Bilateral Trading and Renegotiation

Robert Evans[1] and Sönje Reiche[2]

August 2011

Abstract

For a bilateral trade model with a privately informed buyer we characterize trading rules which are implementable when the seller cannot commit not to renegotiate. Let $R(v)$ be the seller's ex ante expected payoff if she offers and commits to a price of v. If R is concave, a rule can be implemented without commitment if and only if (i) it can be implemented with commitment, and (ii) each type's expected quantity is at least as high as if there were no mechanism. If R is strictly concave, the direct mechanism which implements such a rule with commitment will also implement it in any equilibrium without commitment, so the standard mechanism is robust to renegotiation. In the non-concave case there must in addition be bunching.

Keywords: Bilateral Trade, Renegotiation, Mechanism Design

[1]University of Cambridge, UK. robert.evans@econ.cam.ac.uk
[2]University of Cambridge, UK. skr29@cam.ac.uk

1 Introduction

Suppose that two or more parties, some of whom have private information, use a mechanism, or contract, to govern their relationship. It may be that, after the mechanism has delivered an outcome, there is a different outcome which all the parties would prefer, given the information revealed to them by the play of the mechanism. That is, the mechanism may not be renegotiation-proof. It has been argued by many authors that renegotiation-proofness is a robustness property that it is desirable for a mechanism to possess.

In this paper we study mechanism design with renegotiation and incomplete information. Specifically, we consider a bilateral trading model with one-sided asymmetric information (the seller's cost is common knowledge, but the buyer's value is private information). A third party designs a mechanism to govern their trading relationship. In this mechanism the buyer sends a message to the seller, resulting in some trade and payment. Then, if some of the good remains unsold, the mechanism designer cannot prevent the seller from making a take-it-or-leave-it price demand to the buyer for the remaining stock. Of course, the seller's demand at the renegotiation stage will depend on what she has learned from the mechanism, so we cannot assume that the buyer's message reveals his type because the seller, knowing the truth, would subsequently extract all the remaining surplus; this in turn would give the buyer an incentive to understate his value.

For this setting we characterize the implementable quantity and utility schedules, that is, functions mapping the buyer's type to, respectively, expected quantity of trade and expected utility, taking renegotiation into account. As in the commitment case, once the implementable expected quantity schedules have been determined, the implementable expected utility schedules can be derived by integration: utility is the integral of quantity.

Define a function R on the buyer's types by letting $R(p)$ be the seller's expected payoff if she can commit (before learning anything about the buyer's type) to a take-it-or-leave-it price of p. The most relevant case for applications is the case in which R is concave. For that class of models, it turns out (Theorem 4) that the implementable quantity schedules are essentially those which both (i) can be implemented in the standard commitment setting, when renegotiation can be prevented; and (ii) have, for every type of buyer, an expected quantity which is at least as great as it would be if there were a null mechanism followed by renegotiation. Moreover, one mechanism which implements a particular schedule is simply the same direct revelation mechanism which would implement it in the commitment case, although the equilibrium is very different. In equilibrium, rather than tell the truth with probability 1, the buyer uses a mixed strategy, as does the seller. A type v of the buyer randomizes over all messages up to v, so that the seller, given announcement v, has a posterior belief distributed over types v and above. The seller's renegotiation demand is then distributed over values v and higher. Nevertheless, strikingly, each type gets the same utility as if they told the truth and were committed to the outcome. Perhaps our most surprising result (Theorem 5) is for the case of strictly concave R. In that case, the implementation is strong. Take a quantity schedule, and associated utility schedule, as just described, and the standard direct revelation mechanism which implements it in the commitment case. The mechanism designer can give this mechanism to the parties and know that even if renegotiation cannot be prevented, every type will get the desired expected quantity and utility in any equilibrium of the game.

For the case in which R is non-concave, bunching must take place for low buyer types. Let $\tilde{R}$ be the least concave function greater than or equal to R, i.e., the least concave majorant of R. Then in any interval on which $\tilde{R} > R$, all types must get the same expected quantity. Subject to this restriction, the implementable schedules are the same as in the concave case, and the mechanism and equilibrium have the same features. If the seller is restricted to play a pure strategy then (Theorem 1) the quantity schedule can take at most two values, i.e., low types trade $q \leq 1$ in expectation, while high types trade 1.

These results can be regarded as contributing to the bargaining literature as well as to the mechanism design literature. Given a fixed bargaining game of incomplete information, one can ask: in what ways is it possible for an uninformed outsider to alter the outcome of the game by obliging the parties to sign a contract beforehand? Our framework can also be interpreted from this point of view.

Related Literature

Various notions of renegotiation-proofness for mechanisms have been proposed. In the incomplete information case, much of the literature concerns interim renegotiation, i.e., the parties have an opportunity to renegotiate before they play the mechanism. For example, Holmström and Myerson (1983) define a decision rule (or mechanism) M as *durable* if, given any type profile, and any alternative mechanism $\tilde{M}$, the players would not vote unanimously to replace M by $\tilde{M}$ if a neutral third party were to propose it to them (see also Crawford (1985), Palfrey and Srivastava (1991) and Lagunoff (1995)). Ex post renegotiation has been studied by Green and Laffont (1987), Forges (1994), and Neeman and Pavlov (2007). In these contributions the concepts employed are variations on the principle that a mechanism is (ex post) renegotiation-proof if, for any outcome x of the mechanism and any alternative outcome y, the players would not vote unanimously for y in preference to x if a neutral third party were to propose it to them. Such definitions of renegotiation-proofness have the merit that, if a given mechanism satisfies it, the mechanism is robust against all possible alternative outcomes. However, it also has the drawback that the implied renegotiation process does not have a non-cooperative character. Under an alternative modeling of this process, a renegotiation proposal would be made by one of the parties to the mechanism, or, more generally, the players would play an exogenously given non-cooperative bargaining game after the mechanism is completed.

In this paper we use the latter notion of renegotiation. This is closer to the one generally used for the complete information case (Maskin and Moore (1999), Segal and Whinston (2002)), in which, for any inefficient outcome of the mechanism, there is a single renegotiation outcome, which can be predicted by the players. It also corresponds to the approach used in the literature on contract renegotiation (e.g. Dewatripont and Maskin (1990), Hart and Tirole (1988), Laffont and Tirole (1988,1990)) in which a trading opportunity is repeated a number of times and the focus is on comparing the outcomes of long-term contracts, sequences of short-term contracts, and long-term contracts which can be renegotiated (i.e., in the two-period case, the parties are committed for one period, but in the second period there is an opportunity to change the contract). The contract renegotiation literature is concerned with analyzing the optimal mechanism from the point of view of the principal (one of the two parties to the contract). The same applies to Skreta (2006), who considers a similar model to ours, but with T periods and discounting, and shows that it is optimal for the principal to offer a price in each period. Our paper differs from this literature in that we are concerned with characterizing the set of outcome functions which could in principle be implemented by a third party.

Our analysis is also related to the literature on incomplete information bargaining beginning with Fudenberg and Tirole (1983). Firstly, one interpretation of a mechanism is that it is a device for understanding what can be achieved by non-cooperative bargaining games and secondly, as noted above, our analysis can be understood as a characterization of what can be achieved by imposing a contract on two bargainers before they begin an exogenous non-cooperative bargaining game.

Outline

Section 2 sets out the model, which has a continuum of types for the buyer. Section 3 analyzes a version of the model with three types so as to illustrate some of the arguments and results which follow. Section 4 derives necessary conditions for implementation in the continuum case. Section 5 derives sufficient conditions. Section 6 establishes the strong implementation result for the strictly concave case. Section 7 concludes. The majority of proofs are in the Appendix.

2 The Model

There are two players, a buyer (B) and a seller (S). They may trade up to one unit of a divisible good. The seller has a zero cost of production (this is a normalization) and the buyer's value v, which is privately known to the buyer, is distributed according to a distribution F on the interval $V = [\underline{v}, \overline{v}]$, where $\underline{v} > 0$. We assume that F has a differentiable density $f > 0$. Both players are risk-neutral and have quasi-linear utility for money. If S produces and trades a quantity $q \leq 1$ of the good for payment p, then S's payoff is p and B's payoff is $vq - p$. The seller's expected profit function is denoted by $R(p) = p(1 - F(p))$. This is S's expected payoff, given F, if she commits to price p. Let $p^* = \max\{\arg\max_p R(p)\}$.

The players' interaction is governed by a mechanism which is intended to determine how much to trade and at what price. A mechanism is a set of messages M and a pair of functions $q : M \to [0,1]$ and $t : M \to \Re$. B chooses a message $m \in M$. When message m is sent, $q(m)$ is the contracted quantity to be supplied by S and $t(m)$ is the contracted payment to be paid by B to S.

Renegotiation

We assume, however, that the two parties are not able to commit to the outcome of the mechanism. Specifically, if the outcome of the mechanism gives quantity $q < 1$, it is common knowledge that this outcome is inefficient. In such a situation, we would expect the parties, if they can, to try to negotiate a Pareto-superior outcome. We consider perhaps the simplest bargaining game, in which S makes a single take-it-or-leave it demand. For example, we may think of the seller as a retail outlet belonging to a large firm, which, in order to further its own wider objectives, has chosen a mechanism (price schedule) for all its retailers to employ. When a customer arrives at a shop and places an order which is less than the maximum, the salesman has an incentive to say 'you can have the rest as well for price p'. This, of course, potentially undermines the mechanism because the customer, knowing that the salesman's offer will depend on the customer's choice, will have an incentive to alter her behaviour.

More specifically, the following sequence of events takes place. S observes the outcome of the mechanism (M, q, t) and the message m which B sent. (That is, we restrict attention to direct, as opposed to mediated, communication). If $q = 1$ then the game is over. If $q < 1$ then S demands a new pair (q', t') and B either accepts or rejects it. If B accepts, then they trade q' and B pays t' to S (and they ignore the mechanism outcome). If B rejects, then the mechanism outcome is implemented, i.e. they trade q and B pays t.

We will assume that B accepts the renegotiation demand if he is indifferent. It is easy to see that S must demand $q' = 1$ and that the demand will be accepted by all types $v \geq v'$ for some threshold value v'. Equivalently, we can assume that S produces q and transfers it for payment t (as stipulated by the outcome of the mechanism), but then makes a take-it-or-leave-it unit price demand p for the remaining quantity $1 - q$, which will be accepted by all types $v \geq p$. S's payoff is then t if the demand is rejected and $t + (1 - q)p$ if it is accepted, while type v of B gets payoff $vq - t$ if $v < p$ and $v - t - (1 - q)p$ if $v \geq p$. If, as will turn out to be the case, the buyer's type is not fully revealed by the play of the mechanism, then the outcome of the post-mechanism bargaining need not be efficient since the bargaining game is one of incomplete information. This corresponds to a situation in which, whatever mechanism the parties play, there will always be time to renegotiate after it ends, but from that point there is only a finite amount of time in which to reach agreement (hence there may be unrealized gains).

Equilibrium

A mechanism (M, q, t) and the post-mechanism bargaining together define a game of incomplete information. We follow Milgrom and Weber (1985) by considering Bayes-Nash equilibria in distributional strategies (see also Crawford and Sobel (1982)). We take the message space M to be an arbitrary metric space and we define the players' strategies as measures on the space $M \times V$, where the measurable structure is given by the Borel sets. Given a measure μ on $M \times V$, we denote by μ_M and μ_V the marginal measures over M and V respectively; that is, for a Borel set $E \subseteq M$, $\mu_M(E) = \mu(E \times V)$ and similarly for μ_V.

A distributional strategy for the buyer, μ^B, is a measure over $M \times V$ such that $\mu^B_V = \eta_F$, where η_F is the measure over V corresponding to the distribution function F governing B's types. Note that this does not specify B's renegotiation strategy: as noted above, we assume that if he is type v he accepts a renegotiation demand p

if and only if $p \leq v$. Given μ^B, a distributional strategy for S is a measure μ^S over $M \times V$ which satisfies $\mu^S_M = \mu^B_M$. Corresponding to μ^B there exists, for each $v \in V$, a regular conditional probability measure $\mu^B(\cdot|v)$ over M , which we interpret as B's distribution (effectively, mixed strategy) over the messages in M if he has type v. Similarly, for each $m \in M$, there are conditional measures $\mu^S(\cdot|m)$ and $\mu^B(\cdot|m)$ over V . We interpret $\mu^S(\cdot|m)$ as S's distribution over renegotiation demands[3] in V if she receives message m and we will take $\mu^B(\cdot|m)$ to be S's belief about B's type. In the case in which $M = V$ we will refer to S's belief as $\beta(\cdot|m)$; that is, $\beta(\cdot|m) = \mu^B(\cdot|m)$. This is to avoid confusion between the seller's belief and the buyer's mixed strategy. A *strategy profile* is a pair (μ^B, μ^S) of strategies which are related as just described.

