ISSN 1745-8587

Birkbeck
UNIVERSITY OF LONDON

School of Economics, Mathematics and Statistics

BWPEF 0717

Unforeseen Contingency and Renegotiation with Asymmetric Information

Jihong Lee
Birkbeck,University of London

October 2007

Birkbeck Working Papers in Economics & Finance

Unforeseen Contingency and Renegotiation with Asymmetric Information

Jihong Lee*

(Birkbeck College, Univertity of London)†

*This paper grew out of the first chapter of my PhD thesis. Above all, I thank my supervisors, Luca Anderlini and Hamid Sabourian, for their guidance and kindness. I am also grateful to Leonardo Felli, Jim Mirrlees and, especially, an anonymous referee for their insightful comments.

†email: J.Lee@ems.bbk.ac.uk

Abstract

This paper considers a buyer-seller contracting model in which the seller possesses private information about all relevant aspects of the state of nature, including how much each action is worth to the buyer. We argue that, given asymmetric information, the buyer may not entirely dismiss an unforeseen contingency claim by the seller. Then, if the buyer lacks the foresight/awareness to "expect the unexpected", the model admits an equilibrium in which a seemingly complete contract is written and then renegotiated along its outcome path to generate inefficiency ex post.

JEL Classification: C79, D82, D86, L14

1 Introduction

Consider an owner of a building who wishes to install a new lift in her building. A potential seller tells the owner that two types of lifts are possible, A_1 or A_2, but she will find out which is the right choice only after conducting some initial investigation. The contract, however, needs to be written now.

As is common with many such procurement contracts, it is only the seller who observes, and then tells the buyer, which alternative best suits the buyer's needs. The uncertainty here concerns not merely the seller's own costs of installing different lifts; rather it involves all relevant aspects of the environment, including how much each lift should be worth to the buyer. The buyer and the seller therefore write a contract that postulates the terms of trade contingent on the seller's recommendation (as to which lift should be installed).

Upon finishing the initial investigation, the seller however makes an unexpected announcement: the situation turns out to be unexpectedly complex, and as a result, she cannot tell for sure which of the two lifts, A_1 or A_2, is the correct choice. The seller then recommends another lift, say G, which is a better option under these unforeseen circumstances.

What is crucial about this rather common story of renegotiation is asymmetric information. After all, there is nothing abnormal about not being able to foresee all relevant future contingencies, and moreover, the buyer has no way of verifying whether or not the claim is true. Then, how can the buyer dismiss for sure that the seller is not lying? Renegotiation may therefore happen even when the seller is making a false claim.

This paper formalises this kind of contracting situations which feature asymmetric information and the possibility of unforeseen contingency. We find that, modelled in the way described above, asymmetric information coupled with a plausible behavioural assumption on the uninformed player's foresight/awareness of the environment can result in a seemingly complete contract being written and then renegotiated along its outcome path to generate inefficiency *ex post*.

We therefore propose a fresh perspective on the role of renegotiation, based on asymmetric information and unforeseen contingency. Our approach departs from the traditional (complete information) contracting literature which views renegotiation as a constraint on contracting and explains its occurrence through some form of contractual incompleteness that it creates (Maskin and Moore, 1999; Maskin and Tirole, 1999; Hart and Moore, 1999; Segal, 1999; Che and Hausch, 1999). Also, in this literature, renegotiation is assumed to be always efficient, and hence, the source of inefficiency is often associated with the "hold up" problem *ex ante* (Hart, 1995).

To concretely illustrate our arguments, let us revisit the above example with more details. There are two possible states of nature, and the state-contingent payoff consequences of the three lifts are summarised in the table below.

	State 1			State 2		
	A_1	A_2	G	A_1	A_2	G
Buyer's valuation	10	10	6	10	10	16
Seller's cost	5	15	4	15	5	14
Net surplus	5	-5	2	-5	5	2

Thus, in one contingency, A_1 embodies the highest tradeable surplus and A_2 the lowest; in the other contingency, A_2 embodies the highest tradeable surplus and A_1 the lowest; G embodies an intermediate, non-contingent surplus. Suppose that ex ante both parties (who are risk-neutral) believe that the two states of nature are equally likely and that the seller will perfectly observe the state when revealed.

The sequence of events is as follows. First, the seller offers a contract which the buyer can accept or reject. The buyer's reservation payoff, which she obtains upon rejecting a contract offer, is zero. Second, the seller receives a private signal about the state of nature and then sends a message to the buyer. Third, the seller may opt to make a take-it-or-leave-it renegotiation offer to the buyer. If the seller opts not to make a renegotiation offer, or if she does make an offer which is rejected, the original contract is enforced.

Consider the following message contract: if the seller recommends A_2, trade A_2 at price 10; otherwise, trade A_1 at price 10. If the seller announces the truth in each state, this contract clearly delivers the first-best and the maximum possible payoff to the seller.

Suppose then that the seller encounters an unforeseen contingency and does not in fact observe the realised state, and hence, still believes that the two states are equally likely. The seller informs this to the buyer and asks to renegotiate, trading G at price $11 - \epsilon$ for some $\epsilon \in (0, 2)$ instead of trading A_1 at price 10 as stipulated by the contract (the seller has not recommended A_2). Note that if the claimed unforeseen contingency were to occur the seller would indeed prefer to renegotiate as such: under the original contract her expected payoff is $\frac{(10-5)+(10-15)}{2} = 0$, while under the terms of the renegotiation the expected payoff is $11 - \epsilon - \frac{14+4}{2} = 2 - \epsilon > 0$.

Given this, and since it is not possible to verify the seller's claim, let us assume that the buyer takes the (unexpected) message at its face value.[1] She then accepts the renegotiation offer since her expected payoff from the contracted trade is 0 while she obtains from the renegotiation $\frac{6+16}{2} - (11-\epsilon) = \epsilon > 0$.

Moreover, the following is also true here. If the buyer were to respond in such a way to the unforeseen contingency claim, then the seller would want to renegotiate even if she were to observe state 1. To see this, note that in this state the seller's payoff under the original contract is $10 - 5 = 5$, while with the renegotiated deal it is $(11 - \epsilon) - 4 = 7 - \epsilon > 5$ since $\epsilon \in (0, 2)$. It can be readily verified that, if the seller observes state 2, she is better off staying with the contract.

