



Signalling and Renegotiation in Contractual Relationships

Paul Beaudry; Michel Poitevin

Econometrica, Vol. 61, No. 4 (Jul., 1993), 745-782.

Stable URL:

<http://links.jstor.org/sici?sici=0012-9682%28199307%2961%3A4%3C745%3ASARICR%3E2.0.CO%3B2-O>

Econometrica is currently published by The Econometric Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/econosoc.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

SIGNALLING AND RENEGOTIATION IN CONTRACTUAL RELATIONSHIPS

BY PAUL BEAUDRY AND MICHEL POITEVIN¹

This paper examines how the possibility of renegotiation affects contractual outcomes in environments in which adverse selection is a problem. The game setup is an extension of the one-shot signalling game in which an infinite number of rounds of renegotiation are permitted before contracted actions are in fact executed. The main results of the paper are (1) executed contracts may still contain distortions although players can never commit not to renegotiate, (2) the popular “efficient” separating-equilibrium outcome of one-shot signalling games is never an equilibrium outcome when an infinite number of rounds of renegotiation are permitted, (3) standard incentive-compatibility constraints can be easily generalized to incorporate situations that allow for an infinite number of rounds of renegotiation, (4) equilibrium outcomes can be separating and nevertheless depend on the uninformed player’s priors as informed types pool in the first stage and use the renegotiation stages to separate, (5) renegotiation in signalling games may lead to outcomes similar to equilibrium outcomes of screening games in which multiple contract purchases are allowed.

KEYWORDS: Informed principal, contracts, signalling, renegotiation.

1. INTRODUCTION

SIGNALLING GAMES ARE OFTEN USED to analyze outcomes in environments with adverse selection. One of the areas in which these games have been especially useful is in the determination of contractual outcomes when a privately informed contractor faces a market of uninformed contractees.² The major insight of the signalling literature is that equilibrium behavior is generally characterized by the informed player revealing his information by the choice of a self-selecting distortion relative to the symmetric-information outcome; however, when this ex post inefficiency is simply written into a contract, the contracting parties have an interest in renegotiating the contract immediately after it is signed.

In this paper, we examine the determinants of contractual outcomes when the informed party cannot commit to abstain from proposing renegotiations. Contract renegotiation in signalling games, or more precisely in informed-principal settings, has been previously examined by Maskin and Tirole (1992). These authors focused on the particular issue of renegotiation of decision rules or menus of contracts, and impose that parties be committed to not renegotiate once an allocation within the menu has been chosen. In their framework equilibrium allocations are often characterized by ex post inefficiencies and thus

¹ We would like to thank Marcel Boyer, Dominique Desruelle, Martin Hellwig, Michel Moreaux, Roger Myerson, Michael Peters, Michael Riordan, Jacques Robert, Daniel Vincent, Andy Weiss, and two anonymous referees for many helpful comments. We would also like to acknowledge financial support from the Social Sciences and Humanities Research Council of Canada and the Fonds pour la formation de chercheurs et l’aide à la recherche du Québec.

² Among others, papers by Brennan and Kraus (1987), Giammarino and Lewis (1989), Heinkel (1982), Milde and Riley (1988), and Poitevin (1989, 1990) are all examples of contract signalling within the financial market.

there are further incentives for renegotiation.³ Therefore to complement their work, this paper examines the implications of allowing the informed party to propose renegotiations after a specific allocation has been agreed upon. The issue of renegotiating allocations is of particular interest because most signalling models consider situations in which the informed party offers a contract specifying a single allocation as opposed to a menu of allocations, and therefore only this later stage of renegotiation is potentially relevant.

Other papers on renegotiation that are closely related to the current paper are Hart and Tirole (1988), Laffont and Tirole (1990), and Dewatripont (1989).⁴ These papers are concerned with the implementation of long-term contracts when parties cannot commit to not renegotiate the contract as it unfolds. Our approach differs from these papers in essentially three features. First, the renegotiation process we consider includes a potentially infinite number of rounds.⁵ Secondly, renegotiation occurs before any part of the contract is actually executed, but after a precise allocation has been contracted. Thirdly, it is the informed party that initiates the renegotiation at each round. The first feature is adopted in order to rule out situations where the last stage of renegotiation is used by the players to commit to sorting distortions (which would simply displace our previous criticism to the last stage of the renegotiation process). The second feature is introduced to consider the effect of allowing agents to potentially renegotiate away ex post inefficiencies characterizing the allocations specified in a contract. The third feature reflects the idea that it is the informed agent which may have the bargaining power in many market situations, as is the case in many financial-market examples.

In order to present our analysis of contractual outcomes in the presence of renegotiation, we have structured the paper as follows. In Section 2, we describe the adverse-selection environment we are considering, and we motivate the need for studying renegotiation in such environment. Section 3 presents the game setup and gives the characterization of the Perfect Bayesian Equilibria (PBE). The game is essentially an infinitely repeated version of the one-shot signalling game except that each new signalling stage game takes as the status quo position the previously agreed upon contract. The set of PBE is characterized and compared to the set of PBE of the one-shot signalling game. Section 4 refines the set of PBE and characterizes the equilibria that satisfy an extended version of Banks and Sobel's (1987) notion of Divinity. Concluding remarks follow in Section 5. All proofs are relegated to Appendix A.

The main results of the paper are (1) executed contracts may still contain distortions although neither party can commit to never renegotiate a contract,

³ These different stages at which incentives for renegotiation arise are similar in spirit to those to which Holmström and Myerson (1983) referred as issues of durability versus posterior durability of decision rules.

⁴ Other papers on renegotiation include Hart and Moore (1988), Aghion, Dewatripont, and Rey (1989), Dewatripont (1988), Fudenberg and Tirole (1990), and Nosal (1991).

⁵ Hosios and Peters (1988) study a renegotiation game with a potentially infinite number of rounds in the case of a monopoly insurer. In their paper, the uninformed agent proposes the renegotiations.

but the popular “efficient” separating-equilibrium outcome is never an equilibrium when renegotiation is permitted, (2) standard incentive-compatibility constraints can be easily generalized to incorporate situations that allow for an infinite number of rounds of renegotiation, (3) equilibrium outcomes can be separating and nevertheless depend on priors as the informed types pool in the first stage and use the renegotiation stages to separate, (4) renegotiation in signalling games may lead to outcomes similar to equilibrium outcomes of screening games in which multiple contract purchases are allowed.

2. THE ENVIRONMENT AND THE MOTIVATION FOR RENEGOTIATION

Let us consider a situation with two agents, one informed agent and one uninformed agent. These two agents can contract on a vector of actions $m = (m_a, m_b)$ belonging to a set $M \subseteq \mathbf{R}_+^2$. We assume that the null contract $m = (0, 0)$ is an element of M . The preferences of these agents over the contract m are represented by $U(m, t)$ for the informed agent and $V(m, t)$ for the uninformed agent, where $t \in \{L, H\}$ represents the informed agent’s private information or type.⁶ The utility levels $U(0, t)$ and $V(0, H) = V(0, L) =: V(0)$ give the agents’ reservation utility. We will call the contract $m'(V(0))$, defined as the solution to

$$\left\{ \max_m U(m, t) \text{ subject to } V(m, t) \geq V(0) \text{ and } U(m, t) \geq U(0, t) \right\},$$

the symmetric-information contract. In general, $m'(\cdot)$ will be called type t ’s contract curve. We also make the following assumptions about agents’ preferences.

ASSUMPTION 1: $U(\cdot, t)$ and $V^P(\cdot, \mu) := \mu V(\cdot, H) + (1 - \mu)V(\cdot, L)$ are quasi-concave (one strictly), monotonic (one strictly), continuous and differentiable in m , for all $0 \leq \mu \leq 1$. Also, when $V(m, t)$ and $U(m, t)$ are both strictly monotonic in m we assume that $\text{sign } \nabla_m V(m, t') = -\text{sign } \nabla_m U(m, t)$ for $t, t' = L, H$.

ASSUMPTION 2: $V(m, H) > V(m, L)$ for all $m \in M \setminus (0, 0)$; $V(0, 0, H) = V(0, 0, L)$.

ASSUMPTION 3A:

$$-\frac{V_a(m, L)}{V_b(m, L)} \geq -\frac{V_a(m, H)}{V_b(m, h)} \quad \text{for all } m \in M.$$

⁶ The setup we consider is one of pure adverse selection; however it can be trivially generalized to include some types of moral-hazard problems. For example, the utility functions $U(\cdot)$ and $V(\cdot)$ can be considered as indirect utility functions where an unobserved action has been replaced by the optimal action as a function of the contract and the type, or similarly, the type may represent a past choice of an unobservable action as in Fudenberg and Tirole (1990).

ASSUMPTION 3B:

$$\begin{aligned} \text{Case RS: } & -\frac{U_a(m, H)}{U_b(m, H)} < -\frac{U_a(m, L)}{U_b(m, L)} \quad \text{for all } m \in M, \\ \text{Case S: } & -\frac{U_a(m, H)}{U_b(m, H)} > -\frac{U_a(m, L)}{U_b(m, L)} \quad \text{for all } m \in M. \end{aligned}$$

ASSUMPTION 4: $U(0, t) < U(m^L(V(0)), t) < U(m^H(V(0)), t)$ for $t = L, H$.

Assumption 1 ensures that, for $0 \leq \mu \leq 1$, there always exists a solution to $\{\max_m U(m, t) \text{ subject to } V^P(m, \mu) \geq \bar{V}\}$,⁷ and that, when both agents have strictly monotonic preferences in M they have opposite preferences over each element of the contract. Assumption 2 provides a simple ordering of types, that is, type H is the “good” type since uninformed agents would always prefer to contract with him than with the L type (for any given contract except the null contract).

The third assumption imposes the single-crossing property on both the informed and uninformed agents’ preferences. For the informed agent’s single-crossing property, Case RS represents situations where the marginal trade-offs are ranked differently for the informed and the uninformed player. This situation corresponds, for example, to the Rothschild and Stiglitz’s (1976) insurance model, where insurance is marginally less costly but also marginally less valuable for the high type. In contrast, Case S represents the situation where the ranking of the marginal trade-offs between types is similar for both the informed and uninformed players. This situation corresponds, for example, to Spence’s (1973) education model, where education is both marginally more productive and less costly for the high type. Note that the important difference between Cases S and RS is that in Case S the high type generally wants to separate by overinvesting in the signal compared to the symmetric-information outcome, while in Case RS the high type wants to separate by underinvesting in the signal. Assumption 3 is general enough to include the vast majority of signalling games discussed in the literature.⁸

Finally, Assumption 4 implies that there are potential gains from trade associated with the informed agent agreeing to a contract with an uninformed agent, that is, $m'(V(0)) \neq (0, 0)$, and that the first-best solution is not incentive-compatible.

In the absence of renegotiation, the analysis of contractual outcomes for this type of environment has proceeded mainly along two lines: screening games and signalling games. Both approaches assume that the informed agent can take advantage of the potential competition between uninformed agents and therefore represent attempts to characterize competitive outcomes in markets with

⁷ Note that $V^P(\cdot, 0) = V(\cdot, L)$ and that $V^P(\cdot, 1) = V(\cdot, H)$.

⁸ Notice that the elements of M can be arbitrarily renamed and therefore Assumption 3 also encompasses the cases where all signs are reversed.

adverse selection. The most common prediction for either one of these games is that if a pure-strategy equilibrium exists, it is characterized by type L negotiating his symmetric-information contract, and the high type negotiating a contract which is distorted relative to the symmetric-information case. Nevertheless, the type H informed agent receives the highest possible level of utility consistent with an incentive-compatible separation. We will denote this well-known “efficient” separating outcome by $\{m_L^s, m_H^s\}$.

The potential problem with considering this pair of contracts as a prediction of the contractual outcome in situations with adverse selection is that, conditional on having signed the contract m_H^s , there exist renegotiations that would improve the uninformed agent’s payoff regardless of his beliefs concerning the type of the informed agent that has signed the contract m_H^s , and would also improve the informed agent’s payoff. Therefore, unless Pareto-improving renegotiations can be ruled out, the type L agent will foresee the possibility of renegotiation and therefore will prefer initially m_H^s to m_L^s which causes the proposed equilibrium to break down. Consequently, the need to examine the case where renegotiations cannot be ruled out seems to be relevant.

3. THE RENEGOTIATION GAME AND THE EQUILIBRIUM CHARACTERIZATION

Although the need to extend adverse-selection models to include renegotiation may be quite obvious, the best way of modeling it is surely not. We have chosen to extend the signalling framework to the following sequence of moves in the hope of capturing the essential elements of an environment where an informed contractor has the bargaining power and neither party can commit to abstain from renegotiating the original contract. Moreover, we choose to examine the case where the informed player makes all the offers since we believe that understanding situations where the choice of contracts conveys information is highly relevant and is a first step towards understanding more complex environments with bilateral private information.⁹

The environment we consider can be related to the following situation. Suppose an informed seller faces an uninformed buyer. The seller and the buyer must sign a contract before the buyer can take delivery of the good. The seller’s private information affects both his valuation of a contract and that of the buyer’s. For example, a contract may specify a quantity and a price for the good to be delivered, while the seller’s private information may be about the quality of the good. The game is infinite and time is divided in subperiods of length Δ . The relationship evolves over time as follows.

(1) At $i = 0$, nature chooses the seller’s type.

(2) At $i = 1$, the seller announces a contract $m_1 \in M$.¹⁰ The buyer can then decide to accept the offer. If the offer is not accepted, the game ends. If the

⁹ Most previous papers on renegotiation examine the case where all offers are made by an uninformed agent and therefore contract offers do not convey information.

¹⁰ The main results of this paper are not changed if the seller is allowed to offer menus of contracts.

offer is accepted, the buyer must specify a "delivery date" at which the contract must be fulfilled. The minimum delivery delay is one period such that the delivery date σ_1 must be greater than 1, that is, $\sigma_1 > 1$.¹¹

(3) At $i = 2$, the seller can propose a renegotiation m_2 . If the buyer accepts it, the previous contract is void and the buyer must specify a new delivery date $\sigma_2 > 2$. If the renegotiation is refused by the buyer and $\sigma_1 = 2$, the buyer takes delivery of the good at the conditions specified by contract m_1 , and the game ends; if $\sigma_1 > 2$, the game continues.

(4) All other periods $i \geq 3$ of the game are similar to period 2. The seller can propose a renegotiation m_i which can be accepted or rejected by the buyer. A contract acceptance must be followed by the specification of a new delivery date $\sigma_i > i$ by the buyer. The game ends at period i if and only if the proposed renegotiation is rejected and the current delivery date is i . In this case, the last accepted contract is fulfilled.

The main elements of this game are (1) the seller makes both the initial contract offer and all following propositions of renegotiation, while the buyer accepts or refuses offers, (2) the buyer sets a delivery date for the good which always allows at least one round of renegotiation and therefore neither the buyer or the seller can commit not to renegotiate, and (3) the renegotiation process only ends when the delivery date comes due and the last proposed renegotiation is refused, and therefore the renegotiation process can potentially last an infinite number of rounds. This last element implies that the delivery date can be set such that many renegotiation rounds can take place following a rejected offer.¹²

It is natural to define payoffs for this game as the discounted utilities associated with the implemented contract where the discounting is related to the time used in the renegotiation process. The possibility of renegotiation implies that delay only occurs if delivery takes place after period 2 and therefore the number of periods of delay is given by $n = \sigma_i - 2$ where σ_i is the date at which the good was delivered. The type t seller's payoff associated with a contract m delivered with delay n is then $\exp\{-r_s \Delta n\}U(m, t)$, while the corresponding buyer's payoff is $\exp\{-r_b \Delta n\}V(m, t)$, where r_s and r_b are respectively the seller's and buyer's instantaneous discount rates, and Δn is the total length of delay.