Given a mechanism (M, q, t), a strategy profile (μ^B, μ^S) and $(m, v, r) \in M \times V^2$, define expected payoff functions as follows (to simplify notation we do not include the mechanism as an argument of the functions):

$$U^S(m, r, \mu^B) = t(m) + [1 - q(m)]r\mu^B([r, \overline{v}]|m)$$

and

$$U^B(m, v, \mu^S) = vq(m) - t(m) + [1 - q(m)] \int_{\underline{v}}^{v} (v - r)d\mu^S(r|m).$$

The seller's expected payoff, given message m and renegotiation demand r, consists of the revenue delivered by the mechanism ($t(m)$) plus the expected revenue obtained by then demanding r, where the expectation is with respect to the correct belief, given m and B's strategy, about B's type. The buyer's expected payoff, similarly, consists of the payoff given by the mechanism plus her expected surplus generated by the renegotiation, with respect to S's renegotiation strategy μ^S.

Given $v \in V$ and μ^S, let $M^B(v, \mu^S) \subseteq M$ be the set of maximizers of $U^B(m, v, \mu^S)$. Given $m \in M$ and μ^B, let $V^S(m, \mu^B) \subseteq V$ be the set of maximizers of $U^S(m, r, \mu^B)$.

Definition 1: A *renegotiation equilibrium* is a profile (μ^B, μ^S) of strategies such that

[3]Clearly, we can ignore demands outside V.

(i) for each $v \in V$, $\mu^B(M^B(v, \mu^S)|v) = 1$, and

(ii) for each $m \in M$, $\mu^S(V^S(m, \mu^B)|m) = 1$.

Note that sequential rationality is built in to the equilibrium concept because the seller maximizes her expected payoff for each message, using the correct updated belief (i.e. a version of the conditional distribution) given μ^B.

Given message $m \in M$, $v \in V$ and strategy profile (μ^B, μ^S), let $\tilde{q}(m, v, \mu^S)$ be the expected quantity traded if S uses strategy μ^S, message m is sent and renegotiation demands less than or equal to v are accepted. That is,

$$\tilde{q}(m, v, \mu^S) = q(m) + [1 - q(m)]\mu^S([\underline{v}, v]|m).$$

Let $\bar{q}(v, \mu^B, \mu^S)$ be the expected quantity for type v. So

$$\bar{q}(v, \mu^B, \mu^S) = \int_M \tilde{q}(m, v, \mu^S) d\mu^B(m|v).$$

Similarly, let $\bar{u}(v, \mu^B, \mu^S)$ be the expected utility of buyer type v. That is,

$$\bar{u}(v, \mu^B, \mu^S) = \int_M U^B(m, v, \mu^S) d\mu^B(m|v).$$

We will be interested in the ways that expected utility and trade quantity may vary with the buyer's type.

Definition 2: A function $U : V \to \Re_+$ is a *r-implementable utility schedule* if there exists a mechanism (M, q, t) and a renegotiation equilibrium (μ^B, μ^S) of the mechanism such that, for all $v \in V$, $U(v) = \bar{u}(v, \mu^B, \mu^S)$. It is *strongly r-implementable* if there exists a mechanism such that, for all $v \in V$, $U(v) = \bar{u}(v, \tilde{\mu}^B, \tilde{\mu}^S)$ for every renegotiation equilibrium $(\tilde{\mu}^B, \tilde{\mu}^S)$ of the mechanism.

Definition 3: A function $Q : V \to [0, 1]$ is a *r-implementable quantity schedule* if there exists a mechanism (M, q, t) and a renegotiation equilibrium (μ^B, μ^S) of the mechanism such that, for all $v \in V$, $Q(v) = \bar{q}(v, \mu^B, \mu^S)$.

The fact that U must be non-negative reflects the fact that B's outside utility has been normalized to zero and we implicitly allow him not to participate in the mechanism. We refer to a utility schedule or quantity schedule as *c-implementable* if it can be implemented in the case in which the players can be committed to the mechanism. By standard results (see Fudenberg and Tirole (1993), Milgrom and Segal (2002)) Q is c-implementable if and only if $Q(\cdot)$ is non-decreasing, and $U \geq 0$ is c-implementable if and only if, for all $v \in V$, $U(v) - U(\underline{v}) = \int_{\underline{v}}^{v} Q(u)du$ for some non-decreasing function $Q : V \rightarrow [0, 1]$. A c-implementable U is absolutely continuous and a.e. differentiable. r-implementability and c-implementability are related as follows.

Proposition 1 *If U (resp. Q) is r-implementable then U (resp. Q) is c-implementable.*

Proof Consider, for the model with commitment, a direct revelation mechanism which, given announced type v, reproduces the random outcome of the equilibrium (μ^B, μ^S) of the mechanism which r-implements U or Q, including the buyer's acceptance strategy, assuming he is of type v. Clearly it is optimal for the buyer to tell the truth in this mechanism. QED

Although the profit-maximizing mechanism for the seller is simply a take-it-or-leave-it price offer there are several reasons why we may be interested in implementing other quantity schedules. For example, as in the hold-up literature, there may be a prior investment stage. Suppose, for example, that the buyer first chooses a level of costly unverifiable investment and that higher investment will lead, on average, to a higher value for the buyer. It can be shown that in that case the optimal quantity schedule for the seller (taking into account the need to give investment incentives to the buyer) can be strictly increasing over a range of type values. For a second example, consider a case in which the seller is a division of a bigger firm. The division wants to maximize profits but the headquarters imposes constraints; in particular, a constraint that everybody must be served with some probability, perhaps increasing in willingness-to-pay. It may be, for example, that if consumers are shut out of the market then they are less likely to buy in other markets served by the firm, or in the same market in the future. Alternatively, it may be that there are learning effects: the firm may want to serve low-value customers because they may discover a taste for the good and then become high-value types in the future. If consumers and firm divisions have a one-period perspective but the headquarters takes long-term considerations into account then the headquarters may want to impose an increasing quantity schedule on the division.

3 An Example with Three Types

In this section we examine a version of our model with three types. It turns out that this simple model can illustrate some of the results and arguments that apply to the continuum model set out in the previous section. The three types are v_0, v_1 and v_2, with $0 < v_0 < v_1 < v_2$. The prior probability of type v_i is denoted by $p_i > 0$ $(i = 0, 1, 2)$. Let $R(v_i) = v_i \sum_{j \geq i} p_j$. $R(v_i)$ is the seller's expected revenue if she demands a price of v_i. We assume that R is strictly increasing and strictly concave. This implies that the unique optimal demand for the seller, in the absence of any mechanism, is v_2. Strict concavity of R is equivalent to the following condition (see expression (7) below).

$$\frac{p_1}{v_2 - v_1} > \frac{v_0}{v_2}(\frac{p_0}{v_1 - v_0}). \tag{1}$$

Take any pair of numbers q_0, q_1 such that $0 \leq q_0 < q_1 < 1$. Suppose that we want to implement the quantity schedule in which type v_i $(i = 0, 1)$ gets quantity q_i and type v_2 gets $q_2 = 1$. Suppose also that we want the utilities of types v_1 and v_2 to be the lowest consistent with this, given the utility of type v_0, (i.e., we want the downward incentive constraints to bind). Then, if we can commit to a mechanism, we can use an incentive-compatible direct revelation mechanism $(q_i, t_i)_{i=0,1,2}$, in which the message set is $\{0, 1, 2\}$ and the contracted quantity and payment, given message i, are q_i and t_i. Let U_i be the expected utility of type v_i in this mechanism. Then $U_0 = v_0 q_0 - t_0$, $U_1 = U_0 + q_0(v_1 - v_0)$ and $U_2 = U_1 + q_1(v_2 - v_1)$. Let $U_i(j)$ denote the payoff of type v_i if he sends message j. Then, if $j > i$, $U_i(j) < U_i(i) = U_i$. Also $U_2(0) < U_2$.

Suppose now that, after the mechanism is played, if the contracted quantity is less than 1, i.e. if the message i was either 0 or 1, the seller makes a take-it-or-leave-it price offer for the remaining quantity $1 - q_i$. The combination of communication stage and subsequent renegotiation defines a dynamic game of incomplete information. Clearly there will be no perfect Bayesian equilibrium in which the buyer tells the truth with probability 1. For example, if there were such an equilibrium, type v_1 could then send message 0, the seller would offer a renegotiation price of v_0, and the buyer's payoff would be $v_1 q_0 - t_0 + v_1(1 - q_0) - v_0(1 - q_0)$. This is strictly greater than $v_1 q_0 - t_0 = U_1$. If he were to tell the truth (announce 1) he would get no surplus in the renegotiation (the seller would offer v_1) and so his payoff would be U_1.

Nevertheless, in *any* perfect Bayesian equilibrium of this game, the expected utilities and quantities are the same, for each type of buyer, as in the truthful equilibrium of the commitment case.

First, we construct an equilibrium which has this property.[4] The message strategies of the three types of buyer are as follows. Each type puts strictly positive probability on its true type and also on each lower type, but none on any higher type. Thus, v_0 sends message 0 with probability 1. v_1 sends 0 with probability α and 1 with probability $1 - \alpha$, where

$$\alpha = \frac{p_0 v_0 (v_2 - v_1)}{p_1 v_2 (v_1 - v_0)}.$$

v_2 sends 0 with probability β_0, 1 with probability β_1 and 2 with probability $1 - \beta_0 - \beta_1$, where

$$\beta_0 = \frac{p_0 v_0 v_1}{p_2 v_2 (v_1 - v_0)}$$

and

$$\beta_1 = \frac{v_1 p_1}{p_2 (v_2 - v_1)} - \beta_0.$$

These strategies are well-defined: it follows from (1) that $\alpha < 1$ and that $\beta_1 > 0$. The

[4]The equilibrium is not unique. There is a double continuum of other equilibria: replace α, β_0 and β_1 with $\hat{\alpha}$, $\hat{\beta}_0$ and $\hat{\beta}_1$ such that $\hat{\alpha} \in [\alpha, 1]$, $\hat{\beta}_0 = v_1 \hat{\alpha} p_1 [p_2(v_2 - v_1)]^{-1}$ and $\hat{\beta}_1 \in [\beta_0 + \beta_1 - \hat{\beta}_0, 1]$. But, as we will show, all these equilibria are payoff-equivalent.

fact that $\beta_0 + \beta_1 < 1$ follows from the fact that $R(v_2) > R(v_1)$. Given a renegotiation offer p, type v_i accepts if and only if $p \le v_i$. The seller's renegotiation strategy is as follows. After message 0, she randomizes over offers v_1 and v_2: she offers v_1 with probability $(q_1 - q_0)(1 - q_0)^{-1}$ and she offers v_2 with the complementary probability, $(1 - q_1)(1 - q_0)^{-1}$. After message 1 she offers v_2 with probability 1. (After message 2 there is no action to take, since $q_2 = 1$).

Let $p(i|j)$ denote the seller's probability, given message j, that she is facing type v_i. Since each message has strictly positive probability these probabilities are given by Bayes' Rule. It is straightforward to check that $v_0 = v_1(p(1|0)+p(2|0)) = v_2 p(2|0)$ and that $v_1 = v_2 p(2|1)$. Therefore, after message 0 the seller's optimal offers are v_0, v_1 and v_2; and, after message 1, her optimal offers are v_1 and v_2. This shows that the seller's strategy is optimal.

To see that the buyer's strategy is optimal, first note that type v_0 cannot get any renegotiation surplus no matter what message he sends, since the renegotiation offer will be at least v_0. Therefore his payoff from message j is $U_0(j)$, which is maximized by message 0, giving payoff U_0. Similarly, since there will be no offer below v_1, type v_1 gets $U_1(j)$ from message j. Since $U_1(0) = U_1(1) = U_1 > U_1(2)$, his optimal messages are 0 and 1. If type v_2 sends message 2 his payoff is U_2. If he sends message 1 he will get no renegotiation surplus so his payoff will be $U_2(1) = U_2$. If he sends message 0 he will get quantity 1, taking renegotiation into account. His expected payment is

$$t_0 + \frac{(q_1 - q_0)}{(1 - q_0)} v_1(1 - q_0) + \frac{(1 - q_1)}{(1 - q_0)} v_2(1 - q_0) = t_0 + (q_1 - q_0)v_1 + (1 - q_1)v_2,$$

which, since $v_1 q_0 - t_0 = v_1 q_1 - t_1$ and $v_2 q_1 - t_1 = v_2 - t_2$, is equal to t_2. Therefore his payoff is $v_2 - t_2 = U_2$. This shows that the above is an equilibrium.

Now suppose that there is some other equilibrium, in which type v_i has equilibrium payoff $\tilde{U}_i$ and expected quantity $\tilde{q}_i$. Let the seller's beliefs be $\tilde{p}(i|j)$. It is easy to see that $\tilde{U}_i \ge U_i$ for each i because it is always possible to tell the truth and then decline to renegotiate. It is also clear that, in this equilibrium, type v_0 cannot get any renegotiation surplus, so it must be that he sends message 0 with probability 1. Hence $\tilde{U}_0 = U_0$.