At the time of contracting, of course, it is *not* expected that the seller

[1]This is of course an extreme assumption. In the main analysis below, we consider the general case in which the buyer ex post believes the seller's unforeseen contingency claim to be true with any (arbitrarily small) probability.

will be unable to learn the identity of the optimal trading opportunity: this is an unforeseen contingency. Therefore, the buyer may well not anticipate that the seller could actually send such a message. This being the case, she will agree to the original contract if offered, and the seller will indeed offer it, knowing that, while she would stick with the contract in one of the two foreseen contingencies, she would be able to induce renegotiation in the other.

The net effect is that we may well see a seemingly complete contract being written and then renegotiated on the path of play, and moreover, the renegotiation leads to an inefficient trade. The seller obtains a greater payoff this way, compared to enforcing the contract and making the efficient trade in each state, at the expense of the buyer whose actual expected payoff is below her reservation payoff.

This insight is driven by the assumptions that (i) the buyer is unable to verify the truth of a message when reported and (ii) the buyer lacks the foresight to "expect the unexpected" and anticipate the possibility that the seller claims not to have learned the realised state.

While the first asymmetric information assumption is prevalent, the second assumption, which has a "behavioural" flavour, is also plausible because of asymmetric information. What this assumption captures is the likely fact that the uninformed buyer possesses limited knowledge of the environment. For example, she may not be aware of the third alternative G in the first instance and only later learns about its availability, presumably through the seller's "newly arrived" information. The seller, on the other hand, is an experienced repeat player in the business, and therefore, should know more about the environment. It is also plausible that she should have better understanding of her customer's behaviour (i.e. that the customer does not expect an unforeseen contingency claim at the contract writing stage) than the customer herself.

The rest of the paper presents a more general treatment of the above ideas and results. In Section 2, we describe the basic model (in the spirit of the "widget" models of Segal, 1999, Hart and Moore, 1999, and others), without the possibility of unforeseen contingency and renegotiation, in which asymmetric information poses no restriction to the scope of contracting and efficiency. Section 3 then introduces our two key assumptions on the buyer's beliefs in response to an unforeseen contingency claim by the seller and on the buyer's limited awareness/foresight at the contract writing stage. The extensive form concerning renegotiation and the equilibrium concept are also spelled out in this section. In Section 4, the main results are presented. We demonstrate that, even if the buyer believes an unforeseen contingency claim to be true with an arbitrarily small probability, it is possible to find an equilibrium in which a contract is written and then renegotiated along the outcome path to generate an inefficient trade. Finally, Section 5 relates our results to the recent literature on incomplete contracts and renegotiation in closer detail. In particular, we argue that our analysis offers an explanation for why individuals may indeed be reluctant to commit not to renegotiate.

2 The Basic Model

Consider a situation in which two risk neutral players, buyer (B) and seller (S), face a trade opportunity. Let $j = B, S$ index a player. B requires one unit of a widget from S. There are three widgets the players can choose from; a_1 and a_2 are *special* widgets and g is the *generic* widget. Let us denote the set of available widgets by $A = \{a_1, a_2, g\}$ and its element by a.

There are two possible states of nature, indexed by θ, drawn from the set $\Theta = \{A_1, A_2\}$ with equal probabilities. The two special widgets both yield the same fixed valuation v to B in either state of nature. S's costs of producing these widgets, on the other hand, are determined by the realisation of uncertainty. Let $c_i(\theta)$ denote the cost of producing widget a_i in state θ (for $i = 1, 2$). For expositional simplicity, we impose a symmetric structure over these parameters such that

$$c_1(A_1) = c_2(A_2) = c$$

$$c_1(A_2) = c_2(A_1) = \hat{c}\ ,$$

where $v > c$ and $\hat{c} > c > 0$. Thus, widget a_1 is "special" only in state A_1, while it is widget a_2 which offers the high-surplus generating opportunity in state A_2.

Widget g is "generic" in that it yields the same surplus in all possible states of nature. We shall denote the value of this surplus by G. Although the surplus is fixed, g's valuation and cost are state-dependent. We assume that widget g yields both higher valuation and cost in A_2 than in A_1 (think of the "gold-plated" and "cheap" imitation widgets suggested by Segal, 1999).

Specifically, we assume

$$v_g^1 - c_g^1 = v_g^2 - c_g^2 \equiv G > 0 \tag{1}$$

and

$$(v_g^2 - v_g^1 =) \; c_g^2 - c_g^1 = \hat{c} - c \equiv d > 0 \; , \tag{2}$$

where v_g^i and c_g^i are the valuation and cost of the generic widget in state A_i (for $i = 1, 2$).

Notice here that (2) assumes the cross-state difference in the generic widget's payoff consequences to be exactly equal to that in the costs of producing a special widget. This simplifies the analysis. As we shall see below, these two parametric features of the model are critical determinants of the results, but the central insights can be most economically conveyed with them treated as one parameter.

We make two further assumptions about the nature of the problem. First, a special widget, if traded in the right state, yields a greater surplus than the generic widget.

Assumption 1 $\quad v - c > G.$

Second, the generic widget surplus has the following lower bound.

Assumption 2 $\quad G > v - \frac{1}{2}(c + \hat{c}).$

This inequality implies that the generic widget performs better than a special widget under uncertainty; that is, the fixed surplus G is greater than the *expected* surplus from trading either special widget. We shall henceforth refer to this property of the generic widget as *risk dominance*.[2]

Assumptions 1 and 2 make the informational content a critical aspect of the model. We shall invoke asymmetric information. We assume that the realisation of uncertainty is observed only by S. In particular, the private information postulated here involves all relevant aspects of the environment, that is, *all* payoff consequences of *all* widgets. It is only S who learns these from nature's move.

A *trade rule* in our model is a mapping $\pi : \Theta \to 2^A \backslash \{\emptyset\}$. Let $a(\theta)$ denote the efficient widget in state θ. The *first-best* trade rule, denoted by π^*, is such that $\pi^*(\theta) = \{a(\theta)\}$ for each θ, thereby yielding the (maximum possible) total surplus of $v - c$.

