Contract renegotiation: a simple framework and implications for organization theory

PAUL BEAUDRY University of British Columbia
MICHEL POITEVIN Université de Montréal

Abstract. This paper provides a unifying framework for studying renegotiation of contracts in the presence of asymmetric information. We show that interim renegotiation does not constrain the set of contracts attainable with full commitment, regardless of whether renegotiation offers are made by the informed or the uninformed agent. Ex post renegotiation, however, does constrain the set of attainable contracts. These constraints depend on the identity of the agent making the renegotiation offer. We then show how the theory of contract renegotiation can provide insights for organization theory. Specifically, we show how decentralization of decision making can be an optimal response to the threat of ex post renegotiation. Finally, we show that our framework can be used to analyse the trade-off between internal and external markets.

La renégociation de contrat: un cadre conceptuel simple et les implications pour la théorie des organisations. Ce mémoire fournit un cadre conceptuel unificateur pour étudier le renégociation de contrats quand l'information est asymétrique. On montre qu'une renégotiation intérimaire ne contraint pas l'ensemble des contrats réalisables avec un plein engagement, que les offres de renégociation soient faites par l'agent le mieux ou le moins bien informé. La renégociation ex post contraint cependant l'ensemble des contrats réalisables. Ces contraintes dépendent de l'identité de l'agent qui fait l'offre de renégociation. Les auteurs montrent aussi comment la théorie de la renégociation de contrats peut améliorer notre compréhension de certains aspects de la théorie des organisations. En particulier, on montre comment la décentralisation de la prise de décision peut être une réponse optimale à la menace de renégociation ex post. Finalement, les auteurs montrent que leur cadre conceptuel peut être utilisé pour analyser la relation d'équivalence entre marché interne et marché externe.

I. INTRODUCTION

The last two decades have seen the development of the economics of information

We would like to thank Michael Peters for detailed and perceptive comments on an earlier version. We would like to acknowledge financial support from SSHRCC and FCAR.

Canadian Journal of Economics Revue canadienne d'Economique, XXVIII, No. 2
May mai 1995. Printed in Canada Imprimé au Canada

and its application to the realm of contracts. In particular, in the presence of asymmetric information the role played by contracts in coordinating activities expands: not only do contracts govern the terms of exchanges, but they also become vehicles for transmitting information. This dual role has led to new and important predictions: the presence of asymmetric information may lead contracts to dictate inefficient outcomes, since the information-transmitting role of contracts may hinder their traditional role as a means of attaining allocative efficiency. More recently, the standard framework for analysing contractual relations has been questioned. In particular, most of the predictions of contract theory have been developed in environments where agents are assumed to be fully committed to the terms of a contract. In many relevant economic environments, however, full commitment to the terms of a contract may be an unrealistic assumption. Therefore, recent research has explored how the predictions of contract theory are affected when the assumption of full commitment is relaxed.

There are three dimensions in which the issue of commitment can be and has been examined. First, there is the possibility that parties to a contract may not be able to commit to all types of actions or to actions in the distant future. The analysis of this possibility has given rise to the literature on incomplete contracts. Second, there is the possibility that parties to a contract may not be able to make a commitment to obey the contract at any point during the relationship. This has given rise to the literature on self-enforcing contracts. Third, there is the possibility that parties to a contract may not be able to make a commitment not to change the terms of a contract, even though in certain circumstances it may be in their mutual benefit to do so. The third issue, which is referred to as the possibility of renegotiation, is the subject of this paper. At first glance, it may seem surprising that renegotiation is even an issue in contract theory, since optimal contracts should leave no room for mutually beneficial modifications. Although this intuition is right when the contract is examined at the time of original signing, it no longer holds once the contract is being carried out and information is revealed. In effect, once information is revealed, the objectives of the contracting parties may become more precise, and therefore mutually beneficial changes to the contract may exist. If these changes to the contract are initially expected to occur after information is actually transmitted, however, they may change the initial incentive vis-à-vis the manner in which information is transmitted.

The approach we favour for analysing renegotiation draws heavily on the work of Holmström and Myerson (1983) in that we search to isolate how the possibility of renegotiating a contract after the arrival of information affects attainable outcomes in private-information environments. This contrasts slightly with mainstream literature on renegotiation that has examined primarily situations where there are repeated interactions between information revelation and implemented actions (see, e.g., Dewatripont 1988; Hart and Tirole 1988; Laffont and Tirole 1990). The advantage of our approach is that it sufficiently simplifies the problem to allow for a thorough examination of different renegotiation processes. Moreover, once this more elementary framework is understood, it becomes easier to analyse the multiperiod problems that are predominant in the literature.

This paper is structured as follows: in section II we present a general framework for analysing the effects of renegotiation in a hidden-information environment. The simplicity of the framework permits us to examine how the timing of renegotiation as well as the identity of the proposer affects the equilibrium outcome. This section concludes with a discussion of the extension of our results to adverse selection, moral hazard, and multi-period problems. Section III, which is more suggestive, discusses how the theory of renegotiation can be insightful for the theory of institutional design. In particular, we argue that the issue of renegotiation offers new and important insights with respect to questions of centralization versus decentralization in organizations. Finally, section IV offers concluding comments.

II. THE THEORY

Let us consider a situation where two individuals want to coordinate a set of actions. The actions of interest are represented by a vector $a = \{a_1, a_2\} \in \mathcal{A}$, where $\mathcal{A}$ is a compact and convex subset of R^2 and where action a_i is associated with agent $i = 1, 2$. The environment is uncertain and the possible states of nature are indexed by $t \in \mathcal{T} = \{1, \ldots, T\}$ with the prior probability of state t's being denoted $p_0(t)$. The preferences of each player depend on the actions a and the state of nature t and are represented by the utility function $U(a,t)$ for player 1, and $V(a,t)$ for player 2. The reservation utility for players 1 and 2 is $U(0,t)$ and $V(0,t)$, respectively.

In order to fix ideas, it is helpful to think of the set-up as one that examines contingent trade relationships where one of the two elements of a is a transfer payment. For example, (a) in a labour market relationship, a_1 may be the wage paid, a_2 the number of hours worked, and t the state of demand; (b) in a goods market relationship a_1 may be the price paid for the good, a_2 the quantity of the good that is transacted, and t the quality of the good; and (c) in a financial market relationship, a_1 may be the level of investment into a project, a_2 the amount paid to the financier when the project is successful, and t the probability that the project is successful.

In order to ensure that the contractual problem is well behaved, we assume that the sign of V_{a_i} is opposite to the sign of U_{a_i}, that $V(\cdot, t)$ and $U(\cdot, t)$ are continuously differentiable and concave in a for all t, and that for all $t > t'$ and $a \neq \{0, 0\}$, $V(a,t) > V(a,t')$. The first assumption specifies that the two players have opposite preferences with respect to the vector of actions; that is, if player 1 prefers more a_i, then player 2 prefers less; the first two assumptions imply that Pareto optimal trades exist for any state of nature. Finally, the last assumption implies that a higher index for t corresponds to a better state of nature for player 2.

In this environment, an allocation μ is defined as a vector of action-pairs with one element for each state of nature, namely, $\mu = \{\mu_t\}_{t=1}^{T}$ with $\mu_t \in \mathcal{A}$. Our interest is in characterizing allocations that arise when the two individuals can write contracts knowing that only one player, specifically player 1, observes the state of nature. The allocations that the players will be able to obtain depend on the type of contract that can be written and on the process by which the contract is chosen and carried out. In order to maintain a balance between ease of presentation and generality, we let contracts belong to the following class.

DEFINITION 1. *A contract c (or mechanism) is defined by*
1. *A menu of actions* $m(c) = \{a^n\}_{n=1}^N$, *where* $a \in \mathcal{A}$ *from which player 1 (the informed agent) is required to choose after she learns her private information.*
2. *The penalties imposed on each individual in the case where the selected actions are not carried out. In general, we assume that these penalties are infinite.*

A contract is therefore a game form to be played by the two players. The game form has three important features. First, we have allowed for mechanisms other than direct revelation mechanisms, since in the presence of renegotiation the latter may be restrictive. We do, however, restrict the space of messages that can be sent during the relationship. Specifically, we assume that only player 1 can send a message, since it is she who has the private information. Second, we have restricted our attention to contracts that specify only choices of deterministic outcomes. Third, we allow contracts to be specific about the pay-offs associated with different actions before the actions are actually carried out, namely, the specification of infinite penalties imposed on an agent choosing an action different from that prescribed by the contract. This effectively makes the contract enforceable. In summary, a contract specifies a set of possible actions that can be undertaken, some stage of communication between the two players in which they coordinate on a certain course of action, and finally the implication of executing different actions. The purpose of the analysis is to characterize allocations that result under various assumptions regarding the commitment possibilities available to the two players when they are executing a contract.

1. Complete commitment contracts

Before introducing renegotiation, it is useful to review the benchmark case in which both players can make a commitment to honour the terms of a contract and therefore cannot renegotiate it under any circumstances.[1] A contract is said to be a complete commitment contract if it is supported along the equilibrium path of the following game called the commitment game. Renegotiation will later be introduced through modifications of this basic game.

1. In the first stage, player 1 proposes a contract c_0 to player 2.
2. In the second stage, player 2 accepts or rejects the contract offer. If it is rejected, the game ends and both players receive their reservation utility.
3. In the third stage (if reached), player 1 observes the state of nature (her type) t.
4. In the fourth stage, the contract is carried out; that is, player 1 selects an element $s_0 \in m(c_0)$, and then both players choose their actions as prescribed by the element $m(c_0)$.

1 For a more thorough introduction to this class of problem, see Hart (1983). Note, however, that our set-up is slightly different from that examined by Hart, since we allow for that state of nature to affect both players' preferences; that is, we examine the case of common values.

This commitment game has two important features. First, the environment we have chosen is one of *hidden information*: the contract is signed with the two agents' having the same information structure, but it is carried out just after player 1 has privately observed the state of nature. Below, we shall discuss how the case of pre-contractual private information (adverse selection) can be handled. Second, once a contract c_0 has been chosen, there is no possibility of modifying it.

The strategy of player 1 is represented by a tuple $\sigma_1 = \{\tilde{c}_0, \tilde{s}_0(c_0, t)\}$, where $\tilde{c}_0$ is an initial contract offer and $\tilde{s}_0(\cdot)$ represents player 1's decision rule regarding the choice of an element in $m(c_0)$. The decision by player 1 not to make an offer is denoted by $\emptyset$. The strategy of player 2, σ_2, is represented by the function $\tilde{d}_0(c_0) \in \{0, 1\}$, which represents the decision rule concerning the acceptance or rejection of the initial contract offer with $\tilde{d}_0(c_0) = 1$ if the contract c_0 is accepted and 0 otherwise.

Given this game, a *Perfect Bayesian Equilibrium* (PBE) is a pair of strategies σ_1 and σ_2 that are best replies to one another given beliefs in every contingency in which agents are forced to make a choice, and a pair of beliefs that are updated using Bayes rule whenever possible.[2]

The following proposition provides a characterization of equilibrium allocations of the commitment game. In the proof of the proposition, we give the equilibrium contract as well as strategies and beliefs that support a characterized allocation as a PBE outcome. (Note: all proofs are provided in the appendix.)

PROPOSITION 1. *An allocation, $\mu^c = \{\mu_1^c, \mu_2^c, \ldots, \mu_T^c\}$, is an equilibrium allocation of the commitment game if and only if it is a solution to the following maximization problem.*

$$\max_{\{\mu_t\}_{t=1}^T} \sum_{t=1}^{T} p_0(t) U(\mu_t, t)$$

$$\text{s/t} \quad (i) \quad \sum_{t=1}^{T} p_0(t) V(\mu_t, t) \geqq \sum_{t=1}^{T} p_0(t) V(0, t)$$

$$(ii) \quad U(\mu_t, t) \geqq U(\mu_t, t) \; \forall t, t' \in \mathcal{T}. \tag{1}$$

Proposition 1 states the equivalence between equilibrium allocations and the solutions to a well-defined maximization problem.[3] The equilibrium characterization corresponds to player 1's preferred allocation among the set of allocations satisfying her incentive-compatibility constraints (ii) and player 2's participtation constraint (i). An important property of the allocation μ^c is that it is interim efficient; that is, no other allocation can increase the utility of one type of player

2 See Fudenberg and Tirole (1991) for a precise definition of a Perfect Bayesian Equilibrium.
3 Under our assumptions the constraint set is closed and therefore there exists a solution to this maximization problem.