Suppose that $\tilde{U}_1 > U_1$. If v_1 sends $i > 0$ with strictly positive probability in equilibrium the seller will then demand at least v_1 after message i (since type v_0 does not send i). Therefore he gets $U_1(2) < U_1$ if $i = 2$ and he gets U_1 if $i = 1$. Hence he must send 0 with probability 1. Since, by hypothesis, he gets more than U_1 in equilibrium, he must get some renegotiation surplus, so the seller must offer v_0 with strictly positive probability after message 0. This implies that, in particular,

$$v_0 \geq v_1(\tilde{p}(1|0) + \tilde{p}(2|0)).$$

Suppose that v_2 is also an optimal offer for the seller. Then

$$v_2\tilde{p}(2|0) \geq v_1(\tilde{p}(1|0) + \tilde{p}(2|0)).$$

These two inequalities imply (see (6) below) that

$$\tilde{p}(1|0) \leq \frac{(v_2 - v_1)v_0}{(v_1 - v_0)v_2}\tilde{p}(0|0),$$

which, since types v_0 and v_1 pool on message 0, implies that

$$p_1 \leq \frac{(v_2 - v_1)v_0}{(v_1 - v_0)v_2}p_0,$$

which contradicts (1). This shows that the seller will not offer v_2 after message 0. Therefore $\tilde{q}_1 = 1$. Clearly $\tilde{q}_2 = 1$, so

$$\tilde{U}_2 = \tilde{U}_1 + (v_2 - v_1) > U_1 + (v_2 - v_1)$$

$$> U_1 + q_1(v_2 - v_1) = U_2.$$

This shows that $\tilde{U}_2 > U_2$. This in turn implies that v_2 does not send message 2 (which would give payoff U_2). Nor can he send message 1 with positive probability in equilibrium - if he did, the renegotiation offer would be v_2 (since neither v_0 nor v_1 sends 1) and so his payoff would be $U_2(1) = U_2 < \tilde{U}_2$. Therefore v_2 must send 0 with probability 1. However, if all three types pool on message 0, $p(i|0) = p_i$ for $i = 0, 1, 2$, and so, since $R(v_2) > R(v_0)$, v_0 cannot be an optimal offer for the seller after message 0. This contradiction establishes that $\tilde{U_1} = U_1$.

Suppose then that $\tilde{U}_0 = U_0$, $\tilde{U}_1 = U_1$ and $\tilde{U}_2 > U_2$. After message 0, the seller does not offer v_0 because this would give a strictly positive renegotiation surplus to v_1, contradicting the fact that $\tilde{U}_1 = U_1$. $\tilde{U}_2 > U_2$ implies that v_2 does not send message 2 in equilibrium. Therefore message 2 has zero probability. If message 1 has zero probability then all three types pool on message 0, implying that the renegotiation offer is v_2, giving v_2 strictly less than U_2. Hence message 1 has strictly positive probability. If, for $i = 0, 1$,

$$v_1(\tilde{p}(1|i) + \tilde{p}(2|i)) < v_2\tilde{p}(2|i) \tag{2}$$

then, after message i, the seller offers v_2 with probability 1. This implies that v_2 gets no renegotiation surplus after message i, so, since $\tilde{U}_2 > U_2$, v_2 does not send message i in equilibrium. Therefore $\tilde{p}(2|i) = 0$, contradicting (2). Therefore, for $i = 0, 1$, that is, for each positive probability message,

$$v_1(\tilde{p}(1|i) + \tilde{p}(2|i)) \geq v_2\tilde{p}(2|i),$$

which implies that $v_1(p_1 + p_2) \geq v_2 p_2$, contradicting the fact that $R(v_2) > R(v_1)$. This shows that $\tilde{U}_2 = U_2$.

Therefore, in any perfect Bayesian equilibrium, the utility of each type is the same as in the commitment case. It can also be shown that the expected quantity bought by type v_i is q_i, the same as in the truthful equilibrium of the commitment case. We show in the following sections that this generalizes to the continuum model set out in Section 2. Specifically, suppose the profit function R is strictly concave and consider any non-decreasing quantity schedule which gives each type at least as much as he would get if there were no mechanism (simply a bargaining stage). Then the incentive-compatible direct revelation mechanism which implements this quantity (and utility) schedule if the parties can commit will also implement them, in any equilibrium, if they cannot commit.

4 Implementable Schedules: Necessary Conditions

In this section and the next we ask which quantity and utility schedules can be implemented in our general model of Section 2, when renegotiation cannot be prevented. Here we derive necessary conditions for r-implementability.

Fix an arbitrary mechanism (M, q, t) and an arbitrary renegotiation equilibrium (μ^B, μ^S) of the mechanism. Henceforth we will typically drop the arguments μ^B and μ^S, for example in $M^B(v, \mu^S)$, $V^S(m, \mu^B)$ and $\tilde{q}(m, v, \mu^S)$. Several results below are true for "almost all" messages; that is, there is a subset N of M which has measure zero according to μ^B_M such that each of these results holds for all messages outside N.

Our first Lemma shows that, as in the no-renegotiation case, expected quantity must be monotonically non-decreasing in type both on average (which follows from the fact that expected quantity, being r-implementable, is also c-implementable) and also message by message. Recall that $\tilde{q}(m, v)$ is type v's expected quantity after message m:

$$\tilde{q}(m, v) = q(m) + (1 - q(m))\mu^S([\underline{v}, v]|m). \tag{3}$$

Let $\tilde{q}(m, v_-)$ be the expected quantity of buyer type v if he sends message m and accepts a renegotiation offer r if and only if $r < v$ (which is an optimal acceptance strategy, although we assume that in equilibrium B accepts if indifferent).

Lemma 1 (Monotonicity)

(i) Take any v and v' such that $v' > v$. If $m \in M^B(v)$ and $m' \in M^B(v')$, i.e. m is optimal for v and m' is optimal for v', (a) $\tilde{q}(m, v) \leq \tilde{q}(m', v')$ and (b) $\tilde{q}(m, v) \leq \tilde{q}(m', v'_-)$.

(ii) If $v' > v$, $\bar{q}(v', \mu^B, \mu^S) \geq \bar{q}(v, \mu^B, \mu^S)$.

In equilibrium a given type v may randomize over many messages, but the next Lemma shows that each of these messages, taking renegotiation into account, must lead to the same expected trade quantity.

Lemma 2 *For all $v \in V$ and almost all messages m and m' with $m, m' \in M^B(v)$, $\tilde{q}(m,v) = \tilde{q}(m',v)$.*

To see why Lemma 2 is true, suppose that $\tilde{q}(m',v) < \tilde{q}(m,v)$. Then $\tilde{q}(m',v) < 1$, so higher types than v also find it optimal to send m' (otherwise S would never demand more than v after m', so v's quantity would be 1). There are two possibilities: (a) it is optimal for types arbitrarily close to v, and above v, to send m'; (b) there exists $v' > v$ such that m' is optimal for v', but not for any type in (v, v'), hence S doesn't make any demand in this interval after message m'. In case (a), since expected quantity given m' is continuous in v to the right, there are types to the right of v whose expected quantity is close to $\tilde{q}(m',v) < \tilde{q}(m,v)$, which violates monotonicity (Lemma 1). In case (b), it is optimal for v' to send m' and accept only renegotiation demands strictly less than v'. Since S will not demand prices in (v, v'), this gives expected quantity $\tilde{q}(m',v) < \tilde{q}(m,v)$, which, since $v' > v$, violates monotonicity.

Definition 4: A type $v \in int(V)$ is a point of increase of $\bar{q}$ if, for all $v_1 < v < v_2$,

$$\bar{q}(v_1) < \bar{q}(v_2).$$

The next Lemma is the key to characterizing the implementable quantity schedules.

Lemma 3 *Suppose that $v \in V$ is a point of increase of $\bar{q}$. Then for almost all $m \in \cup_{\tilde{v}<v} M^B(\tilde{v})$,*

(i) $v \in V^S(m)$ and $m \in M^B(v)$;

(ii) $\mu^S((v-\varepsilon, v+\varepsilon]|m) > 0$ for all $\varepsilon > 0$.

That is, suppose that v is a point of increase of the quantity schedule, and take any message m which is optimal for some type lower than v. Then m must also be optimal for v and S's renegotiation strategy must, after m, put strictly positive probability on any neighbourhood of v. Essentially, the support of v's message strategy must include all the supports of the strategies of lower types. A very crude intuition for this is that if low types send different messages from high types then after renegotiation the low types must trade for sure, and so there can be no higher types which are points of increase of the quantity schedule. A somewhat more precise argument is as follows. Suppose $\bar{q}(v) > \bar{q}(\tilde{v})$ for all $\tilde{v} < v$ and suppose that it is suboptimal for S to demand v after message m in the support of some type lower than v. Let $v' < v$ be the highest type below v which sends m. If no type above v' sends m then v' trades for sure after sending m and so, by Lemma 2, $\bar{q}(v') = 1$, which contradicts the fact that $\bar{q}(v) > \bar{q}(v')$. Let $v'' > v$ be the lowest type above v which sends m. Then, after m, S will not demand prices in (v', v''). As argued after the statement of Lemma 2, this violates monotonicity since it is optimal for v'' to send m and accept only demands strictly below v'', giving expected quantity $\bar{q}(v') < \bar{q}(v)$.

Lemma 3 implies that there must be some bunching of types; that is, some types must, taking renegotiation into account, trade the same amount as other types. This follows immediately from Lemma 3*(ii)* in the case in which μ^S is pure, i.e., if, for any m, there exists $v \in V$ such that $\mu^S(\{v\}|m) = 1$. In fact, in that case, the types must be grouped into at most two intervals; in other words, there can only be one point of increase of $\bar{q}$. Suppose that, for some equilibrium (μ^B, μ^S) in which μ^S is pure, the function $\bar{q}(v)$ has a point of increase at $v_i > \underline{v}$ and another at $v_j > v_i$. Then, after any message in the union of the supports of types strictly below v_i, Lemma 3 implies that μ^S puts strictly positive probability in the neighbourhood of v_i and also that of v_j, which contradicts the fact that μ^S is pure. Therefore, the quantity schedule must be a step function with at most two steps: quantity is equal to $\hat{q} \leq 1$ up to some critical type, and equal to 1 for all higher types.

Theorem 1   For any renegotiation equilibrium (μ^B, μ^S) of any mechanism (M, q, t), if μ^S is a pure strategy then, for some $\hat{q} \leq 1$ and $\hat{v} \leq p^*$, $\bar{q}(v)$ takes the form $\bar{q}(v) = \hat{q}$ for $v < \hat{v}$; $\bar{q}(v) = 1$ for $v > \hat{v}$.

Our aim is to characterize the implementable schedules in the case in which the seller may use a mixed strategy. It turns out that in this case too, bunching must take place. We show below that all types above p^* (the highest optimal renegotiation demand when the belief is F) must trade quantity 1. But for lower types too there may be bunching, depending on the shape of the profit function R.

To show this, we take a v which is a point of increase of the quantity schedule $\bar{q}$, and we consider the renegotiation demand decision of S after a message m which is optimal for some arbitrary type $\tilde{v} < v$; that is, $m \in \cup_{\tilde{v}<v} M^B(\tilde{v})$. Fix a pair (v_1, v_2) such that $v_1 < v < v_2$. By Lemma 3, it is optimal for S to demand v rather than v_1 or v_2, so we obtain the following inequalities.

$$v\mu^B([v, \bar{v}]|m) \geq v_1\mu^B([v_1, \bar{v}]|m),$$

and

$$v\mu^B([v, \bar{v}]|m) \geq v_2\mu^B([v_2, \bar{v}]|m).$$

Rearranging:

$$(v - v_1)\mu^B([v, \bar{v}]|m) \geq v_1\mu^B([v_1, v)|m) \tag{4}$$

and

$$v\mu^B([v, v_2)|m) \geq (v_2 - v)\mu^B([v_2, \bar{v}]|m). \tag{5}$$

Using (4) we can write

$$\begin{aligned}\mu^B([v_2, \bar{v}]|m) &= \mu^B([v, \bar{v}]|m) - \mu^B([v, v_2)|m) \\ &\geq \frac{v_1}{v - v_1}\mu^B([v_1, v)|m) - \mu^B([v, v_2)|m).\end{aligned}$$

Then, replacing $\mu^B([v_2, \bar{v}]|m)$ in (5) and simplifying, we obtain

$$\mu^B([v, v_2)|m) \geq \frac{(v_2 - v)v_1}{(v - v_1)v_2}\mu^B([v_1, v)|m). \tag{6}$$

Integrating over $M_v = \cup_{\tilde{v}<v} M^B(\tilde{v})$ with respect to μ_M^B gives

$$\mu^B([v, v_2)|M_v) \geq \frac{(v_2 - v)v_1}{(v - v_1)v_2}\mu^B([v_1, v)|M_v),$$

which, using Bayes' Rule, implies that

$$\mu^B(M_v|[v,v_2))\eta_F([v,v_2)) \geq \frac{(v_2-v)v_1}{(v-v_1)v_2}\mu^B(M_v|[v_1,v))\eta_F([v_1,v)).$$

Since $\mu^B(M_v|[v_1,v)) = 1 \geq \mu^B(M_v|[v,v_2))$, it follows that

$$\frac{F(v_2)-F(v)}{(v_2-v)} \geq \frac{v_1}{v_2}\frac{F(v)-F(v_1)}{(v-v_1)}. \tag{7}$$

Since

$$v_2F(v_2) - v_2F(v) = R(v) - R(v_2) + (v_2-v)[1-F(v)]$$

and

$$v_1F(v) - v_1F(v_1) = R(v_1) - R(v) + (v-v_1)[1-F(v)]$$

we have

$$\frac{R(v_2)-R(v)}{(v_2-v)} \leq \frac{R(v)-R(v_1)}{(v-v_1)} \tag{8}$$

and, therefore,

Lemma 4 *Take any v which is a point of increase of $\bar{q}$ for some equilibrium (μ^B,μ^S) of some mechanism. Then, for all v_1, $v_2 \in V$* with $v_1 < v < v_2$

$$R(v) \geq \frac{v_2-v}{v_2-v_1}R(v_1) + \frac{v-v_1}{v_2-v_1}R(v_2). \tag{9}$$

That is, the chord between $(v_1, R(v_1))$ and $(v_2, R(v_2))$ must lie weakly below the graph of R at v. It follows from Lemma 4 that we can characterize the points at which the schedule *cannot* increase in terms of the least concave majorant of R, i.e., the least concave function that lies above R.