The main objective of the present paper is to investigate what can be achieved by an *ex ante* contract. For this purpose, we make no explicit description of the game that may be played in the absence of a contract, such as some ex post bargaining game. We simply assume that, should the players fail to agree on a contract, they proceed to obtain some reservation payoffs, which are normalized to zero. There are always gains from engaging in a trade with each other, but the opportunity arises only if the players enter into a contractual relationship.[3]

[2] Notice here that from B's perspective trading the generic widget under uncertainty is risky while trading either special widget is not. This is a simplification. The analysis is unaltered by making B's valuation of each special widget also state-contingent.

[3] Many procurement contracts are consistent with this description. Another way to justify the need for a contract is to incorporate ex ante relationship-specific investment

Under the revelation principle, we focus on revelation contracts in which S reports her private information. Let us define $t = (a, p)$ as a *trade* specifying a widget to be traded, $a \in A$, and its price, $p \in \mathbb{R}$. Let T denote the set of all such trades. A revelation contract, z, is then defined as a function $z : \Theta \to T$ such that $z(\theta) \in T$ for each θ, specifying the terms of trade for each message.

The extensive form of the model is as follows. There are two stages. In Stage 1, the players negotiate a contract. In particular, we assume that S decides whether or not to make a take-it-or-leave-it offer of a contract which B then either accepts or rejects. If a contract is agreed, we move to Stage 2. If a contract is offered and rejected, or if S chooses not to offer a contract, the game ends with the players getting their reservation payoffs. The sequence of events in Stage 2 is as follows. At the beginning of the stage, nature draws a state; then S sends a message $\theta \in \Theta$ followed by production and trade taking place according to the terms of trade stipulated by the contract.[4] We assume for now that the players can commit not to renegotiate.

Let $E_z(\theta) \subseteq A$ be the set of equilibrium outcomes of contract z in Stage 2 when the realised state is θ. (We shall later lay out more details about the equilibrium concept when the possibility of unforeseen contingency and renegotiation are introduced. The problem is trivial at this point.) We then say that a contract z *implements* a trade rule π if $E_z(\theta) \subseteq \pi(\theta)$ for all θ.

---

such that "hold-up" is a feature of null contract and ex post (re)negotiation. The trade opportunities arise only if a sufficient level of investment is made up front but null contract induces under-investment. The contracting problem is confined to the problem of *ex post* incentives as it suffices to convey the idea of the paper.

[4] S's production capacity is just one unit of a widget.

Let z^p define a revelation contract such that $z^p(\theta) = (a(\theta), p)$ for each θ. The following result is trivial.

Proposition 1 *For any p, z^p implements π^*. (S always reveals her private information truthfully.)*

Under the bargaining structure assumed for the contract negotiation stage, an *equilibrium contract* of the model is one that gives S the maximum expected payoff while giving B her reservation payoff.[5] It then follows that:

Corollary 1 *z^v is an equilibrium contract of the model.*

What we thus have here is a simple contracting model in which asymmetric information does not limit the scope of contracting in any way. There are substantial *common interests* here, and therefore, it is straightforward to make the informed player always want to trade the correct widget. Notice also that the generic widget plays no role whatsoever. As we shall shortly see, however, the presence of this alternative plays an important role when we introduce the possibility of unforeseen contingency and renegotiation. Having such a simple, efficient benchmark will serve to highlight the impact of the added features.

[5]A contract may of course generate multiple equilibrium outcomes with different payoffs. In what follows, however, the equilibrium contract generates a unique payoff to S.

3 Unforeseen Contingency and Renegotiation

We now introduce the possibility of unforeseen contingency and renegotiation to the basic model described above.

3.1 Behavioural Assumptions

In Stage 1, the players anticipate that S will perfectly learn the state of nature when revealed. Consider now that, in Stage 2, having written a contract, S claims an unforeseen contingency that she has *not* in fact been able to observe the realised state, and therefore, is still equally uncertain as to whether it is A_1 or A_2. How should B respond?

Although ex ante B does not foresee the possibility of such an event, we take the stance that, ex post when S claims an unforeseen contingency, B does not entirely dismiss the possibility that S may be telling the truth. After all, it is a common practice not to anticipate all relevant contingencies, and moreover, there is asymmetric information; B cannot actually verify S's claim.

In order to formally model this idea, let us first introduce some notation. With slight abuse of notation, define $\Theta' = \{A_1, A_2, \phi\}$ as the set of seller "types", where A_1 (A_2) is the anticipated type who observes nature's draw and learns the state to be A_1 (A_2) and ϕ the unanticipated type who does not observe nature's draw. For expositional clarity, in what follows, θ will index a type and σ will refer to a message. Also, define q as a probability distribution on Θ', and let Q denote the set of all such distributions.

Then, we define B's *prior beliefs* over the possible seller types, denoted by ψ, as a mapping

$$\psi : \emptyset \cup \Theta' \to Q$$

such that $\psi(\sigma) \in Q$ for any $\sigma \in \emptyset \cup \Theta'$, where $\emptyset$ refers to no message. We shall sometimes refer to $\psi(\emptyset)$ as B's *ex ante* priors (held in Stage 1) and $\psi(\sigma)$ for any $\sigma \in \Theta'$ as her *ex post* priors (held in Stage 2).

We now present our key assumption below.

Assumption 3

1. *For any $\sigma \neq \phi$, $\psi(\sigma) = q^*$ such that $q^*(\phi) = 0$ and $q^*(A_1) = q^*(A_2) = \frac{1}{2}$; and*
2. *For $\sigma = \phi$, $\psi(\sigma) = q^\epsilon$ such that $q^\epsilon(\phi) = \epsilon$ and $q^\epsilon(A_1) = q^\epsilon(A_2) = \frac{1-\epsilon}{2}$ for some $\epsilon \in (0, 1]$.*

In other words, the players have common ex ante priors such that types A_1 and A_2 are equally likely and type ϕ is impossible, but B's ex post priors change if S claims an unforeseen contingency that she has not observed the realised state; in this case B forms a prior belief that there is a probability $\epsilon > 0$ with which the unforseen contingency may actually take place. Note that B does not realise the possibility of unforeseen contingency unless prompted; the transformation in B's priors is triggered only by message ϕ and not by messages A_1 or A_2. Also, Assumption 3 concerns *prior* beliefs; as we shall clarify below, B's *posterior* beliefs will be formulated in the usual Bayesian procedure from such priors and equilibrium strategies.