Before discussing the equilibria of the game it is necessary to introduce "some" notation. Let us define a potential history realized before the uninformed or informed agent moves in the i th period by a vector $\mathcal{H}^V(i) = \{m_1, \sigma_1, m_2, \sigma_2, \dots, m_i\}$, or $\mathcal{H}^U(i) = \mathcal{H}^V(i-1) \cup \{\sigma_{i-1}\}$ respectively, where m_1 is the initial contract offer, m_k represents the renegotiation proposed at date k ,

¹¹ None of our results would be changed if the seller could offer contracts contingent on the delivery date. For ease of presentation, we have opted for an extensive form in which the buyer specifies the delivery date.

¹² We call our game a "signalling-cum-renegotiation" game and we refer to a "one-shot signalling game" as the game that corresponds only to the first two stages of our game. A "one-shot signalling game" corresponds to what Cho and Kreps (1987) call simply a "signalling game."

and σ_k the delivery date associated with any newly accepted contract in period k . At date $i = 1$, there is no relevant history for the informed agent, so $\mathcal{H}^U(1) = \emptyset$. The collection of possible histories is $\Xi^V(i)$ and $\Xi^U(i)$ for the uninformed and informed agent respectively. Given this notation for histories, we can define the agents' strategies and beliefs.¹³

A *strategy* Ω for the informed agent is a sequence of type-contingent contracts $\Omega' = \{\Omega'_1, \Omega'_2, \Omega'_3, \dots\}$ where Ω'_i is chosen conditionally on $\mathcal{H}^U(i)$ and $\Omega'_i \in M \cup \emptyset$ represents the initial contract offer, if $i = 1$, or the $(i - 1)$ th renegotiation offer if $i \geq 2$. We denote by $\Omega'_i = \emptyset$ the informed agent's option not to offer a renegotiation in period i . We assume, without loss of generality, that the informed agent must offer a contract in period 1, that is, $m_1 \neq \emptyset$.

A *strategy* σ for the uninformed agent is a sequence of delivery dates $\sigma = \{\sigma_1, \sigma_2, \dots\}$ where σ_i is chosen conditionally on $\mathcal{H}^V(i)$ and $\sigma_i \in \{0, i + 1, i + 2, \dots\}$. Upon acceptance of an offer, the uninformed agent must specify a delivery date, and therefore his acceptance decision can be represented by the choice of a delivery date. Hence, if the offer m_i is accepted, $\sigma_i > i$ represents the delivery date of contract m_i . We define the decision to reject offer m_i by $\sigma_i = 0$. If $m_i = \emptyset$, then $\sigma_i = 0$, that is, the uninformed agent cannot change the delivery date if no new contract is offered. Finally, we assume that the delivery date be finite.

The *beliefs* μ for the uninformed agent is a sequence $\mu = \{\mu_0, \mu_1, \dots\}$ where (1) $\mu_0 \in (0, 1)$ are the priors held by the uninformed agent at the beginning of the game about the probability that the informed agent is of type H ; (2) $\mu_i \in [0, 1]$ are the beliefs held after a history $\mathcal{H}^V(i) \in \Xi^V(i)$.

There are two other concepts that are useful to define before we discuss the equilibrium concept, that is, the notion of continuation strategies and an outcome function.

A *continuation strategy* $\sigma_+(i)$ or $\Omega'_+(i)$ is simply a truncated strategy. For example, $\sigma_+(i) = \{\sigma_i, \sigma_{i+1}, \sigma_{i+2}, \dots\}$.

The *contract-outcome function* f gives the executed contract for any compatible history and pair of continuation strategies: $f(\mathcal{H}^V(i), \sigma_+(i), \Omega'_+(i + 1)) \in M$ and $f(\mathcal{H}^U(i), \sigma_+(i), \Omega'_+(i)) \in M$. The executed contract corresponds to that described in the game setup for the given history and continuation strategies.

We are now in a position to discuss the equilibria of this game. A minimal requirement is to consider the Perfect Bayesian Equilibria (PBE) of the game. Recall that a set of strategies and beliefs $\{\Omega, \sigma, \mu\}$ forms a PBE if at every information set of the game (1) the strategies for the remainder of the game are Nash given beliefs, (2) beliefs are consistent with observed behavior and presumed strategies.¹⁴

In this general setup, self-selection of the different types of the informed player can occur through either his choices of contract proposals or through

¹³ Note that we restrict our attention to pure strategies. This assumption is discussed at the end of the section.

¹⁴ See Fudenberg and Tirole (1991) for a precise definition of a PBE.

delay. In the main body of the paper we choose to abstract from the signalling role of delay and to concentrate only on the signalling role of contract proposals. We choose such a focus for two main reasons. First, it is important to understand the precise effect of renegotiation of contracts that may include distortions as a consequence of self-selection. Second, in many negotiation situations, the cost of delay may be perceived by the players as being irrelevant (of second order), and therefore emphasizing only the signalling role of contract proposals may be warranted.¹⁵ Consequently, in what follows we examine the limit case where discount rates for both the seller and the buyer are set equal to zero, which de facto ensures that delay cannot be used as a signal. Nevertheless, to emphasize the relevance of the analysis, Appendix B shows how many of our results can be extended to the case of positive discount rates, and we discuss how our characterizations can be thought of as the limit of the game with positive discount rates in which discount rates approach zero. The only problem with the limit game is that it leaves undefined the payoffs for the case when the renegotiation process continues to infinity. In this case, we assume that both players receive their reservation utility as if no relationship had ever begun.¹⁶

Proposition 1 offers a characterization of the set of PBE outcomes for our signalling-cum-renegotiation game when discount rates are zero. An outcome of a PBE is a set of contracts $\{m_L, m_H\}$ and a set of delays $\{n_L, n_H\}$, where the implemented contract m_t is delivered in period $n_t + 2$ if the informed player's type is t . However, as argued above, when both discount rates are zero, delay plays no role in the agents' payoffs, and therefore Proposition 1 describes only restrictions imposed on contractual outcomes $\{m_L, m_H\}$.

PROPOSITION 1: *Any pair of contracts $\{m_L, m_H\}$ can be supported as an outcome of a Perfect Bayesian Equilibrium if and only if the contracts satisfy the following conditions:*

- (i) $\mu_0 V(m_H, H) + (1 - \mu_0) V(m_L, L) \geq V(0)$,
- (ii) $U(m_L, L) \geq \left\{ \max_m U(m, L) \text{ subject to } V(m, L) \geq V(0) \right\}$,
- (iii) $U(m_H, H) \geq \left\{ \max_m U(m, H) \text{ subject to } V(m, L) \geq V(0) \right\}$,
- (iv) $U(m_H, H) \geq \left\{ \begin{array}{l} \max_m U(m, H) \text{ subject to } V(m, H) \geq V(m_L, H) \\ V(m, L) \geq V(m_L, L) \end{array} \right\}$,
- (v) $U(m_L, L) \geq \left\{ \begin{array}{l} \max_m U(m, L) \text{ subject to } V(m, H) \geq V(m_H, H) \\ V(m, L) \geq V(m_H, L) \end{array} \right\}$,

¹⁵ Rubinstein (1991), among others, argues in favor of directly examining games with zero discounting when the per-period cost of delay is likely to be perceived by players to be negligible.

¹⁶ This is equivalent to assuming that any finite number of rounds of renegotiation is costless, but that an infinite number of rounds eliminates all gains from trade.

$$(vi) \quad U(m_H, H) \geq \left\{ \begin{array}{ll} \max_m U(m, H) & \text{subject to } V(m, H) \geq V(m_H, H) \\ & V(m, L) \geq V(m_H, L) \end{array} \right\},$$

$$(vii) \quad U(m_L, L) \geq \left\{ \begin{array}{ll} \max_m U(m, L) & \text{subject to } V(m, H) \geq V(m_L, H) \\ & V(m, L) \geq V(m_L, L) \end{array} \right\},$$

(viii) if $m_L \neq m_L^s$ and $m_L \neq m_H$, then for $t = L, H$,

$$U(m_t, t) \geq \left\{ \begin{array}{ll} \max_m U(m, t) & \text{subject to } V(m, H) \geq V(m_P, H) \\ & V(m, L) \geq V(m_P, L) \\ & V(m_P, H) = V(m_H, H) \\ & V(m_P, L) = V(m_L, L) \end{array} \right\}.$$

Many of the conditions defining PBE contractual outcomes of the signalling-cum-renegotiation game are identical to those which define the PBE outcomes of one-shot signalling games. In particular, condition (i) is the uninformed player's individual-rationality constraint, while conditions (ii) and (iii) reflect the constraint that the informed player can be no worse off in equilibrium than if the uninformed player thinks he is a type L for sure. However, the other conditions clearly reflect the addition of the renegotiation process.

The main factor to notice behind conditions (iv)–(vii) is that starting from any contract \hat{m} , the uninformed player will always eventually accept a renegotiation m_i , if

$$m_i \in \{m/V(m, L) > V(\hat{m}, L) \text{ and } V(m, H) > V(\hat{m}, H)\},$$

that is, if, regardless of his beliefs, the uninformed player cannot lose from accepting the renegotiation. For example, conditions (vi) and (vii) indicate that all executed contracts must be such that it is not in the interest of the informed player, starting from his equilibrium contract, to propose a renegotiation that would be accepted for sure. Similarly, conditions (iv) and (v) reflect the same argument that one type cannot increase its utility by proposing an assured renegotiation, but now starting from the other type's equilibrium contract. In fact, conditions (iv) and (v) represent a very simple and intuitive generalization of the standard incentive-compatibility constraints to a situation where the informed player cannot commit not to renegotiate a contract. We call these constraints "renegotiation-induced incentive-compatibility constraints".¹⁷

Condition (viii) is an attainability condition on all separating outcomes that involve subsidization across types, which is the case when $m_L \neq m_L^s$ and $m_L \neq$

¹⁷ Conditions (iv)–(vii) would also represent necessary conditions on contractual outcomes for hidden-information models with renegotiation, that is, situations where the private information is known after the initial contract is signed but with the same type of renegotiation possibilities after the information is revealed.


m_H . In this case, separation can only be achieved after a pooling stage. Therefore there must exist a pooling contract, m_p , from which: first, the uninformed player would accept, starting from this contract, the separating renegotiation necessary to eventually achieve the executed contracts; secondly, also starting from this pooling contract, the informed player must be no better off by offering a “surely-acceptable” renegotiation than renegotiating towards the proposed equilibrium contract.¹⁸

The equilibrium path for the different outcomes can be constructed with all equilibrium contracts being attained by an initial step of pooling followed by separation. Consider any equilibrium outcome $\{m_L, m_H\}$ and consider the pooling contract m_p defined by $V(m_p, L) = V(m_L, L)$ and $V(m_p, H) = V(m_H, H)$. The informed player first makes the offer m_p regardless of his type. The uninformed player accepts this contract and sets the delivery date for next period. The informed player then proposes a renegotiated contract m_L if his type is L , and m_H if his type is H . Again these contracts are accepted with delivery required for period 3. Finally, each type of informed player does not make a new offer, and the game ends with the execution of the contract m_L or m_H . Off the equilibrium path, the uninformed player accepts renegotiations only if they are weakly preferred to the outstanding contract regardless of the informed player's type. This last condition states that the only accepted renegotiations m_i , starting from a contract \hat{m} , must belong to the set $\{m / V(m, L) \geq V(\hat{m}, L) \text{ and } V(m, H) \geq V(\hat{m}, H)\}$.

The construction of these equilibrium paths may lead one to believe that all stages of renegotiation beyond the first are redundant, and therefore a simpler game with only one stage of potential renegotiation (as opposed to an infinite number) would deliver the same characterization; however, this reasoning is very misleading: in the case where only one round of renegotiation is allowed, all the separating equilibria of the one-shot signalling game can be attained by initially offering the null contract, followed by the appropriate renegotiation proposal. In contrast, Proposition 1 implies that many equilibria of the one-shot signalling game cannot be supported when further renegotiations are possible. For example, as suggested by Weiss (1983) and Admati and Perry (1987) and as stated in Corollary 1, the “efficient” separating-equilibrium outcome can never be supported as a PBE of our signalling-cum-renegotiation game.¹⁹

¹⁸ These conditions can easily be generalized for an arbitrary finite number of types. For example, suppose we have a finite set of types $\{L, \dots, H\}$ which all satisfy a suitable generalization of Assumptions 1–4. Then, a condition such as condition (iv) would have to be satisfied for all types and all equilibrium contracts, with the same two constraints on the right-hand side of the inequality representing the set of surely-acceptable renegotiations. The addition of these new constraints would then define the set of necessary and sufficient conditions for the PBE outcomes of the game with multiple types.

¹⁹ Gale and Stiglitz (1989) examine a signalling game with only one round of renegotiation. Although they focus on the Pareto-optimal equilibrium outcome, the standard “efficient” separating-equilibrium outcome is an equilibrium outcome of their game whenever there is no information exogenously revealed between contract offers.



≡ Set of PBE outcomes for type H of the one-shot signalling game
Set of PBE outcomes for type H of the signalling-cum-renegotiation game

FIGURE 1.—Comparison of equilibrium sets.

COROLLARY 1: The contractual outcome²⁰ $\{m_L^s, m_H^s\}$, defined by $m_L^s = m^L(V(0))$, $U(m_L^s, L) = U(m_H^s, L)$, $U(m_L^s, H) \leq U(m_H^s, H)$ and $V(m_H^s, H) = V(0)$, can never be supported as a PBE of the signalling-cum-renegotiation game.

Despite Corollary 1, some of the separating-equilibrium outcomes of the one-shot signalling game can still be supported. In order to illustrate the restrictions implied by the introduction of renegotiation, Figure 1 illustrates the set of PBE outcomes for type H for the one-shot signalling game (Case S of Assumption 3B) and for the signalling-cum-renegotiation game when $m_L = m_L^s$. The effects of conditions (v) and (vi) can be clearly seen. Condition (v) limits the extent to which the low type would prefer mimicking the high type than

²⁰ This contract corresponds to the “efficient” separating-equilibrium outcome in one-shot signalling games.

separating, given the assured-renegotiation possibilities.²¹ The effect of condition (vi) is to delimit a set of contractual outcomes which are immune from renegotiations from the high type.

Figure 1 may give the impression that the introduction of renegotiation into contractual one-shot signalling games strictly reduces the set of attainable outcomes. Although renegotiation limits the set of attainable fully separating (equilibria where separation is achieved at the initial move) or fully pooling outcomes, it does permit the informed player to attain, through a history of initial pooling followed by separating renegotiations, a whole set of outcomes not attainable in one-shot signalling games. Furthermore, some of these outcomes may dominate *ex ante*, at least from the informed player's perspective, all PBE outcomes of one-shot signalling games. This possibility may at first glance appear odd, since the removal of a power to commit *ex post* usually hinders a player *ex ante*;²² however, in one-shot signalling games, the set of attainable outcomes does not generally contain the set of constrained Pareto optima.²³ Therefore, because informed players may, through renegotiation, get closer to some constrained Pareto optimum, they may prefer environments in which they cannot commit to abstain from renegotiating.

