multi-period game. A recursive application of the Cho-Kreps (1987) intuitive criterion is defined. The original game is modified to allow for stochastically changing types (but with arbitrarily small probabilities of changes). Theorem 2 illustrates that, in this game, the equilibrium path described in Theorem 1 is the only one to survive this restriction.

4. A Refinement of Perfect Bayesian Equilibrium

The equilibrium characterized in Theorem 1 has the feature that the private information of the buyer is revealed in the first period of the game. This stark learning behavior points to an intriguing and controversial feature of the equilibrium. If the seller observes an equilibrium low quantity in an early period of the game, she must believe she is facing a low type with probability one. Furthermore, she must continue to believe she is facing a low type throughout the rest of the game. Even so, it is not sequentially rational for her to revert to the optimal monopoly price, $a_L/2$ against a low type. Instead, since the low type is demanding off a lower demand curve, a_{Lt} - p_t, the seller's sequentially rational price falls to $a_{Lt}/2$. Despite the equilibrium generated knowledge that the low type is in the game, behavior continues to be distorted along the equilibrium path.

Is this apparently counterintuitive feature simply a curiosum of the large size of the set of perfect Bayesian equilibria of repeated signalling games? One way to approach this question is to determine whether the equilibrium would survive the application of a common refinement of sequential equilibrium. A difficulty arises however, in the attempt to extend the definitions of these refinements to multistage games with full separation. After separation occurs, the seller will believe with probability zero that the buyer is of a particular type. But many belief-based refinements require the comparison of the value to various types of potential continuation paths following a deviation. In this environment, such a comparison would be vacuous if we did not allow the seller to consider the possibility of changing her belief from probability zero to a positive probability.

The equilibrium characterized in Theorem 1 exhibits this phenomenon of increasing supports off the equilibrium path. Seller updating includes the feature that for high-quantity deviations the seller will change her belief that the buyer is a high-type with probability zero to a belief that the buyer is a high-type with some positive probability. Note that perfect Bayesian equilibria (and sequential equilibria) of signalling games not only allows for this type of updating, in some games, it is required to ensure the existence of sequential equilibrium. (See, Madrigal, Tan and Werlang, 1987 and van Damme and Noldeke, 1990.) Beaudry and Poitevin also use a dynamic extension of a standard refinement to analyze one-shot signalling game with the possibility of later renegotiation. Their model also yields equilibria where the uninformed agent's beliefs feature increasing supports off the equilibrium path.

In this section, I skirt the issue by extending the model to ensure that the uninformed agent can never believe she is facing any given type with probability one. The extended model introduces the possibility that informed agents' types may change exogenously in every period. Bayesian updating by the seller must take this possibility into account when she formulates her new beliefs. Specifically, let a_{t+1} be the buyer's type in period t+1 and assume that in any period, buyer types follow a stationary Markovian process of the following form:

(3) $$Prob[a_t = a_H | a_{t+1} = a_i] = \epsilon_i, i = L,H, 0 \leq \epsilon_L \leq \epsilon_H \leq 1, \quad for\ \ t = 1,2,...,T-1.$$

Thus, conditional on being of type i in period t+1, the buyer is relatively more likely to be of type i in period t. I focus on the limiting case where ϵ_H approaches one and ϵ_L approaches zero, but the model could, of course, be interpreted literally as a description of a game where there is a significant probability that types change over the course of the game. Consistent with full rationality, both the buyer and the seller factor in the possibility of future type changes in the determination of current optimal strategies. The operative role of this modification is that for any $\epsilon_L > 0$, in no period, does the seller believe with probability zero, she is facing a high-type buyer and therefore, the issue of non-increasing supports does not arise. This feature is used to describe the extension of the refinement.

Before defining the refinement formally, it may be helpful to walk through a short example to illustrate its application in a three period game with $a_H = 1$, $\delta = 1$, $a_L = 3/4$ and $\epsilon_H = 1$.[7] Consider the subform of the game beginning at the second stage of the second last period with the seller's price, p_2, already set and with some current seller belief, μ_2, *strictly* above 0. Sequential rationality imposed on the buyer in the last period implies that we can represent buyer types's preferences in (q_2,p_1) space by the following induced utility function:

$$V(q_2,p_1;p_2,a_i)=(a_i-p_2)^2-(a_i-p_2-q_2)^2+\epsilon_i(1-p_1)^2+(1-\epsilon_i)(3/4-p_1)^2,\quad i=L,H.$$

Figure 2 illustrates these preferences in (q_2,p_1) space. It can be shown (Lemma 1) that these indifference curves are concave and that the slope of the high type indifference curves are steeper than those of the low type. Seller behavior in the last period is determined for any belief μ_1, and since buyer behavior in the last period is a function only of seller price, p_1, a two-period signalling game emerges from the overlap of the second stage of period two and the first stage of the last period. The buyer's second period quantity demand acts as a signal which the seller observes and generates a response which is the final period price.