Let $C(R)$ be the set of pairs $(v,y) \in V \times \Re$ such that there exist $v_1 \in V, v_2 \in V$ and $\lambda \in [0,1]$ for which $(v,y) = \lambda(v_1, R(v_1)) + (1-\lambda)(v_2, R(v_2))$. This is the set consisting of the union of the graph of R and all the chords between points on the graph. Let $\tilde{R}$ be the function whose graph is the upper boundary of $C(R)$. It is straightforward to show that $\tilde{R}$ is the least concave majorant of R. By continuity of R, the set of values of v such that $\tilde{R}(v) > R(v)$, if non-empty, is a union of disjoint intervals $I = \cup_{\gamma \in A}(a_\gamma, b_\gamma)$. The quantity schedule must be flat on any such interval.

Lemma 5 *If $\tilde{R}(v) > R(v)$, then v cannot be a point of increase of the quantity schedule $\bar{q}$; that is, $\bar{q}$ is constant on each interval $(a_\gamma, b_\gamma)_{\gamma \in A}$.*

A sufficient condition for R to be concave (in fact strictly concave) on $[\underline{v}, p^*)$ is that F satisfies the increasing hazard rate assumption,[5] i.e., $\frac{d}{dv}\{f(v)[1-F(v)]^{-1}\} \geq 0$. If this is satisfied then there can be no v such that $R'(v) = 0$ and $R''(v) \geq 0$, so R has a unique maximum p^* and is strictly increasing below p^*. Furthermore, for $v < p^*$, since $R'(v) > 0$, we have $f'(v) > -\frac{f(v)}{v}$, so $R''(v) < 0$.

Even in the case in which R is a concave function, and so $\tilde{R} = R$, there must exist an interval on which the quantity schedule is flat since equilibrium quantity must equal 1 for types above p^*.

Lemma 6 *$\bar{q}(v) = 1$ for $v > p^*$.*

The idea of the proof of Lemma 6 is the following. Suppose that in some equilibrium (μ^B, μ^S) of some mechanism there is a point of increase $\hat{v} > p^*$. Then, after any message sent by a type lower than $\hat{v}$, one optimal demand is $\hat{v}$, by Lemma 3. Suppose we change the strategies of types $\hat{v}$ and above so that they are all the same (the average strategy of these types under μ^B), leaving the strategies of lower types unchanged. Call this new strategy $\tilde{\mu}^B$. Then, conditional on $v \geq \hat{v}$, S's belief under $\tilde{\mu}^B$ is a truncation of F, which implies that $\hat{v}$ is the optimal demand (the least concave majorant $\tilde{R}$ must be decreasing above p^*, and if $\hat{v}$ is a point of increase, it must be that $R(\hat{v}) = \tilde{R}(\hat{v})$, so $R(\hat{v}) \geq R(v)$ for all $v \geq \hat{v}$). $\hat{v}$ remains optimal, given $\tilde{\mu}^B$, after messages sent by types below $\hat{v}$, since the change does not affect the probability of acceptance of any demand $v \leq \hat{v}$, while demands $v \geq \hat{v}$ will only be accepted by types $\hat{v}$ and higher and, conditional on these types, $\hat{v}$ is optimal, as just argued. Therefore $\hat{v}$

[5]cf., for the classic adverse selection model, the analysis of bunching in the principal's optimal contract (e.g. Mussa and Rosen (1978)).

is optimal regardless of the message sent, hence optimal ex ante. But this contradicts the assumption that p^* is the maximal ex ante optimal demand. It follows that $\hat{v}$ cannot be a point of increase, so the quantity schedule must reach 1 at p^* or below.

Lemma $A.2$ in the Appendix shows that $\bar{q}$ must be right-continuous on $(\underline{v}, \bar{v}]$. In summary,

Theorem 2 (Necessary conditions for r-implementability) *If* $Q : V \to [0, 1]$ *is a r-implementable quantity schedule then* Q *is non-decreasing, right-continuous on* $(\underline{v}, \bar{v}]$ *and constant on each interval* $(a_\gamma, b_\gamma)_{\gamma \in A}$*, and* $Q(v) = 1$ *for* $v \geq p^*$*.*

5 Implementable Schedules: Sufficient Conditions

Here we show that any quantity schedule Q which satisfies the necessary conditions of Theorem 2 can be r-implemented.

Take such a schedule Q and assume first that Q is right-continuous at $\underline{v}$ as well as on the rest of V. Since Q is non-decreasing and right-continuous and $Q(\bar{v}) = 1$, Q can be regarded as a cumulative distribution function of a random variable taking values in V. We define a direct revelation mechanism $\hat{\Gamma}(Q) = (V, Q(.), t(.))$ such that, for all $v \in V$, $t(v) = \int_{\underline{v}}^{v} u dQ(u)$. In other words, the unit price paid by B if he claims to be type v is the expectation of his value conditional on his true type being less than or equal to v, where the expectation is taken with respect to the distribution Q. We will show that Q is r-implemented by $\hat{\Gamma}$. It turns out in fact that $\hat{\Gamma}$ is incentive-compatible and so is the same as the direct revelation mechanism which implements Q in the standard commitment case. Of course, in our setting, B does not necessarily tell the truth. The equilibrium is a generalization of the one we derived for the three-type example of Section 3. A buyer of type $v \leq p^*$ randomizes over all announcements up to v in such a way that S is indifferent, given message m, between all renegotiation demands in $[m, p^*]$, excluding those in I. The distribution of S's demands is then Q conditional on B's type being at least m.

First we show, by construction, that there exists a strategy μ^B for the buyer with the above property; namely, given announcement v, S's optimal demands are $[v, p^*]/I$. We do this here for the case in which I is empty; the other case is considered in the Appendix.

If I is empty then $\tilde{R} = R$ and so R is concave. p^* maximizes $R(v) = v[1 - F(v)]$, so

$$1 - F(p^*) = p^* f(p^*). \tag{10}$$

Define a function $\phi : V \rightarrow \Re$ by $\phi(v) = v^2 f'(v) + 2vf(v)$. R is concave if and only if $\phi(v) \geq 0$ for all $v \in V$.

Suppose that the buyer's message strategy is as follows. If his value is $v \leq p^*$ then he randomizes over all messages up to and including his true type v according to distribution function $(f(m)m^2)(f(v)v^2)^{-1}$. That is, for $m \leq v$,

$$\mu^B([\underline{v}, m]|v) = \frac{f(m)m^2}{f(v)v^2}$$

and, for $m > v$, $\mu^B([\underline{v}, m]|v) = 1$. If his value is $v > p^*$ then he has the same strategy as type p^*. Each type therefore puts a mass point on $\underline{v}$; the strategy of type $v \leq p^*$ has density $f^B(m|v) = \phi(m)[f(v)v^2]^{-1}$ for $m \leq v$ and that of a type above p^* has density $f^B(m|v) = \phi(m)[f(p^*)(p^*)^2]^{-1}$ for $m \leq p^*$. These distributions are well-defined because $\phi(m) \geq 0$.

After any message $m \leq p^*$, S puts zero probability on types below m since such types would not have sent message m. Therefore demands below m are suboptimal.

Consider first case (a), in which $\underline{v} < m \leq p^*$. Let the density of S's updated belief given m be denoted by $g^B(v|m)$. By Bayes' Rule, for $v_1 \in [m, \bar{v}]$,

$$g^B(v_1|m) = \frac{f^B(m|v_1)f(v_1)}{\int_m^{\bar{v}} f^B(m|v)f(v)dv}.$$

The denominator is equal to

$$\int_m^{p^*} \frac{\phi(m)}{v^2} dv + \int_{p^*}^{\bar{v}} \frac{\phi(m)f(v)}{f(p^*)(p^*)^2} dv = \phi(m)[\frac{1}{m} - \frac{1}{p^*} + \frac{1 - F(p^*)}{f(p^*)(p^*)^2}]$$

and, hence, using (10),

$$\int_m^{\bar{v}} f^B(m|v)f(v)dv = \frac{\phi(m)}{m}.$$

Therefore, for $m \leq v_1 \leq p^*$, $g^B(v_1|m) = m/(v_1)^2$, and so,[6] integrating over v_1, $\beta([\underline{v}, v)|m) = 1 - mv^{-1}$ for $v \in [m, p^*]$. Hence $v[\beta([v, \bar{v}]|m)] = m$. This implies that S, after message m, is indifferent between all offers in $[m, p^*]$. For $v_1 > p^*$,

$$g^B(v_1|m) = \frac{mf(v_1)}{f(p^*)(p^*)^2}$$

and so, integrating,

$$\beta([\underline{v}, v)|m) = 1 - \frac{m}{p^*} + \int_{p^*}^{v} \frac{mf(v_1)}{f(p^*)(p^*)^2} dv_1.$$

Hence, using (10), for $v > p^*$,

$$\beta([v, \bar{v}]|m) = \frac{m[1 - F(v)]}{(p^*)^2 f(p^*)}.$$

However, $p^*[1 - F(p^*)] > v[1 - F(v)]$, so, using (10), $(p^*)^2 f(p^*) > v[1 - F(v)]$. Hence $v\beta([v, \bar{v}]|m) < m$, which implies that demanding $v > p^*$ is inferior to demanding m. Hence the set of optimal demands is $[m, p^*]$.

Case (b), in which $m = \underline{v}$, is similar, and treated in the Appendix.

To summarize,

Lemma 7 Given any mechanism with message set V there exists a strategy μ^B such that $\mu^B_M([\underline{v}, p^*]) = 1$ and such that for all messages $m \leq p^*$, $V^S(m, \mu^B) = [m, p^*]/I$.

Let B's strategy μ^B be as given by Lemma 7. Define S's strategy as follows. For any $m \in V$,

$$\mu^S([\underline{v}, m]|m) = \mu^S((p^*, \bar{v}]|m) = 0$$

[6] Recall that $\beta(\cdot|m)$ denotes S's belief about B's type after message m.

and, for $m < v_1 \leq v_2 \leq p^*$,

$$\mu^S([v_1, v_2]|m) = \frac{Q(v_2) - Q(v_1)}{1 - Q(m)}.$$

To check that the strategy profile (μ^B, μ^S) gives expected quantity $Q(v)$ for each type v, note that if type $v \leq p^*$ sends message $m \leq v$ her expected volume of trade is

$$\begin{aligned} & q(m) + [1 - q(m)]\mu^S([\underline{v}, v]|m) \\ = {} & Q(m) + [1 - Q(m)]\frac{Q(v) - Q(m)}{1 - Q(m)} \\ = {} & Q(v). \end{aligned}$$

By Lemma 7, $V^S(m, \mu^B) = [m, p^*]/I$; by construction $\mu^S([m, p^*]|m) = \frac{Q(p^*)-Q(m)}{1-Q(m)}$ $= 1$ and, since Q is flat on each interval in I, $\mu^S(I|m) = 0$. Therefore condition (ii) of Definition 1 is fulfilled, i.e. S's strategy is optimal.

It remains to show that B's strategy is optimal, that is, $\mu^B(M^B(v, \mu^S)|v) = 1$ for all $v \in V$. To see this, take a type $v \leq p^*$. Note first that all messages $m \leq v$ give the same expected volume of trade for v, namely $Q(v)$. Also, expected payments are the same, since the expected payment given message m is

$$t(m) + [1 - Q(m)] \int_m^v v' \frac{dQ(v')}{1 - Q(m)} = t(v).$$

Therefore, $U^B(m, v, \mu^S) = U^B(v, v, \mu^S)$ for all $m \leq v$. Second, suppose that v sends message $m > v$. Then her expected payoff is $vQ(m) - t(m)$, whereas if she sends message v her payoff is $vQ(v) - t(v)$. The latter payoff is higher if $t(m) - t(v) \geq v(Q(m) - Q(v))$, i.e., if

$$\int_v^m v' dQ(v') \geq \int_v^m v dQ(v'),$$

which is true since $v' \geq v$ for all $v' \in [v, m]$. Therefore, $U^B(v, v, \mu^S) \geq U^B(m, v, \mu^S)$ for all $m > v$, so $\mu^B(M^B(v, \mu^S)|v) = 1$ for all $v \leq p^*$. Finally, any type $v > p^*$ gets the same payment and trade volume from any message m as type p^* does (since $\mu^S([\underline{v}, p^*]|m) = 1$ for all $m \in V$), so for those types too we have $\mu^B(M^B(v, \mu^S)|v) = 1$.

The case in which Q is not right-continuous at $\underline{v}$ is treated in the Appendix.