It is natural to treat ϵ as taking a small value. If it is large, or indeed $\epsilon = 1$, Assumption 3 amounts to saying that B takes message ϕ at (near) its face value, which is less plausible and conveys more of a bounded rationality flavour.[6] However, we do not rule our this extreme possibility. As will be shown below, the magnitude of ϵ does not affect the substantive contents of the results.

Next, we make an additional assumption about the extent of the uninformed player's forecasting capability/awareness of the environment.

Assumption 4 *B does not "expect the unexpected"; that is, in Stage 1, B anticipates only $\sigma = A_1$ or $\sigma = A_2$ in Stage 2. S knows this in Stage 1.*

As argued earlier, we view this to be a plausible behavioural assumption, particularly in light of the type of asymmetric information modelled. S can be thought of as an experienced seller in the business possessing superior knowledge of the environment than her customers who only seldom enter the market. For instance, B may initially know only about the availability of the two special widgets, while S knows that the generic widget is also available. In such a relationship, it is also likely that the seller knows more about how a customer behaves than the customer herself.[7]

[6]It is possible to interpret Assumption 3 as positing information processing errors. However, we do not emphasise this view; rather, Assumption 3 seems a realistic description of how an individual would respond to an unforeseen contingency claim that cannot be verified. From the bounded rationality perspective, this paper may be considered related to some recent papers that study the impact of information processing errors explicitly (for example, Crawford, 2003; Eyster and Rabin, 2005; Ettinger and Jehiel, 2007), and also, to the literature on *non-partitional* information structures (for example, Geanakoplos, 1989; Brandenburger *et al.*, 1992; Lipman, 1995; Rubinstein, 1998).

[7]Behavioural assumptions of similar flavour can be found in the recent literature on

Finally, in Stage 1, both players maximise their expected payoffs in terms of the common ex ante priors in which ϕ is an unforeseen contingency but S takes account of the facts that her opponent's ex post priors will change if she claims an unforeseen contingency (Assumption 3) *and* that B does not expect the unexpected (Assumption 4), while B only expects the anticipated types A_1 or A_2. We clarify below how the players choose their behaviour in Stage 2.

3.2 Contract, Renegotiation and Equilibrium

The assumptions asserted above naturally motivate renegotiation. In particular, we are interested in whether the informed player can exploit the uninformed player's imperfect awareness of the environment by writing a contract that allows her to later pretend that the unforeseen contingency has taken place and subsequently induce renegotiation to the inefficient generic widget. Since the generic widget risk-dominates the special widgets, it is the efficient widget to trade if S is telling the truth when announcing ϕ.

A contract is now defined by $z : \{A_1, A_2, \phi\} \rightarrow T$. Since type ϕ is not expected by B (Assumption 4), a natural way to interpret ϕ as a contract clause here is to think of it as representing *non*-announcement of either A_1 or A_2 (assuming that this is itself verifiable). A contract then offers a complete

---

contracting with dynamically inconsistent preferences where the agents are assumed to be "naive" and imperfectly aware of their future selves at the contract writing stage while the principal correctly knows them. See, for example, O'Donoghue and Rabin (1999), Dellavigna and Malmendier (2004) and Eliaz and Spiegler (2006).

coverage of all possible contingencies: A_1, A_2 or else.

Moreover, a contract is subject to renegotiation. The following defines the *contract game*. In Stage 2 (after a contract has been agreed and the state of nature revealed), S first sends a message; subsequently B has an option to propose a fresh take-it-or-leave-it offer of a trade which S can either accept (Y) or reject (N). If B makes a renegotiation offer and S rejects it, or if B makes no renegotiation offer, the contract is enforced.[8]

Let us now define strategies and discuss the equilibrium notion for the contract game with the possibility of unforeseen contingency and renegotiation. We look for perfect Bayesian equilibria assuming that, within this stage, B has full awareness of the environment. At the beginning of Stage 2 (after agreement of some contract), the players have common knowledge about the possibility of unforeseen contingency and B's beliefs in response to message ϕ as in Assumption 3, compute an equilibrium of the contract game and choose the strategies accordingly. Since B may now perceive that type ϕ is a possibility, we must consider what that type would do in equilibrium in order to pin down B's post-ϕ behaviour and beliefs. B fully behaves as an Bayesian updater.

We need to introduce some extra notation at this juncture. First, define the set of all (partial) histories within the contract game relevant for actions

[8] A shaper result (in terms of uniqueness of the equilibrium set) is obtained if we consider the alternative renegotiation process in which S gets to make an offer (see the next section). However, since the proposed renegotiation works against B's interests, we shall present the case in which B has the bargaining power at renegotiation.

as

$$H = \{\emptyset, (\sigma), (\sigma, t)\}_{\sigma \in \{A_1, A_2, \phi\}, t \in T} \ ,$$

where, with some abuse of notation, $\emptyset$ implies the beginning of the contract game, (σ) refers to the partial history of message σ, and finally, (σ, t) refers to the partial history of message σ followed by renegotiation offer t. Let h index an element of this set.

Also, let us define

$$H_j = \{h \in H | \text{ it is } j\text{'s turn to play after } h\}.$$

We thus have

$$H_S = \{\emptyset, (\sigma, t)\}_{\sigma \in \{A_1, A_2, \phi\}, t \in T} \ \text{ and } \ H_B = \{(\sigma)\}_{\sigma \in \{A_1, A_2, \phi\}}.$$

We can then define S's strategy as a mapping

$$f_S : \Theta' \times H_S \to \{A_1, A_2, \phi\} \cup Y \cup N$$

such that $f_S(\theta, h) \in \{A_1, A_2, \phi\} \cup Y \cup N$ for any $(\theta, h) \in \Theta' \times H_S$. We define B's strategy as a mapping

$$f_B : H_B \to T \cup \emptyset$$

such that $f_B(h) \in T \cup \emptyset$ for any $h \in H_B$, where, with some further abuse of notation, $\emptyset$ refers to no renegotiation offer.