[FIGURE 2 ABOUT HERE.]

In a two-period signalling game, the Cho-Kreps (1987) Intuitive Criterion (what will be refinement R_2 in the multi-period game) rules out candidate *pBe* outcomes such as point *A* in Figure 2 as follows.[8] Trace the indifference curve of the high-type downward and to the left until it crosses the $p_1 = 3/8$ line at q'. Suppose a deviant quantity, $q' - \Delta$ is demanded at this stage. There is no seller belief and subsequent best response that could yield an outcome for the high type buyer that the buyer prefers to the candidate outcome, *A*. On the other hand, if the seller updates her beliefs following the deviation by putting zero weight on the high type and responds with the price $p_1 = 3/8$, for small enough Δ, the low type gains a strictly higher payoff than from *A*. In a simple two-stage signalling game, the Cho-Kreps criterion would imply that the seller then should believe that only a low type would demand such a low quantity and therefore update with a belief $\mu_1 = \epsilon_L$ and respond with a price in the last period of $p_1 = 3/8$.

This argument can be used to eliminate all pooling outcomes. The only outcome that survives the restriction to this type of seller updating rule is the separating outcome B for the high type and C for the low type shown also in Figure 2. The high type receives his full information equilibrium payoff while the low type underdemands just enough so as to dissuade imitation from the high type. This last condition yields the quantity $q_{L2} = a_{L2} - p_2 = 5/8 - p_2$ for any price offer p_2. Thus a lower "effective" demand curve determines the low type's behavior in the second to last period.

Given the necessary buyer behavior in period 2, the seller's best response in period 2 is again a simple function of her beliefs: $p(\mu_2) = (\mu_2 + (1-\mu_2)*5/8)/2$. The fact that the buyer types separate for any sequentially rational price offer of the seller implies that there is no informational (and therefore no dynamic) role of prices. Instead, the seller is again in a situation analogous to that of the static monopolist facing one of two possible demand curves. This time, though, the low demand curve is lower than before because of the low type's desire to separate from the high type.

Observe that, except for the determination of the seller's price, p_2, this argument is made independent of the actual value of the seller's current prior, μ_2. What is required is that, fixing any candidate equilibrium, for `low enough' deviant quantities, the seller will believe that a low type made the demand. This requirement, in turn, requires that to support the only equilibrium outcome satisfying this feature, we be able to place a high enough probability that a high type makes a deviant quantity above the low type's quantity, q_{L2}. If $T = 2$, then these requirements are satisfied as long as $a_L > a_H/2$ and $\mu_T = \mu_2 > 0$. In a three-period game, the issue becomes more delicate if $\epsilon_L = 0$, because then it is possible that along the equilibrium path, $\mu_2 = 0$.

Given that buyers expect to separate in this manner for all prices in the second stage of the middle period, the utility that each buyer type expects from any continuation path at the beginning of period 2 will be a function only of the seller price offer, p_2, independent of the history of the game. Furthermore, if we move forward in the game to the second stage of the initial period, with a seller price, p_3, outstanding, it can be shown that the induced buyer preferences over the quantity they demand given this price and the subsequent p_2 that this generates from the seller are qualitatively similar to those in Figure 2. The equilibrium path isolated in Theorem 1 is obtained by applying the intuitive criterion (R_3) after replacing the continuation paths of the game with the expected payoffs (which depend only on the seller beliefs and through them on subsequent seller price offers). An argument similar to that for the overlap of period 2 and 1 applies here as well and yields complete separation again in the first period, this time with the low type buyer demanding a quantity, $q_{L3} = 1-.5((1-5/8)(3-5/8))^{.5} = .528 - p_t$.

The equilibrium price path follows one of two patterns. Independent of the buyer type, the initial price is given by $(\mu_3+(1-\mu_3)q_{L3})/2$. If the buyer is a high type, the quantity demanded is relatively high, and subsequent prices move immediately to 1/2 in each of the remaining two periods. If the buyer is a low type, the initial demand is lower than the low type's static demand, the seller's next period price falls to a price below her static monopoly price against a low type and then rises as the game continues.

The characterization of the formal refinement requires some more notation and definitions. Let h_t denote a history of a game up to the end of period t and (h_t,p_{t-1}) denote the history to the middle of period t-1.[9] Since a strategy determines the continuation play of the game for any given history, we can compute expected payoffs from t-1 onwards for buyer and seller given a history and a strategy, σ.