1 without decreasing the utility of another type or player 2 and/or violating the incentive-compatibility constraints.[4]

When the two players can commit to the terms of the contract, the equilibrium contract results in an interim-efficient allocation. In general interim-efficient allocations are not ex post efficient, since they involve carrying out pairs of actions that are Pareto dominated conditional on the state of nature.[5] This can easily seen from the following first-order condition associated with maximization (1).

$$\frac{\{p_0(t) + \Sigma_{t' \neq t}\gamma(t, t')\}U_{a_1}(\mu_t^c, t) - \Sigma_{t' \neq t}\gamma(t', t)U_{a_1}(\mu_t^c, t')}{\{p_0(t) + \Sigma_{t' \neq t}\gamma(t, t')\}U_{a_2}(\mu_t^c, t) - \Sigma_{t' \neq t}\gamma(t', t)U_{a_2}(\mu_t^c, t')} = \frac{V_{a_1}(\mu_t^c, t)}{V_{a_2}(\mu_t^c, t)}.$$

In the above first-order condition, $\gamma(t', t)$ is the multiplier associated with the incentive-compatibility constraint stating that in state t' player 1 must prefer the allocation $\mu_{t'}^c$ to the allocation μ_t^c. In most interesting applications, at least one multiplier γ will be different from zero and the ex post efficient condition will not be satisfied for at least one state.

As the above argument illustrates, equilibrium allocations of the commitment game generally prescribe ex post distortions as a result of the self-selection constraints. The allocative role of contracts is therefore impeded by its role as a vehicle for transmitting information. The presence of these distortions suggests that parties may want to renegotiate the contract to attain a mutually preferred allocation once the private information has been revealed. Therefore the prediction that the allocation μ^c describes the behaviour of players in a contracting relationship is valid only if the environment allows players to make a commitment never to renegotiate a contract once it is signed. This is a fairly strong requirement, and therefore it seems relevant also to determine allocations that are likely to arise when such commitment is not possible and players cannot prevent renegotiations.

In relation to our commitment game, there are two potential instances in which players may want to renegotiate a contract. First, the simple arrival of information may create some opportunity for renegotiation. Players may therefore want to renegotiate immediately after player 1 has observed the state of nature but *before* she chooses which element of the menu she it to take. In this case renegotiation would occur after stage 3 but before stage 4. Second, the actual selection by player 1 of an element in the menu of the outstanding contract may also create some opportunity for renegotiation. In this case players would renegotiate *after* player 1 has selected an element from the menu but before the actions are actually executed, namely, after stage 4 but before stage 5. The first type of renegotiation will be referred to as interim renegotation, while we call the second type ex post renegotiation. The following two subsections study the implications of each of these two potential occurrences of renegotiation.

4 See Holmström and Myerson (1983) or Maskin and Tirole (1992) for a precise definition of interim (incentive) efficiency.

5 An action-pair μ_t is ex post efficient in state t if $U_{a_1}(\mu_t, t)/U_{a_2}(\mu_t, t) = V_{a_1}(\mu_t, t)/V_{a_2}(\mu_t, t)$.

2. Interim renegotiation-proof contracts

Suppose the two players have signed a contract c_0. As shown above, this contract generally prescribes ex post distortions to resolve ex ante incentives; once player 1 has privately observed the state of nature, however, ex ante incentives are no longer relevant, and the two players may wish to eliminate ex post distortions to improve their utility conditional on the realization of the state of nature. Therefore, if the two players are not committed to the initial contract, they may renegotiate it.

Interim renegotiation-proof contracts are characterized as contracts that can be supported along the equilibrium path of the following game called the interim-renegotiation game. Interim renegotiation is introduced by modifying our benchmark game as follows.

1. In the first stage, player 1 proposes a contract c_0 to player 2.
2. In the second stage, player 2 accepts or rejects the contract offer. If it is rejected, the game ends and both players receive their reservation utility.
3. In the third stage (if reached), player 1 observes the state of nature t.
   - 3.1 In stage 3.1, player i proposes a contract c_i to player j.
   - 3.2 In stage 3.2, player j accepts or rejects the contract offer. If it is rejected, the contract c_0 is the outstanding contract; if it is accepted, the contract c_i becomes the outstanding contract;
4. In the fourth stage, the outstanding contract c is carried out; that is, player 1 selects an element $s \in m(c)$, and then both players choose their actions as prescribed by the element $m(c)$.

We refer to this new game as the interim-renegotiation game. This game has two important aspects. First, it allows for either player to make the renegotiation offer. We shall study in turn the cases in which either player 1 or player 2 is making the renegotiation offer. Second, the interim-renegotiation game allows for the possibility of renegotiation just after player 1 has privately observed the state of nature. Therefore, regardless of the identity of the player offering the renegotiation, it takes place under asymmetric information.

In the case in which player 1 makes the renegotiation offer, the strategy of player 1 is $\Omega_1 = \{\tilde{c}_0, \tilde{c}_1(c_0, t), \tilde{s}(c_0, c_1, d_1, t)\}$, where $\tilde{c}_0$ is the initial contract offer, $\tilde{c}_1(\cdot)$ is her decision rule regarding the renegotiation offer, and $\tilde{s}(\cdot)$ has the same interpretation as it does in the commitment game but is contingent on the complete history of the game; the strategy of player 2 is $\Omega_2 = \{\tilde{d}_0(c_0), \tilde{d}_1(c_0, c_1)\}$, where $\tilde{d}_1(\cdot)$ is his decision rule concerning his acceptance decision of player 1's renegotiation proposal, and $\tilde{d}_0(\cdot)$ has the same interpretation as it does in the commitment game but is contingent on the complete history of the game. When player 2 makes the renegotiation offer, the strategy of player 1 is $\Omega_1 = \{\tilde{c}_0, \tilde{d}_0(c_0, c_2, t), \tilde{s}(c_0, c_2, d_2, t)\}$, where $\tilde{d}_2(\cdot)$ is her decision rule concerning her acceptance decision of player 2's renegotiation proposal; the strategy of player 2 is $\Omega_2 = \{\tilde{d}_0(c_0), \tilde{c}_2(c_0)\}$, where $\tilde{c}_2(\cdot)$ is his decision rule regarding the renegotiation offer.

If player 1 makes the renegotiation offer, the beliefs of player 2 are updated after stage 3.1 and are denoted $p_2(t|c_0, c_1)$. If player 2 makes the renegotiation offer, the beliefs of player 2 remain constant throughout the game.

To characterize equilibrium allocations that arise when interim renegotiation is possible, it is unfortunately not interesting or appropriate simply to characterize the whole set of equilibrium allocations of the interim-renegotiation game. For example, any equilibrium allocation of the commitment game supported by a contract c^c can be supported as an equilibrium allocation of the interim-renegotiation game when player 1 makes the renegotiation offer as follows: (a) player 1 initially offers a contract c_0 that specifies the trivial null menu ($m(c_0) = \{(0,0)\}$); (b) the contract is accepted by player 2; (c) player 1 then proposes a renegotiation consisting of the contract $c_1 = c^c$; (d) the proposed renegotiation is again accepted by player 2. In this example, renegotiation has no effect on equilibrium allocations, since players use the last period of the game to commit effectively to ex post distortions as they do in the commitment case. Therefore, in order to characterize allocations that are robust to renegotiation, it is desirable to restrict attention to the set of equilibrium allocations that can be supported by equilibrium strategies that do not involve any renegotiation. Such a selection of equilibria ensures that the last stage of the game is not used arbitrarily to support ex post distortions. This 'equilibrium selection' approach to examining the implications of renegotiation is similar to Holmström and Myerson's (1983) work on durable mechanisms.[6] Allocations that are robust to the introduction of interim renegotiation are called interim-renegotiation-proof and are defined as follows.

DEFINITION 2

- *An equilibrium allocation μ_1^{ir} is* I1-renegotiation-proof *if (1) it is an equilibrium allocation of the interim-renegotiation game in which it is player 1 that proposes the renegotiation in stage 3.1, and (2) it can be supported by strategies where player 1 makes no attempt to renegotiate the initial contract along the equilibrium path.*
- *An equilibrium allocation μ_2^{ir} is* I2-renegotiation-proof *if (1) it is an equilibrium allocation of the interim-renegotiation game in which it is player 2 that proposes the renegotiation in stage 3.1, and (2) it can be supported by strategies where player 2 makes no attempt to renegotiate the initial contract along the equilibrium path.*

This definition asserts that an allocation μ_i^{ir} is Ii-renegotiation-proof if it can be supported by a contract c_i^{ir}, where $m(c_i^{ir}) = \{\mu_{it}^{ir}\}_{t=1}^{T}$, and c_i^{ir} is being offered and accepted in stages 1 and 2, respectively, but is not being renegotiation in stage 3. Hence for any Ii-renegotiation-proof allocation, there must exist equilibrium strategies and beliefs for the interim-renegotiation game such that, when its supporting contract c_i^{ir} is offered and accepted in stages 1 and 2, players will not renegotiate it, even though it is possible to do so in stage 3.1. This definition eliminates as

6 This approach can be seen as a simple alternative to formally modelling the renegotiation process as an infinite game in which there is never a last stage that can be used to make a commitment not to renegotiate. At the end of this section we discuss at greater length the parallel between our proposed approach and that of an infinite game.

interim-renegotiation-proof allocations those that can be obtained only by being offered in the renegotiation stage 3.1 following the trivial contract offer in stage 1.

Note that this definition includes two types of interim-renegotiation-proofness, since allocations satisfying this property may potentially depend on the identity of the player allowed to make the renegotiation offer. The following proposition characterizes the equilibrium interim-renegotiation-proof allocations of the interim-renegotiation game for the cases in which player 1 or player 2 is making the renegotiation offer.

PROPOSITION 2. *Regardless of the identity of the player making the renegotiation offer in stage 3.1, an allocation* μ_i^{ir} *is an equilibrium Ii-renegotiation-proof allocation if and only if it is attainable with full commitment.*

Proposition 2 captures the main forces behind the results obtained by Homström and Myerson (1983), Dewatripont (1988), Maskin and Tirole (1992), and Nosal (1991); that is, the possibility of interim renegotiation does not affect equilibrium allocations. In fact, these authors have shown, using slightly different approaches, that interim-renegotiation-proof allocations are interim efficient. Since interim-efficient allocations are those chosen in the absence of renegotiation, it follows that interim renegotiation does not affect the set of allocations that are chosen.[7]

3. Ex-post-renegotiation-proof contracts

In the last subsection we showed that allowing for players to renegotiate a contract before player 1 chooses an item in the menu of the outstanding contract has no effect on the equilibrium allocations. Hence, despite interim renegotiation, equilibrium allocations still exhibit some distortions. This type of renegotiation assumes that players are committed to execute the action prescribed by player 1's choice in the menu; however, the possibility that players can renegotiate just after player 1 has chosen an item in the menu can possibly change equilibrium allocations. The next step therefore is to consider whether allowing players to renegotiate after player 1 has communicated her choice to player 2 has any effect on equilibrium allocations.

As with interim renegotiation, the effect of ex post renegotiation can be assessed by introducing an appropriate modification of the commitment game. We now describe the ex post renegotiation game.

1. In the first stage, player 1 proposes a contract c_0 to player 2.
2. In the second stage, player 2 accepts or rejects the contract offer. If it is rejected, the game ends and both players receive their reservation utility.
3. In the third stage (if reached), player 1 observes the state of nature t.
4. In the fourth stage, the contract c_0 is carried out; that is, player 1 selects an element $s_0 \in m(c_0)$.

   4.1 In stage 4.1, player i proposes a contract c_i to player j.

7 The basic idea behind this result was first uncovered by Milgrom and Stokey (1981). We thank Michael Peters for pointing out this fact.

4.2 In stage 4.2, player j accepts or rejects the contract offer. If it is rejected, the contract c_0 remains the outstanding contract;

4.3 If c_i is accepted, it becomes the outstanding contract and player 1 selects an element $s_i \in m(c_i)$, and then both players choose their actions as prescribed by the element $m(c_i)$.

As with the interim-renegotiation game, the ex post renegotiation game allows for either player to make the renegotiation offer, and we shall study the two cases in turn; also, renegotiation is taking place under asymmetric information, since player 1 has privately observed the state of nature before the renegotiation offer. There is a major difference between the interim-renegotiation game and the ex post-renegotiation game, however, and it is related to the status quo position following the rejection of a renegotiation. In the latter game, the informed player can take a costly action (a choice from a menu restricting her choice set) which can be used to signal her private information before she offers to renegotiate or responds to an offer to renegotiate. With interim renegotiation, the offer to renegotiate is simply cheap talk.