One implication of Corollary 1 is that the standard separating-equilibrium outcome of a Spence-type education game cannot be supported when renegotiation is possible and education is contracted simultaneously with the wage. This result is in direct contrast with a recent result by Noldeke and van Damme (1990a), and therefore a comparison is in order. In response to the critique about the efficient separating equilibrium put forth by Weiss (1983) and Admati and Perry (1987), Noldeke and van Damme looked at non-Markovian mixed strategy equilibria of a Spence-type education game where employment offers by uninformed players can be made throughout the education process. They show that, for small discount rates, the efficient separating-equilibrium outcome of the static Spence game is approximately the same as the unique equilibrium outcome of the dynamic version of the game which satisfies the Never-a-Weak-Best-Response criterion. One may think of this extended game as virtually allowing employers to "renegotiate" the education-contingent employment offers of the static game, once they have actually observed the attainment of different levels of education; however, this notion of renegotiation refers to the

²¹ In particular, the "Miyazaki" outcome (Miyazaki (1985)), defined as the solution to

$$\left\{ \begin{array}{ll} \max_{m_L, m_H} U(m_H, H) & \text{subject to } \mu V(m_H, H) + (1 - \mu)V(m_L, L) \geq V(0) \\ & U(m_H, H) \geq U(m_L, H) \\ & U(m_L, L) \geq U(m_H, L) \end{array} \right\}$$

is never a PBE outcome of our game since it never satisfies condition (v).

²² This negative effect of limiting commitment has been generally confirmed in the renegotiation literature. In fact, Dewatripont and Maskin (1990) consider this to be one of the important insights of the renegotiation literature.

²³ In the case where the one-shot signalling game allows a menu of contracts to be offered by the informed player, constrained Pareto optima are achievable, although they do not satisfy the Intuitive Criterion (Maskin and Tirole (1992)).

renegotiation of *implicit* contracts, that is, if one views the static Spence game as a contract game where a worker and an employer initially agree on an education-contingent wage contract, the result of Noldeke and van Damme implies that the separating-equilibrium contract is self-enforcing (in an ex-post competitive environment). In comparison, the result of this paper refers to the renegotiation of *explicit* contracts. The difference between the renegotiation of implicit versus explicit contracts is the fallback position when a renegotiation is refused. In the case of explicit contracts, if a renegotiation is refused, the worst outcome for any player is that the initial contract is eventually implemented. In the case of implicit contracts, if a renegotiation is refused, this might lead to a subgame where a player may be worse off than in the initial contract. It is this type of off-the-equilibrium-path punishment that allows the efficient separating-equilibrium outcome to be supported in Noldeke and van Damme's analysis. In our analysis, the type of punishment that can be imposed on a low type who tries to mimic a high type, and renegotiate, is limited by the explicit contract, and therefore it is impossible to support the efficient separating outcome.

We now turn to a brief discussion of our assumption to rule out stochastic contract offers. Previous papers on renegotiation, in particular Laffont and Tirole (1990), have emphasized the importance of mixed strategies on the part of the informed player as a means of hiding information. It therefore seems important to examine the extent to which Proposition 1 would need to be modified if mixed strategies on the part of the informed player were allowed. In fact, allowing the informed agent to make stochastic contract offers would not enlarge the set of PBE outcomes characterized by Proposition 1 beyond mixes over elements which satisfy constraints (ii)–(viii) and satisfy constraint (i) in expectations. To see this, suppose it was not true. This would imply that one of the seven conditions would not be satisfied by one possible realization of a PBE outcome. First, conditions (ii)–(iii) must clearly be satisfied by any realization since otherwise it would involve a dominated strategy; secondly condition (vii) must also be satisfied, otherwise the equilibrium outcome would not be attainable even with stochastic offers. So, suppose one of the remaining conditions (iv)–(vi) is not satisfied by a contract \tilde{m} reached with positive probability in equilibrium. This means that there exists a type that could increase his utility by mimicking another type's (or his own) strategy with probability one until \tilde{m} is signed, and then offering a surely-acceptable renegotiation that would make him better off compared to his equilibrium contract. So \tilde{m} could not be an equilibrium outcome. Therefore, since each possible realization of a PBE with stochastic contract offers must satisfy constraints (iv)–(vi), the standard efficient separating-equilibrium outcome could not even be an equilibrium outcome when stochastic contract offers are allowed. The reason mixed strategies on the part of the informed player do not seem very important in our game, in relation to the game analyzed by Laffont and Tirole (1990), is that in our game the informed player is making the offers and therefore can hide information by offering a pooling contract. In Laffont and Tirole, the informed agent only accepts or rejects offers made by the uninformed agent and therefore informa-

tion can be hidden only by randomized acceptance decisions by the informed agent.²⁴

It would be appropriate to verify the existence of PBE for our game; however, existence will be proven for a subset of the PBE in the following section, and therefore we do not prove it here.

4. EQUILIBRIUM REFINEMENT

In the previous section we characterized the set of PBE outcomes of a signalling-cum-renegotiation contract game and we compared this set with the set of PBE outcomes of a one-shot signalling contract game. This comparison provided some insights about the effects of renegotiation in markets with adverse selection; however, given the multiplicity of equilibria, our characterization has only little positive content. In fact, many equilibria were supported by out-of-equilibrium beliefs that were such that only surely-acceptable renegotiations were accepted by the uninformed agent. In many cases such beliefs imply that a renegotiation directly benefiting one type was rejected on the beliefs that it was offered by the other type. To rule out such beliefs, we now discuss the possibility of refining the equilibrium concept.

A common way to reduce the set of equilibria in one-shot signalling games is to place restrictions on out-of-equilibrium beliefs such as those suggested by the Intuitive Criterion (Cho and Kreps (1987), Cho (1987)), Divinity (Banks and Sobel (1987)) or D1 (Cho and Kreps (1987)). These three criteria are based on the examination of a subset of possible Best Responses of the uninformed player that would make the informed player better off following a deviation. We proceed along the same lines by defining a general notion of a Best Response that incorporates the fact that the deviating player will be called again to play following his deviation. We follow Cho (1987) in defining the set of Best Responses.

A response $\tilde{\sigma}_i$ belongs to the *set of Best Responses* $BR(\mathcal{H}^V(i))$ following a history $\mathcal{H}^V(i)$ if there exist continuation strategies $\tilde{\Omega}'_+(i+1)$, for $t = L, H$, $\tilde{\sigma}_+(i+1)$ for the uninformed agent, and beliefs $0 \leq \tilde{\mu} \leq 1$ such that

$$\begin{aligned} & \tilde{\mu}V\left(f\left(\mathcal{H}^V(i), \tilde{\sigma}_i, \tilde{\sigma}_+(i+1), \tilde{\Omega}^H_+(i+1)\right), \dot{H}\right) \\ & + (1 - \tilde{\mu})V\left(f\left(\mathcal{H}^V(i), \tilde{\sigma}_i, \tilde{\sigma}_+(i+1), \tilde{\Omega}^L_+(i+1)\right), L\right) \\ & \geq \tilde{\mu}V\left(f\left(\mathcal{H}^V(i), \sigma'_i, \tilde{\sigma}_+(i+1), \tilde{\Omega}^H_+(i+1)\right), H\right) \\ & + (1 - \tilde{\mu})V\left(f\left(\mathcal{H}^V(i), \sigma_i, \tilde{\sigma}_+(i+1), \tilde{\Omega}^L_+(i+1)\right), L\right) \end{aligned}$$

for all σ_i .

²⁴ Bolton (1990) emphasizes that when it is the uninformed party that proposes the renegotiations, the speed of information revelation is one of the most important elements in the characterization of equilibrium outcomes. In contrast, when it is the informed agent that proposes the renegotiations, the speed of information revelation is not necessarily important.

A response is a Best Response if there exist beliefs and continuation strategies for all players (including the uninformed player himself) such that the uninformed player does at least as well playing the Best Response as any other available strategy. This definition does not put any restriction on continuation strategies of any player.

Out of these Best Responses we must now define the notion of Preferred Responses for a given type of the informed agent, following an out-of-equilibrium contract offer. The set of Preferred Responses should include all “plausible” continuation strategies following the out-of-equilibrium contract offer. We also adopt Cho’s (1987) definition of Preferred Responses. Call $\Pi_{k=i}^{\infty} BR(\mathcal{H}^V(k))$ the set of all continuation strategies for the uninformed player consisting of possible Best Responses at all information sets following $\mathcal{H}^V(i)$. This allows us to define the sets of Preferred and Strictly Preferred Responses.

The set of Preferred and Strictly Preferred Responses of a given type t are

$$\begin{aligned} \Lambda(t/(\mathcal{H}^V(i))) \\ &:= \left\{ \tilde{\sigma}_+(i) \in \prod_{k=i}^{\infty} BR(\mathcal{H}^V(k)) \middle/ \right. \\ &\quad \left. \max_{\Omega'_+(i+1)} U(f(\mathcal{H}^V(i), \tilde{\sigma}_+(i), \Omega'_+(i+1)), t) \geq U^*(t) \right\}, \\ \Lambda^+(t/(\mathcal{H}^V(i))) \\ &:= \left\{ \tilde{\sigma}_+(i) \in \prod_{k=i}^{\infty} BR(\mathcal{H}^V(k)) \middle/ \right. \\ &\quad \left. \max_{\Omega'_+(i+1)} U(f(\mathcal{H}^V(i), \tilde{\sigma}_+(i), \Omega'_+(i+1)), t) > U^*(t) \right\}, \end{aligned}$$

where $U^*(t)$ is type t ’s equilibrium payoff.

The set $\Lambda(t/\mathcal{H})(\Lambda^+(t/\mathcal{H}))$ is the set of all sequences of Best Responses by the uninformed agent following history \mathcal{H} that can make type t weakly (strictly) better off relative to his equilibrium payoff.

Given this definition of Preferred Responses, it is now quite straightforward to extend any of the previous refinement criteria defined for one-shot signalling games to multistage games. We choose to work with the Divinity Criterion for two reasons. First, given that the set of Best Responses is potentially very large in our game, an extended version of the Intuitive Criterion, such as the one proposed by Cho (1987), does not have much bite. Secondly, a stronger criterion such as D1 has much less intuitive appeal than Divinity. The following definition extends the Divinity criterion in what we believe to be the most direct fashion.

Beliefs satisfy the *Extended-Divinity (XD) Criterion* if, following an equilibrium-path history $\mathcal{H}^U(i)$ and an out-of-equilibrium contract offer \tilde{m}_i , the

following two conditions hold, where $\mathcal{H}^V(i) = \mathcal{H}^U(i) \cup \{\tilde{m}_i\}$:

(1) $\Lambda^+(t/\mathcal{H}^V(i)) \subseteq \Lambda^+(t'/\mathcal{H}^V(i))$ implies that $\mu_i \leq \mu_0$ if $t' = L$ and $\mu_i \geq \mu_0$ if $t' = H$;

(2) if $\Lambda(t/\mathcal{H}^V(i))$ is empty and $\Lambda^+(t'/\mathcal{H}^V(i))$ is nonempty, then $\mu_i = 0$ if $t' = L$ and $\mu_i = 1$ if $t' = H$.

The first part of the XD Criterion says that if, following a first deviation from the equilibrium path, the set of Strictly Preferred Responses of type t is included in the set of Strictly Preferred Responses of type t' , then the uninformed player should put posterior beliefs no higher than the priors on type t if this deviation is observed.²⁵ This criterion therefore leaves unrestricted beliefs at information sets reached after more than one deviation from the equilibrium path. The second part of the criterion says that posterior beliefs should put no weight on a type that has no incentive to deviate regardless of the uninformed agent's sequence of Best Responses following the deviation. That part captures the idea behind the Intuitive Criterion. Finally, we can define the equilibrium concept we will use to restrict the set of PBE in the analysis of this game.

An *XD-equilibrium* for the game is a Perfect Bayesian Equilibrium with beliefs that satisfy the Extended-Divinity Criterion.

We will now use that definition to characterize all PBE outcomes that satisfy the XD Criterion. Proposition 2 introduces three necessary conditions for a PBE outcome to be an XD-equilibrium outcome. Subsections 4.1 and 4.2 show that, in Cases RS and S respectively, these conditions are also sufficient to characterize all XD-equilibrium outcomes.

PROPOSITION 2: *Suppose that $m^* = \{m_L^*, m_H^*\}$ is the contractual outcome of an XD-equilibrium; then the following three conditions must hold:*

$$\begin{aligned}
 \text{(iii')} \quad & U(m_H^*, H) \geq \left\{ \max_m U(m, H) \quad \text{subject to} \quad V^P(m, \mu_0) \geq V(0) \right\}, \\
 \text{(iv')} \quad & U(m_H^*, H) \geq \left\{ \max_m U(m, H) \quad \text{subject to} \quad \begin{aligned} & V(m, H) \geq V(m_L^*, H) \\ & V^P(m, \mu_0) \geq V^P(m_L^*, \mu_0) \end{aligned} \right\}, \\
 \text{(v')} \quad & U(m_L^*, L) \geq \left\{ \max_m U(m, L) \quad \text{subject to} \quad \begin{aligned} & V(m, L) \geq V(m_H^*, L) \\ & V^P(m, \mu_0) \geq V^P(m_H^*, \mu_0) \end{aligned} \right\}.
 \end{aligned}$$

²⁵ We believe that using initial priors to update beliefs following any deviation is warranted in a multistage game even though the deviation may occur late in the game. In the refinement literature (for example, Cho (1987)), the benefit of a deviation for a given type is assessed by comparing its potential payoff relative to the equilibrium payoff, that is, a deviation is not considered to be a tremble but is rather considered as a deliberate move by the informed player. Therefore, a deviation must be interpreted as a move that was thought out from the beginning of the game but only observed in period i . In this case, it is more reasonable to use the initial priors as the basis of updating since the deviation was really planned at the initial stage. This argument is akin to the reasoning in favor of allowing for switches in the support of beliefs.

Proposition 2 presents the implications of the XD criterion in a manner that highlights its relationship with Proposition 1. The three conditions (iii'), (iv'), and (v') are obvious counterparts to conditions (iii), (iv), and (v). The difference between each of these conditions is only within the constraint set of each maximization problem.

Conditions (iv) and (v) were based on the notion that some renegotiations were always accepted regardless of the uninformed agent's beliefs and that only those renegotiations could be considered as assured renegotiations. Conditions (iv') and (v') are very similar with the exception that the XD criterion enlarges the set of assured renegotiations starting from the alternative type's equilibrium contract. As a result, the renegotiation-induced incentive-compatibility constraints become more stringent. For example, if (v') is not satisfied, the type L agent can play type H 's equilibrium strategy until m_H^* is reached and then offer a renegotiation which, if accepted, is always profitable to him. Furthermore, this renegotiation can be chosen such that for any belief in $[0, \mu_0]$, the uninformed agent should accept it. Since type L always gains from the deviation, the belief imposed by the XD Criterion on that deviation should put no more weight than μ_0 on the type H having played it. Hence the renegotiation should be accepted, thus upsetting the equilibrium. This implies that (v') must be satisfied in any XD-equilibrium.²⁶ A second implication of the XD criterion is given by condition (iii'). This condition reflects the fact that the set of initial contract offers that will be accepted for sure by the uninformed agent is also enlarged by the XD criterion.

Conditions (iii'), (iv'), and (v') are stronger than their counterpart of Proposition 1 and therefore the set of XD-equilibrium outcomes is necessarily smaller than the set of PBE outcomes. In the next two subsections we show that these three conditions, in conjunction with the other conditions of Proposition 1, are also sufficient to characterize all XD-equilibrium outcomes. Since the precise characterization of the XD-equilibrium outcomes of the game depends on whether we are in Case RS or S of Assumption 3B, the sufficiency part of the proof is separated for each case.