Denote the buyer's expected payoff by $v^{\sigma}_{t}(h_t,p_{t-1},a_j)$ when the history is (h_t,p_{t-1}), the buyer type in period t is a_j and the strategies from t onwards are determined by σ. The seller's expected payoff from the same stage in the game onward, conditional on the buyer type in period t being a_j is $u^{\sigma}_{t}(h_t,p_{t-1},a_j)$. Given a history, (h_t,p_{t-1}) and a profile of pure strategies,[10] σ, let $q^{\sigma}_{t-1}(h_t,p_{t-1},a_j)$ be the strategy choice of a buyer who is of type a_j in period t-1. Similarly, $p_{t-1}{}^{\sigma}(h_t)$ is the corresponding seller price choice fixed by σ.

A *pBe* also characterizes seller beliefs after any history. Let $\mu_{t-1}{}^{\sigma}(h_t)$ denote the seller's interim probability that the buyer is a high type at the beginning of period t-1 (following the random move by nature at the beginning of period t-1). Recall that I consider only *pBe* such that beliefs may change only after buyer deviations.[11]

Definition 1: A subset of perfect Bayesian equilibrium strategies, Σ, satisfies condition C_t if for all strategy profiles, $\sigma \in \Sigma$, for all $i \leq t$, for all h_{i+1},h'_{i+1}, for all $p_i \leq a_H/2$, $q^{\sigma}_{i}(h_{i+1},p_i,a_i) = q^{\sigma}_{i}(h'_{i+1},p_i,a_i)$, $a_i \in \{a_L,a_H\}$.

Definition 1 characterizes a class of strategies which exhibit a strong type of stationarity from some period t to the end of the game. Since for any *pBe*, $q^{\sigma}_{1}(h_2,p_1,a_1) = a_1 - p_1$, for all histories, the set of all pBe strategies satisfies C_1.

Definition 2: For any set of *pBe* strategies, Σ, satisfying C_t, define

$$BR^{t}_{\Sigma}(\mu)=\{p_t \mid p_t \in argmax(\mu u^{\sigma}_{t}(h_{t+1},p_t,a_H)+(1-\mu)u^{\sigma}_{t}(h_{t+1},p_t,a_L) \quad for\ some\quad \sigma\in\Sigma\}.$$

BR^{t}_{Σ} is not necessarily the set of seller equilibrium strategies since any given seller belief μ may never arise in a *pBe*. However, in order to apply the refinement, I want the ability to conduct thought experiments that range over all possible seller beliefs following a deviation. This device allows that flexibility. Notice that u^σ is defined as a function of the strategy profile alone, not seller beliefs. The variation of beliefs, μ, in the current period is not assumed to affect future play of the game.[12]

I now define a refinement of a subset of *pBe*.

Definition 3: If Σ is a subset of *pBe* of the *T* period game satisfying $C_{t\text{-}1}$,

$$\begin{gathered} R_t(\Sigma)=\{\sigma\in\Sigma\,|\,\forall h_{t+1},\forall p_t\le a_H/2, \textit{for any } \tilde{q}\notin\{q_t^\sigma(h_{t+1},p_t,a_L),q_t^\sigma(h_{t+1},p_t,a_H)\} \\ \textit{such that if} \\ v_t^\sigma(h_{t+1},p_t,\tilde{q},p,a_j)\le v_t^\sigma(h_{t+1},p_t,a_j)\ \ \forall p\in BR_\Sigma^t(\mu)\forall\mu\in[\epsilon_L,\epsilon_H], \\ \textit{and} \\ v_t^\sigma(h_{t+1},p_t,\tilde{q},p,a_k)>v_t^\sigma(h_{t+1},p_t,a_k),\ \textit{for some } p\in BR_\Sigma^t(\epsilon_k),j,k\in\{L,H\},k\ne j, \\ \textit{then}, \mu_{t-1}^\sigma(h_t,p_t,\tilde{q})=\epsilon_k\}. \end{gathered}$$

If Σ does not satisfy $C_{t\text{-}1}$, then $R_t(\Sigma) = \Sigma$. The refinement is generated by applying R_t iteratively. Let Σ_1 be the full set of *pBe*. The *T*-fold application of the refinement yields the subset of *pBe*, $\Sigma_T = R_T(R_{T\text{-}1}(...R_1(\Sigma_1)...)) \equiv R^T(\Sigma_1)$.