In the case in which player 1 makes the renegotiation offer, the strategy of player 1 is $\Phi_1 = \{\tilde{c}_0, \tilde{s}_0(c_0, t), \tilde{c}_1(c_0, s_0, t), \tilde{s}_1(c_0, s_0, c_1, t)\}$, where $\tilde{c}_0$ is the initial contract offer, $\tilde{c}_1(\cdot)$ is her decision rule regarding the renegotiation offer, and $\tilde{s}_0(\cdot)$ and $\tilde{s}_1(\cdot)$ are her decision rules concerning her choice in the menu of c_0 and c_1 respectively; the strategy of player 2 is $\Phi_2 = \{\tilde{d}_0(c_0), \tilde{d}_1(c_0, s_0, c_1)\}$, where $\tilde{d}_1(\cdot)$ is his decision rule concerning his acceptance decision of player 1's renegotiation proposal, and $\tilde{d}_0(\cdot)$ has the same interpretation as it does in the commitment game but is contingent on the complete history of the game. When player 2 makes the renegotiation offer, the strategy of player 1 is $\Phi_1 = \{\tilde{c}_0, \tilde{s}_0(c_0, t), \tilde{d}_2(c_0, s_0, c_2, t), \tilde{s}_2(c_0, s_0, c_2, t)\}$, where $\tilde{d}_2(\cdot)$ is her decision rule concerning her acceptance decision of player 2's renegotiation proposal; the strategy of player 2 is $\Phi_2 = \{\tilde{d}_0(c_0), \tilde{c}_2(c_0, s_0)\}$, where $\tilde{c}_2(\cdot)$ is his decision rule regarding the renegotiation offer.

If player 1 makes the renegotiation offer, the beliefs of player 2 are updated after stage 4.1 and are denoted $p_2(t|c_0, s_0, c_1)$. If player 2 makes the renegotiation offer, the beliefs of player 2 are updated after he observes player 1's choice $s_0 \in m(c_0)$ in stage 4 and are denoted by $p_2(t|c_0, s_0)$.

Once again, the full effect of ex post renegotiation cannot be properly understood by only looking at the equilibrium allocations of the ex-post-renegotiation game. As with interim renegotiation, players can always use the last stage of offers to commit to ex post inefficiencies that would potentially be renegotiated away if another round of renegotiation was allowed. Therefore, we restrict attention to equilibrium allocations that are ex-post-renegotiation-proof in that they can be supported by a contract that is not renegotiated along the equilibrium path in stage 4.1.

DEFINITION 3

- *An equilibrium allocation* μ_1^{pr} *is* P1-renegotiation-proof *if (1) it is an equilibrium allocation of the ex-post-renegotiation game where it is player 1 that proposes* *the renegotiation at stage 4.1; (2) the renegotiation offer consists of a contract whose menu has a single element; and (3) it can be supported by strategies wherein player 1 does not attempt to renegotiation the initial contract along the equilibrium path.*

- *An allocation μ_2^{pr} is supported by a status quo contract offered and accepted in stages 1 and 2 respectively is* P2-renegotiation-proof *if it can be supported by strategies wherein player 2 does not attempt to renegotiate the initial contract along the equilibrium path of the subgame of the ex-post-renegotiation game starting with stage 3.*

This definition asserts that an allocation μ_i^{pr} is Pi-renegotiation-proof if there exist equilibrium strategies and beliefs for the ex-post-renegotiation game such that, when its supporting contract c_i^{pr} is offered and accepted in stages 1 and 2, players will not renegotiate it, even though it is possible to do so.

Our approach favours the use of a finite renegotiation game to capture the effects of renegotiation; however, this is not as straightforward an exercise as it may seem. To avoid the difficulties associated with a finite game (and the commitment possibilities it allows) we introduce some particularities in the two definitions of renegotiation-proofness. We now discuss these particularities.

The definition of P1-renegotiation-proofness includes the additional provision that player 1 be restricted to only offer contracts whose menu consists of a single element. This condition is necessary to guarantee existence of P1-renegotiation-proof allocations. At the end of this section we motivate this assumption by showing how the definition of P1-renegotiation-proofness is equilvalent to modelling an infinite renegotiation game where a menu of contracts can be offered at each stage.

The definition of P2-renegotiation-proofness only considers the subgames of the ex-post-renegotiation game starting with stage 3. Each such subgame is characterized by a (status quo) contract that has been offered and accepted in stages 1 and 2, respectively. Given a status quo contract, we are interested in whether this contract can be supported by strategies and beliefs of the continuation subgame such that it is not renegotiated along the equilibrium path. We cannot define P2-renegotiation-proof allocations as equilibrium allocations of the ex-post-renegotiation game, since it may be shown that the allocation μ^c is the only equilibrium allocation of that game. This result is a direct consequence of the ex-post-renegotiation game's being finite, however, and is therefore not very interesting. This explains why we adopt the alternative approach of considering only the renegotiation subgames (starting with stage 3).[8]

There are two concepts of ex-post-renegotiation-proofness, since allocations satisfying this property may potentially depend on the identity of the player allowed to make the renegotiation offer. Propositions 3 and 4 give necessary and sufficient conditions, respectively, for P1-renegotiation-proofness, while proposition 5 does the same for P2-renegotiation-proofness.

8 The same approach has been used by Maskin and Tirole (1992).

PROPOSITION 3. *An equilibrium P1-renegotiation-proof allocation, μ_1^{pr}, must satisfy the following conditions.*

$$(i)\ \sum_{t=1}^{T} p_0(t)V(\mu_{1t}^{pr}, t) \geqq \sum_{t=1}^{T} p_0(t)V(0, t)$$

$$(ii)\ U(\mu_{1t}^{pr}, t) \geqq \max_{\mu} \{U(\mu, t),\ \text{s/t}\ V(\mu, t'') \geqq V(\mu_{1t'}^{pr}, t'') \qquad \forall t'' \in \mathcal{T}\}$$

$$\forall t, t' \in \mathcal{T}.$$

Proposition 3 provides a set of necessary conditions for an allocation to be P1-renegotiation-proof. In proposition 4 we characterize one such allocation as an equilibrium allocation of the ex-post-renegotiation game, namely, the allocation that yields player 1 the highest expected utility.

PROPOSITION 4. *If an allocation μ_1^{pr} is a solution to the following maximization problem, then it is an equilibrium P1-renegotiation-proof allocation.*

$$\max_{\{\mu_{1t}\}_{t=1}^{T}} \sum_{t=1}^{T} p_0(t)U(\mu_{1t}, t)$$

$$\text{s/t}\quad (i)\ \sum_{t=1}^{T} p_0(t)V(\mu_{1t}, t) \geqq \sum_{t=1}^{T} p_0(t)V(0, t)$$

$$(ii)\ U(\mu_{1t}, t) \geqq \max_{\mu} \{U(\mu, t),\ \text{s/t}\ V(\mu, t'') \geqq V(\mu_{1t'}, t'') \qquad \forall t'' \in \mathcal{T}\}$$

$$\forall t, t' \in \mathcal{T}. \tag{2}$$

The set of constraints (ii) in propositions 3 and 4 clearly illustrates the effect of ex post renegotiation on the equilibrium contract. These constraints are more stringent than the usual incentive-compatibility constraints, and therefore they represent generalized incentive-compatibility contraints that incorportate the possibility of ex post renegotiation. The set of constraints in the maximization problem of each constraint (ii) implies that, given a status quo position $\mu_{1t'}$, player 2 will only accept those renegotiation offers that increase his utility regardless of his beliefs. Suppose constraint (ii) is satisfied at a status quo position $\mu_{1t'}$. For any offer that player 1 prefers to $\mu_{1t'}$, there exists a belief for player 2 such that he is worse off under the new offer than under the status quo position. When assigned with this belief, player 2 simply rejects the offer of player 1. The renegotiation offers satisfying the constraint set (ii) are referred to as 'surely-acceptable renegotiations.' With this interpretation, the constraints (ii) state that it cannot be possible for any type t to increase her utility by selecting the equilibrium element of any type t' and then offering a surely acceptable renegotiation. It is in this sense that the constraints (ii) represent generalized incentive-compatibility constraints.[9] Therefore ex post renegotiation, as opposed to interim renegotiation, generally affects the equilibrium allocations and therefore reduces player 1's expected utility. Ex post distortions can still arise in the presence of ex post renegotiation, however, when it is the informed player that proposes the renegotiation.[10]

Before analysing the case of P2-renegotiation-proofness, we would like to compare the above results with results that have obtained when the renegotiation process is formally modelled as an infinite game. Beaudry and Poitevin (1993) study a contracting model with adverse selection. The renegotiation process is an infinite repetition of stages 4.1 to 4.3 of the ex-post-renegotiation game. The game basically ends when a renegotiation offer is rejected by player 2. It is shown that the effects of renegotiation are captured entirely by the constraints (ii) of proposition 3 in the sense that all individually rational allocations that satisfy these constraints can be supported as PBE outcomes of the infinite-renegotiation game. The basic intuition for this result is that it is not rational for player 2 to reject a renegotiation offer that increases his utility with respect to the status quo outcome regardless of his beliefs. Hence constraints (ii) of proposition 3 must be satisfied. Therefore the approach we have adopted in this article yields similar results to those of an infinite-renegotiation game.

We can now explain at length why, to guarantee existence, the definition of P1-renegotiation-proofness must restrict player 1's renegotiation contract offer to a single element. Consider a two-state example in which a candidate equilibrium allocation μ^* is ex post efficient for the two states and satisfies constraints (i) and (ii) of proposition 3. For some preference configurations, it is the case that (1) state 1's outcome, μ_1^*, is not ex post efficient conditional on state 2 – that is, $U_{a_1}(\mu_1^*, 2)/U_{a_2}(\mu_1^*, 2) \neq V_{a_1}(\mu_1^*, 2)/V_{a_2}(\mu_1^*, 2)$; (2) there exists an outcome $\hat{\mu}$ such that $U(\hat{\mu}, 2) > U(\mu_1^*, 2)$, $V(\hat{\mu}, 2) > V(\mu_1^*, 2)$, and $U(\hat{\mu}, 1) < U(\mu_1^*, 1)$. Suppose that player 1 selects the element μ_1^* of the initial contract and that she follows with offering a contract $\hat{c}_1$ with $m(\hat{c}_1) = \{\mu_1^*, \hat{\mu}\}$. It is a weakly dominant strategy for player 2 to accept this offer, since he cannot lose regardless of his beliefs and he strictly gains for all beliefs putting a positive weight on state 2. Player 2 then accepts this offer with the consequence that the allocation μ^* is not P1-renegotiation-proof. A similar argument could be applied to all candidate P1-renegotiation-proof allocations, and thus none would exist. But this argument is valid only in a finite game. In an infinite game, accepting the contract $\hat{c}_1$ may not be a weakly dominant strategy for player 2, since his pay-off from doing so depends on the resolution of the future stages of the game. This problem disappears if player 1 is restricted to offering a contract whose menu consists of a single element. In the example above, any contract yielding the outcome $\hat{\mu}$ can be rejected in the

9 Note that, under our assumptions, the constraint set of the maximization problem (2) is closed and therefore existence of P1-renegotiation-proof allocations is guaranteed.

10 In the special case of private values, that is, when player 2's preferences are independent of t, it can easily be seen from constraint (ii) that P1-renegotiation-proofness implies ex post efficiency.

belief that it was offered by a type 1.[11] This motivates our assumption regarding the contracts that constitute a valid renegotiation offer by player 1.

We now turn to the analysis of P2-renegotiation-proofness.

PROPOSITION 5. *An allocation μ_2^{pr} is an equilibrium P2-renegotiation-proof allocation if and only if it satisfies the following conditions.*

$$(i)\quad U(\mu_{2t}, t) \geqq U(\mu_{2t'}, t)\ \forall t, t' \in \mathcal{T}$$

$$(ii)\quad \Sigma_{\tau\in\mathcal{T}(\mu_{2t})}\, p_0(\tau)V(\mu_{2t}, \tau) \geqq$$

$$\left\{ \max_{\{\mu_\tau\}_{\tau\in\mathcal{T}(\mu_{2t})}} \Sigma_{\tau\in\mathcal{T}(\mu_{2t})} p_0(\tau)V(\mu_\tau, \tau) \right.$$

$$\text{s/t}\quad \left.\begin{array}{l} U(\mu_\tau, t') \geqq U(\mu_{2t}, t')\ \forall t' \in \mathcal{T}(\mu_{2t}) \\ U(\mu_\tau, \tau) \geqq U(\mu_{\tau'}, \tau)\ \forall \tau, \tau' \in \mathcal{T}(\mu_{2t}) \end{array}\right\} \forall t \in \mathcal{T},$$

where $\mathcal{T}(\mu_{2t}) = \{\tau \in \mathcal{T}\,|\mu_{2\tau} = \mu_{2t}\}$.