We also want to define B's posterior beliefs. Let $\mu(\theta|h)$ denote her belief after $h \in H_B$ that S is of type $\theta \in \Theta'$. Define $\mu \equiv \{\mu(\theta|h)|\theta \in \Theta', h \in H_B\}$ as B's belief system.

Let E_z be the set of perfect Bayesian equilibria of the contract game induced by contract z. Also let $e = (f_S, f_B, \mu)$ refer to a single perfect Bayesian equilibrium (PBE). To constitute an equilibrium, f_B has to be such that, at every $h \in H_B$, B chooses the optimal renegotiation offer given the opponent strategies f_S and her posterior (Bayesian) beliefs μ, and f_S has to be optimal for each type $\theta \in \Theta'$ at every $h \in H_S$ against f_B and μ.

B's equilibrium beliefs μ are formulated by the Bayes' rule and the opponent equilibrium strategies f_S while being consistent with Assumption 3. Since $\psi(A_1) = \psi(A_2) = q^*$ such that $q^*(\phi) = 0$, it must be that $\mu(\phi|\sigma) = 0$ for $\sigma = A_1, A_2$. On the other hand, since $\psi(\phi) = q^\epsilon$ such that $q^\epsilon(\phi) = \epsilon$, $\mu(\phi|\phi)$ may be positive depending on the equilibrium (and the contract). For example, in a *pooling* equilibrium such that $f_S(\theta, \emptyset) = \phi$ for all $\theta \in \Theta'$, we must have that $\mu(\phi|\phi) = \epsilon$ and $\mu(A_1|\phi) = \mu(A_2|\phi) = \frac{1-\epsilon}{2}$; in a *separating* equilibrium such that $f_S(\theta, \emptyset) = \theta$ for each $\theta \in \Theta'$, we must have that $\mu(\phi|\phi) = 1$.

Finally, if we from now on say that a contract constitutes an equilibrium contract of the model then we mean a contract that, under the additional assumptions made in the previous sub-section, maximises S's expected payoff among the contracts that will be accepted if offered.

4 Main Result

Consider the following question: For any given $\epsilon \in (0, 1]$, does there exist a contract which (i) B will accept if offered and (ii) induce an equilibrium in which renegotiation occurs on its outcome path and S expects a payoff of at least $v - c$ (the maximum possible surplus S can guarantee herself by implementing the first-best trade rule π^*)?

Let $\hat{z}$ define the following contract:

$$\hat{z} : \begin{cases} \sigma = A_1 & \rightarrow \quad (a_1, v) \\ \sigma = A_2 & \rightarrow \quad (a_2, v) \\ \sigma = \phi & \rightarrow \quad (a_1, v) \ . \end{cases}$$

An alternative, and more convincing, interpretation of this contract is:

> CONTRACT: *If S sends message A_2, B and S trade a_2 at price v; otherwise, B and S trade a_1 at price v.*

It is straightforward to establish that in the absence of Assumption 3 this contract will implement π^*. Then, since B does not expect the unexpected in Stage 1, her ex ante *perceived* payoff from this contract is equal to her reservation payoff. Thus, we can make the following statement.

Remark 1 *B accepts $\hat{z}$ if offered in Stage 1.*

We now characterise the set of PBEs of contract $\hat{z}$ in Stage 2. In particular, it will be shown that, under certain parametric configurations, S has incentives to induce renegotiation in state A_1, but not in state A_2, such that the inefficient generic widget ends up getting traded; moreover, the resulting surplus is distributed such that, despite sub-optimality of the total surplus generated, S does better than (or at least as well as) what she would otherwise obtain from always telling the truth and trading the efficient widget.

Let us first illustrate this in the special case where $\epsilon = 1$.

Proposition 2 *Suppose that $\epsilon = 1$. Then, there exists $\hat{e} \in E_{\hat{z}}$ in which (i) renegotiation occurs in state A_1 such that S reports ϕ followed by trade of g and (ii) the contract is enforced with trade of a_2 in state A_2. Also, in this equilibrium, S's expected payoff is $v - c$ and B's expected payoff is less than zero.*

Proof. Define $u_\theta(t)$ as type θ seller's payoff from making trade t. Let us consider the following profile $\hat{e} = (\hat{f}_S, \hat{f}_B, \mu)$ where

1.

$$\hat{f}_S(\theta, \emptyset) = \begin{cases} \phi & \text{if } \theta = A_1 \text{ or } \phi \\ A_2 & \text{if } \theta = A_2 \end{cases}$$

$$\hat{f}_S(\theta, (\sigma, t)) = \begin{cases} Y & \text{if } \theta = A_1 \text{ or } A_2 \text{ and } u_\theta(t) \geq v - c \\ Y & \text{if } \theta = \phi \text{ and } u_\theta(t) \geq v - \frac{1}{2}(c + \hat{c}) \\ N & \text{otherwise} \end{cases}$$

2.

$$\hat{f}_B(\sigma) = \begin{cases} \emptyset & \text{if } \sigma = A_1 \text{ or } A_2 \\ (g, \hat{p}) & \text{if } \sigma = \phi \end{cases}$$

$$\text{where} \quad \hat{p} = v - c + c_g^1$$

3.

$$\mu(A_1|\sigma) = \begin{cases} 1 & \text{if } \sigma = A_1 \\ 0 & \text{if } \sigma = A_2 \\ 0 & \text{if } \sigma = \phi \end{cases}$$

$$\mu(A_2|\sigma) = \begin{cases} 0 & \text{if } \sigma = A_1 \\ 1 & \text{if } \sigma = A_2 \\ 0 & \text{if } \sigma = \phi \end{cases}$$

$$\mu(\phi|\sigma) = \begin{cases} 0 & \text{if } \sigma = A_1 \\ 0 & \text{if } \sigma = A_2 \\ 1 & \text{if } \sigma = \phi \end{cases} .$$

First, it is straightforward to check that the beliefs are consistent with Assumption 3 (for $\epsilon = 1$) and the Bayes' rule on the equilibrium path.