In words, the *t*'th refinement states the following. Suppose that all the *pBe* under consideration, $\Sigma_{t\text{-}1}$, have the feature that buyer behavior is stationary in all periods following period *t* and, so, continuation payoffs will depend only on the seller's subsequent price offers, which in turn depend only on her subsequent beliefs. Then, if a deviant demand occurs in period *t* with the characteristic that, given the hypothesized equilibrium continuation, one type does worse for any sequentially rational seller price offer while there is a sequentially rational seller price offer for which the other type does strictly better, then the seller must believe that it was the latter type who deviated in period t. Theorem 2 shows that for any $\epsilon_L > 0$, the equilibrium path described in Theorem 1 is the only one to survive the T-fold application of this refinement.

Theorem 2: *Let Σ_1 be the set of pBe and define $\Sigma_T = R^T(\Sigma_1)$. Assume **A1** and suppose $\epsilon_L > 0$. If $\sigma \in \Sigma_T$, then σ generates the equilibrium path described in Theorem 1 with a_{Lt} defined as*

$$a_{Lt}=\min\{a_L,a_H-\sqrt{\delta(\epsilon_H-\epsilon_L)(3(\epsilon_H a_H+(1-\epsilon_H)a_{Lt-1})-(\epsilon_L a_H+(1-\epsilon_L)a_{Lt-1})}/2\}$$

Proof: The proof of Theorem 2 is found in the Appendix.

The direct application of refinement, R^T puts restrictions on how the seller can update when she observes *lower* than expected quantities demanded. In a way, it allows the low type the opportunity to destroy any pooling equilibrium by signalling his type with a low quantity demand. The richness of the preferences implies that the high type is not willing to sacrifice high quantity consumption now for lower prices in the future. However, this restriction eliminates many pBe and as a consequence it also forces restrictions on how she may update when *higher* than expected quantities are observed. Even though information is fully revealed, screening costs are incurred throughout the game. After the first period, both the buyer and the seller know all the relevant information for the rest of the game. Nevertheless, the equilibrium strategy of the low type is to underdemand for the remainder of the game (except the final period). He acts as if he had a demand curve with a strictly lower intercept. Given this behavior, the seller can do no better than to post a lower price. The buyer who is informed that he is of a low type signals this to the seller in the first period but is `forced' to continue to convince the seller throughout the game. There is a sense in which, although, all the information is revealed immediately, complete separation has really not occurred until the game is fully over.

Observe that a condition of the theorem is that $\epsilon_L > 0$. If $\epsilon_L = 0$, there are other *pBe* satisfying R^T. For example, in the three-period game[13], a *pBe* exists with the following characteristics. In the first period, for any price, p_3, the low type demands a low enough quantity, q_3, that even if the high type demanded q_3, was offered $p_L = a_L/2$ for the rest of the game and the buyer was able to demand a_H-p_L in the last two periods, the high type still prefers to demand a_H-p_3 and reveal himself. In this equilibrium, the seller believes $\mu_2 = 0$ ($\mu_2 = 1$) if she sees the low (high) quantity in the first period and never changes it for the rest of the game. In the subsequent periods, she offers the full information static price, $a_L/2$ or $a_H/2$. The equilibrium is not eliminated by R^3. It may seem odd that if the seller sees first a low quantity and then the quantity a_H - p_2, she never wavers from her belief, $\mu = 0$. However, the demand $a_H - p_2 = q_2^o(h_3,p_2,a_H)$ and therefore, according to Definition 3, there is no restriction implied for how she should update.

This type of equilibrium does not result in the peculiar phenomenon of reaching a point in the game where it is common knowledge that the buyer is a low type and yet underdemanding persists and, perhaps, is attractive for that very reason. On the other hand, it implies a sort of dogmatic belief formation by the seller. She forms her belief in the first period of the game and never changes her mind thereafter. Not surprisingly, for very long games, it is much harder to support complete separation with this type of equilibria since the temptation for the high type to deviate in the first period and get a low price for many periods in spite of high demands is very strong. To support these dogmatic equilibria for games of arbitrarily many periods and for any $a_L > a_H/2$ requires a discount factor below 4/9 which is much lower than the 4/5 bound in the equilibrium characterized in Theorem 1 (see the comment following assumption **A1**.) Of course, there may also be other separating equilibria with less dogmatic behavior by the seller. For example she may change her mind only after observing some fixed number of deviations, but then some assessment would have to be made concerning what is a reasonable number of deviations before the seller should switch her beliefs.

No equilibrium of this type survives the application of the refinement with $\epsilon_L > 0$ because $\mu_2 > 0$ after any history leading to period 2 and, by Bayes' rule, if the seller observes a high demand in period 2 she must believe it came from a high type and respond with a high price in period 1. The definition of q_3 would then fail to satisfy the incentive compatibility constraint on the high type. Of course, the *pBe* described in Theorem 1 also survives in the limit as ϵ_L goes to zero.