The important element in proposition 5 is the set of constraints (ii), which captures the effects of ex post renegotiation by player 2.[12] These constraints imply that, conditional on the information revealed by the choice of an element in the menu of the initial contract offer, player 2 cannot increase its expected utility without reducing the utility of one type of player 1 in the support of his revised beliefs.[13] In particular, these constraints state that separating allocations must be ex post efficient to be P2-renegotiation-proof; that is, any state that is uniquely identified in equilibrium must correspond to an outcome that is ex post efficient. We should note that renegotiation lead by the uninformed agent can result in an allocation that fails to separate the different types, that is, a pooling allocation. This will be the case if no ex post efficient allocation is incentive compatible.[14]

We know from the commitment case that separation with ex post distortion generally characterizes optimal contracts, and therefore it is clear that ex post renegotiation by the uninformed party can reduce the ex ante potential gains from trade. The reason why ex post renegotiation led by the uninformed player imposes efficiency on separating outcomes is that there is nothing that stops player 2 in stage 4.1 from offering an efficient solution once a state is identified. In contrast, when the ex post renegotiation is led by the informed player, it is still possible for a separating allocation to be ex post inefficient. In this case player 1 cannot

11 Note that $V(\hat{\mu}, 1) < V(\mu_1^*, 1)$, since μ^* satisfies constraints (ii) of proposition 3.

12 Forges (1990) gives a definition of ex-post-renegotiation-proofness that is similar to P2-renegotiation-proofness.

13 Under our assumptions the constraint set is closed, and therefore there exists a solution to this maximization problem.

14 It is also worth noting that in the case of private values, P2-renegotiation-proof allocations are always separating and therefore ex post efficient.

try to renegotiate to an efficient outcome knowing that any such offer would be interpreted as signal of a different state of nature and would therefore be rejected.

Proposition 5 gives no prediction as to which allocation will emerge in an economic environment in which player 2 can make renegotiation offers. For completeness and further discussions we give a characterization of one allocation that is likely to emerge when player 1 has all the initial bargaining power. This is the P2-renegotiation-proof allocation that maximizes the ex ante welfare of player 1 subject to a participation constraint for player 2. This allocation is the solution to the following maximization problem.

$$\max_{\{\mu_{2t}\}_{t=1}^{T}} \sum_{t=1}^{T} p_0(t)U(\mu_{2t}, t)$$

$$(i) \quad \sum_{t=1}^{T} p_0(t)V(\mu_{2t}, t) \geqq \sum_{t=1}^{T} p_0(t)V(0, t)$$

$$(ii) \quad U(\mu_{2t}, t) \geqq U(\mu_{2t'}, t)\ \forall t, t' \in \mathcal{T}$$

$$(iii) \quad \sum_{\tau \in \mathcal{T}(\mu_{2t})} p_0(\tau)V(\mu_{2t}, \tau)$$

$$\geqq \left\{ \max_{\{\mu_\tau\}_{\tau\in\mathcal{T}(\mu_{2t})}} \sum_{\tau\in\mathcal{T}(\mu_{2t})} p_0(\tau)V(\mu_\tau, \tau) \right.$$

$$\text{s/t} \quad \left. \begin{array}{l} U(\mu_\tau, t') \geqq U(\mu_{2t}, t')\ \forall t' \in \mathcal{T}(\mu_{2t}) \\ U(\mu_\tau, \tau) \geqq U(\mu_{\tau'}, \tau)\ \forall \tau, \tau' \in \mathcal{T}(\mu_{2t}) \end{array} \right\} \forall t \in \mathcal{T}, \qquad (3)$$

where $\mathcal{T}(\mu_{2t}) = \{\tau \in \mathcal{T} \,|\, \mu_{2\tau} = \mu_{2t}\}$.

It is especially important to note that the pay-offs obtained by player 1 in maximizations (2) and (3) cannot be ranked; that is, the constraints under (2) and (3) are not subsets of each other.

In summary, propositions 2 to 5 demonstrate how renegotiation can effect equilibrium allocations in hidden information environment. There are three results of this analysis that we believe are especially relevant. First, the nature of the renegotiation process matters for describing renegotiation-proof allocations; that is, there is no unique notion of renegotiation-proofness: the restrictions imposed by renegotiation depend on both the timing of renegotiation and the identity of the proposer. Renegotiation-proofness is therefore a property of processes rather than a property of allocations (as incentive compatibility is). The game-theoretic methods described in this paper show how to translate the effects of the renegotiation process into an additional set of constraints on the contract. Second, allowing for contracts to be renegotiated does not in general imply that outcomes will be ex post efficient.

In fact, it is only when renegotiation is lead by the uninformed player after a specific action-pair has been agreed upon that ex post efficiency is generally expected. Finally, the distinction between the interim and ex post cases illustrates that renegotiation is not a consequence of the arrival of new information. That new information is perfectly anticipated in a probabilistic sense and is contracted away by both parties. Rather, renegotiation occurs because only part of the contract can be implemented at one time. Once a part of the contract has been fulfilled, it is impossible for the contracting parties not to let bygones be bygones.

It is worth mentioning that the current analysis of renegotiation can easily be extended to the cases of adverse selection and moral hazard. The different definitions of renegotiation-proofness involve restrictions on the equilibrium strategies played in the renegotiation stage given an initial contract offer. These definitions are contingent only on the contract initially offered, not on the information structure under which it was offered.

The adverse-selection equivalence to our interim-renegotiation game has been examined by Maskin and Tirole (1992) in their analysis of informed-principal problems. The only difference in our game is that player 1 knows the state of nature before offering the initial contract in stage 1, that is, stage 3 becomes stage 0, and player 1 makes the renegotiations offer. They show that interim-renegotiation-proof allocations of this modified game are interim efficient. This result is the analog to proposition 2 in the hidden-information framework.

In the simplest moral hazard problem, the information problem consists of an unobservable action that affects a random outcome. Suppose now that player 1 chooses the level of the unobserved action, known as effort, in stage 3 and that the two players could renegotiate the contract before the realization of the random outcome. At the time of renegotiation, the effort of player 1 is not observed by player 2, and therefore renegotiation takes place under asymmetric information. In this case the analysis of renegotiation in the moral-hazard case is similar to that of the cases we previously examined, since player 1's effort level can be interpreted as the state of nature or her type. This problem has been formally analysed by Fudenberg and Tirole (1990).

We conclude this section by comparing our analysis of renegotiation in a static environment with that in dynamic environments. As we mentioned in the introduction, much of the literature on renegotiation, for example, Hart and Tirole (1988), Laffont and Tirole (1990), and Dewatripont (1989), has been concerned with the effect of renegotiation in multiperiod environments. The starting point of this literature was the observation that optimal long-term contracts are in general time inconsistent. This observation implies that optimal long-term contracts are reasonable only in environments in which both players can commit not to renegotiate the contract in between periods.

Most of the renegotiation literature that examines multi-period problems has limited itself to cases where it is the uninformed player who proposes the renegotiations. The concept of renegotiation-proofness used in this literature is a hybrid between our concept of interim and ex-post-renegotiation-proofness, since the complete contract is a menu of submenus. In each period, player 1 first chooses a submenu within a larger menu, then the possibility of renegotiation arises in the form of offering submenus, and finally player 1 chooses a particular action vector to undertake for the current period. Therefore, in this set-up two forces are at play: in each period, the renegotiation looks like the interim renegotiation game, but the overall renegotiation game has some aspects of ex post renegotiation, since certain submenus are chosen before renegotiation occurs.

Until now, we have examined how renegotiation affects outcomes within different contractual environments. One of the results that we have emphasized is that there is no general notion of renegotiation-proofness, but instead that the effects of renegotiation can be assessed only within the context of a precise renegotiation process. Such a result may seem disheartening, since it suggests that the theory lacks a strong positive content. In the next section, however, we argue that renegotiation theory offers important insights into the theory of organizations exactly because different renegotiation processes have different implications.

III. RENEGOTIATION AS AN ELEMENT IN A THEORY OF TRADING INSTITUTIONS

In market economies, the exchange of goods and services can take place either within firms or between firms. A firm has often been described as an internal market, while Walrasian or external markets describe the trading place for different firms. Economists have always been preoccupied by the differences between these trading institutions, and early on, they realized that internal markets were better designed to achieve the necessary coordination of different groups of agents, that is, any trade on external markets can be replicated on internal markets. At this point two important questions were central to the preoccupations of organizational theorists. First, if internal markets seem to dominate external markets, why do we observe such a great proportion of trade on external markets, or equivalently, what are the limits to the size of internal markets (the firm)? And second, if one examines more closely the functioning of internal markets, how should decision making take place; that is, should decision making be centralized or decentralized? These two questions are central to the understanding of trading institutions in modern economies. A first answer to these questions is that, in a world of perfect and costless information, institutions are a matter of indifference: there is equivalence between internal and external markets. This suggests that informational imperfections may be an important factor in understanding the emergence and structure of trading institutions.

Recent developments in incentive theory have provided a framework for discussing the rationale for and the structure of trading institutions in the presence of informational imperfections. In particular, contract theory has shown that informational imperfections can create incentives for parties to coordinate activities through contracts that are richer (or more complex) than the standard spot contract which specifies only a transfer payment for the exchange of one good (as in a Walrasian market). In several situations, these richer contracts can be associated with internal markets, since they dictate rules of behaviour, communication, and compensation for a defined group of individuals. Therefore, contract theory can be interpreted as predicting that a laissez-faire economy will develop an institutional structure that involves internal markets which are much more complex and diversified than simple Walrasian (external) markets.

Even though standard contract theory has formally explained the emergence of internal markets, in many dimensions, it remains a very incomplete theory of institutions, since it fails to answer the two central questions stated above. First, contract theory confirms the intuition of early economists to the effect that trade with informational imperfections can best be handled through internal markets instead of external markets, but the theory does not provide any limits regarding the size of field of activity of firms. Second, contract theory has no predictive power regarding the structure of decision making within internal markets, that is, whether or not decisions are centralized or decentralized. In effect, the revelation principle shows that any decentralized structure can always be replicated by an appropriate centralized structure; that is, centralizing all information and decisions is regarded as always being an optimal organizational structure.

In this section, we use the results of the previous section to indicate how the introduction of renegotiation into standard contract theory can shed some light on the two central questions about the rationale for and structure of institutions. Our discussion will be rather informal and is meant mainly to suggest how the theory of renegotiation can be used to improve our understanding of institutional design. A thorough formalization of the link between renegotiation and institutional design is the subject of our ongoing research.

Our analysis of institutions will be broken down into two subsections. First, we shall study the functioning of internal markets when there are informational imperfections and renegotiation is possible. We shall show that, because of renegotiation, a decentralized decision-making process may dominate a more centralized structure. Second, we shall use these results to study the scope of internal markets. We shall also consider situations where, because of renegotiation, informational imperfections may best be handled by external markets instead of internal markets. In both cases, we shall argue that the preferred trading institution can be viewed as a means of avoiding certain types of renegotiation. Taken together, these two cases highlight ways in which the theory of renegotiation can be used to advance contract theory as a theory of trading institutions.

1. Decision making in internal markets: centralization versus decentralization

There is a widespread belief that organizations (firms) often gain by delegating decisions instead of centralizing them; as mentioned above, however, standard contract theory does not offer any support to this view. For example, Myerson's (1979) proof of the revelation principle shows that for any decentralized structure there exists a centralization scheme that results in the same outcome. Therefore, within the framework of contract theory, one must invoke something like transaction costs as the reason behind the common use of delegation. The transaction costs explanation is in itself not completely satisfactory, however, since it is not very precise as to the nature of these costs and therefore offers only limited predictions about when delegation is going to dominate centralization. In this subsection we argue that, because renegotiation can undo ex ante incentives, adopting a decentralized decision-making process may be viewed as an optimal commitment to ex ante desirable distortions even in environments where transaction costs alone would suggest centralization.