4.1. Case RS

In this subsection we characterize XD-equilibrium outcomes for the case in which single-crossing properties of the informed and uninformed player have the same sign, that is,

$$-\frac{U_a(m, L)}{U_b(m, L)} > -\frac{U_a(m, H)}{U_b(m, H)} \quad \text{and} \quad -\frac{V_a(m, L)}{V_b(m, L)} \geq -\frac{V_a(m, H)}{V_b(m, H)} \quad \text{for all } m \in M.$$

This corresponds to Case RS in Assumption 3B.

²⁶ A similar argument for type H shows that (iv') is also necessary for a PBE outcome to satisfy the XD Criterion.

PROPOSITION 3: *In Case RS, the three necessary conditions of Proposition 2 are sufficient to characterize the unique XD-equilibrium contractual outcome $\{m_L^{RS}, m_H^{RS}\}$. This outcome has the following properties:*

- (1) $V^P(m_H^{RS}, \mu_0) = V(0)$,
- (2) $m_H^{RS} = \arg \left\{ \max_m U(m, H) \quad \text{subject to} \quad V^P(m, \mu_0) \geq V(0) \right\}$,
- (3) $m_L^{RS} = \arg \left\{ \max_m U(m, L) \quad \text{subject to} \quad \begin{array}{l} V(m, L) \geq V(m_H^{RS}, L) \\ V^P(m, \mu_0) \geq V^P(m_H^{RS}, \mu_0) \end{array} \right\}$.

The first thing to note from Proposition 3 is that, in the RS case, the XD Criterion is strong enough to yield a unique equilibrium contractual outcome, and this unique outcome depends on priors. The equilibrium outcome for Case RS can be constructed in the following manner: the high type's contract is first located at the tangency of his utility function and the zero-rent pooling line, and then the low type's contract is determined by the renegotiation-induced incentive-compatibility constraint (v').

We now briefly describe one equilibrium path supporting this outcome, referring the reader to Appendix A for a complete description of the out-of-equilibrium play. Along the equilibrium path, each type offers m_H^{RS} in the first stage of the game which is accepted with delivery being set for next period. In the second stage, the high type does not make any offer, therefore ending the game with the fulfillment of contract m_H^{RS} . The low type offers his equilibrium contract m_L^{RS} which is accepted with delivery being set for next period. In the third stage, the low type does not make any offer, thus imposing the fulfillment of contract m_L^{RS} .

The following intuition is suggestive of why renegotiation leads to this particular characterization. In the RS case, the high type wants to signal by underinvesting in the signal; however, it is quite easy for the low type to mimic such a move. The low type knows that if he mimics the high type he will eventually be able to renegotiate and thereby invest more in the signal. This problem of separation is in fact so severe that the high type cannot find any means of separation and hence both types offer the same initial contract; however, once the low type has benefited from cross-subsidization, he can decide whether or not any additional investment in the signal that is expected to be accepted by the informed agent is beneficial, and if so he renegotiates.

The precise circumstances when the low type decides to renegotiate, that is, whether the final outcome is separating or pooling, depends on the relative magnitude of the single-crossing property for the informed and uninformed agents' preferences and the priors μ_0 . This can be shown using the properties of




FIGURE 2.—Case RS with a separating equilibrium.

Proposition 3. The second property indicates that type H 's contract is at the tangency of its utility function and the zero-profit pooling line. The third property characterizes type L 's contract. The relative slope of type L 's indifference curve and the uninformed agent's isoprofit curve when evaluated at the contract m_H^{RS} determines the configuration of the equilibrium. If the slope of the informed agent's indifference curve is greater than that of the uninformed agent's isoprofit curve, type L wants to invest further in the signal at the marginal cost related to its type. This additional investment would not be profitable to type H , and therefore type L can propose a surely-acceptable renegotiation. Consequently the equilibrium outcome is separating. If it is smaller, type L wants to disinvest from the signal at the marginal cost related to its type; however this would also be profitable to type H which would then mimic type L to profit from the high "selling price." Therefore type L cannot propose an acceptable renegotiation from the point of view of the uninformed agent and hence the equilibrium outcome is pooling. These two cases are illustrated in Figures 2 and 3 respectively.

In a somewhat different setting, and for the parameterization corresponding exactly to the Rothschild and Stiglitz's (1976) insurance game,²⁷ the equilibrium configuration identified in Proposition 3 was first singled out by Jaynes (1978) who studied competitive insurance markets in which information sharing between insurance companies about consumer purchases is derived endogenously

²⁷ The exact parameterization of a Rothschild and Stiglitz (1976) insurance game is a special case of our RS configuration when contract curves are independent of type.




FIGURE 3.—Case RS with a pooling equilibrium.

(see also Hellwig (1988)). Gale (1991),²⁸ within the context of the securities market, has provided an alternative interpretation or foundation to the “Jaynes” outcome. He looks at a model of sequential competitive markets à la Bertrand in which offers are always made by uninformed agents and in each period, only one contract can be purchased. He shows that the “Jaynes” outcome arises as the equilibrium outcome of his game. In this respect, the work of Gale is closely related to the analysis of this paper since it highlights some of the implications of sequential trading of contracts, which in fact is very similar to renegotiation.²⁹

4.2. Case S

We now briefly characterize XD-equilibrium outcomes for case S. In case S, the single-crossing properties of the informed and uninformed agents go in opposite directions. Formally this is written as

$$-\frac{U_a(m, L)}{U_b(m, L)} < -\frac{U_a(m, H)}{U_b(m, H)} \quad \text{and} \quad -\frac{V_a(m, L)}{V_b(m, L)} \geq -\frac{V_a(m, H)}{V_b(m, H)} \\ \text{for all } m \in M.$$

The Spence signalling game corresponds to this configuration.

²⁸ Gale (1991) builds on Gale (1986).

²⁹ Beaudry and Poitevin (1990) study a similar model to that of Gale (1991) in which the entrepreneur can issue sequentially multiple securities on his project.

In case S the three necessary conditions of Proposition 2 are also sufficient to characterize all the XD-equilibrium outcomes. We do not have a uniqueness result as in the RS configuration, but all XD-equilibrium outcomes have the same characterization except that they differ in the level of rents left to the uninformed agent.

PROPOSITION 4: *In Case S, the three necessary conditions of Proposition 2 are sufficient to characterize all XD-equilibrium contractual outcomes. These outcomes $\{m_H^S, m_L^S\}$ have the following properties.*³⁰

- $$\begin{aligned}
 (1) \quad & V^P(m_L^S, \mu_0) = V^P(m_H^S, \mu_0), \\
 (2) \quad & m_t^S = \arg \left\{ \max_m U(m, t) \quad \text{subject to} \quad \begin{aligned} & V(m, t) \geq V(m_{t'}^S, t) \\ & V^P(m, \mu_0) \geq V^P(m_{t'}^S, \mu_0) \end{aligned} \right\} \\
 & \text{for } t, t' = L, H \text{ and } t \neq t', \\
 (3) \quad & V^P(m_H^S, \mu_0) \leq V(0), \\
 (4) \quad & \mu_0 V(m_H^S, H) + (1 - \mu_0) V(m_L^S, L) \geq V(0).
 \end{aligned}$$

Conditions 1 and 2 of Proposition 4 imply that, in Case S, all equilibrium outcomes are separating and depend on priors. They are characterized by the fact that both types' renegotiation-induced incentive-compatibility constraints are binding. This is in contrast to one-shot signalling games in which only one type's incentive-compatibility constraint binds in equilibrium. Conditions (iv') and (v') can then only be satisfied if both types' equilibrium contracts are at the tangency of their respective utility functions and the same pooling line. There is a continuum of XD-equilibrium outcomes, most of which yield positive rents to the uninformed agent. Conditions 3 and 4 give an implicit upper bound to the maximal level of rents that an uninformed agent can earn in any XD-equilibrium. This level is by far lower than in many other PBE. Thus, even if our criterion does not yield uniqueness it is still strong enough to shrink significantly the set of admissible PBE.

We would like to underline three properties of the application of the XD criterion that apply to both Case RS and Case S. First, it must be noted that part of the force of the XD criterion in the current game is related to the fact

³⁰ Play proceeds similarly to Case RS. We therefore refer the reader to the proof of Proposition 4 for the detailed description of strategies and beliefs.

that we consider only pure strategies on the part of the uninformed agent.³¹ This can be best understood by considering the problem of nesting sets of Strictly Preferred Responses for each player, which is the central element for the application of the XD criterion. In the case of pure strategies, this nesting reduces to a comparison of whether or not a type gains on the current contract proposal. In contrast, with mixed strategies it is necessary to compare possibly infinite sequences of probabilities, which obviously makes the use of the XD criterion very difficult. Although we have not characterized the implications of the XD criterion for the case of mixed acceptance rules, we nevertheless believe that our focus on pure strategies is economically reasonable.

Secondly, the inefficiencies characterizing the separating equilibrium outcomes underline a general property of our separating equilibria. In these equilibria, the support of beliefs may change in adjacent information sets. For example, in any separating equilibrium outcome, the uninformed agent eventually knows for sure with which type of agent he is contracting; however, at that point, if one type offers a distortion-reducing renegotiation from its equilibrium contract, it is rejected by the uninformed agent on the belief that this renegotiation may have been offered by either type. At first glance, this belief may appear contradictory. Why is the uninformed agent "changing his mind" about the type that offered the equilibrium contract once a Pareto-improving renegotiation for that type is offered from there? The answer may lie in the fact that any such renegotiation relaxes the other type's renegotiation-induced incentive-compatibility constraint. In fact, it is only by notseeing another proposal that the uninformed agent can reasonably maintain the belief that the player is of the conjectured type. Therefore, in order for one type to keep his signal of being that type, he must refrain from offering Pareto-improving renegotiations. A failure to do so would destroy one type's investment in the signal.³² In the light of this discussion, the property of a moving support of beliefs in adjacent information sets may not be as unintuitive as it appears.³³ We also note that it may be easier to think about this property on the normal form of the game in which both players are seen as choosing simultaneously their strategy at the beginning of the game and therefore any deviation really arises when the relevant beliefs are the priors.

Finally, Propositions 3 and 4 are derived only for the case of two types,³⁴ and it is therefore useful to briefly indicate the insights that are robust to the

³¹ Allowing for mixed strategies on the part of the informed agent does not change the implications of the XD criterion.

³² This argument is presented more formally in Noldeke and van Damme (1990b).

³³ Madrigal, Tan, and da Costa Werlang (1987) and Noldeke and van Damme (1990b) present a game of incomplete information in which the unique Nash equilibrium is supported by a changing support of beliefs. Note that most equilibrium concepts allow for the support of beliefs to change. See Vincent (1989) for an application with this property.

³⁴ Although we can provide some results for the case of more than two types, we have not been able to derive a characterization as complete as that found in Propositions 1–4.

inclusion of more types.³⁵ The first implication, which is common to both Case S and Case RS, is that all XD-equilibria are attained through some initial stage of pooling. This result remains valid with multiple types since a restriction analogous to (iii') (see Proposition 2) would conflict with any completely separating equilibria. Therefore, regardless of the number of types, in every XD-equilibria there is some cross-subsidization from the best types towards the worst types. A second result that is common to both cases, and that remains true regardless of the number of types, is that the XD criterion does not single out distortion-free equilibrium outcomes. In particular, in both the RS and the S cases, the best type always chooses to underinvest in the signal, in contrast to the one-shot signalling game. The only important element of Proposition 3 that is not robust to the inclusion of multiple types is the uniqueness result of the equilibrium outcome.³⁶

5. CONCLUSION

One of the main reasons for introducing renegotiation in models of contract determination is to understand which previous characterizations or results are robust and which are not. In the case of one-shot signalling games, the two central results in the literature seem to be that (1) the informed party generally uses self-selecting distortions to separate and thereby convey his information, (2) separating equilibria do not depend on the uninformed party's priors. In this paper we find that, when the contracting parties' power to commit not to renegotiate is eliminated, equilibrium outcomes still generally contain distortions, are usually separating, but nevertheless depend on priors.

The presence of distortions in the equilibrium of the signalling-cum-renegotiation game is that much more surprising once it is realized that the renegotiation process allows for a potentially infinite number of rounds of renegotiation. In particular, every accepted renegotiation proposal is always followed by at least one round of renegotiation. Consequently, our results suggest that one should be very cautious before adopting any axiomatic approach to renegotiation under asymmetric information that exogenously imposes restrictions on distortions.

³⁵ In the case of more than two types, the XD criterion can be quite easily extended along the same lines as Divinity (Banks and Sobel (1987)). The main idea of the extended version of the criterion would be that, following a deviation, the relative posteriors associated with the deviation coming from a type t versus coming from a type t' would need to be greater than the relative priors, whenever all continuation strategies which would make type t' deviate would also make type t deviate.

³⁶ For Case RS with more than two types, Jaynes (1978) constructed an outcome iteratively as follows: (1) maximize the best type's utility along the zero-profit pooling line; this is the final outcome for the best type and the first contract for the other types; (2) eliminate the best type from the pool and repeat the corresponding maximization for the best type within the remaining pool (conditional on having the first contract); the addition of the two contracts is this second type's outcome; (3) repeat this procedure until no one remains in the pool. This outcome can be supported as an XD-equilibrium of our renegotiation game, but it is not the unique outcome.

Our second important finding is that equilibrium outcomes depend on the uninformed party's priors even though executed contracts are separating. This result arises because final separating outcomes are achieved through an initial stage of pooling followed by separating renegotiations. On the one hand, this result implies that Akerlof's (1970) Lemons Problem may emerge, whereby trade breaks down when priors are sufficiently pessimistic, even in environments with a "rich" contracting space. On the other hand, as the prior probability that the informed agent is of the bad type approaches zero, the good type approaches his first-best outcome.

Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215, U.S.A.

and

Département de Sciences Économiques and C.R.D.E., Université de Montréal, C.P. 6128, Succursale A, Montréal, QC, Canada H3C 3J7.

Manuscript received July, 1990; final revision received October, 1992.

APPENDIX A

PROOF OF PROPOSITION 1: It will first be shown, by construction, that every pair of contracts satisfying the eight conditions of the proposition can be supported by a PBE. Then, it will be shown that these eight conditions are in fact necessary. We first introduce some notation. Define by \hat{m} , the outstanding contract, and by \hat{m}_k the k th accepted contract. Finally define by $\mathbf{B}(m_i, t) := \{m \in M / V(m, t) \geq V(m_i, t)\}$, and by $\mathbf{B}(m_i) := \mathbf{B}(m_i, L) \cap \mathbf{B}(m_i, H)$.