5. Implications of Equilibrium

The equilibrium exhibits some intuitive comparative statics. As long as we continue to assume that **A1** holds, the low type's demand falls as δ rises. He must distort his demand even further the more important the future becomes. Similarly, as T becomes large, the more the low type must underdemand in early periods since longer games offer greater rewards to a high type who successfully mimics a low type. Finally, holding $\epsilon_H+\epsilon_L$ fixed, demand rises as ϵ_H-ϵ_L falls. The closer the ϵ's, the less valuable is current information, and therefore the less costly it is for the low type to separate from the high type.

The equilibrium characterized in Section 3 exhibits some additional noteworthy characteristics. A sort of ratchet effect is still present although in a different sense than in the nonlinear model. There, the principal is forced to offer a more generous scheme in order to induce information revelation. In the case of the linear contracting game, the informed agents will often reveal following *any* price offer of the seller. However, it is this behavior which forces the seller's price offer to be lower than in a static price setting problem. For any given belief of the seller, her optimal price is higher with the same belief as the game nears the final period.

Since, in equilibrium, information is completely revealed in each period, it is interesting to compare the results here with those of a similar model where the buyer has the same preferences but acts non-strategically. Both types can benefit from the strategic behavior. To see this, note that the price offer of the seller is typically lower in the strategic game. If the true state is high, the lower price is a straight gain to the buyer. When the true state is low, the lower price is a benefit even though the buyer is also forced to underdemand relative to his true demand curve. For $a_{Lt} > a_L/3$, a condition implied by **A1**, the low type buyer is made strictly better off by the lower price.

Consider the simple $T = 2$ game. In a game where there is no possibility of a high type, subgame perfection forces the buyer to choose $q = a_L - p$ in every period and, therefore, the equilibrium price path is just $p = a_L/2$ in each period. In the game with a small initial probability of a high type, the initial price is (close to) $a_{L2}/2 < a_L/2$ and then reverts to $a_L/2$ in the last period if the type is in fact low. Even with the lower quantity demanded in the first period, the low type does better than in the game with no possibility of a high type.[14] This result is noteworthy since it represents a situation in which informed types are in a position in which they are forced to separate but the separation can benefit both types. In a typical screening model, some types of the informed players are usually forced to incur screening costs to separate themselves from other types. Here the screening environment can bestow an advantage. In a static or non-strategic monopoly pricing game, a low type buyer would prefer it if, by committing to demand a lower quantity, he could convince the seller to offer a lower price. In general, the technology for such a commitment is lacking. Here, though, the low type has a credible concern that the seller will mistake him for a high type in the remainder of the game. The concern serves as a commitment device and allows him to induce a lower price from the seller to shift some of the surplus from the trade in his direction. This differs from results in standard two-stage signalling games because of the addition of at least one earlier stage where the actions of the uninformed agent (the seller's initial pricing stage) is affected by the later signalling concerns of the informed agent. In addition, the strategy space of the uninformed agent is rich enough to allow for strategy choices that both types strictly prefer to the strategy choices in the complete information game. For example, in Kreps and Wilson's (1982) multi-period signalling game, a version of the chainstore paradox, the uninformed entrant can only decide whether to enter or not. In that environment, never enter is the outcome in a complete information game with the strong incumbent and there are no other strategies which can benefit both types as in this game where both types prefer the lower price. As a result, the possibility for this beneficial commitment feature did not arise in their model.

6. Conclusion

The ability of a seller to extract information depends on her strategic power. The more powerful the seller, the more dangerous it is for an informed buyer to reveal his private information and in long games the ratchet effect shows us that the result is very little information transfer. If the strategic power is more equally shared, though, as in linear contracting games, the likelihood of information revelation rises dramatically. In this class of bilateral monopoly games, the signalling space is rich enough for informed buyers to separate in each period. However, even though this separation conveys a great deal of information, it does not relieve the players of the burden of separation in subsequent play of the game. Some transactions are not consummated even if the players are virtually certain that they should be. The disinterested partner shows his true colors in the first period and proves it over and over for the rest of the relationship by demanding less of his partner than he truly desires. Playing hard to get results in the persistence of an under-requited love.