This establishes

Theorem 3 (Sufficient conditions for r-implementability) *A function* $Q : V \to [0,1]$ *is an r-implementable quantity schedule if it is non-decreasing, right-continuous on* $(\underline{v}, \bar{v}]$ *and constant on each interval* $(a_\gamma, b_\gamma)_{\gamma \in A}$*, and* $Q(v) = 1$ *for* $v \geq p^*$*.*

In the renegotiation equilibrium constructed above, if type $v < p^*$ tells the truth, S only offers renegotiation demands of v or higher, so no type of B gets a renegotiation surplus after telling the truth. It follows that each type of B must, if he tells the truth, get the same expected payoff as he would get if he told the truth when committed to the mechanism $\hat{\Gamma}(Q)$. Furthermore, one optimal strategy for the buyer, of any type, is to tell the truth.[7] Therefore, a fortiori, it is optimal to tell the truth in the commitment case. Thus, $\hat{\Gamma}(Q)$ is incentive-compatible and, in fact, is the same as the direct revelation mechanism which implements Q in the standard commitment case (with payoff $\underline{v}Q(\underline{v})$ for the lowest type).

If there were no mechanism then S would simply make a take-it-or-leave-it offer p which maximizes $R(p) = p[1 - F(p)]$. Let the quantity schedule which results from demanding the highest such maximizer, p^*, be denoted by Q^0. that is, $Q^0(v) = 0$ if $v < p^*$ and $Q^0(v) = 1$ if $v \geq p^*$. For the case in which R is concave (the most relevant case for applications) the following Theorem gives a simple characterization of r-implementable schedules in terms of Q^0. The result follows straightforwardly from Proposition 1, Theorems 2 and 3, and the fact that, for any c-implementable U and associated Q, $U(v) = U(\underline{v}) + \int_{\underline{v}}^{v} Q(u)du$.

Theorem 4 *Suppose that* R *is concave.*

(i) $Q : V \to [0,1]$ *is r-implementable if and only if (*Q *is c-implementable and right-continuous on* $(\underline{v}, \bar{v}]$*, and* $Q \geq Q^0$*).*

[7] This is related to the result of Bester and Strausz (2001). They showed that, when the principal cannot commit, any incentive-efficient outcome can be implemented by a direct revelation mechanism in which it is optimal for each type to tell the truth (but each lies with positive probability). In our case this property applies to any implementable outcome, not only incentive-efficient ones.

(ii) $U : V \to \Re_+$ *is r-implementable if and only if (*U *is c-implementable and* $U'(v) \geq Q^0(v)$ *wherever the derivative is defined).*

6 Strong Implementability

The previous section shows that any quantity schedule satisfying the conditions of Theorem 3, and its associated utility schedule, can be implemented in the renegotiation case by the same direct revelation mechanism $\hat{\Gamma}$ which implements it in the commitment case. However, for the case in which R is strictly concave[8], we can establish a significantly stronger result: the equilibrium quantity and utility schedules are the same for *any* renegotiation equilibrium of $\hat{\Gamma}$. In other words, applying Theorem 4, the mechanism designer who wants to implement an incentive-compatible quantity schedule Q which gives each type at least as much as he would get if there were no mechanism does not need to be concerned about whether renegotiation might or might not be possible - the same mechanism will work, in a strong sense, in either case.

Theorem 5 Suppose that R *is strictly concave. Suppose that* $U : V \to \Re_+$ *is r-implementable. Let* $\Gamma(U) = [V, Q(.), t(.)]$ *be a direct revelation mechanism which c-implements* U*. Then* $\Gamma(U)$ *strongly r-implements* U*.*

Proof $Q : V \to [0, 1]$ is the quantity schedule associated with U, so Q is non-decreasing and, for some $\hat{p} \leq p^*$, $Q(v) < 1$ if $v < \hat{p}$ and $Q(v) = 1$ if $v \geq \hat{p}$. There is no loss of generality in assuming that Q is right-continuous at $\underline{v}$, so Q is right-continuous.

Let $(\tilde{\mu}^B, \tilde{\mu}^S)$ be an arbitrary renegotiation equilibrium of $\Gamma(U)$. Let $\tilde{U}$ be the expected utility schedule and $\tilde{Q}$ be the expected quantity schedule given by this equilibrium. That is,

$$\tilde{Q}(v) = \bar{q}(v, \tilde{\mu}_B, \tilde{\mu}_S)$$

[8]In fact, we only require R to be locally strictly concave below a unique maximizer p^*, which, as shown above, is guaranteed by the increasing hazard rate assumption.

and

$$\tilde{U}(v) = vQ(m) - t(m) + [1 - Q(m)] \int_{\underline{v}}^{v} (v - r) d\tilde{\mu}^S(r|m)$$

where $m \in M^B(v, \tilde{\mu}^S)$.

Let $U(v, m)$ denote the expected utility of type v if, in mechanism $\Gamma(U)$ in the commitment case, he sends message $m \in V$.

We need to show that, for all v, $\tilde{U}(v) = U(v)$ and $\tilde{Q}(v) = Q(v)$. Clearly each type v gets at least $U(v)$ since he can always tell the truth and then refuse to renegotiate. So $\tilde{U}(v) \geq U(v)$ for all $v \in V$.

Henceforth, for brevity, we omit the strategies $\tilde{\mu}^S$ and $\tilde{\mu}^B$ from terms such as $M^B(v, \tilde{\mu}^S)$. Recall that there is possibly a set of messages $N \subseteq V = M$, which has zero $\tilde{\mu}^B_M$-measure, for which the statements in Lemmas 2 and 3 (and Lemma A.1 in the Appendix) do not apply.

We proceed via several intermediate claims.

Claim 1 *Given $v \in V$, suppose that there exists $m \in M^B(v)/N$ such that $m \notin M^B(\tilde{v})$ for all $\tilde{v} < v$. Then $U(v, m) = U(v) = \tilde{U}(v)$.*

Proof Take $m \in M^B(v)/N$ such that $m \notin M^B(\tilde{v})$ for all $\tilde{v} < v$. It follows from Lemma $A.1(ii)$ in the Appendix that, after message m, S will not offer less than v, i.e. $\tilde{\mu}^S([\underline{v}, v)|m) = 0$. Hence type v gets no renegotiation surplus if he sends m. Therefore his expected payoff is $U(v, m) \leq U(v) \leq \tilde{U}(v)$. However, m is optimal for him, so $U(v, m) = U(v) = \tilde{U}(v)$. This proves the Claim. QED.

The proofs of the next three Claims are in the Appendix.

Claim 2 *If $\tilde{U}(v) = U(v)$ then (i) $\tilde{Q}(v) \geq Q(v)$ and (ii) if $v > \underline{v}$, $lim_{v' \uparrow v} \tilde{Q}(v') \leq lim_{v' \uparrow v} Q(v')$.*

Claim 3 *If, for some $v_1 \in V$, $\tilde{Q}(v_1) > Q(v_1)$ and $\tilde{U}(v_1) = U(v_1)$ then (i) if $v_1 > \underline{v}$, v_1 is a point of increase of $\tilde{Q}$; (ii) $M^B(v_1)/N = \cup_{v' \leq v_1} M^B(v')/N$; and (iii) for all $m \in M^B(v_1)/N$, $v_1 \in V^S(m)$, i.e., v_1 is an optimal demand for S after message m.*

Claim 4 Suppose that $v_1 < v_2$ and, for all $v \in [v_1, v_2)$, $M^B(v)/N \subseteq \cup_{v' \leq v_1} M^B(v')$. Then there exists at most one point v in $[v_1, v_2]$ such that $v \in V^S(m)$ for all $m \in \cup_{v' \leq v_1} M^B(v')/N$.

Take v such that there exists $m \in M^B(v)/N$ such that $m \notin M^B(\tilde{v})$ for all $\tilde{v} < v$. By Claim 1, $\tilde{U}(v) = U(v)$.

In order to show that $\tilde{U}(v) = U(v)$ for all other types, assume that there exists $\hat{v} < \hat{p}$ such that $\tilde{U}(\hat{v}) > U(\hat{v})$. Let

$$v_1 = sup\{v < \hat{v} | \tilde{U}(v) = U(v)\}$$

and let

$$v_2 = inf\{v > \hat{v} | \tilde{U}(v) = U(v)\}.$$

v_1 exists because $\tilde{U}(\underline{v}) = U(\underline{v})$, since the lowest type can get no renegotiation surplus. By continuity, $\tilde{U}(v_1) = U(v_1)$.

To show that v_2 exists, assume that it does not. Then $\tilde{U}(v) > U(v)$ for all $v > v_1$. Hence, applying Claim 1, $M^B(v)/N \subseteq \cup_{v' \leq v_1} M^B(v')$ for all $v \in (v_1, \bar{v}]$. So $\cup_{v' \in V} M^B(v')/N = \cup_{v' \leq v_1} M^B(v')/N$. This implies that there is no point of increase v of $\tilde{Q}$ in (v_1, p^*): if there were, it would be optimal for S, by Lemma 3, to demand v after any $m \in \cup_{v' \in V} M^B(v')/N$, hence optimal to demand v ex ante, which contradicts the fact that $v < p^*$ (recall that p^* is the unique maximizer of R, by strict concavity of R). Suppose that $\tilde{Q}(v_1) > Q(v_1)$. Then, by Claim 3, v_1 is optimal for S after any message in $\cup_{v' \leq v_1} M^B(v')/N$, hence after any message in $\cup_{v' \in V} M^B(v')/N$. Hence, as above, we have a contradiction. Therefore $\tilde{Q}(v_1) \leq Q(v_1)$. We then have, since there is no point of increase in (v_1, p^*), $\tilde{Q}(v) = \tilde{Q}(v_1) \leq Q(v)$ for all $v \in [v_1, p^*)$. Therefore, since $\tilde{U}(v) = \tilde{U}(v_1) + \int_{v_1}^{v} \tilde{Q}(u)du$ and $U(v) = U(v_1) + \int_{v_1}^{v} Q(u)du$, we have $\tilde{U}(v) \leq U(v)$ for all $v \in [v_1, p^*)$, which contradicts the definition of v_1. This shows that v_2 exists.

v_1 and v_2 satisfy the hypotheses of Claim 4, so, by Lemma 3, there is at most one point of increase in $(v_1, v_2]$. Hence on some neighbourhood $[v_1, v_1 + \delta)$ $\tilde{Q}$ is equal to a constant q' and on some neighbourhood $(v_2 - \delta, v_2)$ $\tilde{Q}$ is equal to a constant q''. For $v' \in (v_1, v_1 + \delta)$, $\tilde{U}(v') = \tilde{U}(v_1) + \int_{v_1}^{v'} q' du$, $U(v') = U(v_1) + \int_{v_1}^{v'} Q(u)du$ and $\tilde{U}(v') > U(v')$. Hence $q' > Q(v_1)$, i.e. $\tilde{Q}(v_1) > Q(v_1)$. Therefore v_1 satisfies the hypotheses of Claim 3, which implies that $v_1 \in V^S(m)$ for all $m \in \cup_{v' \leq v_1} M^B(v')/N$.

By continuity, $\tilde{U}(v_2) = U(v_2)$. Therefore, by Claim 2, $\tilde{Q}(v_2) \geq Q(v_2)$. For $v' \in (v_2 - \delta, v_2)$, $\tilde{U}(v') = \tilde{U}(v_2) - \int_{v'}^{v_2} q'' du$, $U(v') = U(v_2) - \int_{v'}^{v_2} Q(u)du$ and $U(v') < \tilde{U}(v')$. Therefore $lim_{v \uparrow v_2} \tilde{Q}(v) = q'' < lim_{v \uparrow v_2} Q(v) \leq \tilde{Q}(v_2)$. Hence $\tilde{Q}(v_2) > lim_{v \uparrow v_2} \tilde{Q}(v)$, so v_2 is a point of increase of $\tilde{Q}$. Hence, by Lemma 3, $v_2 \in V^S(m)$ for all $m \in \cup_{v' \leq v_1} M^B(v')/N$. Since the same is true for v_1 this contradicts the fact, by Claim 4, that there is at most one such point in $[v_1, v_2]$.

This shows that there does not exist $\hat{v} < \hat{p}$ such that $\tilde{U}(\hat{v}) > U(\hat{v})$. Hence $\tilde{U}(v) = U(v)$ for all $v \in [\underline{v}, \hat{p}]$. Therefore, since, almost everywhere, $\tilde{U}'(v) = \tilde{Q}(v)$ and $U'(v) = Q(v)$, and $\tilde{Q}$ and Q are right-continuous, $\tilde{Q}(v) = Q(v)$ for all $v \in [\underline{v}, \hat{p})$. Claim 2 implies that $\tilde{Q}(\hat{p}) \geq Q(\hat{p}) = 1$. Therefore $\tilde{Q}(v) = Q(v) = 1$ for all $v \geq \hat{p}$, so $\tilde{U}(v) = U(v)$ for all $v \geq \hat{p}$. This completes the proof. QED.

7 Conclusion

In this paper we have analyzed the impact of ex-post renegotiation on the set of implementable outcomes in a bilateral trade problem. When full commitment is possible, any increasing trading rule can be implemented by using a direct revelation mechanism that is designed to elicit the truth from privately informed parties. When commitment is not possible, the set of implementable trading rules is restricted because a direct revelation mechanism cannot fully extract all information from the parties. Nevertheless, we have shown that the restriction takes a very simple form - no type's expected trade volume can be reduced by the mechanism. Furthermore, the direct revelation mechanism which is appropriate for the commitment case uniquely implements the desired outcome in the non-commitment case.