(i) For the case when $m_L \neq m_L^s$, and $m_H \neq m_L$, let us consider the following strategies. We first define the contract m_p as the unique solution to $\{V(m_p, H) = V(m_H, H) \text{ and } V(m_p, L) = V(m_L, L)\}$ (uniqueness is due to the single-crossing property of Assumption 3b). We now define subsets of the histories that will facilitate the writing of the strategies.

$$\begin{aligned} \hat{h}_p &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) / \hat{m}_k = m_p \forall k \geq 1 \}, \\ \hat{h} &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h}_p / \hat{m}_1 = m_p \text{ and for the first } \hat{m}_{k'} \neq m_p \\ &\quad (1) \hat{m}_{k'} \in \{m_i\} \cup \mathbf{B}(m_p) \text{ and } (2) \hat{m}_j \in \mathbf{B}(\hat{m}_{j-1}) \forall j > k' \}, \\ \hat{h}_L &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h}_p \cup \hat{h} / \hat{m}_1 = m_p \text{ and for the smallest } k' \text{ and } t' \in \{L, H\} \\ &\quad \text{such that } \hat{m}_{k'} \notin \mathbf{B}(\hat{m}_{k'-1}, t') \text{ with } \{\hat{m}_{k'-1}, \hat{m}_{k'}\} \neq \{m_p, m_i\}, t' = L\} \\ &\quad \text{or } \{\hat{m}_1 \neq m_p\} \}, \\ \hat{h}_H &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h}_p \cup \hat{h} \cup \hat{h}_L \}. \end{aligned}$$

The set \hat{h}_p regroups all histories in which only m_p has been accepted. The set \hat{h} regroups all histories starting with m_p which is then followed either by m_i and surely-acceptable contract acceptances, or only by surely-acceptable acceptances. The set \hat{h}_L regroups all histories starting with m_p that are eventually followed by an out-of-equilibrium contract over which the uninformed agent first makes losses when that contract is offered by type L , and also all histories starting with $\hat{m}_1 \neq m_p$. The set \hat{h}_H is the complement of these three sets, that is, the collection of histories for

which the uninformed agent first makes losses on an out-of-equilibrium contract with type H .

$$\Omega' = \begin{cases} \Omega'_1 = m_p; \\ \text{for all histories } \mathcal{H}^U(i) = \{m_1, \sigma_1, \dots, \sigma_{i-1}\}, i \geq 2, \\ \left\{ \begin{array}{l} m_i \text{ if } \mathcal{H}^U(i) \in \hat{h}_p, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}) \right\} \\ \text{if } \mathcal{H}^U(i) \in \hat{h} \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, L) \right\} \\ \text{if } \mathcal{H}^U(i) \in \hat{h}_L \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, H) \right\} \\ \text{if } \mathcal{H}^U(i) \in \hat{h}_H \text{ and } m_i \neq \hat{m}, \\ \emptyset \text{ otherwise;} \end{array} \right\} \end{cases}$$

$$\sigma = \begin{cases} \sigma_1 = \begin{cases} 2 & \text{if } \mathcal{H}^V(1) \in \{\{m_p\} \cup \{m_1/V(m_1, L) \geq V(0)\}\}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \sigma_i = \begin{cases} i+1 & \text{if } \mathcal{H}^V(i) \in \left\{ \begin{array}{l} \{\hat{h}_p/m_i \in \{m_i\} \cup \mathbf{B}(m_p)\} \cup \{\hat{h}/m_i \in \mathbf{B}(\hat{m})\} \\ \cup \{\hat{h}_L/m_i \in \mathbf{B}(\hat{m}, L)\} \cup \{\hat{h}_H/m_i \in \mathbf{B}(\hat{m}, H)\} \end{array} \right\}, \\ 0 & \text{otherwise;} \end{cases} \end{cases}$$

$$\mu = \begin{cases} \mu_1 = \begin{cases} \mu_0 & \text{if } \mathcal{H}^V(1) \in \{m_p\}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \mu_i = \begin{cases} 1 & \text{if } \mathcal{H}^V(i) \in \left\{ \begin{array}{l} \{\hat{h}_p/m_i \in \{m_H\} \cup \mathbf{B}(m_p, L) \setminus \{m_L\}\} \\ \cup \{\hat{h}/m_i \in \mathbf{B}(\hat{m}, L)\} \cup \hat{h}_H \end{array} \right\}, \\ \mu_{i-1} & \text{if } m_i = \emptyset, \\ 0 & \text{otherwise.} \end{cases} \end{cases}$$

In order to verify that the proposed strategies do constitute an equilibrium, let us first consider whether or not the informed player's strategy is optimal given σ . If either type deviates in the first contract offer, the history will be in \hat{h}_L and they cannot get utility higher than $(\max_m U(m, t) \text{ subject to } V(m, L) \geq V(0))$, since these are the only other contracts that will be accepted in the initial stage of contract offers, and since the following rounds of renegotiation cannot relax this constraint set. Therefore, by conditions (ii) and (iii), they would not want to deviate. At the first renegotiation proposal ($i = 2$), since conditions (iv) and (v) are satisfied, each type prefers his proposed strategy to that of following the other's proposed strategy and renegotiating in the latter rounds. Furthermore, by condition (viii), both types prefer offering the proposed renegotiation to proposing any alternative renegotiation that would be accepted. In the second round of renegotiation ($i = 3$), the history is in \hat{h} and, due to conditions (vi) and (vii), the informed player can do no better than to force the uninformed player to execute the contract by not making any new offer. Finally, in all other out-of-equilibrium subgames the informed player simply offers his preferred renegotiation within the set that will be accepted. The informed player never gains by deviating and not making an offer since this never induces the uninformed agent to revise his beliefs, and therefore the same offers that will be accepted later will also be accepted now. Consequently, Ω is sequentially optimal.

Let us now consider whether or not the uninformed player's strategy is sequentially optimal given Ω and μ . Along the equilibrium path, given condition (i) is satisfied, it is optimal for the uninformed player to accept the initial contract offer as well as each respective first-round renegotiation, and set

the delivery date one period later. At the second and subsequent rounds of renegotiation, the uninformed agent has no choice of strategy since the informed player chooses not to make new offers until the delivery date. Off the equilibrium path, the uninformed party only accepts initial contracts for which $V(m, L) \geq V(0)$ since he believes such deviations come from the type L player. Similarly, at any out-of-equilibrium renegotiation from contract \hat{m} , if the uninformed agent has only accepted surely-acceptable renegotiation in the past, he accepts renegotiations where $m_i \in \mathbf{B}(\hat{m})$ since he believes that the proposal is from a type H if $m_i \in \mathbf{B}(\hat{m}, L)$. This is supported by the fact that if the uninformed ever loses on one type in one round of renegotiation, the continuation strategy for either type in further rounds of renegotiation is to always offer a proposal for which the uninformed party is just indifferent and therefore he will not be able to recoup the losses. Consequently, it is optimal in stage i to only accept offers which are acceptable to the uninformed agent regardless of his beliefs. Upon acceptance, the uninformed agent is indifferent within the set of possible delivery dates. The informed agent always offers his preferred contract among the set of those that will be accepted. Once he has signed his preferred contract in that set, he stops making offers until the delivery date. Since, upon acceptance, the informed agent only needs one period to offer a new contract, picking a far away delivery date cannot be profitable and the informed agent chooses the delivery date for next period. Therefore, σ is sequentially optimal.

Finally, beliefs are consistent with Bayes rule along the equilibrium path since, after an initial pooling move of m_P , beliefs are the priors, and after the first round of separating renegotiations beliefs are respectively updated to $\mu_1 = 1$ if $m_2 = m_H$ and $\mu_1 = 0$ if $m_2 = m_L$.

For the case where $m_L = m_L^*$ and $m_L \neq m_H$, it is left to the reader to verify that the following set of strategies supports the outcome $\{m_L, m_H\}$ (if this outcome satisfies the conditions of the proposition). Moreover, if $m_L = m_H$, the following strategies also support this type of outcome, except that beliefs along the equilibrium path need trivial modifications for it to form a PBE. We again define some subsets of histories that have a similar interpretation as those previously defined.

$$\begin{aligned}
 \hat{h} &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) / \hat{m}_1 = m_H \text{ and } \hat{m}_j \in \mathbf{B}(\hat{m}_{j-1}) \ \forall j > 1 \}, \\
 \hat{h}_L &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h} / \hat{m}_1 = m_H \text{ and for the smallest } k' \text{ and } t' \in \{L, H\} \\
 &\quad \text{such that } \hat{m}_{k'} \notin \mathbf{B}(\hat{m}_{k'-1}, t'), t' = L \text{ or } \{\hat{m}_1 \neq m_H\} \}, \\
 \hat{h}_H &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h} \cup \hat{h}_L \}, \\
 \Omega' &= \left\{ \begin{array}{l} \Omega'_1 = m_i; \\ \text{for all histories } \mathcal{H}^U(i) = \{m_1, \sigma_1, \dots, \sigma_{i-1}\}, i \geq 2, \\ \Omega'_i = \left\{ \begin{array}{l} \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h} \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, L) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h}_L \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, H) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h}_H \text{ and } m_i \neq \hat{m}, \\ \emptyset, \text{ otherwise;} \end{array} \right. \end{array} \right. \\
 \sigma &= \left\{ \begin{array}{l} \sigma_1 = \begin{cases} 2 & \text{if } \mathcal{H}^V(1) \in \{ \{m_L\} \cup \{m_H\} \cup \{m_1 / V(m_1, L) \geq V(0) \} \}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \sigma_i = \begin{cases} i+1 & \text{if } \mathcal{H}^V(i) \in \left\{ \begin{array}{l} \{ \hat{h} / m_i \in \mathbf{B}(\hat{m}) \} \cup \{ \hat{h}_L / m_i \in \mathbf{B}(\hat{m}, L) \} \\ \cup \{ \hat{h}_H / m_i \in \mathbf{B}(\hat{m}, H) \} \end{array} \right\}, \\ 0 & \text{otherwise;} \end{cases} \end{array} \right.
 \end{aligned}$$

$$\mu = \begin{cases} \mu_1 = \begin{cases} 1 & \text{if } \mathcal{H}^V(1) = \{m_H\}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \mu_i = \begin{cases} 1 & \text{if } \mathcal{H}^V(i) \in \{\hat{h}/m_i \in \mathbf{B}(\hat{m}, L)\} \cup \hat{h}_H, \\ \mu_{i-1} & \text{if } m_i = \emptyset, \\ 0 & \text{otherwise.} \end{cases} \end{cases}$$

(ii) The second step of the proof consists in showing that any outcome that can be supported by a PBE must satisfy conditions (i)–(viii).

First notice that for any subgame where a contract \hat{m} has been accepted and a contract m_i has been proposed, where $V(m_i, L) > V(\hat{m}, L)$ and $V(m_i, H) > V(\hat{m}, H)$, m_i must be accepted by the uninformed player. Suppose the informed agent does not make any offer until the delivery date and then offers m_i . Rejection implies the fulfillment of \hat{m} . Therefore, in this period accepting m_i is a Best Response for the uninformed agent regardless of beliefs and of the informed player's continuation strategy. Conditions (iv) to (vii) follow directly from the following fact. For any pair of strategies that implements contracts that do not satisfy one of these constraints, the informed player could follow, either his strategy (if it is condition (vi) or (vii) that is not satisfied), or the other type's strategy (if it is condition (iv) or (v) that is not satisfied) until the proposed contract is accepted, not make any offer until the delivery date, and then propose the renegotiation he prefers within the set $\{m_i/V(m_i, L) > V(\hat{m}, L) \text{ and } V(m_i, H) > V(\hat{m}, H)\}$, which is accepted for sure since rejection signifies the delivery of the good at the conditions specified by \hat{m} . The informed player can then stop making offers in order to end the game.

Condition (i) obviously has to be satisfied, since any strategies that support an outcome that does not satisfy this condition would not even constitute a Nash equilibrium, given that the uninformed player would simply prefer to refuse the initial contract offer than play the proposed strategy. Similarly, conditions (ii) and (iii) follow from the fact that the informed parties must attain at least the described level of utility, since any initial offer within the set $\{m/V(m, L) > V(0)\}$ must, by sequential rationality, be accepted by the uninformed player regardless of beliefs.

Finally, condition (viii) must be satisfied for an outcome to be attainable through an initial stage of pooling followed by a separating renegotiation, which is the case for any outcome $m_L \neq m_L^s$ and $m_H \neq m_H^s$.³⁷ Since, in these cases, separation must be achieved through renegotiating from an initially pooled situation, say m_p , it must be the case that there exists a contract m_p where:

(a) $m_H \in \mathbf{B}(m_p, H)$ and $m_L \in \mathbf{B}(m_p, L)$, for the separating renegotiations to be acceptable to the uninformed party, and where

(b) $U(m_i, t) \geq \{\max_m U(m, t) \text{ subject to } m \in \mathbf{B}(m_p)\}$, for the informed party to prefer to follow his proposed strategy than to simply offer a contract within the set $\mathbf{B}(m_p)$, which is always accepted when the informed agent offers it at the delivery date. But for all the potential contracts that satisfy (a), the one most likely to satisfy (b) is m_p as defined in condition (viii). Any other contract satisfying (a) would simply make the constraint set in (b) larger, and therefore harder to satisfy. Consequently, condition (viii) is the existence condition of a pooling contract through which the proposed outcome can be achieved. Q.E.D.

PROOF OF COROLLARY 1: The outcome $\{m_L^s, m_H^s\}$ never satisfies condition (v) of Proposition 1 since m_L^s is strictly interior to the constraint set of this condition, that is,

(a) $V(m_L^s, L) > V(m_H^s, L)$,

(b) $V(m_L^s, H) > V(m_H^s, H)$.

Inequality (a) holds since, by Assumption 1, m_L^s is the unique solution to $\{\max_m U(m, L) \text{ subject to } V(m, L) \geq V(0)\}$; consequently for all contracts $m (\neq m_L^s)$ such that $U(m, L) = U(m_L^s, L)$ we have $V(m, L) < V(0)$. Inequality (b) holds because $V(m, L) < V(m, H)$ for any contract (Assumption 2), and $V(m_H^s, H) = V(m_L^s, L) = V(0)$ by the definition of the proposed outcome. Q.E.D.

PROOF OF PROPOSITION 2: We will show alternatively that each of these conditions must hold for m^* to be an XD-equilibrium outcome.

³⁷ This condition would not be necessary if we had allowed the informed player to initially offer a menu of contracts as in Maskin and Tirole (1992).

(i) Suppose that condition (v') is not true. Then it is possible to construct the contract m' that has the following properties.

- (1) $U(m', L) > U(m_L^*, L)$,
- (2) $V^P(m', \mu_0) > V^P(m_H^*, \mu_0)$,
- (3) $V(m', L) > V(m_H^*, L)$.

Given that m^* is a PBE and that (v') is not satisfied, it is easy to show that the contract m' exists. Suppose that player L follows type H 's strategy until m_H^* is offered. Since it is along the equilibrium path, it is always accepted by the uninformed agent. Then player L proposes to renegotiate at m' . If it is rejected, both types are not better playing the deviation m' since they get $U(m', t) \leq U(m_t^*, t)$. If it is accepted, in any continuation, player L can always secure at least as much as $U(m', L) > U(m_L^*, L)$ (property (1)). Therefore player L always gains from the deviation if it is accepted. Player H may or may not gain from it, depending on the relation between $U(m', H)$ and $U(m_H^*, H)$, and the continuation game. There are Preferred Responses that would make player H better off and some that may make him worse off. In any case, $\Lambda^+(H/\{\dots, m_H^*, \sigma_i, m'\}) \subseteq \Lambda^+(L/\{\dots, m_H^*, \sigma_i, m'\})$. By the XD criterion, upon being offered m' , beliefs μ' should be $\mu' \leq \mu_{i-1}$ (assuming that m' has been offered in period i). For any such belief, m' should therefore be accepted (properties (2) and (3)). But acceptance of m' breaks the equilibrium. Hence condition (v') must be satisfied in any XD-equilibrium.

(ii) A similar proof, in which the roles of L and H are inverted, shows that condition (iv') must also be satisfied in any XD-equilibrium.