Department of Economics
University of Western Ontario
London, Ontario   N6A 5C2
Canada

APPENDIX

Proof of Theorem 2: If $a_{Lt} = a_L$, then the proof is similar but simpler so I focus on the harder case with $a_{Lt} < a_L$. Let Σ_{t-1} be the subset of *pBe* such that buyer behavior satisfies $\mathbf{H1}_{t-1}$ from Theorem 1. Note that this stationary and separating buyer behavior implies that for $\sigma \in \Sigma_{t-1}$, after any history h_{i+1} resulting in a seller belief, μ_i, $i \leq t$-1, the seller's unique sequentially rational response must be $p_i(\mu_i) = (\mu_i a_H + (1-\mu_i)a_{Li})/2$. Since Σ_{t-1} satisfies C_{t-1}, $BR_{\Sigma}{}^i(\mu) = p_i(\mu)$. Also note that the set of all *pBe* equals Σ_1 and satisfies $\Sigma_1 = R_1(\Sigma_1)$ since any seller posterior μ_0 is irrelevant because the game has ended. By Theorem 1, $\Sigma_T = R_T(R_{T-1}(...R_1(\Sigma_1)..))$ is non-empty. If we can show that for $\Sigma_{t-1} = R_{t-1}(R_{t-2}(...R_1(\Sigma_1)))...)$, $\sigma \in R_t(\Sigma_{t-1})$ implies that buyer behavior according to σ satisfies $\mathbf{H1}_t$ then Theorem 2 follows by induction. Therefore, suppose that $\Sigma_{t-1} = R_{t-1}(R_{t-2}(...R_1(\Sigma_1)))$. Let $\sigma \in R_t(\Sigma_{t-1})$ and let p_t be any price offered after any history. Since $\mathbf{H1}_{t-1}$ fixes buyer behavior in the following periods, for any quantity, q, chosen in this period, and price p_{t-1} offered in the next period, the continuation utility of a buyer of type L in period t, given σ is

$$V(q,p;a_j,p_t) = (2(a_j - p_t) - q)q + \delta\{\epsilon_j(a_H - p)^2 + (1-\epsilon_j)((a_L - p)^2-(a_L-a_{Lt-1})^2)\}+K'_{jt}$$

$$= (a_j-p_t)^2-(a_j - p_t - q)^2 + \delta\{\epsilon_j(a_H - p)^2 + (1-\epsilon_j)(a_L - p)^2\}+K_{jt},$$

for constants K_{jt}, $j = L,H$. By $\mathbf{H1}_{t-1}$ and sequential rationality on the seller, p_{t-1} must lie between $p_{t-1}(\epsilon_L) = p_{Lt-1}$ and $p_{t-1}(\epsilon_H) = p_{Ht-1}$. The worst that can happen to the buyers in the next period is the highest of these prices, therefore a buyer who is a high type in period t has continuation utility that is bounded by what he would get if he chose the current period utility maximizing quantity, $a_H - p_t$ and received the highest price in the next period, p_{Ht-1}:

(5) $$V(q,p;a_H,p_t) \leq (a_H - p_t)^2 + \delta\{\epsilon_H(a_H - p_{Ht-1})^2 + (1-\epsilon_H)(a_L - p_{Ht-1})^2\}+K_{Ht}.$$

We use the following result concerning the continuation function in the proof.

Lemma 1: For any ϵ_H, ϵ_L, with $1 \geq \epsilon_H \geq \epsilon_L \geq 0$, if (q,p) satisfy (5), the indifference curves generated by V are concave and the slope of the H curve is greater than that of the L curve.

Proof: Let $\tilde{a}_H = \epsilon_H a_H + (1-\epsilon_H) a_L$ and $\tilde{a}_L = \epsilon_L a_H + (1-\epsilon_L) a_L$. In what follows I focus on the case, $q \leq a_L - p$. If $q > a_L - p_t$, then the indifference curves slope in opposite directions and the single crossing property follows in a similar but simpler manner. Observe that in (q,p) space, the line $p = \tilde{a}_H - a_H + p_t + q$ lies below the line $p = \tilde{a}_L - a_L + p_t + q$ and so $p \leq \tilde{a}_H - a_H + p_t + q$ implies $p \leq \tilde{a}_L - a_L + p_t + q$.

For $q \leq a_j - p_t$, the level sets of V are given by

$$\frac{dp}{dq}\bigg|_j = \frac{a_j - p_t - q}{\delta(\tilde{a}_j - p)}.$$

If (q,p) satisfy $p = \tilde{a}_j - a_j + p_t + q$, the indifference curve has a slope of $1/\delta \geq 1$. Since $p_{Ht\text{-}1} \leq \tilde{a}_H/2$, the point $((a_H - p_t), p_{Ht\text{-}1})$ lies strictly below the line $p = \tilde{a}_H - a_H + p_t + q$ which has a slope of one. An indifference curve which passes through $((a_H - p_t), p_{Ht\text{-}1})$ cannot cross that line, since at the point of crossing, it would have a slope that exceeds the slope of the line. Therefore the set defined by (5) lies below the line $p = \tilde{a}_H - a_H + p_t + q$. The following results rely on this fact.