There are several dimensions on which it would be desirable to generalize the analysis: for example, to other bargaining games (such as the infinite horizon game with discounting in which the uninformed party makes all the price offers) and to other contracting situations. Our conjecture is that for some of these cases analogous results will hold, but this is left for future work.

Appendix

Given $m \in M$, let $V^B(m) = \{v \in V | m \in M^B(v, \mu^S)\}$, i.e., the set of types for which m is an optimal message.

The following Lemma collects a number of useful properties.

Lemma A.1 *For almost all messages* m:
(i) $\mu^B[(V^B(m))^c|m] = 0$,
(ii) if $[v, v'] \subseteq \left(V^B(m)\right)^c$ *then* $[v, v') \subseteq \left(V^S(m)\right)^c$,
(iii) the function $U^B(m, \cdot, \mu^S) : V \to \Re$ *is continuous,*
(iv) $(V^B(m))^c$ *and* $(V^S(m))^c$ *are open sets, and*
(v) if $v \in V^S(m)$ *then* $m \in M^B(v)$.

Proof *(i)* Let $N \subseteq M \times V$ be defined by

$$N = \{(m, v) \in M \times V | m \notin M^B(v, \mu^S)\},$$

i.e., the message-type pairs such that the message is suboptimal for the type. Let $E \subseteq M$ be a set such that $\mu^B_M(E) > 0$ and such that $\mu^B[(V^B(m))^c|m] > 0$ for all $m \in E$. Then $\int_E \mu^B[(V^B(m))^c|m] d\mu^B_M > 0$. However,

$$\int_E \mu^B[(V^B(m))^c|m] d\mu^B_M = \mu^B((E \times V) \cap N) \leq \mu^B(N) = 0,$$

where the last equality follows from

$$\mu^B(N) = \int_V \mu^B[(M^B(v))^c|v] d\mu^B_V = 0.$$

This gives a contradiction.

(ii) Let $\tilde{U}^S(m, v, \mu^B) = (U^S(m, v, \mu^B) - t(m))(1 - q(m))^{-1}$. Given m such that $q(m) < 1$, maximizing $\tilde{U}^S$ is equivalent to maximizing U^S. For all $\hat{v} \in [v, v')$,

$$\begin{aligned}\tilde{U}^S(m, \hat{v}, \mu^B) &= \hat{v}\mu^B([\hat{v}, v') \mid m) + \hat{v}\mu^B([v', \bar{v}] \mid m) \\ &= \hat{v}\mu^B([v', \bar{v}] \mid m)\end{aligned}$$

for all m satisfying (i).

If $\mu^B([v', \bar{v}]|m) = 0$ then $0 = \tilde{U}^S(m, \hat{v}, \mu^B) < \tilde{U}^S(m, \underline{v}, \mu^B) = \underline{v}$.

If $\mu^B([v', \bar{v}]|m) > 0$ then $\tilde{U}^S(m, \hat{v}, \mu^B) = \hat{v}\mu^B([v', \bar{v}]|m) < v'\mu^B([v', \bar{v}]|m)$. The latter expression is $\tilde{U}^S(m, v', \mu^B)$. Hence $\hat{v}$ is suboptimal for S.

(iii) Take a sequence $\{v_1, v_2, v_3, ...\} \subseteq V$ converging to $v \in V$ from below (the argument is similar for convergence from above).

$$\begin{aligned}U^B(m, v, \mu^S) - U^B(m, v_i, \mu^S) &= (v - v_i)q(m) + (1 - q(m)) \int_{v_i}^{v} (v - r) d\mu^S(r|m) \\ &\quad + (v - v_i)(1 - q(m))\mu^S([\underline{v}, v_i)|m).\end{aligned}$$

The first and third terms converge to zero for $i \to \infty$ and, by the Dominated Convergence Theorem, the second term also converges to zero.

(iv) If $v \in (V^B(m))^c$ then $U^B(m', v, \mu^S) > U^B(m, v, \mu^S)$ for some m'. By continuity, the same inequality is true for any v_i close enough to v. Hence $v_i \in (V^B(m))^c$. Therefore $(V^B(m))^c$ is open.

Suppose that $v \in (V^S(m))^c$. Then there exists $\hat{v} \in V$ with $\tilde{U}^S(m, \hat{v}, \mu^B) > \tilde{U}^S(m, v, \mu^B)$. We show first that $\tilde{U}^S(m, \cdot, \mu^B)$ is left-continuous. This will imply that for ε small enough $\tilde{U}^S(m, \hat{v}, \mu^B) > \tilde{U}^S(m, v - \varepsilon, \mu^B)$, so $v - \varepsilon \in (V^S(m))^c$.

$$\begin{aligned}&\tilde{U}^S(m, v, \mu^B) - \tilde{U}^S(m, v - \varepsilon, \mu^B) \\ &= v\mu^B([v, \bar{v}]|m) - (v - \varepsilon)\mu^B([v - \varepsilon, \bar{v}]|m) \\ &= -v\mu^B([v - \varepsilon, v)|m) + \varepsilon\mu^B([v - \varepsilon, \bar{v}]|m) \to 0\end{aligned}$$

as $\varepsilon \to 0$ by the Monotone Convergence Theorem. Hence $\tilde{U}^S(m, \cdot, \mu^B)$ is left-continuous. Also,

$$
\begin{aligned}
& \tilde{U}^S(m, v, \mu^B) - \tilde{U}^S(m, v+\varepsilon, \mu^B) \\
= \; & v\mu^B([v, \bar{v}]|m) - (v+\varepsilon)\mu^B([v+\varepsilon, \bar{v}]|m) \\
= \; & v\mu^B([v, v+\varepsilon)|m) - \varepsilon\mu^B([v+\varepsilon, \bar{v}]|m) \to v\mu^B(\{v\}|m).
\end{aligned}
$$

So

$$
\tilde{U}^S(m, \hat{v}, \mu^B) > \tilde{U}^S(m, v, \mu^B) \geq \tilde{U}^S(m, v+\varepsilon, \mu^B)
$$

for ε small enough, so $v+\varepsilon \in (V^S(m))^c$ and therefore $(V^S(m))^c$ is open.

(v) Suppose that $m \notin M^B(v)$, i.e., $v \notin V^B(m)$. Then, by *(iv)*, there is a neighbourhood $[v-\varepsilon, v+\varepsilon] \subseteq (V^B(m))^c$. By *(ii)*, $[v-\varepsilon, v+\varepsilon) \subseteq (V^S(m))^c$. Hence $v \notin V^S(m)$. Contradiction. QED

Proof of Lemma 1. Let $\tilde{t}(m, v)$ be the expected amount paid, including the renegotiation price, if message m is sent, and renegotiation demands less than or equal to v are accepted. That is,

$$
\tilde{t}(m, v) = t(m) + [1 - q(m)] \int_{\underline{v}}^{v} r d\mu^S(r|m).
$$

Let $\bar{t}(v)$ be the expected amount paid by type v, i.e.

$$
\bar{t}(v) = \int_M \tilde{t}(m, v) d\mu^B(m|v).
$$

Similarly, let $\tilde{t}(m, v_-)$ be the expected amount paid, including the renegotiation price, if message m is sent and only renegotiation offers strictly less than v are accepted, i.e., $\tilde{t}(m, v_-) = \tilde{t}(m, v) - (1 - q(m))v\mu^S(\{v\}|m)$.

(i) If v sends m and accepts renegotiation demands less than or equal to v (which is an optimal strategy for v), v gets $v\tilde{q}(m, v) - \tilde{t}(m, v)$. If v sends m' and accepts renegotiation demands less than or equal to v' (which is not necessarily optimal), v

gets $v\tilde{q}(m', v') - \tilde{t}(m', v')$. So

$$v\tilde{q}(m, v) - \tilde{t}(m, v) \geq v\tilde{q}(m', v') - \tilde{t}(m', v').$$

Similarly

$$v'\tilde{q}(m', v') - \tilde{t}(m', v') \geq v'\tilde{q}(m, v) - \tilde{t}(m, v).$$

Hence, if $v' > v$, it follows that $\tilde{q}(m, v) \leq \tilde{q}(m', v')$. This proves (a).

Similarly, since it is also optimal for v' to send m' and only to accept renegotiation demands strictly less than v', (b) follows.

(ii) follows from (a) since $\bar{q}(v, \mu^B, \mu^S) = \int_M \tilde{q}(m, v, \mu^S) d\mu^B(m|v)$. QED

Proof of Lemma 2 Suppose $\tilde{q}(m, v) > \tilde{q}(m', v)$ for $m, m' \in M^B(v)$ for some $v \in V$. Then $\tilde{q}(m', v) < 1$ and therefore $\mu^S((v, \bar{v}]|m') > 0$ by (3). This implies $(v, \bar{v}] \cap V^S(m') \neq \emptyset$ and so there is a $\hat{v} \in (v, \bar{v}]$ such that $\mu^B([\hat{v}, \bar{v}]|m') > 0$ (for an offer $\hat{v}$ to be a maximizer of U^S, necessarily $\mu^B([\hat{v}, \bar{v}]|m') > 0$, since the offer $\underline{v}$ has a payoff of $\underline{v} > 0$). Therefore, by Lemma *A*.1*(i)*, $V' := (v, \bar{v}] \cap V^B(m') \neq \emptyset$. Let $v' = \inf V'$.

(i) Assume $v' > v$. By Lemma $A.1(iv)$ $V^B(m')$ is closed, hence $v' \in V^B(m')$, so $m' \in M^B(v')$. Summarizing, we have $m' \in M^B(v) \cap M^B(v')$ and $m' \notin M^B(\tilde{v})$ for all $\tilde{v} \in (v, v')$. Then, by Lemma $A.1(ii)$ $\mu^S((v, v')|m') = 0$ and so

$$\tilde{q}(m', v'_-) = q(m') + (1 - q(m'))\mu^S([\underline{v}, v')|m') = \tilde{q}(m', v) < \bar{q}(m, v),$$

which contradicts monotonicity (Lemma $1(i)(b)$).

(ii) Assume $v' = v$. Then there exists a sequence $\{v + \varepsilon_i\} \subseteq V^B(m')$, with $\varepsilon_i \to 0$
.

$$\tilde{q}(m', v + \varepsilon_i) - \tilde{q}(m', v) = (1 - q(m'))\mu^S((v, v + \varepsilon_i]|m').$$

Therefore, by the Monotone Convergence Theorem, $\tilde{q}(m', v+\varepsilon_i) \to \tilde{q}(m', v) < \tilde{q}(m, v)$. Hence, there exists $v + \varepsilon > v$ such that m is optimal for v, m' is optimal for $v + \varepsilon$, and $\tilde{q}(m', v + \varepsilon) < \tilde{q}(m, v)$. This contradicts monotonicity. QED.

Proof of Lemma 3 *(i)* Suppose that there exist $v', v'' \in V$ such that $v' < v < v''$, $m \in M^B(v')$, $m \in M^B(v'')$ and $(v', v'') \subseteq (V^S(m))^c$, so that $\mu^S((v', v'')|m) = 0$. Then, since v is a point of increase, Lemma 2 implies that $\tilde{q}(m, v') < \bar{q}(\tilde{v})$ for all $\tilde{v}$ with $v < \tilde{v} < v''$. But $\tilde{q}(m, v''_-) = \tilde{q}(m, v')$, so $\tilde{q}(m, v''_-) < \bar{q}(\tilde{v})$ for some $\tilde{v} < v''$, which violates monotonicity (Lemma 1).

To show that, if $v \notin V^S(m)$, we can choose v' and v'' with the required properties, hence obtain a contradiction, first assume (a) that $V^S(m) \cap [\underline{v}, v] \neq \emptyset$. Let $v' = max\{\tilde{v}|\tilde{v} \in V^S(m) \cap [\underline{v}, v]\}$. This exists because $V^S(m)$ is closed by Lemma $A.1$*(iv)*. Also $v' < v$, $(v', v] \subseteq (V^S(m))^c$ and, by Lemma $A.1$*(v)*, $m \in M^B(v')$. Now suppose (b) that $V^S(m) \cap [\underline{v}, v] = \emptyset$. Take any $v' < v$ such that $m \in M^B(v')$. Such a v' exists since $m \in \cup_{\tilde{v}<v} M^B(\tilde{v})$. This defines v'. To define v'', suppose first that $V^S(m) \cap [v, \bar{v}] = \emptyset$. Then $\mu^S((v', \bar{v}]|m) = 0$. Therefore $\tilde{q}(m, v') = 1$ and so, by Lemma 2, $\bar{q}(v') = 1$, which contradicts the fact that v is a point of increase. Hence $V^S(m) \cap [v, \bar{v}] \neq \emptyset$. Define $v'' = min\{\tilde{v}|\tilde{v} \in V^S(m) \cap [v, \bar{v}]\}$. This exists by Lemma $A.1$*(iv)*. Also $v'' > v$, $[v, v'') \subseteq (V^S(m))^c$ and, by Lemma $A.1$*(v)*, $m \in M^B(v'')$. Therefore $v \in V^S(m)$. $m \in M^B(v)$ by Lemma $A.1(v)$.