(iii) Suppose condition (iii') is not satisfied. Then there must exist a contract m' that satisfies

- (1) $U(m', H) > U(m_H^*, H)$,
- (2) $V^P(m', \mu_0) > V(0)$.

Suppose player H offers m' in the first period. Since $V^P(m', \mu_0) > V(0)$, we have that $V(m', H) > V(0)$. If m' is accepted, player H always gains from that deviation (property (1)), while type L may or may not gain. If it is rejected both types lose from the deviation. This implies that $\Lambda^+(L/m') \subseteq \Lambda^+(H/m')$. Therefore the XD criterion imposes that, following the offer m' , $\mu_1 \geq \mu_0$. But these beliefs induce the uninformed player to accept m' . Therefore player H offers m' , thus upsetting the equilibrium. Hence condition (iii') must be satisfied in any XD-equilibrium. Q.E.D.

PROOF OF PROPOSITION 3: We first derive the characterization of the unique PBE outcome satisfying the conditions of Proposition 2. We then show that this outcome exists, and that it satisfies all conditions of Proposition 1. Finally, we construct strategies and beliefs that satisfy the XD criterion and that support this unique outcome m^{RS} as an XD-equilibrium. This proves the sufficiency part of the proof.

(i) Suppose that m_H^{RS} is not defined by (2). Since (iii') is satisfied by any XD-equilibrium outcome, we must have $V^P(m_H^{RS}, \mu_0) < V(0)$. This implies that $V(m_L^{RS}, L) > V(m_H^{RS}, L)$ for the uninformed agent's individual-rationality constraint to be satisfied. But in that case, because $U(\cdot, L)$ is steeper than $U(\cdot, H)$, either (iv) or (v) is not satisfied. Hence $V(m_L^{RS}, L) \leq V(m_H^{RS}, L)$. But this implies that the uninformed player's individual-rationality constraint (i) is not satisfied. Therefore m_H^{RS} must be the tangency point between $U(\cdot, H)$ and $V^P(\cdot, \mu_0) = V(0)$.

(ii) By the first part of the proof, we know that $U(\cdot, H) = U(m_H^{RS}, H)$ is tangent to $V^P(\cdot, \mu_0) = V^P(m_H^{RS}, \mu_0) = V(0)$. This implies that $V(m_L^{RS}, L) \geq V(m_H^{RS}, L)$ for the uninformed player's individual-rationality constraint to be satisfied. But if $V(m_L^{RS}, L) > V(m_H^{RS}, L)$, either (iv) or (v) is not satisfied. Therefore we must have $V(m_L^{RS}, L) = V(m_H^{RS}, L)$. Hence m_L^{RS} must be the solution of the right-hand side of (v') for conditions (i) and (v') to be satisfied.

(iii) The outcome $\{m_L^{RS}, m_H^{RS}\}$ characterized above clearly exists. The contracts m_H^{RS} and m_L^{RS} are each the solution to a well-defined maximization problem which has a unique solution by Assumption 1. Hence the XD-equilibrium outcome m^{RS} exists and is unique.

(iv) We now show that this outcome satisfies all conditions of Proposition 1. By construction, condition (i) is satisfied: it is easy to show that m_P as defined in Proposition 1 is equal to m_H^{RS} . Hence $V^P(m_P, \mu_0) = V^P(m_H^{RS}, \mu_0) = V(0)$. We assume that condition (ii) is satisfied for the problem to be interesting. Conditions (iii') and (v') are satisfied by construction. From the characterization of the outcome, $V(m_H^{RS}, H) \leq V(m_L^{RS}, H)$ and $V^P(m_H^{RS}, \mu_0) \leq V^P(m_L^{RS}, \mu_0)$. This implies that condition (iv') is satisfied. Because $U(\cdot, H) = U(m_H^{RS}, H)$ is tangent to $V^P(\cdot, \mu_0) = V^P(m_H^{RS}, \mu_0)$, condition (vi) is satisfied. By construction, either $m_L^{RS} = m_H^{RS}$, or m_L^{RS} is on his contract curve. In both cases condition (vii) is satisfied. Finally, since $m_P = m_H^{RS}$, condition (viii) is also satisfied. Hence the outcome m^{RS} is the outcome of a PBE.

(v) Finally we have to show that there exist strategies and beliefs that satisfy the XD criterion and that support the outcome m^{RS} as an XD-equilibrium.

First define $\mathbf{B}(m_i, \mu_0) := \{m \in M / V^P(m, \mu_0) \geq V^P(m_i, \mu_0)\}$. Then define the following subsets of histories.

$$\begin{aligned}\hat{h}_{\mu_0} &:= \{\mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) / \{m_1 \notin \mathbf{B}(m_H^{RS}, \mu_0)\} \text{ or } \{m_1 \in \mathbf{B}(m_H^{RS}, \mu_0) \cap \mathbf{B}(m_H^{RS}, L)\} \\ &\quad \text{and for the smallest } k' \text{ and } t' \in \{L, \mu_0, H\} \text{ with } \hat{m}_{k'} \notin \mathbf{B}(\hat{m}_{k'-1}, t'), t' = \mu_0\}\}, \\ \hat{h} &:= \{\mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) / \{m_1 \in \mathbf{B}(m_H^{RS}, \mu_0) \cap \mathbf{B}(m_H^{RS}, L)\} \\ &\quad \text{and } \{\hat{m}_j \in \mathbf{B}(\hat{m}_{j-1}, \mu_0) \cap \mathbf{B}(\hat{m}_{j-1}, L) \forall j > 1\}\}, \\ \hat{h}_L &:= \{\mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h}_{\mu_0} \cup \hat{h}\}.\end{aligned}$$

We now write the strategies and beliefs that support the equilibrium outcome $\{m_H^{RS}, m_L^{RS}\}$.³⁸

$$\begin{aligned}\Omega^t &= \begin{cases} \Omega_1^t = m_H^{RS}; \\ \text{for all histories } \mathcal{H}^U(i) = \{m_1, \sigma_1, \dots, \sigma_{i-1}\}, i \geq 2, \\ \Omega_i^t = \begin{cases} \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, \mu_0) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h}_{\mu_0} \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, \mu_0) \cap \mathbf{B}(\hat{m}, L) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h} \text{ and } m_i \neq \hat{m}, \\ \arg \left\{ \max_m U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, L) \right\} \\ \quad \text{if } \mathcal{H}^U(i) \in \hat{h}_L \text{ and } m_i \neq \hat{m}, \\ \emptyset \quad \text{otherwise;} \end{cases} \end{cases} \\ \sigma &= \begin{cases} \sigma_1 = \begin{cases} 2 & \text{if } \mathcal{H}^V(1) \in \hat{h}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \sigma_i = \begin{cases} i+1 & \text{if } \mathcal{H}^V(i) \in \left\{ \begin{aligned} &\hat{h}/m_i \in \mathbf{B}(\hat{m}, \mu_0) \cap \mathbf{B}(\hat{m}, L) \\ &\cup \left\{ \hat{h}_{\mu_0}/m_i \in \mathbf{B}(\hat{m}, \mu_0) \right\} \cup \left\{ \hat{h}_L/m_i \in \mathbf{B}(\hat{m}, L) \right\} \end{aligned} \right\} \\ 0 & \text{otherwise;} \end{cases} \end{cases} \\ \mu &= \begin{cases} \mu_1 = \begin{cases} \mu_0 & \text{if } \mathcal{H}^V(1) \in \hat{h} \cup \hat{h}_{\mu_0}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \mu_i = \begin{cases} \mu_0 & \text{if } \mathcal{H}^V(i) \in \left\{ \hat{h}/m_i \in \mathbf{B}(\hat{m}, L) \setminus \{m_L^{RS}\} \right\} \cup \hat{h}_{\mu_0}, \\ \mu_{i-1} & \text{if } m_i = \emptyset, \\ 0 & \text{otherwise.} \end{cases} \end{cases}\end{aligned}$$

We now describe the separating equilibrium. Along the equilibrium path, each type offers m_H^{RS} in the first stage of the game which is accepted with delivery next period. In the second stage, the

³⁸ These strategies and beliefs support the separating equilibrium outcome. They would need trivial modifications to support the pooling equilibrium outcome.

high type does not make any offer, therefore ending the game with the fulfillment of the contract. The low type offers his equilibrium contract m_L^{RS} which is accepted with delivery next period. In the third stage, the low type does not make any offer, thus imposing the fulfillment of the contract.

Off the equilibrium path, the informed agent always maximizes against the set of offers accepted by the uninformed agent. In the first stage, the uninformed agent accepts offers $m_1 \in \mathbf{B}(m_H^{RS}, \mu_0) \cap \mathbf{B}(m_L^{RS}, L)$. Any offer such that $V^P(m_1, \mu_0) < V(0)$ is believed to come from either type and is therefore rejected. Any other offer is believed to come from type L and is therefore rejected. In round i , the uninformed player accepts all renegotiations $m_i \in \mathbf{B}(\hat{m}, \mu_0)$ if the history is in \hat{h}_{μ_0} since he believes that those could come from either type. If the history is in \hat{h} , the uninformed agent accepts all offers in the set $\mathbf{B}(\hat{m}, \mu_0) \cap \mathbf{B}(\hat{m}, L)$. In these cases, any other offer must be rejected since by accepting it the uninformed player loses and cannot recoup the losses in future rounds. For example, if he accepts an offer such that $V(m_i, L) < V(\hat{m}, L)$, he expects to renegotiate in round $i+1$ as if he were facing a low type and therefore will accept all offers such that $m_{i+1} \in \mathbf{B}(m_i, L)$. Hence the uninformed player cannot recoup the losses incurred in stage i . Finally, by the same argument as that presented in the proof of Proposition 1, it can be shown that the uninformed is always indifferent within the set of possible delivery dates.

Because the game is infinite and because the definition of Best Responses is fairly large, each type has a Preferred Response to any deviation. This implies that beliefs do not have to be necessarily concentrated on one type. We now show that the beliefs specified satisfy the XD criterion. We only have to check beliefs following a first deviation from the equilibrium path. After the initial offer, $\mu_1 = \mu_0$ trivially satisfies the XD criterion; for all initial offers $m_1 \in \hat{h}_L$ we have $\mu_1 = 0$. This is justified by the fact that if m_1 is accepted and then renegotiated with beliefs $\mu_1 = 0$ then type L gains from the deviation, while type H may or may not gain depending on subsequent play. In the second round, the only history consistent with m_2 being a first deviation off the equilibrium path is \hat{h} . In this case, if $m_2 \in \mathbf{B}(m_H^{RS}, L) \setminus \{m_L^{RS}\}$, then $\mu_2 = \mu_0$ which is consistent with the XD criterion. For all other $m_2 \in \hat{h}$ not in $\mathbf{B}(m_H^{RS}, L)$, then $\mu_2 = 0$. This is consistent with the XD criterion since if such a m_2 were accepted and renegotiated with beliefs $\mu_2 = 0$ then type L gains from the deviation, while type H may or may not gain depending on subsequent play. In the third stage, equilibrium play follows a history in \hat{h} and beliefs are specified as in the second stage and are therefore consistent with the XD criterion. Hence the candidate equilibrium outcome can be supported as an XD-equilibrium with the above beliefs and strategies. Q.E.D.

PROOF OF PROPOSITION 4: We first show that the equilibrium outcomes must have these four properties. Then we show that these outcomes exist and satisfy the conditions to be a PBE. Finally, we show that these outcomes can be supported by strategies and beliefs that satisfy the XD criterion.

(i) The first thing we show is that conditions (iv') and (v') in Proposition 2 imply property (1). Given that $U(\cdot, H)$ is steeper than $U(\cdot, L)$ and that both incentive-compatibility constraints, (iv') and (v'), must be satisfied, the solution of the right-hand side of condition (v') must lie on $V^P(m_H^S, \mu_0)$. If this solution is such that $U(\cdot, L)$ is tangent to $V^P(\cdot, \mu_0) = V^P(m_H^S, \mu_0)$, then we must have that $V^P(m_L^S, \mu_0) \leq V^P(m_H^S, \mu_0)$. If the solution is at m_H^S and $V^P(m_L^S, \mu_0) > V^P(m_H^S, \mu_0)$, then, because $U(\cdot, H)$ is steeper than $U(\cdot, L)$, $U(m_L^S, H) > U(m_H^S, H)$ and type H 's incentive-compatibility constraint cannot be satisfied. Therefore $V^P(m_L^S, \mu_0) \leq V^P(m_H^S, \mu_0)$. A similar argument using (iv') shows that $V^P(m_L^S, \mu_0) \geq V^P(m_H^S, \mu_0)$. Combining these two results we must have that $V^P(m_L^S, \mu_0) = V^P(m_H^S, \mu_0)$.

(ii) We now show that each type's indifference curve must be tangent to the same $V^P(\cdot, \mu_0)$. Suppose that $U(\cdot, L)$ is not tangent to $V^P(\cdot, \mu_0) = V^P(m_L^S, \mu_0)$. By the previous result, we know that both types are on the same $V^P(\cdot, \mu_0)$ line. Because $U(\cdot, H)$ is steeper than $U(\cdot, L)$, it must be the case that $V(m_H^S, H) > V(m_L^S, H)$ for type H 's incentive-compatibility constraint to be satisfied. But this implies that the solution to (iv') is on $V^P(m_L^S, \mu_0)$. Therefore, $U(\cdot, H) = U(m_H^S, H)$ must be tangent to $V^P(\cdot, \mu_0) = V^P(m_L^S, \mu_0)$. But then (v') is not satisfied. Hence $U(\cdot, L)$ must also be tangent to $V^P(\cdot, \mu_0) = V^P(m_L^S, \mu_0)$. A similar argument shows that $U(\cdot, H)$ must be tangent to $V^P(\cdot, \mu_0) = V^P(m_H^S, \mu_0)$.

(iii) We now give the interval of rents that can be earned by the uninformed agent. The highest pooling line to which both types are tangent is the zero-rent pooling line. Consider a candidate XD-equilibrium outcome m^S such that both types are tangent to a pooling line $V^P(\cdot, \mu_0) > V(0)$. Define a contract offer m' such that

- (1) $U(m', H) > U(m_H^S, H)$,
- (2) $V^P(m', \mu_0) > V(0)$.

This contract clearly exists. Now suppose type H offers m' at the first stage of the game. By construction we know that $\Lambda^+(L/m') \subseteq \Lambda^+(H/m')$. Hence $\mu_1 \geq \mu_0$. Since $V(m', H) > V(0)$ (because $V^P(m', \mu_0) > V(0)$), these restrictions on beliefs imply that m' must be accepted. This breaks the candidate equilibrium. Hence, any XD-equilibrium outcome must be such that the tangent pooling line be at most the zero-rent line.

The lowest pooling line to which the informed player's contracts could be tangent is the line which yields $\mu_0 V(m_H^S, H) + (1 - \mu_0) V(m_L^S, L) = V(0)$. A lower-rent equilibrium outcome would not satisfy condition (i).

Having now characterized the equilibrium outcomes, it remains to show that these outcomes exist, that they constitute a PBE, and that they can be supported by strategies and beliefs that satisfy the XD criterion.

(iv) The contracts are such that both types are tangent to a given pooling line. These contracts are the solution to a well-defined maximization problem. Hence, by Assumption 1, they exist.

(v) We now have to show that the outcomes characterized above satisfy the eight conditions of Proposition 1. By construction, condition (i) is trivially satisfied. We will assume that both types are sufficiently different that condition (ii) is satisfied for all XD-equilibrium outcomes. By construction conditions (iii'), (iv'), and (v') are satisfied. Given that both types are tangent to a pooling line, conditions (vi) and (vii) are satisfied. Finally, since $V^P(m_p, \mu_0) > V^P(m_t^S, \mu_0)$ for $t = L, H$ (where m_p is defined as in condition (viii)), and since conditions (iv) through (vii) are satisfied, condition (viii) is also satisfied.