Taking differences gives,

$$\frac{dp}{dq}\bigg|_L - \frac{dp}{dq}\bigg|_H = \frac{1}{\delta}\left(\frac{a_L - p_t - q}{\tilde{a}_L - p} - \frac{a_H - p_t - q}{\tilde{a}_H - p}\right).$$

Rearranging, gives

$$\frac{1}{\delta}\left(\frac{\tilde{a}_L - p + p - (\tilde{a}_L - a_L + p_t + q)}{\tilde{a}_L - p} - \frac{\tilde{a}_H - p + p - (\tilde{a}_H - a_H + p_t + q)}{\tilde{a}_H - p}\right).$$

Cancelling the one in each term and noting that $\tilde{a}_L$ - $a_L > 0 >$ $\tilde{a}_H$ - a_H we get

$$\frac{dp}{dq}_{|_L} - \frac{dp}{dq}_{|_H} \le \frac{1}{\delta}\left(\frac{p-(\tilde{a}_H - a_H + p_t + q)}{(\tilde{a}_L - p)} - \frac{p-(\tilde{a}_H - a_H + p_t + q)}{\tilde{a}_H - p}\right)$$
$$= \frac{1}{\delta}(p-(\tilde{a}_H - a_H + p_t + q)) * \left(\frac{1}{\tilde{a}_L - p} - \frac{1}{\tilde{a}_H - p}\right),$$

which is less than zero for $p \le \tilde{a}_H - a_H + p_t + q$ and $\epsilon_H \ge \epsilon_L$.

Differentiate along the indifference curve to get

$$\frac{d^2p}{dq^2}_{|_j} = \left(\left(\frac{a_j - p_t - q}{\delta(\tilde{a}_j - p)}\right)^2 - 1\right)/(a_j - p_t - q).$$

which is less than zero for $p \le \tilde{a}_j - a_j + p_t + q$. This yields the concavity of the indifference curves■

Let q_H be a quantity prescribed by σ for the high type in period t and let π be the next period seller price. Since $\epsilon_L > 0$, q_H must occur with positive probability and if H alone demands q_H, we must have $\pi = p_{Ht-1}$. In general, we must have $\pi > p_{Lt-1}$. If $\pi < p_{Ht-1}$, then the low type must also choose q_H with positive probability. Define $\bar{q}(q_H, \pi, p_t)$ by

$$V(\bar{q}, p_{Lt-1}; a_H, p_t) = (a_H - p_t)^2 - (a_H - p_t - q_H)^2 + \delta\{\epsilon_H(a_H - \pi)^2 + (1-\epsilon_H)(a_L - \pi)^2\} + K_{Ht}.$$

(In Figure 2, $\bar{q}$ corresponds to q'.) Substituting in the definitions also yields that if $q_H = a_H - p_t$ and $\pi = p_{Ht-1}$, $\bar{q} = q_{Lt} = q_{Lt} - p_t$ where q_{Lt} is defined in equation (4) (if $q_{Lt} < q_L$.) By assumption **A1**, $\bar{q} \ge q_{Lt}$ is strictly positive.

Suppose that a deviant offer $q = \bar{q} - \Delta$ is demanded. By Lemma 1 and assumption **A1**, there is an Δ small enough that $q > 0$ and the high type strictly prefers q_H to q and its consequent continuation to the deviation and the best possible continuation that he can hope for while the low type strictly prefers demanding q, receiving the subsequent price $p_{Lt\text{-}1} = BR_{\Sigma}^{t\text{-}1}(\epsilon_L)$ in the next period and abiding by the behavior described by $\mathbf{H1}_{t\text{-}1}$ subsequently. By the refinement R_t, then the seller must believe that a low type demanded q and sequential rationality along with the buyer behavior, $\mathbf{H1}_{t\text{-}1}$ requires her to respond with the price, $p_{Lt\text{-}1}$ in the next period. This breaks any pooling behavior in period t. A similar argument follows to show that any separating outcome must satisfy condition $\mathbf{H1}_t$. Therefore, $\Sigma_t = R_t(\Sigma_{t\text{-}1})$ and yields buyer behavior $\mathbf{H1}_t$ and induction yields the result. ■

FOOTNOTES

[1]E-mail -- vincent@sscl.uwo.ca. This paper has benefitted from conversations with Morton Kamien and Alejandro Manelli and the detailed comments of two referees.