In order for us to discuss the issue of delegation, a reinterpretation of the environment presented in section II would be helpful. Assume that player i is producing action a_i, action a_1 is the production level of a good or service, and action a_2 is a transfer payment from player 2 to player 1. Therefore, in this example, player 1 is a producer (employee) and player 2 is the buyer (employer). Because the state of nature affects the utility of both players and it is observed only by player 1, there are potential benefits for both players to coordinate their respective actions through some set of rules. The type of rules that will govern the relationship is influenced by the organizational structure.

An organizational structure that coordinates these actions can be called *centralized* if player 2 has control over action a_1 and player 1 is required to transmit her private information, either directly or indirectly, to player 2 before he makes a decision on the level of action a_1 to be produced. Such a structure is justifiably considered centralized, since it does not delegate authority over decisions to the agent that gathers the relevant information. Alternatively, an organizational structure can be called *decentralized* if control over action a_1 is relinquished from player 2 to player 1 so that player 1 can directly choose the appropriate action level based on her private information.

For the purpose of this example, we assume that, initially, player 2 has control over action a_1 but can delegate his control to player 1 at a cost of $c \in \mathbf{R}$; that is, after paying a cost of c player 1 becomes in control of a_1. The cost c can therefore be interpreted as a transaction cost associated with the delegation of control over action a_1. If $c > 0$, then it is costly to delegate control over a_1 to player 1; if $c < 0$, then it is costly to centralize control over a_1 to player 2.

The contract-theory approach to organizational design predicts that, in the absence of renegotiation and transaction costs, a centralized organization always weakly dominates a delegated structure. Therefore, if the transaction cost c is positive, the theory predicts that the organizational structure should be centralized, while it should be decentralized if the transaction cost c is negative enough. In the absence of renegotiation, transaction costs become an important determinant of organizational structure. In contrast, we want to argue that the threat of renegotiation can modify this result, since delegating authority changes the scope for renegotiation. To see this, it is necessary to compare the type of renegotiation that can arise under a centralized or a decentralized structure and therefore identify the links between these two organizational structures and the results of propositions 2–5.

In the case of centralization, the rules governing the relationship between players 1 and 2 need to specify (1) a menu of state-contingent action-pairs; (2) the obligation by player 1, after having observed her private information, to send a message to player 2 indicating which level of action a_1 (and hence the associated transfer a_2) should be chosen by player 2; and (3) the (large) penalties if the prescribed action a_1 and transfer a_2 are not executed. These rules are equivalent to the contract studied in the previous section. The message sent by player 1 conditions the choice of an action-pair and penalties are imposed if that precise action-pair is not executed. In such an environment, as was emphasized in section II, once player 1 sends her message to player 2, there may be scope for renegotiation before the action is actually carried out. This type of renegotiation was referred to as ex post renegotiation. Therefore, if renegotiation is a possibility, the allocation that is likely to be implemented through a centralized structure should satisfy propositions 4 or 5, since both describe the optimal allocation that can be implemented through a contract that is ex post renegotiation-proof.[15] An implication of these propositions is that ex post renegotiation usually precludes players' achieving an interim-efficient allocation, and therefore a centralized organization faced with renegotiation will not be able to attain an information-constrained Pareto optimum.

In the case of a decentralization, the rules governing the relationship are slightly different. These rules need to specify (1) a menu of state-contingent action-pairs and (2) the (large) penalties if none of the action-pairs in the menu is executed. The distinguishing feature of a decentralized structure over a centralized structure is that the delegation of control over action a_1 to player 1 means that there is no longer the need for player 1 to communicate her private information about the state of nature to player 2. Player 1 can then select one of the action-pair on the basis of her private observation of the state of nature, and the transfer payment from player 2 to player 1 becomes a function of the level of action a_1 that player 1 has selected. Therefore penalties can be imposed only if none of the action-pairs within the menu is carried out. This implies that even if player 1 sends a message to player 2 indicating which action-pair she is going to select, this message is not binding and penalties cannot be imposed if the particular action-pair indicated by player 1 is not carried out. It is precisely the fact that messages from player 1 to player 2 do not bind the players to a course of actions that we take as distinguishing a centralized structure from a decentralized structure.

This simple modification to the nature of the contractual relationship has important implications for the type of renegotiation that can arise. For example, suppose that player 1, after observing the state of nature, sends a message to player 2 indicating the action she plans to undertake and then tries to renegotiate. Since the message sent by player 1 does not contractually bind the players to a specific action-pair, renegotiation takes place with the status quo being player 1's (yet to be chosen) preferred choice among the set of action-pairs in the menu regardless of the message sent. Since this preferred choice is contingent upon the privately

15 Note that the contract supporting the optimal ex-post-renegotiation-proof allocation is also robust to ex ante renegotiation.

observed state of nature, the strategic interaction during renegotiation is closely related to the interim-renegotiation game of section II. In effect, if interim-efficient allocations are characterized by one-to-one mappings between action a_1 and types, [16] then the choice of a_1 by player 1 replaces her message to player 2 and is actually a method for him to commit to a particular element within the menu.[17] Therefore, in these cases, the optimal allocation that can be supported with a decentralized structure is actually equivalent to the interim-renegotiation-proof allocations described in proposition 2, which were shown to be interim efficient and thus information-constrained Pareto optimal.

The above discussion indicates that different organizational structures may be subject to different types of renegotiation. In particular, prescribed allocations in a centralized organization are vulnerable to ex post renegotiation, while those in a decentralized structure are vulnerable only to interim renegotiation. Since we showed in section II that interim renegotiation does not impose any restrictions on the optimality of allocations (beside interim efficiency), while ex post renegotiation generally does, this observation on the interaction between organizational structure and the potential for renegotiation offers new insights on the role of delegation of decision making. In effect, this interpretation of organizations suggests that delegating decision power to the agent who gathers the relevant information will dominate a centralized system whenever the costs of delegating c are not too large, since doing so allows the players to avoid ex post renegotiation. Consequently, even though the delegation of authority may involve certain costs, the theory of renegotiation provides an explanation for why delegation may nevertheless be a preferred organizational structure.

There are similarities between the ideas presented in this section and Milgrom's (1988) notion of influence costs. Milgrom's approach has a similar flavour to ours; although influence activities differ from renegotiation, both approaches stress the costs of ex post communication in preventing optimal allocations from being attained. In both cases a decentralized organizational structure helps to improve efficiency as it acts as a commitment towards eliminating ex post communication.

Before we examine a second implication of renegotiation, it is worthwhile briefly discussing a limit to our current analysis. In particular, our results on the superiority of decentralization over centralization is presented within the context of one-sided asymmetric information; once both sides in a relationship obtain relevant information, however, there are direct gains to centralizing information. Consequently, with bilateral asymmetric information there emerges a trade-off between centralization and decentralization: decentralization avoids renegotiation, while centralization permits a more efficient use of information. Understanding the nature of this trade-off

16 See Melumad and Reichelstein (1987) for a discussion of the conditions under which this configuration arises.

17 Note that any attempt to renegotiate after action a_1 has been chosen is useless, since players have strictly opposing preferences with respect to the transfer a_2. Hence this type of renegotiation does not change the efficiency properties of the allocation attained with a decentralized structure.

and deriving conditions where one form of organization dominates the other seems to be a fruitful area for future research.

2. Internal versus external markets

The previous discussion indicates how the threat of renegotiation can influence organizational design and in particular can favour decentralization over centralization. In this section we want to compare the allocations that arise when trading occurs in internal markets and allocations issued of trading on external markets. We shall show that when a decentralized organization is not a feasible alternative, external markets may be a preferred trading institution over internal markets. We present this result as a potential explanation for the limited scope of institutions.[18]

In order to compare external markets with internal markets, it is necessary first to describe what we mean by the (external) market solution to a relationship with informational imperfections. We have called the organizational (or contractual) solution to our trading problem the situation where player 1, who knows that she will eventually become privately informed, enters into a contract before her private information is actually revealed. In contrast, we shall call the market solution the outcome that arises when player 1 does not immediately begin a contractual relationship but instead (1) waits until her information is revealed, (2) decides on a level for a_1, and (3) offers player 2 a trade between a_1 and a_2. In this case the rules governing the relationship between the two players is reduced to a quid pro quo contract or what we call a market contract. The first point to note from our description of the market solution is that it corresponds exactly to a signalling game. For example, if a_1 represents education, a_2 a wage payment, and t the worker's ability, then the market outcome for this situation is the solution to the Spence (1973) education game. Consequently, we can characterize the market solution by directly exploiting the results from the signalling literature.

The equilibrium allocation of a signalling game in general depends on the equilibrium concept used. Although debates regarding the appropriate equilibrium concept still abound, there exists a rather broad consensus that under our assumptions it is reasonable to describe the outcome of a signalling game as that of the 'efficient' separating equilibrium (see Cho and Kreps 1987; Cho and Sobel 1988). Therefore, we will adopt this convention and take as the market solution to our trading problem the allocation defined by the following maximization.

$$\max_{\{\mu_t\}_{t=1}^{T}} \sum_{t=1}^{T} p_0(t)U(\mu_t, t)$$

$$\text{s/t} \quad (i) \quad V(\mu_t, t) \geqq V(0, t) \qquad \forall t \in \mathcal{T}$$

$$(ii) \quad U(\mu_t, t) \geqq U(\mu_{t'}, t) \qquad \forall t, t' \in \mathcal{T}. \tag{4}$$

18 Dewatripont and Maskin (1989) in a financial market example and Dearden, Ickes and Samuelson (1990) in a problem of innovation adoption compare the relative efficiency of external markets with that of internal markets.

The first point to note about the market solution is that, in ex ante terms, it is dominated by the full-commitment contract described in proposition 1. This can be seen by noting that player 2's participation constraint needs to hold in expected terms only in a contractual relationship, while it has to hold across each state in a market relationship. Therefore, in the absence of renegotiation possibilities, player 1 will find in her interest to enter into a contractual relationship before her information is revealed instead of waiting for the market outcome. In this sense, standard contract theory can be interpreted as suggesting that informational imperfections will best be handled by organizations instead of directly by the market. Therefore, the main drawback of the market solution (as we have defined it) is that it eliminates the risk-sharing possibilities associated with the full-commitment contact by imposing ex-post-participation constraints.

Although the full-commitment contract generally dominates the market solution, the superiority of the organization over the market may not hold once renegotiation is introduced. In particular, consider the case where any feasible organizational arrangement is subject to ex post renegotiation by the informed agent, that is, suppose that the organization can never refrain the informed party from renegotiating after a particular action-pair has been agreed to.[19] In this situation, the organizational arrangement will dominate the market only if the value of the maximization (2) is greater than the value of the maximization (4). In general it is not possible to rank these two different organizational structures, since each has certain advantages in terms of insurance. On the one hand, the ex-post-renegotiation-proof organizational structure requires only that player 2's participation constraint be satisfied in expectation and thereby provides room for explicit insurance, although the effective incentive-compatibility constraints are more stringent. On the other hand, the market solution can support more distortions, which can provide an implicit type of insurance even though explicit cross-subsidization between states is impossible.

The comparison between the organizational and market arrangements implicitly assumes that the organization is vulnerable to ex post renegotiation, while no renegotiation occurs under the market arrangement. We would like to argue that these assumptions arise quite naturally in many economic environments, and therefore the comparison that we set up between an organizational arrangement and the market is in fact relevant.

Consider the following modification to the trading framework we have set up so far and assume that the implementation of the action a_1 takes time. For example, the action a_1 may be an investment level in education that player 1 must make before getting the transfer from player 2. Suppose now that the investment period is partitioned into a large number of small discrete subperiods. In each of these subperiods, players 1 and 2 may potentially renegotiate the outstanding agreement.

In this framework, the organizational arrangement corresponds to the situation in which (1) both players sign a contract, (2) player 1 privately observes the state of nature, (3) she selects an element in the menu of the contract, (4) produces

19 We discuss, below, the relevance of this assumption.

in each subperiod until the desired level of a_1 has been reached, and finally (5) trades her production with player 2 in exchange for the contractually corresponding transfer a_2. Player 1 herself selects her preferred action-pair in the menu, and therefore this organizational arrangement corresponds to a decentralized structure as described in the last subsection. Even though in many cases a decentralized structure may not be vulnerable to ex post renegotiation for reasons explained above, when the implementation of action a_1 takes time it may become so. If player 2 can observe the productive activity of player 1 and if both players can communicate in each subperiod, then a form of ex post renegotiation can effectively occur in each subperiod.