Second, let us check optimality of $\hat{f}_B$. It suffices to establish that $\hat{f}_B(\phi) = (g, \hat{p})$ is optimal. Since $d \equiv \hat{c} - c = c_g^2 - c_g^1$, we have

$$\begin{aligned} u_\phi(g, \hat{p}) &= \hat{p} - \frac{1}{2}(c_g^1 + c_g^2) \\ &= v - c + c_g^1 - c_g^1 - \frac{d}{2} \\ &= v - \frac{1}{2}(c + \hat{c}) , \end{aligned}$$

implying that $\hat{f}_S(\phi, (\phi, g, \hat{p})) = Y$.

Thus, given $\hat{f}_S$ and the beliefs μ ($\epsilon = 1$), renegotiation offer $(g, \hat{p})$ following

ϕ will yield the following expected payoff to B: (using $d \equiv \hat{c} - c = v_g^2 - v_g^1$)

$$\begin{aligned} \frac{1}{2}(v_g^1 + v_g^2) - \hat{p} &= G - (v - c) + \frac{d}{2} \\ &= G - \left[v - \frac{1}{2}(c + \hat{c})\right] \\ &> 0 \, , \end{aligned}$$

where the last inequality follows from Assumption 2. Thus, we establish optimality of $\hat{f}_B(\phi) = (g, \hat{p})$.

Third, let us check optimality of $\hat{f}_S$. For each θ, the second part on the responses to B's renegotiation offers is obviously optimal, as is $\hat{f}_S(\phi, \emptyset) = \phi$. It therefore remains to consider what each of the other types could obtain from mimicking type ϕ.

Consider type A_2. Given $\hat{f}_B(\phi)$ above, announcing message ϕ and accepting the subsequent renegotiation offer gives her a payoff

$$\begin{aligned} \hat{p} - c_g^2 &= v - c + c_g^1 - c_g^2 \\ &= v - c - d \\ &< v - c \, . \end{aligned}$$

Thus, mimicking ϕ is not worthwhile for A_2.

Consider type A_1. Given $\hat{f}_B(\phi)$ above, announcing message ϕ and accepting the subsequent renegotiation offer gives her a payoff

$$\begin{aligned} \hat{p} - c_g^1 &= v - c + c_g^1 - c_g^1 \\ &= v - c \, . \end{aligned}$$

Thus, mimicking ϕ is optimal for A_1. This completes the proof that $\hat{e} \in E_{\hat{z}}$.

Finally, it is straightforward to show that the expected payoffs of the two parties in the above equilibrium match the claim. ‖

The intuition for the above result is straightforward. B takes the face value of the report that an unforeseen contingency has occurred and the world is equally likely to be A_1 or A_2. Also, the generic widget risk-dominates either special widget. Therefore, upon receiving message ϕ, B will make a renegotiation offer to trade the generic widget instead of trading a_1 as stipulated by contract $\hat{z}$. Such renegotiation is worthwhile for S only in state A_1 because the cost of producing the generic widget is higher in state A_2. In fact, the difference in these costs across the states is sufficiently large (since we assume $c_g^2 - c_g^1 = \hat{c} - c$) that the offered renegotiation price makes S indifferent between renegotiating (and trading the inefficient generic widget) and enforcing the contract (and trading the efficient special widget).

We now demonstrate that contract $\hat{z}$ can admit the same type of equilibrium in the general case for any $\epsilon \in (0, 1)$. The additional requirement is that d, which measures both the cross-state difference in the costs/valuations of the generic widget and that in the costs of the special widgets, is sufficiently large. We want the first of these two differences $(c_g^2 - c_g^1)$ to be large for the same reason as in the special case above: in state A_1, S needs to be getting at least as much from renegotiating as from enforcing the contract. The second difference $(\hat{c} - c)$ needs to be large in order to make the renegotiation worthwhile for B, who now believes upon receiving message ϕ that type ϕ is possible with only a small probability.

Proposition 3 *For any $\epsilon \in (0,1)$, there exists $\bar{d} \in (0,\infty)$ such that, if $d \geq \bar{d}$, then there exists $\hat{e} \in E_{\hat{z}}$ in which (i) renegotiation occurs in state A_1 such that S reports ϕ followed by trade of g and (ii) the contract is enforced with trade of a_2 in state A_2. Also, in this equilibrium, S's expected payoff is $v-c$ and B's expected payoff is less than zero.*

Proof. As before, define $u_\theta(t)$ as type θ seller's payoff from making trade t. Fix any $\epsilon \in (0,1)$. Also, define

$$\bar{d} \equiv \frac{(1+\epsilon)}{\epsilon}(v-c-G) \ ,$$

and fix any $d \geq \bar{d}$. (Notice that $\bar{d} > 0$ given Assumption 1.)

Let us consider the following profile $\hat{e} = (\hat{f}_S, \hat{f}_B, \mu)$ where

1.

$$\hat{f}_S(\theta, \emptyset) = \begin{cases} \phi & \text{if } \theta = A_1 \text{ or } \phi \\ A_2 & \text{if } \theta = A_2 \end{cases}$$

$$\hat{f}_S(\theta, (\sigma, t)) = \begin{cases} Y & \text{if } \theta = A_1 \text{ or } A_2 \text{ and } u_\theta(t) \geq v-c \\ Y & \text{if } \theta = \phi \text{ and } u_\theta(t) \geq v - \frac{1}{2}(c+\hat{c}) \\ N & \text{otherwise} \end{cases}$$

2.

$$\hat{f}_B(\sigma) = \begin{cases} \emptyset & \text{if } \sigma = A_1 \text{ or } A_2 \\ (g,\hat{p}) & \text{if } \sigma = \phi \end{cases}$$

$$\text{where} \quad \hat{p} = v - c + c_g^1$$

3.

$$\mu(A_1|\sigma) = \begin{cases} 1 & \text{if } \sigma = A_1 \\ 0 & \text{if } \sigma = A_2 \\ \frac{1-\epsilon}{1+\epsilon} & \text{if } \sigma = \phi \end{cases}$$

$$\mu(A_2|\sigma) = \begin{cases} 0 & \text{if } \sigma = A_1 \\ 1 & \text{if } \sigma = A_2 \\ 0 & \text{if } \sigma = \phi \end{cases}$$

$$\mu(\phi|\sigma) = \begin{cases} 0 & \text{if } \sigma = A_1 \\ 0 & \text{if } \sigma = A_2 \\ \frac{2\epsilon}{1+\epsilon} & \text{if } \sigma = \phi \end{cases} .$$

First, it is straightforward to check that the beliefs are consistent with Assumption 3 and the Bayes' rule on the equilibrium path.