[2]The denotation, period *t*, refers to the period in which *t* periods remain to the end of the game. Thus period 1 is the last period and period *T* is the first.

[3]With the exception of the treatment in Section 4, for most of the paper, I assume that once a buyer's type is chosen by Nature, it remains fixed for the remainder of the game.

[4]The case in which a seller offers non-linear contracts is examined in the two-period case in a slightly different model by Laffont and Tirole (1987).

[5]For a definition of perfect Bayesian equilibrium, see Fudenberg and Tirole (1991). This definition which more closely corresponds to that of sequential equilibrium is slightly more restrictive than that used in earlier applications such as Freixas, Guesnerie and Tirole (1985).

[6]Where it is clear, the τ subscript is dropped.

[7]I also set ϵ_L very small so some of the actual numbers are not precisely correct but are the limits as ϵ_L goes to 0.

[8]Point *A* represents a pooling outcome since any seller offer strictly between 3/8 and 1/2 in the last period can only be generated by a seller belief strictly between 0 and 1.

[9]For buyers, a history includes the realized price offers, demand choices AND the realization of buyer types to that period. For the seller, a history only includes the first two sequences.

[10]This definition does not require pure strategies but i) as long as **A1** is satisfied, the equilibrium will be in pure strategies, and ii) restriction to pure strategies requires less notation so I will refer only to the pure strategy case here.

[11]There is another somewhat technical restriction. The original definition of perfect Bayesian equilibrium (for example, Freixas, Guesnerie, Tirole 1985) placed relatively few restrictions following

out-of-equilibrium histories. Suppose that an out-of-equilibrium history occurs and the *pBe* assigns subsequent beliefs $\{\mu_i\}_{i=t-1}^{1}$ following it. Even if the continuation path follows the prescription of the *pBe* following the out-of-equilibrium history, the original definition did not force μ_i and μ_{i-1} to be consistent with Bayes' rule and the equilibrium strategies. However, in finite games this additional restriction would be implied by sequential equilibrium via the condition of consistency and it seems natural to require it here as well. This restriction requires that beliefs following an out-of-equilibrium move be what Fudenberg and Tirole (1991) term `reasonable' beliefs.

[12]More generally, one might prefer to consider how the continuation path for the rest of the game changes also with changes in the period *t* belief. However, since *pBe* fixes strategies and then appends beliefs, the machinery of *pBe* does not allow us to specify variations of continuation strategies following changes in period *t* beliefs. An alternative approach which could achieve this would be to use the concept of meta-strategy introduced in Grossman and Perry (1986).

[13]In a two-period game, the equilibrium path described in this paragraph and that of Theorem 1 coincide.

[14]Observe that since with $T = 2$, the "dogmatic" equilibrium discussed after Theorem 2 and the equilibrium described in Theorem 1 coincide so this result does not necessarily rely on the refinement.

REFERENCES

Beaudry, P. and M. Poitevin., "Signalling and Renegotiation in Contractual Relationships," *Econometrica* 61 (1993), 745-782.

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

Cho, In Koo and Joel Sobel, "Strategic Stability and Uniqueness in Signalling Games," *Journal of Economic Theory* 50 (1990), 381-413.

Freixas, Xavier, Guesnerie, Roger, and Jean Tirole, "Planning Under Incomplete Information and the Ratchet Effect," *Review of Economic Studies* LII (1985), 173-191.

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

Grossman, Sanford and Motty Perry, "Sequential Bargaining Under Asymmetric Information," *Journal of Economic Theory* 39 (1986), 120-154.

Hart, Oliver and Jean Tirole, "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies* LV(4) 184 (1988), 509-540.

Kreps, David M., and Robert Wilson, "Reputation and Imperfect Information," *Journal of Economic Theory* 27 (1982), 253-279.

Laffont, J. J. and Jean Tirole, "Comparative Statics of the Optimal Dynamic Incentive Contract," *European Economic Review* 31 (1987), 901-926.

Madrigal, V., T. Tan, and S. R. de Costa Werlang, "Support Restrictions and Sequential Equilibria," *Journal of Economic Theory* 43 (1987), 329-334.

Noldeke, G. and E. van Damme, "Switching Away From Probability One Beliefs," Discussion Paper A-304, University of Bonn, 1990.

Rubinstein, Ariel, "Perfect Equilibria in Bargaining Models," *Econometrica* 50 (1982), 97-109.

Vincent, Daniel R., "Repeated Signalling Games and Dynamic Trading Relationships," mimeo, Northwestern University, 1994.