It was argued that a decentralized arrangement was not vulnerable to ex post renegotiation because at the time renegotiation was taking place the status quo position was player 1's (yet to be chosen) preferred allocation in the contractual menu, and therefore this type of arrangement was vulnerable only to interim renegotiation. When production takes time, this argument may no longer be valid. When renegotiation occurs with some units' having been produced, then the status quo position is player 1's preferred allocation among those elements in the menu that specify at least as many units as the number already produced. When a large number of units have been produced, there may not be many elements in the menu that remain attainable, and therefore the status quo position of player 1 becomes much more precise. This type of dynamic renegotiation is quite close to ex post renegotiation, and therefore time-consuming production may make a decentralized structure vulnerable to ex post renegotiation.[20] The above argument implies that in many interesting economic applications ex post renegotiation may be a concern for different types of institutional arrangements (centralized or decentralized).

In the same framework, the market arrangement corresponds to the situation in which (1) player 1 privately observes the state of nature, (2) she produces in each subperiod until the desired level a_1 has been reached, and finally (3) trades her production with player 2 in exchange for a competitively determined transfer a_2. The market arrangement differs from the organizational arrangement in that no contract is signed before or after player 1 observes the state of nature. In fact, the expected competitive resolution of the game provides an implicit contract for player 1. Now suppose that, in any subperiod, player 1 can decide to stop production of a_1 indefinitely and bring the produced units to the market in exchange for a transfer a_2. In some sense, this possibility allows player 1 to renegotiate the implicit contract by bringing her units to the market before they have reached the level prescribed by the equilibrium. In this framework, Noldeke and Van Damme (1990) have shown that the market solution is appropriately described by the static market solution to problem (4). The intuition behind this result is that when player 1 tries to renegotiate the implicit market contract, the status quo position is not dictated by a formally written contract but rather by player 2's best response to the renegotiation

20 This case is examined in Beaudry and Poitevin (1994), where it is shown that the allocation in this dynamic renegotiation problem is identical to the characterization provided here under ex-post-renegotiation-proofness.

offer. This best response depends on player 2's beliefs about the state of nature following player 1's offer. For example, if player 2 has pessimistic beliefs, he will reject most renegotiation offers until player 1 reaches her equilibrium level of a_1. The dependence of the status quo position on player 2's beliefs removes most incentives to renegotiate ex post.

This argument supports the assumption that the market arrangement is not vulnerable to renegotiation, and therefore the comparison that we made above between trading on internal versus external markets is relevant for many interesting economic environments.[21]

IV. CONCLUSION

In this paper we have pursued two main goals. First, we have developed a framework in which different renegotiation processes can be examined. The principal advantage of this framework is that it clarifies several of the conflicting results regarding the effects of renegotiation by emphasizing the difference between interim and ex post renegotiation. Second, we have indicated why and where the theory of renegotiation may be relevant for understanding economic relationships. In particular, we have indicated how the theory of renegotiation can provide insights regarding the structure of institutions, the merits of decentralization, and the value of the market. However, we believe that the latter issue is still in its infancy and deserves further attention.

APPENDIX

Proof of proposition 1

We shall first show by contradiction that the equilibrium allocation must be the solution to problem (1). Sufficiency will then be shown by constructing strategies and beliefs that support this equilibrium allocation as a PBE outcome of the game.

Let $\hat{c}$ represent a candidate equilibrium contract offer and let $\hat{\mu} \neq \mu^c$ be the corresponding equilibrium allocation. Let us also assume that $\hat{\mu}$ is such that $\Sigma_{t=1}^{T} p_0(t) U(\hat{\mu}_t, t) < \Sigma_{t=1}^{T} p_0(t) U(\mu_t^c, t)$. Then by definition there must exist an allocation $\tilde{\mu}$ that satisfies the following set of constraints:

$$\sum_{t=1}^{T} p_0(t) U(\tilde{\mu}_t, t) > \sum_{t=1}^{T} p_0(t) U(\hat{\mu}_t, t)$$

$$\sum_{t=1}^{T} p_0(t) V(\tilde{\mu}_t, t) > \sum_{t=1}^{T} p_0(t) V(0, t), \qquad U(\tilde{\mu}_t, t) > U(\tilde{\mu}_{t'}, t) \qquad \forall t, t' \in \mathcal{T}$$

21 The comparison between dynamic renegotiation in external versus internal markets implies that explicit contracts are far more vulnerable to renegotiation than implicit contracts. This implication follows from the position of their respective status quo.

Given this allocation $\tilde{\mu}$, player 1 will want to deviate by offering a contract $\tilde{c}$ with $m(\tilde{c}) = \tilde{\mu}$, since the only subgame equilibrium of the game induced by such a deviation involves player 2's accepting the contract and player 1's choosing $\tilde{\mu}_t$ if the state of nature is revealed to be t. Consequently, an equilibrium allocation must at least provide the level of utility of player 1 defined in the maximization. Let us now assume that $\Sigma_{t=1}^{T} p_0(t)U(\hat{\mu}_t, t) > \Sigma_{t=1}^{T} p_0(t)U(\mu_t^c, t)$. This inequality implies that either condition (i) or condition (ii) is not satisfied when evaluated at $\hat{\mu}$. However, by a standard dominant strategy argument, this possibility can be ruled out, since it implies that player 2 would gain by simply refusing to play the game or that player 1 would gain by simply choosing his preferred element within the menu. Therefore, an equilibrium allocation must necessarily solve the stated maximization.

The following strategies and beliefs support the equilibrium allocation as a PBE outcome.

$$\sigma_1^c = \begin{cases} \tilde{c}_0 = c^c \text{ with } m(c^c) = \mu^c \\ \tilde{s}_0(c_0, t) = \arg\max_{a^n \in m(c_0)} U(a^n, t) \end{cases}$$

$$\sigma_2^c = \tilde{d}_0(c_0) = \begin{cases} 1 & \text{if } \sum_{t=1}^{T} p_0(t)V(\tilde{s}_0(c_0, t), t) \geqq \sum_{t=1}^{T} p_0(t)V(0, t). \\ 0 & \text{otherwise} \end{cases}$$

It is easy to verify that these strategies and beliefs do in fact constitute a PBE. In stage 4, player 1 selects the element she prefers most in the menu given the outstanding contract c_0. These strategies condition player 2's acceptance decision of the contract d_0 in stage 2: he accepts only contracts c_0 satisfying his participation constraint given the expected resolution of c_0, that is, given the incentive constraints of player 1 and her choice $\tilde{s}_0(c_0, t)$. In the first stage, player 1 offers her most preferred contract in the set of contracts that are acceptable to player 2. Along the equilibrium path we thus have that $\tilde{c}_0 = c^c$ with $m(c^c) = \{\mu_1^c, \ldots, \mu_T^c\}$, $\tilde{d}_0(c^c) = 1$, and $\tilde{s}_0(c^c, t) = \mu_t^c$ for all t. ||

Proof of proposition 2

We shall first show by contradiction that the equilibrium allocation must be the solution to problem (1). Sufficiency will then be shown by constructing strategies and beliefs that support this equilibrium allocation as a PBE outcome of the game.

It is clear that the introduction of interim renegotiation cannot increase the ex ante expected utility of player 1 (otherwise the interim-renegotiation-proof allocation could have been offered intially in the commitment game). Therefore, it is only necessary to show here that, regardless of the identity of the player making the renegotiation, player 1 cannot have in equilibrium less expected utility than that associated with the solution to problem (1). Let $\hat{c}$ represent a candidate equilibrium contract offer and let $\hat{\mu} \neq \mu^{ir}$ be the corresponding equilibrium allocation. Let us also assume that $\hat{\mu}$ is such that $\Sigma_{t=1}^{T} p_0(t)U(\hat{\mu}_t, t) < \Sigma_{t=1}^{T} p_0(t)U(\mu_t^{ir}, t)$. Suppose first that $\Sigma_{t=1}^{T} p_0(t)V(\hat{\mu}_t, t) = \Sigma_{t=1}^{T} p_0(t)V(0, t)$. Then by definition there must exist a type

t' and an outcome $\tilde{\mu}_{t'}$ such that $V(\tilde{\mu}_{t'}, t') > V(\hat{\mu}_{t'}, t')$, $U(\tilde{\mu}_{t'}, t') > U(\hat{\mu}_{t'}, t')$, $U(\tilde{\mu}_{t'}, t') > U(\hat{\mu}_{t'}, t')$ for all $t \neq t'$, and $U(\hat{\mu}_t, t) > U(\tilde{\mu}_{t'}, t)$ for all $t \neq t'$. Consider the contract $\tilde{c}$ whose menu is $m(\tilde{c}) = \{\hat{\mu}_t\}_{t \neq t'} \cup \tilde{\mu}_{t'}$. This allocation is incentive compatible; that is, every type $t \neq t'$ prefers $\hat{\mu}_t$ to any other outcome in the menu, and type t' prefers by construction $\tilde{\mu}_{t'}$ to any other outcome in the menu. Furthermore the contract $\tilde{c}$ yields (weakly) more utility to both players regardless of the realized state. This implies that if $\tilde{c}$ is offered in stage 3.1 by one player, it will be accepted by the other player. The contract $\hat{c}$ can therefore not be Ii-renegotiation-proof in this case. Suppose now that $\Sigma_{t=1}^{T} p_0(t) V(\hat{\mu}_t, t) > \Sigma_{t=1}^{T} p_0(t) V(0, t)$. In this case, player 1 can initially offer the solution to problem (1) with $V(0,t) + \epsilon$ replacing $V(0,t)$. Since player 2 cannot lose in the renegotiation round, he accepts this offer from player 1. This implies that $\hat{c}$ is not Ii-renegotiation-proof in this case either. Therefore, an equilibrium allocation must necessarily solve the stated maximization.

For each case, we provide strategies and beliefs that support the equilibrium allocation as a PBE outcome. Consider first the case in which player 1 makes the renegotiation offer in stage 3.1. The following strategies and beliefs support the allocation μ_1^{ir} as an equilibrium I1-renegotiation-proof allocation.

$$\Omega_1^{ir} = \begin{cases} \tilde{c}_0 = c_1^{ir} \text{ with } m(c_1^{ir}) = \{\mu_{1t}^{ir}\}_{t=1}^{T} \\ \tilde{c}_1(c_0, t) = \begin{cases} \arg\max_{c_1} & \sum_{t=1}^{T} p_0(t) U(\tilde{s}(c_0, c_1, 1, t), t) \\ \quad \text{s/t} & V(\tilde{s}(c_0, c_1, 1, t), t) \geqq V(\tilde{s}(c_0, c_1, 0, t), t) \\ & \hfill \forall t \in \mathcal{T} \\ & \text{if it is different from } c_0 \\ \emptyset & \text{otherwise} \end{cases} \\ \tilde{s}(c_0, c_1, d_1, t) = \arg\max_{a^n \in \{a^n\}_{n=1}^{N}} U(a^n, t), \text{ where } \{a^n\}_{n=1}^{N} \\ \qquad = \begin{cases} m(c_1) & \text{if } d_1 = 1 \\ m(c_0) & \text{if } d_1 = 0 \end{cases} \end{cases}$$

$$\Omega_2^{ir} = \begin{cases} \tilde{d}_0(c_0) = \begin{cases} 1 & \text{if } \sum_{t=1}^{T} p_0(t) V(\tilde{s}(c_0, \tilde{c}_1(c_0, t), \tilde{d}_1(c_0, \tilde{c}_1(c_0, t)), t), t) \\ & \qquad \geqq \sum_{t=1}^{T} p_0(t) V(0, t) \\ 0 & \text{otherwise} \end{cases} \\ \tilde{d}_1(c_0, c_1) = \begin{cases} 1 & \text{if } c_1 \neq c_0 \text{ and } V(\tilde{s}(c_0, c_1, 1, t), t) \\ & \qquad \geqq V(\tilde{s}(c_0, c_1, 0, t), t) \forall t \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases} \end{cases}$$

$$p_2^{ir}(t | c_0, c_1) = \begin{cases} p_0(t) & \text{if } V(\tilde{s}(c_0, c_1, 1, t), t) \geqq V(\tilde{s}(c_0, c_1, 0, t), t) \qquad \forall t \in \mathcal{T} \\ 1 & \text{if } \exists t' \text{ such that } V(\tilde{s}(c_0, c_1, 1, t'), t') < V(\tilde{s}(c_0, c_1, 0, t'), t') \\ & \text{and } t \text{ is the smallest such } t' \\ 0 & \text{otherwise} \end{cases}$$

We shall now argue that these strategies and beliefs do in fact constitute a PBE.