Second, let us check optimality of $\hat{f}_B$. Given $\hat{f}_S$ and μ, $\hat{f}_B(A_1)$ and $\hat{f}_B(A_2)$ are obviously optimal. Let us consider the remaining part of B's strategy. Since $d \equiv \hat{c} - c = c_g^2 - c_g^1$, we have

$$\begin{aligned} u_\phi(g, \hat{p}) &= \hat{p} - \frac{1}{2}(c_g^1 + c_g^2) \\ &= v - c + c_g^1 - c_g^1 - \frac{d}{2} \\ &= v - \frac{1}{2}(c + \hat{c}) \end{aligned} \tag{3}$$

and

$$\begin{aligned} u_{A_1}(g, \hat{p}) &= \hat{p} - c_g^1 \\ &= v - c \ . \end{aligned} \tag{4}$$

This implies that renegotiation offer $(g, \hat{p})$ would be accepted by both type ϕ and type A_1. Notice also that $\hat{f}_S(\phi, (\phi, g, p)) = N$ and $\hat{f}_S(A_1, (\phi, g, p)) = N$ for any $p < \hat{p}$.

Thus, given $\hat{f}_S$ and μ, renegotiation offer $(g, \hat{p})$ following message ϕ will yield the following expected payoff to B: (using $d \equiv \hat{c} - c = v_g^2 - v_g^1$)

$$\begin{aligned}
\mu(A_1|\phi)(v_g^1 - \hat{p}) \quad &+ \quad \mu(A_2|\phi)(v_g^2 - \hat{p}) + \mu(\phi|\phi)\left[\frac{1}{2}(v_g^1 + v_g^2) - \hat{p}\right] \\
&= \quad v_g^1 - \hat{p} + \frac{\epsilon d}{1+\epsilon} \\
&= \quad v_g^1 - v + c - c_g^1 + \frac{\epsilon d}{1+\epsilon} \\
&= \quad G - (v - c) + \frac{\epsilon d}{1+\epsilon} \\
&\geq \quad 0 \; ,
\end{aligned}$$

where the last inequality follows from $d \geq \bar{d}$. Thus, we establish optimality of $\hat{f}_B(\phi) = (g, \hat{p})$.

Third, let us check optimality of $\hat{f}_S$. For each θ, the second part on the responses to B's renegotiation offers is obviously optimal. Let us consider the first part on the message choice. Given $\hat{f}_B$, (3), (4) and μ, it is straightforward to check optimality of $\hat{f}_S(\phi, \emptyset)$ and $\hat{f}_S(A_1, \emptyset)$ as above. For the remainder, it suffices to show that for type A_2 it is not worthwhile to mimic type ϕ. By sending message ϕ and accepting the subsequent renegotiation offer $(g, \hat{p})$, type A_2 would get a payoff

$$\hat{p} - c_g^2 = v - c + c_g^1 - c_g^2 = v - c - d < v - c \; ,$$

which establishes just that. This completes the proof that $\hat{e} \in E_{\hat{z}}$.

Finally, it is straightforward to show that the expected payoffs of the two parties in the above equilibrium match the claim. $\|$

Notice that the equilibrium $\hat{e}$ established above in Propositions 2 and 3 is not the *unique* PBE of contract $\hat{z}$ because, in state A_1, S is actually indifferent between reporting the truth and claiming the unforeseen contingency; both yield a payoff $v - c$. This means that there also exists a fully separating equilibrium in which the contract is enforced to deliver the efficient trade in both states A_1 and A_2. However, it is easily seen that there cannot be any other equilibrium outcome (and payoff to S). The reason is that, given the higher cost of producing the generic widget, type A_2 would never want to mimic type ϕ, and therefore, pooling can only occur between types A_1 and ϕ as in the posited equilibrium.

Nonetheless, we emphasise the reported equilibrium, and renegotiation, for the following reasons. First, the multiplicity of equilibria is no longer an issue if we consider an alternative extensive form for the contract game and allow instead S to make a take-it-or-leave-it renegotiation offer (as in fact assumed for the example discussed in Introduction). In this case, from the proof of Proposition 2 above, it is straightforward to show that, with sufficiently large d, there cannot be a fully separating equilibrium since otherwise type A_1 could become strictly better off by deviating to mimic type ϕ and induce renegotiation. Thus, we can derive a unique PBE which implements the same trades as reported in Propositions 2 and 3 above and gives S an expected payoff strictly greater than $v - c$.

Second, we can also break the indifference and eliminate the fully separating equilibrium by relaxing the parametric assumption that $d \equiv c_g^2 - c_g^1 = \hat{c} - c$, which was made to simplify the analysis. If the two differences ($c_g^2 - c_g^1$ and $\hat{c} - c$) are treated independently, and if we have $c_g^2 - c_g^1 > \hat{c} - c$ and these differences are sufficiently large (given any ϵ), contract $\hat{z}$ has a unique equilibrium which features inefficient renegotiation just as above.

It is obvious that contract $\hat{z}$ generates the highest possible payoff for S among all the contracts that are acceptable to B. Thus, we can put together Remark 2 and Propositions 2 and 3 to state the following Corollary.