(ii) Assume that there exists an $\varepsilon > 0$ such that $\mu^S((v - \varepsilon, v + \varepsilon]|m) = 0$. We show that we can then find v' and v'' as above, which is in contradiction with v being a point of increase. Consider the set $V' := V^S(m) \cap [\underline{v}, v)$. Assume that it is non-empty. If $\sup V' = v$, it follows that there is a sequence in V' converging to v, and we can therefore find a $v' \in (v - \varepsilon, v)$ with $m \in M^B(v')$ and $\mu^S((v', v]|m) = 0$. If $\sup V' < v$, set $v' := \sup V'$. If $V' = \emptyset$, take any $v' < v$ such that $m \in M^B(v')$. This defines v'. To define v'' consider the set $V'' = V^S(m) \cap (v, \bar{v}]$. If $V'' = \emptyset$, we can use the same argument as above to show that this violates the assumption that v is a point of increase. Hence, V'' is non-empty. If $\inf V'' = v$, it follows that there is a sequence in V'' converging to v, and we can therefore find a $v'' \in (v, v + \varepsilon)$ with $m \in M^B(v'')$ and $\mu^S([v, v'')|m) = 0$. If $\inf V'' > v$, let $v'' = \inf V''$. This v'' has the required properties. QED

Proof of Lemma 5 That v cannot be a point of increase follows from Lemma 4 and the fact that $\tilde{R}(v) > R(v)$ implies that the graph of R at v must lie strictly below a chord of R. Take an interval (a, b) (for ease of notation we suppress the subscript γ) with $\tilde{R}(v) > R(v)$ for all $v \in (a, b)$, $\tilde{R}(a) = R(a)$ and $\tilde{R}(b) = R(b)$. We show that $\bar{q}$ must be constant on (a, b). Fix $v \in (a, b)$. Since v is not a point of increase there exists $\tilde{v} < v$ with $\bar{q}(\tilde{v}) = \bar{q}(v)$. Let $\tilde{v}(v) = \inf\{\tilde{v} < v | \bar{q}(\tilde{v}) = \bar{q}(v)\}$. We know that $\bar{q}$ is flat on $(\tilde{v}(v), v]$ since $\bar{q}$ is monotonic. Assume $\tilde{v}(v) > a$. Then, for all $v' < \tilde{v}(v)$ and $\varepsilon > 0$, $\bar{q}(v') < \bar{q}(\tilde{v}(v) + \varepsilon)$. But this implies that $\tilde{v}(v) \in (a, b)$ is a point of increase, which is a contradiction. Hence $\tilde{v}(v) \leq a$. But then $\bar{q}$ is flat on $(a, v]$, hence flat on (a, b). QED

Proof of Lemma 6 We show that there can be no point of increase $\hat{v} > p^*$. Since $\bar{q}(\bar{v}) = 1$, the result follows. Suppose then that there is such a point of increase $\hat{v}$, given equilibrium (μ^B, μ^S) of some mechanism. We claim that there then exists a strategy $\tilde{\mu}^B$ for B such that $\hat{v} \in V^S(M_{\hat{v}}, \tilde{\mu}^B)$ and $\hat{v} \in V^S((M_{\hat{v}})^c, \tilde{\mu}^B)$. The claim implies that $\hat{v}$ maximizes $v\tilde{\mu}^B([v, \bar{v}]|M_{\hat{v}})\tilde{\mu}^B_M(M_{\hat{v}}) + v\tilde{\mu}^B([v, \bar{v}]|(M_{\hat{v}})^c)\tilde{\mu}^B_M((M_{\hat{v}})^c)$ and therefore that it maximizes $v\tilde{\mu}^B([v, \bar{v}]|M) = v\eta_F([v, \bar{v}]) = R(v)$, which contradicts the assumption that p^* is the highest maximizer of R. (Without loss of generality we can ignore messages which have zero probability under $\tilde{\mu}^B$.)

Let $M_{\hat{v}} = \cup_{v<\hat{v}} M^B(v, \mu^B)$ and $(M_{\hat{v}})^c = M \backslash M_{\hat{v}}$. Given any measurable set T in M, let $U^S(T, v, \mu^B) = v\mu^B([v, \bar{v}]|T)$ and $V^S(T, \mu^B) = \{v \in V | v \text{ maximizes } U^S(T, v, \mu^B)\}$.

To prove the claim above, set, for any measurable set T in M, $\tilde{\mu}^B(T|v) = \mu^B(T|v)$ for all $v \leq \hat{v}$ and $\tilde{\mu}^B(T|v) = \mu^B(T|[\hat{v}, \bar{v}])$ for all $v > \hat{v}$.

We first show that for all $v \leq \hat{v}$

$$\hat{v}\tilde{\mu}^B([\hat{v}, \bar{v}]|M_{\hat{v}}) \geq v\tilde{\mu}^B([v, \bar{v}]|M_{\hat{v}}). \tag{11}$$

Using Bayes' Rule, we can show that $\tilde{\mu}^B([v, \bar{v}]|M_{\hat{v}}) = \mu^B([v, \bar{v}]|M_{\hat{v}})$ for all $v < \hat{v}$.

This is because $\tilde{\mu}^B_M(M_{\hat{v}}) = \mu^B_M(M_{\hat{v}})$ and

$$\begin{aligned}
\tilde{\mu}^B(M_{\hat{v}} \times [v, \bar{v}]) &= \int_v^{\hat{v}} \tilde{\mu}^B(M_{\hat{v}}|v')dF(v') + \int_{\hat{v}}^{\bar{v}} \tilde{\mu}^B(M_{\hat{v}}|v')dF(v') \\
&= \int_v^{\hat{v}} \mu^B(M_{\hat{v}}|v')dF(v') + \int_{\hat{v}}^{\bar{v}} \mu^B(M_{\hat{v}}|[\hat{v}, \bar{v}])dF(v') \\
&= \int_v^{\hat{v}} \mu^B(M_{\hat{v}}|v')dF(v') + \int_{\hat{v}}^{\bar{v}} \mu^B(M_{\hat{v}}|v')dF(v') \\
&= \mu^B(M_{\hat{v}} \times [v, \bar{v}])
\end{aligned}$$

Because $\hat{v}$ is a point of increase of $\bar{q}$, we know from Lemma 3 that $\hat{v}\mu^B([\hat{v}, \bar{v}]|M_{\hat{v}}) \geq v\mu^B([v, \bar{v}]|M_{\hat{v}})$ for all v and so (11) follows for all $v \leq \hat{v}$.

We now show that (11) also holds for all $v > \hat{v}$. Since only types above $\hat{v}$ accept offers that are above $\hat{v}$, (11) follows if we can show that for all $v > \hat{v}$

$$\hat{v}\tilde{\mu}^B([\hat{v}, \bar{v}]|M_{\hat{v}} \times [\hat{v}, \bar{v}]) \geq v\tilde{\mu}^B([v, \bar{v}]|M_{\hat{v}} \times [\hat{v}, \bar{v}]). \tag{12}$$

$$\tilde{\mu}^B([v, \bar{v}]|M_{\hat{v}} \times [\hat{v}, \bar{v}]) = \tilde{\mu}^B(M_{\hat{v}} \times [v, \bar{v}])/\tilde{\mu}^B(M_{\hat{v}} \times [\hat{v}, \bar{v}]),$$

noting that $\tilde{\mu}^B(M_{\hat{v}} \times [\hat{v}, \bar{v}]) > 0$ since, if $\mu^B(M_{\hat{v}} \times [\hat{v}, \bar{v}]) = 0$, $\hat{v}$ could not be an optimal demand after a message in $M_{\hat{v}}$, contradicting Lemma 3. Since

$$\tilde{\mu}^B(M_{\hat{v}} \times [v, \bar{v}]) = \int_v^{\bar{v}} \tilde{\mu}^B(M_{\hat{v}}|v')dF(v') = \mu^B(M_{\hat{v}}|[\hat{v}, \bar{v}])\eta_F([v, \bar{v}]),$$

(12) is equivalent to

$$\hat{v}\eta_F([\hat{v}, \bar{v}]) \geq v\eta_F([v, \bar{v}]),$$

i.e., $R(\hat{v}) \geq R(v)$, for all $v > \hat{v}$. This is true since if not, then, since $\hat{v}$ is a point of increase of $\bar{q}$, $\tilde{R}(\hat{v}) = R(\hat{v}) < R(v)$. Also $\tilde{R}(p^*) \geq R(p^*) > R(\hat{v})$ and so $\tilde{R}(p^*) > \tilde{R}(\hat{v})$ and $\tilde{R}(v) > \tilde{R}(\hat{v})$, which contradicts concavity of $\tilde{R}$.

Finally, we need to show that for all $v \in V$

$$\hat{v}\tilde{\mu}^B([\hat{v}, \bar{v}]|(M_{\hat{v}})^c) \geq v\tilde{\mu}^B([v, \bar{v}]|(M_{\hat{v}})^c). \tag{13}$$

For any measurable set T in M set $V^B(T) = \cup_{m \in T} V^B(m)$. From the definition of $(M_{\hat{v}})^c$, we know that $[\underline{v}, \hat{v}) = \left(V^B((M_{\hat{v}})^c)\right)^c$ and so from Lemma $A.1$*(i)* $\mu^B([\underline{v}, \hat{v})|(M_{\hat{v}})^c) = 0$. So, (13) is true for all $v \in V$ if

$$\hat{v}\tilde{\mu}^B([\hat{v}, \bar{v}]|\,(M_{\hat{v}})^c \times [\hat{v}, \bar{v}]) \geq v\tilde{\mu}^B([v, \bar{v}]|\,(M_{\hat{v}})^c \times [\hat{v}, \bar{v}]) \tag{14}$$

for all $v \geq \hat{v}$. This can be shown to hold using the same argument as the one used to prove (12). QED

Lemma A.2 $\bar{q}(\cdot)$ *is right-continuous on* $(\underline{v}, \bar{v}]$.

Proof Suppose not. Then, for some $v \in V$ such that $v > \underline{v}$ and $\bar{q}(v) < 1$, and some $\varepsilon > 0$, $lim_{v_i \downarrow v}\bar{q}(v_i) \geq \bar{q}(v) + \varepsilon$. Then v is a point of increase of $\bar{q}$. Suppose that there exists a sequence $\{v_j\}_{j=1}^{\infty}$, such that $v_j > v$ for all j and $lim_j v_j = v$, and such that each v_j is a point of increase of $\bar{q}$. Take $m \in \cup_{v' < v} M^B(v')$ for which Lemmas 2 and 3 hold. This must exist because the set N of messages for which they do not hold has probability zero, but $\cup_{v' < v} M^B(v')$ has strictly positive probability. By Lemma 3 $m \in M^B(v)$ and $m \in M^B(v_j)$ for all j. By Lemma 2, $\bar{q}(v_j) = \bar{q}(v) + (1 - q(m))\mu^S((v, v_j]|m)$. Therefore, for every v_j, $\mu^S((v, v_j]|m) \geq \varepsilon(1 - q(m))^{-1}$. However, by the Monotone Convergence Theorem, $lim_{j \to \infty}\mu^S((v, v_j]|m) = 0$. This shows that, for some $\eta > 0$, there are no points of increase of $\bar{q}$ in $(v, v + \eta)$. This implies that $\bar{q}$ is constant on $(v, v + \eta)$, equal to say $\hat{q} > \bar{q}(v)$. Let $v' = inf\{\tilde{v} > v | m \in M^B(\tilde{v})\}$, where m is as above; v' exists since $\bar{q}(v) < 1$. If $v' = v$ then, as above, there exists a sequence v_j converging to v from above such that $m \in M^B(v_j)$ for all j, which gives a contradiction. Hence we can assume $v' > v$. By Lemma $A.1(iv)$ $m \in M^B(v')$ and by Lemma $A.1(ii)$ $\mu^S((v, v')|m) = 0$. If v' sends message m and accepts only offers strictly below v' then v' gets $\bar{q}(v) < \hat{q} = \bar{q}(v'')$ for some $v'' \in (v, v')$. This contradicts Lemma 1. Hence $\bar{q}(v) = \hat{q}$, which contradicts the assumption that $lim_{v_i \downarrow v}\bar{q}(v_i) \geq \bar{q}(v) + \varepsilon$. QED

Proof of Lemma 7 Case (b): $m = \underline{v}$. Recall that $\beta(\cdot|m)$ denotes S's belief about

B's type after message m.

$$\beta([\underline{v}, v)|\underline{v}) = \frac{\mu^B([\underline{v}, v) \times \{\underline{v}\})}{\mu^B_M(\{\underline{v}\})}.$$

$$\mu^B_M(\{\underline{v}\}) = \int_{\underline{v}}^{p^*} \frac{f(\underline{v})(\underline{v})^2}{v^2} dv + \int_{p^*}^{\bar{v}} \frac{f(\underline{v})(\underline{v})^2 f(v)}{f(p^*)(p^*)^2} dv = f(\underline{v})(\underline{v})^2 [\frac{1}{\underline{v}} - \frac{1}{p^*} + \frac{1 - F(p^*)}{f(p^*)(p^*)^2}]$$

which equals $f(\underline{v})\underline{v}$.