(vi) We now show that all outcomes in the set characterized above can be supported by strategies and beliefs which satisfy the XD criterion. First define implicitly the following three contracts, $m_p(m'), m_L(m'), m_H(m')$, as a function of an arbitrary contract m' by

- (1) $V^P(m_p(m'), \mu_0) = V^P(m', \mu_0)$,
- (2) $m_t(m') = \arg \max_m U(m, t)$ subject to $V^P(m, \mu_0) \geq V^P(m_t(m'), \mu_0)$,
- (3) $V(m_p(m'), L) = V(m_L(m'), L)$,
- (4) $V(m_p(m'), H) = V(m_H(m'), H)$,
- (5) $\mu_0 V(m_H(m'), H) + (1 - \mu_0) V(m_L(m'), L) = V^P(m', \mu_0)$.

For a given contract m' , $m_p(m')$ is on the same pooling line as m' (by (1)). The contracts $m_t(m')$ are tangent to a pooling line (by (2)). For a given type t , the uninformed agent is indifferent between $m_p(m')$ and $m_t(m')$ (by (3) and (4)). The contracts $m_t(m')$ are such that their expected value is equal to the expected value of contract m' (by (5)). It can be shown that there exists a unique solution to these three contracts for a given m' . Define by \hat{m}_{-k} the k th accepted contract before \hat{m} , and define the following subsets of histories:

$$\begin{aligned} \hat{h} &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) / \hat{m}_k = m_L^S \ \forall k \geq 1 \}, \\ \hat{h}_{\mu_0} &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h} / \{ \hat{m}_1 = m_L^S \} \text{ or } \{ \hat{m}_1 \neq m_L^S \text{ and } V^P(\hat{m}, \mu_0) \geq V(0) \} \} \\ &\quad \text{and } \hat{m}_j \in \mathbf{B}(\hat{m}_{j-1}, \mu_0) \ \forall j > 1 \}, \\ \hat{h}_p &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h} \cup \hat{h}_{\mu_0} / m_t(\hat{m}_{-2}) \neq \hat{m} \neq m_p(\hat{m}_{-1}) \}, \\ \hat{h}_t &:= \{ \mathcal{H}^U(i) \text{ or } \mathcal{H}^V(i) \notin \hat{h} \cup \hat{h}_{\mu_0} / \hat{m} = m_p(\hat{m}_{-1}) \}. \end{aligned}$$

Now consider the following strategies and beliefs:

$$\Omega^t = \begin{cases} \Omega_1^t = m_L^S; \\ \text{for all histories } \mathcal{H}^U(i) = \{m_1, \sigma_1, \dots, m_{i-1}, \sigma_{i-1}\}, i \geq 2, \\ \Omega_i^t = \begin{cases} m_H^S & \text{if } \mathcal{H}^U(i) \in \hat{h} \text{ and } t = H, \\ m_p(\hat{m}) & \text{if } \mathcal{H}^U(i) \in \hat{h}_p, \\ m_t(\hat{m} - 1) & \text{if } \mathcal{H}^U(i) \in \hat{h}_t, \\ \arg \max_m \left\{ \max U(m, t) \text{ subject to } m \in \mathbf{B}(\hat{m}, \mu_0) \right\} & \text{if } \mathcal{H}^U(i) \in \hat{h}_{\mu_0} \text{ and } m_t \neq \hat{m}, \\ \emptyset & \text{otherwise;} \end{cases} \end{cases}$$

$$\sigma = \begin{cases} \sigma_1 = \begin{cases} 2 & \text{if } \mathcal{H}^V(1) \in \{ \{m_L^S\} \cup \{m_1/V^P(m_1, \mu_0) \geq V(0)\} \}, \\ 0 & \text{otherwise;} \end{cases} \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \sigma_i = \begin{cases} i+1 & \text{if } \mathcal{H}^V(i) \in \left\{ \begin{aligned} & \{ \hat{h} \cup \hat{h}_{\mu_0}/m_i \in \mathbf{B}(\hat{m}, \mu_0) \} \\ & \cup \{ \hat{h}_i/m_i \in \{m_i(\hat{m}_{-1})\} \cup \mathbf{B}(\hat{m}, \mu_0) \} \\ & \cup \{ \hat{h}_P/m_i \in \{m_P(\hat{m})\} \cup \mathbf{B}(\hat{m}, \mu_0) \} \end{aligned} \right\}, \\ 0 & \text{otherwise;} \end{cases} \end{cases}$$

$$\mu = \begin{cases} \mu_1 = \mu_0; \\ \text{for all histories } \mathcal{H}^V(i) = \{m_1, \sigma_1, \dots, m_i\}, i \geq 2, \\ \mu_i = \begin{cases} 1 & \text{if } \mathcal{H}^V(i) \in \{ \hat{h}/m_i = m_H^S \} \cup \{ \hat{h}_i/m_i = m_H(\hat{m}_{-1}) \}, \\ 0 & \text{if } \mathcal{H}^V(i) \in \{ \hat{h}_i/m_i = m_L(\hat{m}_{-1}) \}, \\ \mu_{i-1} & \text{if } m_i = \emptyset, \\ \mu_0 & \text{otherwise.} \end{cases} \end{cases}$$

We now describe the equilibrium. Along the equilibrium path, each type offers m_L^S in the first stage which is accepted with delivery next period. In the second stage, type L makes no offer therefore ending the game. Type H offers m_H^S which is accepted with delivery next period. He then ends the game by making no offer in the third period. For both types, ending the game is optimal since after the acceptance of m_L^S , the only accepted renegotiations are such that $m_i \in \mathbf{B}(m_L^S, \mu_0)$.

Off the equilibrium path, the informed agent maximizes against the set of offers that will be accepted by the uninformed agent. In the first stage, the uninformed agent only accepts offers m_1 such that $V^P(m_1, \mu_0) \geq V(0)$. In subsequent stages i , he only accepts offers such that $m_i \in \mathbf{B}(\hat{m}, \mu_0)$. It is sequentially rational to only accept those offers since, if he was accepting any other offer m_i such that $m_i \notin \mathbf{B}(\hat{m}, \mu_0)$, the informed agent would follow by offering $m_P(m_i)$ and then $m_i(m_i)$. These offers would be accepted by the uninformed agent and he would then earn $V^P(m_i, \mu_0) < V^P(\hat{m}, \mu_0)$. By the same argument as that presented in the proof of Proposition 1, it can be shown that the uninformed agent is always indifferent within the set of possible delivery dates.

Because the game is infinite and because the definition of Best Responses is fairly large, each type has a Preferred Response to any deviation. This implies that beliefs do not have to be necessarily concentrated on one type. We now show that the beliefs specified satisfy the XD criterion. We only have to check beliefs following a first deviation from the equilibrium path. After the initial offer, $\mu_1 = \mu_0$ trivially satisfies the XD criterion. In the second stage, any out-of-equilibrium offer following $\hat{m}_1 = m_L^S$ is a history in \hat{h} and has posterior $\mu_2 = \mu_0$. These beliefs trivially satisfy the XD criterion. In the third stage, the equilibrium history is in \hat{h}_{μ_0} and any offer following such a history is then assessed with posterior $\mu_3 = \mu_0$. These beliefs trivially satisfy the XD criterion. Hence all candidate equilibrium outcomes can be supported as XD-equilibria with the above beliefs and strategies. Q.E.D.

APPENDIX B

The object of this appendix is to demonstrate that the equilibrium characterizations presented in the main body of the paper can be thought of as a limit of the set of equilibria of a renegotiation game in which delay is costly. The main interest of such a limiting result is to show that the results of the paper are not pathologies arising only when delay is absolutely costless. Moreover, the exercise has the advantage of highlighting some of the most important assumptions underlying our characterization.

Consider the following modifications of the renegotiation game described in Section 3 in which both players have positive discount rates. The presence of renegotiation after each contract acceptance implies that delay only occurs if delivery takes place after period 2; however, the proofs of the following propositions are considerably simplified when we assume that discounting only

begins after period 3. Therefore, for ease of presentation we set the delay $n = \max(0, \sigma - 3)$ where σ is the delivery date (this assumption is absolutely unessential for proving the following results).

The next proposition characterizes all PBE outcomes of the game in which delay is costly for both players. Corollary 1B then proves two limiting results; one for the case where r_b goes to zero when $r_s = 0$, and one for the case where both r_s and r_b approach zero simultaneously. The reason to discuss these two cases is that the outcomes described in Proposition 1 correspond exactly to those resulting from the first limit, while they represent only a subset of those resulting from the second limit. Proposition 2B, Corollary 2B, and Proposition 3B also extend the results of the paper in a similar manner.

PROPOSITION 1B: *Any pair of contracts $\{m_L, m_H\}$ and pair of delays $\{n_L, n_H\}$ can be supported as part of the outcome of a Perfect Bayesian Equilibrium of the renegotiation game in which discount rates are positive if and only if the contracts and delays satisfy the following conditions:*

$$\begin{aligned}
 (i_B) \quad & \mu_0 \exp(-r_b \Delta n_H) V(m_H, H) + (1 - \mu_0) \exp(-r_b \Delta n_L) V(m_L, L) \geq V(0), \\
 (ii_B) \quad & \exp(-r_s \Delta n_L) U(m_L, L) \geq \left\{ \max_m U(m, L) \text{ subject to } V(m, L) \geq V(0) \right\}, \\
 (iii_B) \quad & \exp(-r_s \Delta n_H) U(m_H, H) \geq \left\{ \max_m U(m, H) \text{ subject to } V(m, L) \geq V(0) \right\}, \\
 (iv_B) \quad & \exp(-r_s \Delta n_H) U(m_H, H) \\
 & \geq \exp(-r_s \Delta n_L) \max \left\{ U(m_L, H), \exp(-r_s \Delta) \right. \\
 & \quad \times \left. \left\{ \max_m U(m, H) \text{ subject to } \begin{aligned} & \exp(-r_b \Delta) V(m, H) \geq V(m_L, H) \\ & \exp(-r_b \Delta) V(m, L) \geq V(m_L, L) \end{aligned} \right\} \right\}, \\
 (v_B) \quad & \exp(-r_s \Delta n_L) U(m_L, L) \\
 & \geq \exp(-r_s \Delta n_H) \max \left\{ U(m_H, L), \exp(-r_s \Delta) \right. \\
 & \quad \times \left. \left\{ \max_m U(m, L) \text{ subject to } \begin{aligned} & \exp(-r_b \Delta) V(m, H) \geq V(m_H, H) \\ & \exp(-r_b \Delta) V(m, L) \geq V(m_H, L) \end{aligned} \right\} \right\}, \\
 (vi_B) \quad & \exp(-r_s \Delta n_H) U(m_H, H) \\
 & \geq \exp(-r_s \Delta n_L) \max \left\{ U(m_H, H), \exp(-r_s \Delta) \right. \\
 & \quad \times \left. \left\{ \max_m U(m, H) \text{ subject to } \begin{aligned} & \exp(-r_b \Delta) V(m, H) \geq V(m_H, H) \\ & \exp(-r_b \Delta) V(m, L) \geq V(m_H, L) \end{aligned} \right\} \right\}, \\
 (vii_B) \quad & \exp(-r_s \Delta n_L) U(m_L, L) \\
 & \geq \exp(-r_s \Delta n_H) \max \left\{ U(m_L, L), \exp(-r_s \Delta) \right. \\
 & \quad \times \left. \left\{ \max_m U(m, L) \text{ subject to } \begin{aligned} & \exp(-r_b \Delta) V(m, H) \geq V(m_L, H) \\ & \exp(-r_b \Delta) V(m, L) \geq V(m_L, L) \end{aligned} \right\} \right\},
 \end{aligned}$$

(viii_B) if $m_L \neq m_L^*$ and $m_L \neq m_H$, then for $t = L, H$,

$$\exp(-r_s \Delta n_t) U(m_t, t)$$

$$\geq \exp(-r_s \Delta \max\{\min\{n_L, n_H\} - 1, 0\}) \max \left\{ U(m_p, t), \right.$$

$$\left. \exp(-r_s \Delta) \left\{ \begin{array}{l} \max_m U(m, t) \text{ subject to} \\ \exp(-r_b \Delta) V(m, H) \geq V(m_p, H) \\ \exp(-r_b \Delta) V(m, L) \geq V(m_p, L) \\ V(m_p, H) = \exp(-r_b \Delta \max\{1, n_H - n_L + 1\}) V(m_H, H) \\ V(m_p, L) = \exp(-r_b \Delta \max\{1, n_H - n_L + 1\}) V(m_L, L) \end{array} \right\} \right\}.$$

PROOF OF PROPOSITION 1B: The proof of this proposition is kept to a minimum since it is very similar to the proof of Proposition 1. We begin by showing that the eight conditions of the proposition are indeed necessary and then we describe the strategies that can be used to support each element of the set.

We first argue that conditions (i_B)–(iii_B) need to hold for any PBE outcome. In the case of condition (i_B), the statement must be satisfied for any pair of equilibrium contracts and delays since, otherwise, the buyer would strictly prefer to reject the initial offer and receive $V(0)$. Conditions (ii_B) and (iii_B) must be satisfied since, otherwise, the seller could always deviate by initially offering his preferred contract within the set $\{m/V(m, L) > V(0)\}$ which is always accepted by the buyer since he can specify a delivery at period 2 (involving no delay) and not accept any proposed renegotiations.

Condition (viii_B) must be satisfied for any equilibrium outcome to be attained. Consider a separating contractual outcome $\{m_L, m_H\}$ in which there is cross-subsidization. Such an outcome can only be attained through some initial stage(s) of pooling (otherwise the buyer would reject the initial equilibrium offer of the subsidized type). Call \hat{m} the last pooling contract before each type's equilibrium offer becomes different from the other type. Also suppose that such separation occurs after \hat{n} of delay. Then it must be the case that, for each type t ,

$$V(\hat{m}, t) \leq \exp\{-r_b \Delta(n_t - \hat{n})\} V(m_t, t)$$

for the buyer to accept the subsequent equilibrium separating offer(s) of each seller's type (otherwise the buyer would take the delivery of the contract \hat{m} after \hat{n} of delay). The necessary condition is then found when the set of contracts that satisfy this condition is the largest. This is the case when the two types pool for the longest possible delay, that is, when $\hat{n} = \min\{n_L, n_H\} - 1$. This argument motivates the definition of m_p in condition (viii_B). It is now easy to see that condition (viii_B) must be satisfied for equilibrium outcomes in which there is cross-subsidization to be attainable.

Conditions (iv_B)–(vii_B) must also be satisfied for any PBE outcome. The reasoning for this claim is identical to that of Proposition 1, that is, if these conditions are not satisfied the seller can always follow the appropriate equilibrium path and offer a contract within the “surely-acceptable” set of offers just before delivery. Since such a deviating proposal would cause one period of delay if accepted, the set of “surely-acceptable” offers includes the consideration of one period of delay. Furthermore standard incentive-compatibility constraints must be included in each condition to ensure that no mimicking occurs in equilibrium, that is, for high discount rates, it is possible that the preferred contract in the set of surely-acceptable offers discounted one period is not preferred to the equilibrium contract of the other type; however, in the limit, when both discount rates become

arbitrarily small, the standard incentive-compatibility constraint will never dominate the constraint over the set of surely-acceptable offers even if discounted one period.