Corollary 2 *For any $\epsilon \in (0, 1]$, there exists $\bar{d} \in [0, \infty)$ such that, if $d \geq \bar{d}$, then $\hat{z}$ is an equilibrium contract of the model.*

Let us summarise our result. S possesses private information not just on her own costs of producing different widgets but also on how much each widget is worth to B. Moreover, B lacks the foresight/awareness to expect the unexpected at the contract writing stage, that S may later claim an unforeseen contingency. These two features allow S to induce agreement of a contract which will be renegotiated along its outcome path. In one of the two (foreseen) states, S falsely claims an unforeseen contingency and recommends a new action, which is in fact inefficient. Since the claim cannot be verified, B does not entirely dismiss it and attaches a small probability that S is telling the truth. Indeed, given the right parameters, there is a price at which such renegotiation is optimal for both parties. Despite inefficiency, this leads to S earning more than (or at least as much as) what she could otherwise obtain from truthfully revealing her information and enforcing the contract to trade the efficient widget. Renegotiation makes B worse off.

5 Commitment Not To Renegotiate

Our results are closely related to the recent literature on incomplete contracts and commitment not to renegotiate. In the traditional contracting/implementation models, renegotiation is treated as a constraint affecting the scope of *off-the-equilibrium* punishments (Maskin and Moore, 1999). In some cases, this goes as far as being the critical factor in driving optimality of null, hence incomplete, contracts.

The debate between Maskin and Tirole (1999), henceforth MT, and Hart and Moore (1999), henceforth HM, is particularly relevant to us. MT define unforeseen contingencies in terms of ex ante *(in)describability* of actions (widgets) and argue that, modelled as such, unforeseen contingencies *per se* need not constrain the scope of contracting as long as the contracting parties can forecast their payoff consequences. Clever message contracts can be devised such that the details of unforeseen contingencies are filled in ex post and correct incentives provided. However, this irrelevance result rests critically on the no-renegotiation assumption. HM show that, when the players cannot commit not to renegotiate, ex ante indescribability of future events can indeed generate optimal incomplete contracts.[9]

[9]Renegotiation is also responsible for incomplete contracts when the contracting environment is sufficiently *complex* (Segal, 1999) and when there are *co-operative* ex ante

But, if renegotiation is such an important source of contractual incompleteness and inefficiency, why is it so prevalent? This observation is puzzling because the theory suggests an ample reason for individuals to bind themselves not to renegotiate when writing a contract, even in the presence of unforeseen contingencies (for example, by involving a third party). We now argue that our analysis offers an explanation for why individuals may indeed be reluctant to rule out renegotiation.

Consider S offering the following contract to B:

> CONTRACT: *If S sends message A_1 (A_2), B and S trade a_1 (a_2) at price v; otherwise, S pays B a large penalty.*

Notice that this contract effectively shuts off any chance of renegotiation (provided that the penalty is large enough) and achieves the first-best.

However, we have shown that this contract may not emerge in equilibrium of the model. In the equilibrium that we identify, the players agree on a contract which is renegotiated along its outcome path; in other words, the players choose *not* to commit not to renegotiate. This results from the seller exploiting her superior informational position and her trading partner's lack of foresight/awareness.

It is also worth mentioning that the aforementioned literature on incomplete contracts and renegotiation assumes that the players have complete information and renegotiation always yields an efficient outcome ex post. Thus, renegotiation is often associated with inefficiency ex ante such as the

relationship-specific investments (Che and Hausch, 1999).

hold up problem. In contrast, notice that the type of renegotiation postulated here generates inefficiency ex post.[10]

Finally, how incomplete is the equilibrium contract highlighted in this paper? On the one hand, it is complete in the sense that it accounts for all the relevant contingencies (message ϕ is part of a contract). At the same time, however, the contract fails to be renegotiation-proof, and this may be viewed as a symptom of incompleteness.

References

Brandenburger, A., Dekel, E. and Geanakoplos, J. (1991). 'Correlated Equilibrium with Generalized Information Structures', *Games and Economic Behavior*, vol. 4, pp. 182-201.

Che, Y-K. and Hausch, D. B. (1999). 'Cooperative Investments and the Value of Contracting', *American Economic Review*, vol. 89, pp. 125-47.

Crawford, V. P. (2003). 'Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions', *American Economic Review*, vol. 93, pp. 133-49.

Dellavigna, S. and Malmendier, U. (2004). 'Contract Design and Self-Control: Theory and Evidence', *Quarterly Journal of Economics*, vol. 119, pp. 353-402.

Eliaz, K. and Spiegler, R. (2006). 'Contracting with Diversely Naive

[10]Few other papers discuss ex post inefficiency of renegotiation in the presence of informational asymmetry. See, for example, Lee (2005).

Agents', *Review of Economic Studies*, vol. 73, pp. 689-714.

Ettinger, D. and Jehiel, P. (2007). 'A Theory of Deception', *mimeo.*

Eyster, E. and Rabin, M. (2005). 'Cursed Equilibrium', *Econometrica*, vol. 73, pp. 1623-72.

Geanakoplos, J. (1989). 'Game Theory without Partitions, and Applications to Speculations and Consensus', *Cowles Foundation Working Paper.*

Hart, O. (1995). *Firms, Contracts, and Financial Structure*, Oxford: Clarendon Press.

Hart, O. and Moore, J. (1999). 'Foundations of Incomplete Contracts', *Review of Economic Studies*, vol. 66, pp. 115-38.

Lee, J. (2005). 'Incomplete Information, Renegotiation and Breach of Contract', *Economics Bulletin*, vol. 3, no. 5.

Lipman, B. (1995). 'Information Processing and Bounded Rationality: A Survey', *Canadian Journal of Economics*, vol. 28, pp. 42-67.

Maskin, E. and Moore, J. (1999). 'Implementation and Renegotiation', *Review of Economic Stidies*, vol. 66, pp. 39-56.

Maskin, E. and Tirole, J. (1999). 'Unforeseen Contingencies and Incomplete Contracts', *Review of Economic Studies*, vol. 66, pp. 115-38.

O'Donoghue, T., and Rabin, M. (1999). 'Incentives for Procrastinators', *Quarterly Journal of Economics*, vol. 114, pp. 769-816.

Rubinstein, A. (1998). *Modelling Bounded Rationality*, Cambridge: MIT Press.

Segal, I. (1999). 'Complexity and Renegotiation: A Foundation for Incomplete Contracts', *Review of Economic Studies*, vol. 66, pp. 57-82.