In stage 4, player 1 selects the element of the menu of the outstanding contract that she prefers given her type. In stage 3.2, player 2 accepts only contracts that result in allocations that improve on his pay-off regardless of the type of player 1. This is supported by the beliefs that if a contract offer in stage 3.1 is expected to result in an allocation for which there exists a smallest t' such that $V(\tilde{s}(c_0, c_1, 1, t'), t') < V(\tilde{s}(c_0, c_1, 0, t'), t')$, then this contract offer must have been offered by type t' and hence it is rejected by player 2. Given this acceptance rule by player 2, player 1 can do no better than offer in stage 3.1 her preferred contract among those accepted by player 2 if this contract is different from the outstanding contract; she offers no contract otherwise. In stage 2, player 2 accepts all contract offers yielding an expected pay-off of $\Sigma_{t=1}^{T} p_0(t) V(0, t)$ given the expected equilibrium resolution of the game following this initial offer. Finally, in stage 1 player 1 offers her preferred contract among those expected to be accepted by player 2.

These strategies and beliefs imply the following equilibrium path. In stage 1, player 1 offers the contract c_1^{ir} which is accepted by player 2 in stage 2. In stage 3.1, player 1, regardless of her type, cannot offer a contract that is expected to be accepted and that improves her welfare. Given the acceptance rule of player 2, the best offer type t player 1 can make is the contract that specifies as a menu the solution to the following maximization problem.

$$\max_{\{\mu_{1t}\}_{t=1}^{T}} U(\mu_{1t}, t)$$

$$\text{s/t} \quad (i)\ V(\mu_{1t}, t) \geqq V(\mu_{1t}^{ir}, t) \qquad \forall t \in \mathcal{T}$$

$$(ii)\ U(\mu_{1t}, t) \geqq U(\mu_{1t'}, t) \qquad \forall t, t' \in \mathcal{T}.$$

Regardless of type, it is easy to see that the first-order conditions of this problem yield the same solution as the first-order conditions of problem (1). Therefore, the solution to this maximization problem is μ_1^{ir} and, regardless of type, player 1 does not make an offer. In stage 4, player 1 selects her preferred element in the menu $m(c_1^{ir})$ and actions are then executed as prescribed by that element.

With these strategies along the equilibrium path it is clear that the allocation μ_1^{ir} is I1-renegotiation-proof.

We now consider the case in which player 2 makes the renegotiation offer in stage 3.1. The following strategies and beliefs support the allocation μ_2^{ir} as an equilibrium I2-renegotiation-proof allocation.

$$\Omega_1^{ir} = \begin{cases} \tilde{c}_0 = c_2^{ir} \text{ with } m(c_2^{ir}) = \{\mu_{2t}^{ir}\}_{t=1}^{T} \\ \tilde{d}_2(c_0, c_2, t) = \begin{cases} 1 & \text{if } c_2 \neq c_0 \text{ and } U(\tilde{s}(c_0, c_2, 1, t), t) \\ & \qquad \geqq U(\tilde{s}(c_0, c_2, 0, t), t) \\ 0 & \text{otherwise} \end{cases} \\ \tilde{s}(c_0, c_2, d_2, t) = \arg\max_{a^n \in \{a^n\}_{n=1}^{N}} U(a^n, t) \text{ where } \{a^n\}_{n=1}^{N} \\ \qquad = \begin{cases} m(c_2) & \text{if } d_2 = 1 \\ m(c_0) & \text{if } d_2 = 0 \end{cases} \end{cases}$$

$$\Omega_2^{ir} = \begin{cases} \tilde{d}_0(c_0) = \begin{cases} 1 & \text{if } \Sigma_{t=1}^{T} p_0(t) V\big(\tilde{s}(c_0,\, \tilde{c}_2(c_0),\, \tilde{d}_2(c_0,\, \tilde{c}_2(c_0),\, t),\, t),\, t\big) \\ & \qquad \geqq \Sigma_{t=1}^{T} p_0(t) V(0,\, t) \\ 0 & \text{otherwise} \end{cases} \\ \tilde{c}_2(c_0) = \begin{cases} \arg\max_{c_2} & \Sigma_{t=1}^{T} p_0(t) V(\tilde{s}(c_0,\, c_2,\, 1,\, t),\, t) \\ \quad \text{s/t} & U(\tilde{s}(c_0,\, c_2,\, 1,\, t),\, t) \geqq U(\tilde{s}(c_0,\, c_2,\, 0,\, t),\, t) \\ & \qquad \forall t \in \mathcal{T} \\ & \text{if it is different from } c_0 \\ \emptyset & \text{otherwise.} \end{cases} \end{cases}$$

We shall now argue that these strategies and beliefs do in fact constitute a PBE.

In stage 4, player 1 selects the element of the menu of the outstanding contract that she prefers given her type. In stage 3.2, player 1 accepts only those contracts leading to an allocation that improves her welfare given her type. In stage 3.1, player 2 offers the best contract that is acceptable to all types of player 1 if it is different from the outstanding contract; he offers no contract otherwise. This is without loss of generality, since player 2 can offer a contract that specifies a menu of elements, one for each type. For example, suppose player 2 offers a contract that is accepted only by a subset of types $\mathcal{T}' \subset \mathcal{T}$. Then all types in $\mathcal{T}'$ prefer an element in $m(c_2)$ to all elements in $m(c_0)$. Furthermore, all types not in $\mathcal{T}'$ prefer an element in $m(c_0)$ to all elements in $m(c_2)$. It would therefore be incentive compatible to include in the menu $m(c_2)$ those elements preferred by all types not in $\mathcal{T}'$. This would not reduce player 2's welfare, since selected elements would be the same under either scheme. Hence the specification of player 2's strategy for his offer of contract c_2 is without loss of generality. In stage 2, player 2 accepts only those contracts that are expected to lead to an equilibrium allocation yielding at least $\Sigma_{t=1}^{T} p_0(t) V(0, t)$ given the resolution of the renegotiation stage. Finally in stage 1, player 1 offers her preferred contract among the set of those expected to be accepted by player 2.

These strategies and beliefs imply the following equilibrium path. In stage 1, player 1 offers the contract c_2^{ir} which is accepted by player 2 in stage 2. In stage 3.1, player 2 cannot offer a contract that is expected to be accepted and that improves her welfare. Given the acceptance rule of player 1, the best offer player 2 can make is the solution to the following maximization problem.

$$\max_{\{\mu_{2t}\}_{t=1}^{T}} \sum_{t=1}^{T} p_0(t) V(\mu_{2t},\, t)$$

$$\begin{aligned} \text{s/t} \quad & (i) \;\; U(\mu_{2t},\, t) \geqq U(\mu_{2t}^{ir},\, t) && \forall t \in \mathcal{T} \\ & (ii) \;\; U(\mu_{2t},\, t) \geqq U(\mu_{2t'},\, t) && \forall t,\, t' \in \mathcal{T}. \end{aligned}$$

Regardless of type, the solution to this maximization problem is μ_2^{ir}. Suppose this was not the case and the solution was $\hat{\mu} \neq \mu_2^{ir}$. By definition, $\hat{\mu}$ satisfies all incentives-compatibility constraints in the maximization problem (1) and improves the expected welfare of player 2 without decreasing player 1's expected welfare. Since player 2's participation constraint is binding in the maximization problem of the statement of the proposition, this contradicts the fact that μ_2^{ir} is a solution to it. It is then clear that, regardless of type, player 2 cannot do better than with c_2^{ir} and therefore he makes no offer. In stage 4 player 1 selects her preferred element in the menu $m(c_2^{ir})$ and actions are then executed as prescribed by that element.

With these strategies along the equilibrium path it is clear that the allocation μ_2^{ir} is I2-renegotiation-proof. ||

Proof of proposition 3

For μ_1^{pr} to be P1-renegotiation-proof, it must be the case that, when it is offered in stage 1 and accepted in stage 2, it is not renegotiation in stage 4.1. If this is the case, we show that the constraints (i) and (ii) must be satisfied. If the allocation μ_1^{pr} is not renegotiated, it will be the implemented allocation and it must therefore satisfy constraint (i) for it to be acceptable to player 2. Constraints (ii) capture the effects of ex post renegotiation by player 1. Suppose there exist $t, t' \in \mathcal{T}$ such that (ii) was not satisfied for a candidate equilibrium P1-renegotiation-proof allocation $\{\tilde{\mu}\}_{t=1}^{T}$ supported by the contract $\tilde{c}$. This implies that there exists an action-pair $\hat{\mu}$ such that $V(\hat{\mu}, t'') > V(\tilde{\mu}_{t'}, t'')$ for all $t'' \in \mathcal{T}$ and $U(\hat{\mu}, t) > U(\tilde{\mu}_t, t)$. Now consider the following strategies. In stage 4, player 1 selects the equilibrium element of type t', that is, $\tilde{s}_0(\tilde{c}, t) = \tilde{\mu}_{t'}$; in stage 4.1, player 1 offers a contract that she knows will be accepted for sure by player 2, since it increases player 2's utility regardless of his beliefs; that is, $\tilde{c}_1(\tilde{c}, \tilde{\mu}_{t'}, t) = \hat{c}_1$ such that $m(\hat{c}_1) = \{\hat{\mu}\}$; in stage 4.2, it is a dominant strategy for player 2 to accept the utility-increasing offer; that is, $\tilde{d}_1(\tilde{c}, \tilde{\mu}_{t'}, \hat{c}_1) = 1$; in stage 4.3, player 1 trivially selects $\hat{\mu} \in m(\hat{c}_1)$. These strategies imply that if $\hat{c}_1$ is offered, then it must be accepted by player 2 regardless of his beliefs. Given the definition of $\hat{\mu}$ and the anticipated response of player 2, player 1 will in fact offer $\hat{c}_1$ after having chosen $\tilde{\mu}_{t'}$, thus breaking the equilibrium supporting the allocation $\tilde{\mu}$ as a P1-renegotiation-proof allocation. This implies that the constraints (ii) are necessary for an allocation to be P1-renegotiation-proof. ||

Proof of proposition 4

We construct strategies and beliefs that support μ_1^{pr} as an equilibrium P1-renegotiation-proof allocation of the ex post-renegotiation game.

Define $\hat{\mu}(a^n, t) = \arg\{\max_{\mu} U(\mu, t) \text{ s/t } V(\mu, t') \geqq V(a^n, t')\ \forall t' \in \mathcal{T}\}$. The set of types that select the element $s_0 \in m(c_0)$ is denoted by $\mathcal{T}(c_0, s_0) = \{t \in \mathcal{T} \mid s_0 = \arg\max_{a^n \in m(c_0)} U(\hat{\mu}(a^n, t), t)\}$.

$$\Phi_1^{pr} = \begin{cases} \tilde{c}_0 = c_1^{pr} \text{ with } m(c_1^{pr}) = \{\mu_{1t}^{pr}\}_{t=1}^{T} \\ \tilde{s}_0(c_0, t) = \arg\max_{a^n \in m(c_0)} U(\hat{\mu}(a^n, t), t) \\ \tilde{c}_1(c_0, s_0, t) = \begin{cases} c_1 \text{ such that } m(c_1) = \{\hat{\mu}(s_0, t)\} & \text{if } \hat{\mu}(s_0, t) \neq s_0 \\ \emptyset & \text{otherwise} \end{cases} \\ \tilde{s}_1(c_0, s_0, c_1, t) = m(c_1) \end{cases}$$

$$\Phi_2^{pr} = \begin{cases} \tilde{d}_0(c_0) = \begin{cases} 1 & \text{if } \Sigma_{t=1}^{T} p_0(t) V(\tilde{s}_1(c_0, \tilde{s}_0(c_0, t), \tilde{c}_1(c_0, \tilde{s}_0(c_0, t)), t), t) \\ & \qquad \geqq \Sigma_{t=1}^{T} p_0(t) V(0, t) \\ 0 & \text{otherwise} \end{cases} \\ \tilde{d}_1(c_0, s_0, c_1) = \begin{cases} 1 & \text{if } m(c_1) \neq \{s_0\} \text{ and } V(m(c_1), t) \geqq V(s_0, t) \\ & \qquad \forall t \in \mathcal{T} \\ 0 & \text{otherwise} \end{cases} \end{cases}$$

$$p_2^{pr}(t|c_0, s_0, c_1) = \begin{cases} p_0(t)/\Sigma_{\tau \in \mathcal{T}(c_0, s_0)} p_0(\tau) & \text{if } \mathcal{T}(c_0, s_0) \neq \emptyset,\ t \in \mathcal{T}(c_0, s_0) \\ & \text{and } V(m(c_1), t) \geqq V(s_0, t)\ \forall t \in \mathcal{T} \\ 1 & \text{if } \exists t' \text{ such that } V(m(c_1), t') < V(s_0, \\ & t') \text{ and } t \text{ is the smallest such } t' \\ 0 & \text{otherwise.} \end{cases}$$

We shall now argue that these strategies and beliefs do in fact constitute a PBE.