In order to show that the conditions of the proposition are also sufficient, we must construct equilibrium strategies and beliefs that can support each of these pairs of contracts and delays. We partition equilibrium outcomes into two classes. In the first class, there is no delay along the equilibrium path, that is, $n_L = n_H = 0$. The second class includes all equilibrium outcomes in which there is delay along the equilibrium path, that is, $n_L > 0$ and/or $n_H > 0$.

When there is no delay along the equilibrium path, the strategies and beliefs described in Proposition 1 need only to be modified slightly to be used in the current proposition. The only necessary changes are that for all periods $i > 3$, the specifications of the strategies that involve the sets $\mathbf{B}(m_i, t)$ and $\mathbf{B}(m_i)$ be redefined so that these sets reflect the cost of delay to the buyer. In particular, for $i > 3$, let $\mathbf{B}(m_i, t) := \mathbf{B}'(m_i, t) := \{m \in M / \exp(-r_b \Delta) V(m, t) \geq V(m_i, t)\}$, and let $\mathbf{B}(m_i)$ be redefined on the new sets $\mathbf{B}'(m_i, t)$, that is, $\mathbf{B}(m_i) := \mathbf{B}'(m_i, L) \cap \mathbf{B}'(m_i, H)$. With these modifications, it is straightforward to verify that all pairs of contracts satisfying the statement of the proposition are actually supported by these strategies. Moreover, all these contracts can be supported by equilibria in which delivery arises either in period 2 or 3.

Equilibria in which there is delay along the equilibrium path are somewhat different from equilibria in which there is no delay. Since these equilibria are not central to our analysis, we describe informally the strategies and beliefs that support them. Define t' as the type which has the smallest equilibrium delay, that is, $t' = \arg \min_t \{n_t\}$. Strategies and beliefs are constructed similarly to those of equilibria involving no delay with the following modifications. We first consider the equilibrium path. The informed agent, regardless of type, initially offers the null contract. After $n_{t'} - 2$ periods of delay, both types offer the contract m_P . After $n_{t'} - 1$ of delay, type t' offers $m_{t'}$; after $n_t - 1$, type t offers m_t . Each equilibrium offer is accepted by the uninformed agent. After accepting the null contract, the uninformed agent sets delivery for $n_{t'} - 2$. Following the acceptance of m_P , he sets delivery at $n_t - 1$. Then each type's equilibrium contract is delivered after its equilibrium delay. Along the path beliefs are updated using Bayes rule. We now consider strategies and beliefs off the equilibrium path. We need only to consider here strategies and beliefs following a deviation by the uninformed agent with respect to the choice of the delivery date (this choice was irrelevant in equilibria with no delay). If delivery following the acceptance of the null contract is set earlier than $n_{t'} - 2$, then it is offered again (and accepted) until $n_{t'} - 3$ with delivery in $n_{t'} - 2$, and the rest of the game then proceeds as along the equilibrium path. If delivery following the acceptance of m_P is set at a date smaller than $n_t - 1$, then renegotiation proceeds in the set $\mathbf{B}(m_P, t)$, which leaves, by condition (viii_B), the buyer at most indifferent between playing its equilibrium strategy and deviating. If delivery of the null contract and m_P is set later, then equilibrium strategies and beliefs are left unchanged. Finally if delivery of equilibrium contracts is set later than the equilibrium dates, then the seller offers in subsequent periods the same equilibrium contract for which delivery is set at its equilibrium date. It is clear that the buyer has no incentive to deviate in choosing delivery dates. Q.E.D.

COROLLARY 1B: (1) For $r_s = 0$ and $r_b > 0$, the set of equilibrium contractual outcomes $\{m_L, m_H\}$ defined in Proposition 1 is the limit as r_b approaches zero of the set of equilibrium contractual outcomes defined in Proposition 1B.

(2) For $r_s > 0$ and $r_b > 0$, the set of equilibrium contractual outcomes $\{m_L, m_H\}$ defined in Proposition 1 is included in the limit as r_b and r_s approach zero of the set of equilibrium contractual outcomes defined in Proposition 1B.

PROOF OF COROLLARY 1B: (1) Set $r_s = 0$ in the eight conditions of Proposition 1B. Conditions (ii_B)–(vii_B) are now independent of equilibrium delays n_L and n_H . There remain conditions (i_B) and (viii_B) that depend on equilibrium delays; however any contractual outcome that satisfies conditions (i_B)–(viii_B) with n_L and/or n_H positive also satisfies them with $n_L = n_H = 0$. Therefore, when $r_s = 0$, there is no loss in generality in setting $n_L = n_H = 0$ when $r_s = 0$ to characterize all equilibrium contractual outcomes. Finally we need to show that the eight conditions of Proposition 1B, when $r_s = 0$ and $n_L = n_H = 0$, converge to the eight conditions of Proposition 1. The right-hand side of conditions (iv_B)–(viii_B) includes standard incentive-compatibility constraints as well as incentive constraints over the set of surely-acceptable offers (discounted one period). It is clear that these last constraints are the binding ones when r_b is small enough. Hence, for small values of r_b , all eight conditions define a continuous correspondence in r_b , and therefore the first part of the corollary follows immediately. (2) By the same argument as above, the set of equilibrium contractual

outcomes $\{m_L, m_H\}$ defined in Proposition 1 is the limit as r_b and r_s approach zero of the set of equilibrium contractual outcomes in which there is no delay ($n_L = n_H = 0$) defined in Proposition 1B. However when both discount rates are positive there can also be equilibria with delay; in the limit these equilibrium outcomes are different from the equilibrium outcomes characterized in Proposition 1. Hence the eight conditions of Proposition 1B define a lower-hemicontinuous correspondence (for contractual outcomes) in both discount rates with the outcomes described in Proposition 1 always being included in the correspondence. *Q.E.D.*

PROPOSITION 2B: *Suppose that $m^* = \{m_L^*, m_H^*\}$ and $\{n_L^*, n_H^*\}$ is a PBE outcome of the renegotiation game with both discount rates being positive, and that this outcome satisfies the XD criterion; then the following three conditions must hold.*

$$\begin{aligned}
 (\text{iii}'_B) \quad & \exp \{-r_s \Delta n_H^*\} U(m_H^*, H) \geq \left\{ \max_m U(m, H) \quad \text{subject to} \quad V^P(m, \mu_0) \geq V(0) \right\}, \\
 (\text{iv}'_B) \quad & \exp \{-r_s \Delta n_H^*\} U(m_H^*, H) \\
 & \geq \exp \{-r_s \Delta n_L^*\} \max \left\{ U(m_L^*, H), \exp -r_s \Delta \right. \\
 & \quad \times \left. \left\{ \max_m U(m, H) \quad \text{subject to} \quad \begin{aligned} & \exp(-r_b \Delta) V(m, H) \geq V(m_L^*, H) \\ & \exp(-r_b \Delta) V^P(m, \mu_0) \geq V^P(m_L^*, \mu_0) \end{aligned} \right\} \right\}, \\
 (\text{v}'_B) \quad & \exp \{-r_s \Delta n_L^*\} U(m_L^*, L) \\
 & \geq \exp \{-r_s \Delta n_H^*\} \max \left\{ U(m_H^*, L), \exp -r_s \Delta \right. \\
 & \quad \times \left. \left\{ \max_m U(m, L) \quad \text{subject to} \quad \begin{aligned} & \exp(-r_b \Delta) V(m, L) \geq V(m_H^*, L) \\ & \exp(-r_b \Delta) V^P(m, \mu_0) \geq V^P(m_H^*, \mu_0) \end{aligned} \right\} \right\}.
 \end{aligned}$$

PROOF OF PROPOSITION 2B: The proof of this proposition is almost identical to that of Proposition 2, and therefore is again kept to a minimum. In particular, conditions (iii)'_B must be satisfied for the exact same reason as condition (iii)', and conditions (v)'_B and (iv)'_B must hold for the same reason as (v'). *Q.E.D.*

COROLLARY 2B: (1) *For $r_s = 0$ and $r_b > 0$, the set of equilibrium contractual outcomes defined in Proposition 2 is the limit as r_b approaches zero of the set of equilibrium contractual outcomes defined in Proposition 2B.*

(2) *For $r_s > 0$ and $r_b > 0$, the set of equilibrium contractual outcomes defined in Proposition 2 is included in the limit as r_b and r_s approach zero of the set of equilibrium contractual outcomes defined in Proposition 2B.*

PROOF OF COROLLARY 2B: (1) Again, for small values of r_b , the standard incentive-compatibility constraints become irrelevant in (iv)'_B and (v)'_B. Then, for $r_s = 0$, each condition in Proposition 2B defines a correspondence that is continuous in r_b , and therefore the first part of the corollary follows immediately.

(2) The second part is shown using a similar argument as in the proof of Corollary 1B. *Q.E.D.*

PROPOSITION 3B: *The contractual outcome $\{m_L^{RS}, m_H^{RS}\}$ described in Proposition 3 can be supported as an XD-equilibrium outcome of the renegotiation game in which both discount rates are arbitrary positive numbers and therefore are included in the limiting set as r_s and r_b go to zero.*

PROOF OF PROPOSITION 3B: The proof of this proposition requires the specification of a set of strategies and equilibrium beliefs which together constitute an XD-equilibrium. Once again the strategies and beliefs described in Proposition 3 can be modified very simply in order to apply to the current proposition. The only necessary modifications to the strategies described in Proposition 3 are that for $i > 3$, the sets $B(m_i, \mu_0)$ and $B(m_i, L)$ be modified to include the cost of delay as follows: $B(m_i, \mu_0) = \{m/\exp(-r_b \Delta)V^P(m, \mu_0) \geq V^P(m_i, \mu_0)\}$ and $B(m_i, L) = \{m/\exp(-r_b \Delta)V(m, L) \geq V(m_i, L)\}$. With these modifications, the equilibrium play of the game is then unchanged but the out-of-equilibrium play of the game is changed to take into account the cost of delay. Q.E.D.

A similar demonstration can be used to show that the XD-equilibrium outcomes specified in Proposition 4 can be supported as XD-equilibrium outcomes of the renegotiation game in which both discount rates are positive, and we therefore omit the proof of this statement.

REFERENCES

- ADMATI, A., AND M. PERRY (1987): "Strategic Delay in Bargaining," *Review of Economic Studies*, 54, 345–364.
- AGHION, P., M. DEWATRIPONT, AND P. REY (1989): "Renegotiation Design under Symmetric Information," M.I.T. Discussion Paper No. 522.
- AKERLOF, G. A. (1970): "The Market of 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 84, 488–500.
- BANKS, J. S., AND J. SOBEL (1987): "Equilibrium Selection in Signalling Games," *Econometrica*, 55, 647–661.
- BEAUDRY, P., AND M. POITEVIN (1990): "Competitive Screening in Financial Markets when Borrowers can Recontract," Cahier de Recherche No. 9035, Département de Sciences Économiques, Université de Montréal.
- BOLTON, P. (1990): "Renegotiation and the Dynamics of Contract Design," *European Economic Review*, 34, 303–310.
- BRENNAN, M., AND A. KRAUS (1987): "Efficient Financing under Asymmetric Information," *Journal of Finance*, 42, 1225–1244.
- CHO, I. (1987): "A Refinement of Sequential Equilibrium," *Econometrica*, 55, 1367–1389.
- CHO, I., AND D. M. KREPS (1987): "Signalling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102, 179–221.
- DEWATRIPONT, M. (1988): "Commitment through Renegotiation-Proof Contracts with Third Parties," *Review of Economic Studies*, 45, 377–390.
- (1989): "Renegotiation and Information Revelation over Time: The Case of Optimal Labor Contracts," *Quarterly Journal of Economics*, 104, 589–619.
- DEWATRIPONT, M., AND E. MASKIN (1990): "Contract Renegotiation in Models of Asymmetric Information," *European Economic Review*, 34, 311–321.
- FUDENBERG, D., AND J. TIROLE (1990): "Moral Hazard and Renegotiation in Agency Contracts," *Econometrica*, 58, 1279–1319.
- (1991): *Game Theory*. Cambridge, Mass.: M.I.T. Press.
- GALE, I. (1986): "Self-Selection and Price Discrimination in a Competitive Dealer Market," Ph.D. Thesis, Chapter 2, Princeton University.
- (1991): "Sequential Trade and Information Revelation," mimeo, University of Wisconsin.
- GALE, I., AND J. E. STIGLITZ (1989): "The Informational Content of Initial Public Offerings," *Journal of Finance*, 44, 469–477.
- GIAMMARINO, R. M., AND T. R. LEWIS (1989): "A Theory of Negotiated Equity Financing," *Review of Financial Studies*, 1, 265–288.
- HART, O., AND J. MOORE (1988): "Incomplete Contracts and Renegotiation," *Econometrica*, 56, 755–785.
- HART, O., AND J. TIROLE (1988): "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55, 509–540.
- HEINKEL, R. (1982): "A Theory of Capital Structure Relevance under Imperfect Information," *Journal of Finance*, 37, 1141–1150.
- HELLWIG, M. (1988): "A Note on the Specification of Interfirm Communication in Insurance Markets with Adverse Selection," *Journal of Economic Theory*, 46, 154–163.

- HOLMSTRÖM, B., AND R. B. MYERSON (1983): "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.
- HOSIOS, A., AND M. PETERS (1988): "Self-Selection with Renegotiation," mimeo, University of Toronto.
- JAYNES, G. D. (1978): "Equilibria in Monopolistically Competitive Insurance Markets," *Journal of Economic Theory*, 19, 394–422.
- LAFFONT, J. J., AND J. TIROLE (1990): "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 57, 597–625.
- MADRIGAL, V., T. TAN, AND S. R. DA COSTA WERLANG (1987): "Support Restrictions and Sequential Equilibria," *Journal of Economic Theory*, 43, 329–334.
- MASKIN, E., AND J. TIROLE (1992): "The Principal-Agent Relationship with an Informed Principal, II: Common Values," *Econometrica*, 60, 1–42.
- MILDE, H., AND J. G. RILEY (1988): "Signaling in Credit Markets," *Quarterly Journal of Economics*, 103, 101–129.
- MIYAZAKI, H. (1985): "The Rat Race and Internal Labour Markets," *Bell Journal of Economics*, 8, 394–418.
- NOLDEKE, G., AND E. VAN DAMME (1990a): "Signalling in a Dynamic Labour Market," *Review of Economic Studies*, 57, 1–23.
- (1990b): "Switching Away from Probability-One Beliefs," Mimeo.
- NOSAL, E. (1991): "Implementing Ex Ante Contracts," Mimeo, University of Waterloo.
- POITEVIN, M. (1989): "Financial Signalling and the 'Deep-Pocket' Argument," *Rand Journal of Economics*, 20, 26–40.
- (1990): "Strategic Financial Signalling," *International Journal of Industrial Organization*, 8, 499–518.
- ROTHSCHILD, M., AND J. E. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, 90, 629–650.
- RUBINSTEIN, A. (1991): "Comments on the Interpretation of Game Theory," *Econometrica*, 59, 909–924.
- SPENCE, M. (1973): "Job Market Signalling," *Quarterly Journal of Economics*, 87, 355–374.
- VINCENT, D. (1989): "Bilateral Monopoly, Non-durable Goods and Dynamic Trading Relationships," mimeo, M.E.D.S., Northwestern University.
- WEISS, A. (1983): "A Sorting-cum-Learning Model of Education," *Journal of Political Economy*, 91, 420–442.