If $v \leq p^*$,

$$\mu^B([\underline{v}, v) \times \{\underline{v}\}) = \int_{\underline{v}}^{v} \frac{f(\underline{v})(\underline{v})^2 f(v_1)}{f(v_1)(v_1)^2} dv_1$$

$$= f(\underline{v})(\underline{v})^2 [\frac{1}{\underline{v}} - \frac{1}{v}]$$

so $\beta([\underline{v}, v)|\underline{v}) = 1 - \underline{v}(v)^{-1}$ as before.

If $v > p^*$,

$$\mu^B([\underline{v}, v) \times \{\underline{v}\}) = \int_{\underline{v}}^{p^*} \frac{f(\underline{v})(\underline{v})^2 f(v_1)}{f(v_1) v_1^2} dv_1 + \int_{p^*}^{v} \frac{f(\underline{v})(\underline{v})^2 f(v_1)}{f(p^*)(p^*)^2} dv_1$$

so

$$\beta([\underline{v}, v)|\underline{v}) = 1 - \frac{\underline{v}}{p^*} + \int_{p^*}^{v} \frac{\underline{v} f(v_1)}{f(p^*)(p^*)^2} dv_1$$

as before. This completes the proof for the case in which I is empty.

Suppose now that I is non-empty. We consider only the case in which I consists of a single interval $(a, b) \subseteq [\underline{v}, p^*]$. The argument extends in an obvious way to the case in which I is an arbitrary union of disjoint intervals.

On $[a, b]$ the graph of $\tilde{R}$ coincides with the chord between $(a, R(a))$ and $(b, R(b))$ and this chord is tangent to the graph of R at $(a, R(a))$ and at $(b, R(b))$. Therefore

$$R(a) + R'(a)(b - a) = R(b) = R(a) + R'(b)(b - a).$$

Hence

$$\frac{af(a)}{b} = \frac{F(b) - F(a)}{b - a} = \frac{bf(b)}{a}$$

from which follow

$$\frac{F(b)-F(a)}{f(a)a^2}=\frac{1}{a}-\frac{1}{b} \tag{15}$$

and

$$f(a)a^2=f(b)b^2. \tag{16}$$

Let the strategies of types in $[\underline{v},a]$ be as defined in the case in which I is empty and let the strategy of all types in (a,b) be the same as the strategy of type a. For $v\geq b$ and $m\notin(a,b)$, let $\mu^B([\underline{v},m]|v)$ be defined as in the case in which I is empty and let $\mu^B((a,b)|v)=0$. That is, if $v\in[b,p^*]$, then, for $m\leq a$,

$$\mu^B([\underline{v},m)|v)=\frac{f(m)m^2}{f(v)v^2},$$

for $m\in(a,b)$,

$$\mu^B([\underline{v},m)|v)=\frac{f(a)a^2}{f(v)v^2},$$

and, for $m\in[b,v]$,

$$\mu^B([\underline{v},m)|v)=\frac{f(m)m^2}{f(v)v^2}.$$

These strategies are well-defined by (16). Types above p^* have the same strategy as type p^*.

Given $m\in[\underline{v},a]$, the probability of acceptance of any demand $\hat{v}\in[\underline{v},a]\cup[b,\bar{v}]$ is the same as before. To see this, note that, with strategies redefined as above, $\int_m^{\bar{v}} f^B(m|v)f(v)dv$ is given by

$$\int_m^a \frac{\phi(m)}{v^2}dv+\int_a^b \frac{\phi(m)}{f(a)a^2}f(v)dv+\int_b^{p^*}\frac{\phi(m)}{v^2}dv+\int_{p^*}^{\bar{v}}\frac{\phi(m)f(v)}{f(p^*)(p^*)^2}dv$$

which, by (15), equals

$$\phi(m)[\frac{1}{m}-\frac{1}{p^*}+\frac{1-F(p^*)}{f(p^*)(p^*)^2}]$$

as before. Other calculations are similar. Therefore S is indifferent between all demands in $[m,a]\cup[b,p^*]$ and demands in $[\underline{v},m)\cup(p^*,\bar{v}]$ are worse. Similarly, after $m\in[b,p^*]$, the optimal demands are $[m,p^*]$. We need to show that, given $m\in[\underline{v},a]$, demands in (a, b) are suboptimal. Suppose then that, after $m \in [\underline{v}, a]$, $\hat{v} \in (a, b)$ is optimal, hence at least as good as a and b. Since all types in $[a, b)$ use the same strategy, the posterior distribution conditional on $[a, b)$ is the same for all messages in the support of these types, i.e., for all $m \in [\underline{v}, a]$. Hence $\hat{v}$ is optimal for S after all such messages, i.e., all messages in the supports of types below b. However, the argument leading to Lemma 4 implies then that

$$R(\hat{v}) \geq \frac{b - \hat{v}}{b - a} R(a) + \frac{\hat{v} - a}{b - a} R(b),$$

which contradicts the assumption that $\hat{v} \in I$.

Messages in (a, b) have zero probability, so, for $m \in (a, b)$, we can choose $\mu^B(\cdot|m)$ arbitrarily. Therefore, take $\mu^B(\cdot|m)$ so that S's optimal offers after m are $[b, p^*]$. QED

Proof of Theorem 3 For the case in which Q is not right-continuous at $\underline{v}$, we adapt the mechanism in the text as follows. Let $\hat{Q}(v) = Q(v)$ for all $v > \underline{v}$ and $\hat{Q}(\underline{v}) = lim_{v \downarrow \underline{v}} Q(v)$. Take the mechanism $\hat{\Gamma}(\hat{Q})$ augmented by adding an extra message, #, with $q(\#) = Q(\underline{v})$ and $t(\#) = \underline{v}(Q(\underline{v}) - \hat{Q}(\underline{v}))$. For all types other than $\underline{v}$ the strategy is the same as defined above, replacing Q by $\hat{Q}$. Type $\underline{v}$ sends # with probability 1. After any message in V, S's strategy is as before, replacing Q by $\hat{Q}$, and after # she believes B is type $\bar{v}$ with probability 1 and demands $\bar{v}$ with probability 1. Since # is a zero-probability message, S's belief is consistent with the definition of equilibrium. Type $\underline{v}$ is indifferent between messages $\underline{v}$ and #, so his strategy is optimal. QED

Proof of Claim 2 For any $\varepsilon > 0$,

$$\tilde{U}(v + \varepsilon) = \tilde{U}(v) + \int_{v}^{v+\varepsilon} \tilde{Q}(u) du$$

and

$$\tilde{U}(v - \varepsilon) = \tilde{U}(v) - \int_{v-\varepsilon}^{v} \tilde{Q}(u) du,$$

and similarly for $U(v + \varepsilon)$ and $U(v - \varepsilon)$. Therefore, since $\tilde{U} \geq U$, $\tilde{U}(v + \varepsilon) \geq U(v + \varepsilon)$

and $\tilde{U}(v-\varepsilon) \geq U(v-\varepsilon)$, so

$$\int_{v}^{v+\varepsilon} [\tilde{Q}(u) - Q(u)]du \geq 0 \tag{17}$$

and

$$\int_{v-\varepsilon}^{v} [\tilde{Q}(u) - Q(u)]du \leq 0. \tag{18}$$

(17) implies, using right-continuity of $\tilde{Q}$ and Q, that $\tilde{Q}(v) \geq Q(v)$, and (18) implies that $lim_{v'\uparrow v}\tilde{Q}(v') \leq lim_{v'\uparrow v}Q(v')$. This proves the Claim.

Proof of Claim 3 By Claim 2, $\tilde{Q}(v_1) > Q(v_1) \geq lim_{v'\uparrow v_1}Q(v') \geq lim_{v'\uparrow v_1}\tilde{Q}(v')$. Therefore $\tilde{Q}(v_1) > lim_{v'\uparrow v}\tilde{Q}(v')$. This shows that v_1 is a point of increase of $\tilde{Q}$.

Lemma 3 then implies that for all $m \in \cup_{v<v_1} M^B(v)/N$, $m \in M^B(v_1)$ and $v_1 \in V^S(m)$. This proves (ii). To prove (iii), we need to show that, for any $m' \in M^B(v_1)/N$ such that $m' \notin \cup_{v<v_1} M^B(v)$, $v_1 \in V^S(m')$. $U(v_1, m') = U(v_1)$, by Claim 1. However,

$$U(v_1, m') = U(m') + (v_1 - m')Q(m')$$

so

$$U(v_1) = U(m') + (v_1 - m')Q(m'),$$

which implies, since U is convex with (right-hand) derivative Q, that U is linear over the interval $[v_1, m']$ or $[m', v_1]$, with either $Q(v_1) = Q(m')$ or $Q(v_1) > lim_{v\uparrow v_1}Q(v) = Q(m')$. In either case, $\tilde{Q}(v_1) > Q(v_1) \geq Q(m')$. This implies (using Lemma 2) that, if type v_1 sends message m', he accepts a renegotiation demand with strictly positive probability. Since v_1 is the lowest type who sends m', this demand must be v_1. Hence, after message m', it must be optimal for S to demand v_1. This proves Claim 3.

Proof of Claim 4 Suppose that v' and $v'' > v'$ are both points in $[v_1, v_2]$ such that v' and v'' are both optimal for S after any $m \in \cup_{v\leq v_1} M^B(v)/N$. Take $v''' \in (v', v'')$. Then

$$v'''\tilde{\mu}^B([v''', \bar{v}]|m) \leq v'\tilde{\mu}^B([v', \bar{v}]|m)$$

and

$$v'''\tilde{\mu}^B([v''',\bar{v}]|m) \leq v''\tilde{\mu}^B([v'',\bar{v}]|m).$$

Therefore, by the argument leading to Lemma 4,

$$\tilde{\mu}^B(\cup_{v\leq v_1}M^B(v)/N|[v''',v''))\eta_F([v''',v'')) \leq \frac{(v''-v''')v'}{(v'''-v')v''}\tilde{\mu}^B(\cup_{v\leq v_1}M^B(v)/N|[v',v'''))\eta_F([v',v''')).$$

$\tilde{\mu}^B(\cup_{v\leq v_1}M^B(v)/N|[v''',v'')) = \tilde{\mu}^B(\cup_{v\leq v_1}M^B(v)/N|[v',v''')) = 1$ (noting that N has zero conditional probability), so

$$R(v''') \leq \frac{v''-v'''}{v''-v}R(v') + \frac{v'''-v'}{v''-v'}R(v''),$$

which contradicts the fact that R is strictly concave. This proves Claim 4.

References

Crawford, V. (1985) "Efficient and Durable Decision Rules: A Reformulation," *Econometrica* 53, 817-835.

Bester , H. and R. Strausz (2001) "Contracting with Imperfect Commitment and the Revelation Principle: The Single-Agent Case," *Econometrica*, 69, 1077-98.

Dewatripont, M. and E. Maskin (1990) "Contract Renegotiation in Models of Asymmetric Information," *European Economic Review* 34, 311-321.

Forges, F. (1993) "Some Thoughts on Efficiency and Information," in *Frontiers of Game Theory*, Ed. K. Binmore, A. Kirman, and P. Tani, MIT Press.

Forges, F. (1994) "Posterior Efficiency", *Games and Economic Behavior* 6, 238-261.

Fudenberg, D. and J. Tirole (1983) "Sequential Bargaining with Incomplete Information", *Review of Economic Studies*, 50, 221-47.

Fudenberg. D. and J. Tirole (1991) *Game Theory.* Cambridge, MA: MIT Press.

Green, J. R. and J.-J. Laffont (1987) “Posterior Implementability in a Two-Person Decision Problem”, *Econometrica* 55, 69-94.

Hart. O and J. Tirole (1988) “Contract Renegotiation and Coasian Dynamics”, *Review of Economic Studies*, 55, 509-40.

Holmström, B. and R. Myerson (1983) “Efficient and Durable Decision Rules with Incomplete Information”, *Econometrica* 51, 1799-1819.

Laffont, J.J. and J. Tirole (1988) “The Dynamics of Incentive Contracts”, *Econometrica* 56, 1153-75.

Laffont, J.J. and J. Tirole (1990) “Adverse Selection and Renegotiation in Procurement”, *Review of Economic Studies* 57, 597-625.

Lagunoff, R. D. (1995) “Resilient Allocation Rules for Bilateral Trade”, *Journal of Economic Theory* 66, 463-487.

Maskin, E. and J. Moore (1999) “Implementation and Renegotiation”, *Review of Economic Studies* 66, 39-56.

Milgrom, P. and I. Segal (2002) “Envelope Theorems for Arbitrary Choice Sets”, *Econometrica*

Milgrom, P. and R . Weber (1985) “Distributional Strategies for Games with Incomplete Information”, *Mathematics of Operations Research* 10, 619-632.

Mussa, M. and S. Rosen (1978) “Monopoly and Product Quality”, *Journal of Economic Theory* 18, 301-17.

Neeman, Z. and G. Pavlov (2007) “Renegotiation-Proof Mechanism Design”, mimeo.

Palfrey, T. R. and S. Srivastava (1991) “Efficient Trading Mechanisms with Pre-Play Communication”, *Journal of Economic Theory* 55, 17-40.

Segal, I., and M. Whinston (2002) “The Mirrlees Approach to Mechanism Design with Renegotiation (with Applications to Hold-Up and Risk Sharing)”, *Econometrica* 70, 1-45.

Skreta, V. (2006) "Sequentially Optimal Mechanisms", *Review of Economic Studies* 73, 1085-1111.