In stage 4.3, if c_1 is accepted, player 1 trivially selects $s_1 = m(c_1)$. In stage 4.2, player 2 accepts the new contract offer c_1 if and only if $m(c_1)$ is preferred to the action s_0 regardless of his beliefs. This is supported by the beliefs that if a contract offer in stage 4.2 yields less than s_0 for a smallest t', then player 2 concentrates his beliefs on t', thus inducing him to reject c_1. Given this acceptance rule by player 2, player 1 can do not better than offer in stage 4.1 her preferred contract among those accepted by player 2. This includes selecting in stage 4 the element s_0 of $m(c_0)$ which gives player 1 the best renegotiation possibility and then offering in stage 4.1 the contract c_1 with the associated menu $m(c_1) = \{\hat{\mu}(s_0, t)\}$. In stage 2, player 2 accepts all contract offers yielding an expected pay-off of $\Sigma_{t=1}^{T} p_0(t) V(0, t)$ given the expected equilibrium resolution of the game following this initial offer. Finally, in stage 1 player 1 offers her preferred contract among those expected to be accepted by player 2.

These strategies and beliefs imply the following equilibrium path. In stage 1, player 1 offers the contract c_1^{pr} which is accepted by player 2 in stage 2. In stage 4, type t of player 1 selects her preferred element μ_{1t}^{pr} in $m(c^{pr})$. In stage 4.1, she makes no offer. Given that μ_1^{pr} satisfies the constraints of the maximization problem (2), the contract c_1^{pr} cannot be renegotiated in stage 4.1 given the equilibrium strategy of player 2.

With these strategies along the equilibrium path, it is clear that the allocation μ_1^{pr} is P1-renegotiation-proof. ||

Proof of proposition 5

We shall first show that constraints (i) and (ii) of the proposition are necessary to describe P2-renegotiation-proof allocations. Then we shall construct strategies and beliefs that support an allocation μ_2^{pr} as an equilibrium P2-renegotiation-proof allocation of the subgame starting with stage 3 of the ex post-renegotiation game.

For μ_2^{pr} to be P2-renegotiation-proof, it must be the case that, when it is offered in stage 1 and accepted in stage 2, it is not renegotiated in stage 4.1. If this is the case, we show that the constraints (i) and (ii) must be satisfied. Constraints (i) represent standard incentive-compatibility constraints that must be satisfied. Constraints (ii) capture the effects of ex post renegotiation by player 2. To show that these constraints are necessary, consider a candidate equilibrium P2-renegotiation-proof allocation $\{\tilde{\mu}\}_{t=1}^{T}$ supported by the contract $\tilde{c}$ with a type $t \in \mathcal{T}$ such that (ii) is not satisfied. This implies that there exists a vector of action-pairs $\{\hat{\mu}_\tau\}_{\tau\in\mathcal{T}(\tilde{\mu}_t)}$ such that $\Sigma_{\tau\in\mathcal{T}(\tilde{\mu}_t)}p_0(\tau)V(\hat{\mu}_\tau,\tau) > \Sigma_{\tau\in\mathcal{T}(\tilde{\mu}_t)}p_0(\tau)V(\mu_{2t}^{pr},\tau)$, $U(\hat{\mu}_\tau,\tau) > U(\hat{\mu}_{\tau'},\tau)$ for all τ, $\tau' \in \mathcal{T}(\tilde{\mu}_t)$, and $U(\hat{\mu}_\tau,\tau) > U(\tilde{\mu}_t,\tau)$ for all $\tau \in \mathcal{T}(\tilde{\mu}_t)$. Now consider the following strategies. In stage 4, player 1 of type $\tau \in \mathcal{T}(\tilde{\mu}_t)$ plays her equilibrium strategy $\tilde{s}_0(\tilde{c},\tau) = \tilde{\mu}_t$; in stage 4.1, the equilibrium dictates that player 2 must believe with probability $p_0(t')/\Sigma_{\tau\in\mathcal{T}(\tilde{\mu}_t)}p_0(\tau)$ that the type of player 1 is t', and consequently he can offer the contract $\tilde{c}_2(\tilde{c},\tilde{\mu}_t) = \hat{c}_2$ such that $m(\hat{c}_2) = \{\hat{\mu}_\tau)_{\tau\in\mathcal{T}(\tilde{\mu}_t)}$; in stage 4.2, type $\tau \in \mathcal{T}(\tilde{\mu}_t)$ of player 1 accepts the utility-increasing renegotiation offer, that is, $\tilde{d}_2(\tilde{c},\tilde{\mu}_t,\hat{c}_2,\tau) = 1$; in stage 4.3, type τ of player 1 selects her preferred element $\hat{\mu}_\tau$ in $m(\hat{c}_2)$. These strategies imply that if $\hat{c}_2$ is offered, then it must be the accepted by all types of player 1 believed by player 2 to have selected $\tilde{\mu}_t$ from the menu of the initial contract offer. Given the definition of $\{\hat{\mu}_\tau\}_{\tau\in\mathcal{T}(\tilde{\mu}_t)}$ and the anticipated response of player 1, player 2 will in fact offer $\hat{c}_2$, after $\tilde{\mu}_t$ has been selected, thus breaking the equilibrium supporting the allocation $\tilde{\mu}$ as a P2-renegotiation-proof allocation. This implies that the constraints (ii) are necessary for an allocation to be P2-renegotiation-proof.

We now show that these constraints are also sufficient by constructing strategies and beliefs that support μ_2^{pr} as a P2-renegotiation-proof allocation. Define c_0 as the contract supporting the allocation μ_2^{pr}; that is, $m(c_0) = \mu_2^{pr}$. Also define

$$\mu^2(s_0) = \arg\left\{ \max_{\{\mu_\tau\}\tau\in\mathcal{T}(s_0)} \sum_{\tau\in\mathcal{T}(s_0)} p_0(\tau)V(\mu_\tau, \tau) \text{ s/t} \quad \left. \begin{array}{ll} U(\mu_\tau, t') \geqq U(s_0, t') & \forall t' \in \mathcal{T}(s_0) \\ U(\mu_\tau, \tau) \geqq U(\mu_{\tau'}, \tau) & \forall \tau, \tau' \in \mathcal{T}(s_0) \end{array} \right\}\right.,$$

where $\mathcal{T}(s_0) = \{\tau \in \mathcal{T} \mid \mu_{2\tau}^{pr} = s_0\}$. Finally define $c^*(s_0)$ as the contract such that $m(c^*(s_0)) = \mu^2(s_0)$.

$$\Phi_1^{pr} = \begin{cases} \tilde{s}_0(c_0, t) = \arg\max_{a^n\in m(c_0)} U(a^n, t) \\ \tilde{d}_2(c_0, s_0, c_2, t) = \begin{cases} 1 & \text{if } \tilde{s}_2(c_0, s_0, c_2, t) \neq s_0 \text{ and} \\ & U(\tilde{s}_2(c_0, s_0, c_2, t), t) \geqq U(s_0, t) \\ 0 & \text{otherwise} \end{cases} \\ \tilde{s}_2(c_0, s_0, c_2, t) = \arg\max_{a^n\in m(c_2)} U(a^n, t) \end{cases}$$

$$\Phi_2^{pr} = \left\{ \tilde{c}_2(c_0, s_0) = \begin{cases} c^*(s_0) & \text{if } c_0 \neq c^*(s_0) \\ \emptyset & \text{otherwise} \end{cases} \right.$$

$$p_2^{pr}(t|c_0, s_0) = \begin{cases} p_0(t)/\Sigma_{\tau\in\mathcal{T}(s_0)}\, p_0(\tau) & \text{if } t \in \mathcal{T}(s_0) \\ 0 & \text{otherwise.} \end{cases}$$

We shall now argue that these strategies and beliefs do in fact constitute a PBE of the renegotiation subgame (starting with stage 3).

In stage 4.3, type t of player 1 selects her preferred element in $m(c_2)$. In stage 4.2, player 1 accepts the new contract offer c_2 if and only if her preferred element in $m(c_2)$ is preferred to s_0. Given this acceptance rule by player 1 and his own beliefs, player 2 can do no better than offer in stage 4.1 his preferred contract among those accepted by player 1. This implies offering the contract $c^*(s_0)$ with the associated menu $\mu_2(s_0)$ if $c_0 \neq c^*(s_0)$. This is supported by the Bayesian revision of player 2's prior when s_0 is selected. No contract is offered if $c_0 = c^*(s_0)$. This is optimal given that constraints (ii) are satisfied by the equilibrium allocation μ_2^{pr}. In stage 4, player 1 selects the element of $m(c_0)$ that yields the highest utility given the expected renegotiation offer by player 2.

These strategies and beliefs imply the following equilibrium path. In stage 4, type t of player 1 selects her preferred element μ_{2t}^{pr} in $m(c_0)$. In stage 4.1, player 2 makes no contract offer. Given that μ_2^{pr} satisfies the constraints (ii), the initial contract c_0 cannot be renegotiated in stage 4.1 given the equilibrium strategy of player 1.

With these strategies along the equilibrium path it is clear that the allocation μ_2^{pr} is P2-renegotiation-proof. ||

REFERENCES

Beaudry, P., and M. Poitevin (1993) 'Signalling and renegotiation in contractual relationship.' *Econometrica* 61, 745–82

— (1994) 'The commitment value of contracts under dynamic renegotiations.' *Rand Journal of Economics* 25, 501–17

Cho, I., and D.M. Kreps (1987) 'Signalling games and stable equilibria.' *Quarterly Journal of Economics* 102, 179–221

Cho, I., and J. Sobel (1988) 'Strategic stability and uniqueness in signalling games' Working Paper No. 87-49, Graduate School of Business, University of Chicago

Dearden, J., B.W. Ickes, and L. Samuelson (1990) 'To innovate or not to innovate: incentives and innovation in hierarchies.' *American Economic Review* 80, 1105–24

Dewatripont, M. (1988) 'Commitment through renegotiation-proof contracts with third parties.' *Review of Economic Studies* 45, 377–90

— (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 (1989) 'Credit and efficiency in centralized and decentralized economies.' Mimeo, Harvard University

Forges, F. (1990) 'Some thoughts on efficiency and information.' Mimeo, CORE

Fudenberg, D., and J. Tirole (1990) 'Moral hazard and renegotiation in agency contracts.' *Econometrica* 58, 1279–319

— (1991) *Game Theory* (Cambridge, MA: MIT Press)

Hart, O. (1983) 'Optimal labour contracts under asymmetric information: an introduction.' *Review of Economic Studies*, 50, 3–35

Hart, O., and J. Tirole (1988) 'Contract renegotiation and coasian dynamics' *Review of Economic Studies* 55, 509–40

Holmström, B., and R.B. Myerson (1983) 'Efficient and durable decision rules with incomplete information.' *Econometrica*, 51, 1799–1819

Laffont, J.J., and J. Tirole (1990) 'Adverse selection and renegotiation in procurement.' *Review of Economic Studies* 57, 597–625
Maskin, E., and J. Tirole (1992) 'The principal-agent relationship with an informed principal, II: common values.' *Econometrica* 60, 1–42
Melumad, N., and S. Reichelstein (1987) 'Centralization versus delegation and the value of communication.' *Journal of Accounting Research* 25, 1–21
Milgrom, P.R. (1988) 'Employment contracts, influence activities, and efficient organization design.' *Journal of Political Economy*, 96, 42–60
Milgrom, P., and N. Stokey (1982) 'Information trade and common knowledge.' *Journal of Economic Theory*, 26, 17–27
Myerson, R.B. (1979) 'Incentive-capatibility and the bargaining problem.' *Econometrica* 47, 61–73
Noldeke, G., and E. van Damme (1990) 'Signalling in a dynamic labour market.' *Review of Economic Studies* 57, 1–23
Nosal, E. (1991) 'Implementing ex ante contracts.' Mimeo, University of Waterloo
Spence, M. (1973) 'Job market signalling.' *Quarterly Journal of Economics* 87, 355–74