

Ex post renegotiation-proof mechanism design*

Zvika Neeman[†]

Gregory Pavlov[‡]

February 4, 2012

Abstract

We study a mechanism design problem under the assumption that ex post renegotiation cannot be prevented. We investigate what kind of equilibria of which mechanisms are ex post renegotiation-proof (EPRP) under a variety of ex post renegotiation procedures, and which social choice functions can be implemented in a way that is EPRP. In complete information environments, we show that the set of EPRP implementable social choice functions contains all ex post efficient allocations when there are at least three agents, but only budget balanced Groves allocations when there are two agents. In incomplete information environments with correlated beliefs and at least three agents, every ex post efficient social choice function is EPRP implementable, but with independent private values only social choice functions that are given by budget balanced “Groves in expectations” mechanisms are EPRP implementable. Examples are provided to illustrate the scope and limitations of our results.

JEL classification: D02; D70; D82; D86.

Keywords: Mechanism design; Implementation; Ex post renegotiation.

1 Introduction

One of the practical concerns of mechanism design theory is that players might often have incentives to change the rules of the game they are playing. Although in some cases the mechanism designer can prevent such changes, in many situations it is impossible or nearly impossible to do so, especially when a change in the rules of the game, contract, or mechanism is mutually beneficial for the agents. Such mutually consensual changes, which are known as *renegotiation*, can occur

*We would like to thank Eddie Dekel, Françoise Forges, Maria Goltsman, Johannes Hörner, Bart Lipman, Dilip Mookherjee, Michel Poitevin, Pasquale Schiraldi and the seminar participants at Arizona State University, Boston University, Hebrew University, Université de Montréal, Tel Aviv University, Toulouse, Vienna, University of Western Ontario, CETC (Montreal, 2007), Decentralization conference (Ann Arbor, 2007), World Congress of the Game Theory Society (Evanston, 2008), and mechanism design conference (Bonn, 2009) for helpful comments and conversations. All remaining errors are ours.

[†]The Eitan Berglas School of Economics, Tel-Aviv University, Tel-Aviv, Israel 69978; Email zvika@post.tau.ac.il, URL <http://www.tau.ac.il/~zvika/>.

[‡]Department of Economics, University of Western Ontario, Social Science Centre, London, Ontario N6A 5C2, Canada; Email gpavlov@uwo.ca, URL <http://economics.uwo.ca/faculty/pavlov/>.

at different stages of the contractual process. *Interim renegotiation* takes place before the mechanism is played and involves a change of the mechanism and the equilibrium the players intend to play. *Ex post renegotiation* takes place after the mechanism is played and involves a change of the outcome or recommendation proposed by the mechanism. The consequences of both interim and ex post renegotiation crucially depend on the details of the renegotiation process: what alternative outcomes or mechanisms are considered? How do the players communicate with each other, and how do they select among the alternative proposals? How is the surplus that is generated by renegotiation shared among the players? Etc.

This paper is devoted to the subject of ex post renegotiation.¹ We study the problem of a mechanism designer who is unable to prevent renegotiation and is ignorant of the exact way in which renegotiation will proceed. More specifically, we are interested in the question of what kind of equilibria of which mechanisms are *renegotiation-proof*, that is, can be expected to be “stable” in the sense of surviving in their original form under a variety of different renegotiation procedures. We also want to know which social choice functions can be implemented in a way that is renegotiation-proof.

The prospect of ex post renegotiation can undermine an equilibrium of a mechanism in two different ways. First, the equilibrium outcome that is reached by the mechanism can be renegotiated. Second, the incentive compatibility of the original equilibrium may be compromised because some agent may find it in his interest to deviate from the original equilibrium when the mechanism is played in anticipation of subsequent renegotiation. This second possibility can be illustrated by the following example. Garratt and Tröger (2006) consider a single object first price auction with two bidders, one that has a positive valuation for the object and one that values the object at zero. They show that if resale is not allowed, then there is a unique equilibrium in which both bidders bid zero and the high value bidder always wins the object. But if resale (or renegotiation) is allowed, then the above equilibrium is no longer incentive compatible, because the second bidder could deviate to a positive bid and try to resell the good to the first bidder. In the game with resale there is a unique equilibrium with resale on the equilibrium path. The former equilibrium is efficient, but the latter is not.

The standard approach to renegotiation that is used in the literature assumes that the mechanism designer can anticipate exactly how any outcome will be renegotiated. The anticipated renegotiation can then be incorporated into the original contract, and (in many cases) without loss of generality attention can then be limited to renegotiation-proof mechanisms.² In contrast, our notion of renegotiation-proofness requires that an equilibrium of a mechanism survives under all plausible renegotiation procedures.³

¹An extension of our approach to interim renegotiation is available in the working paper version of this paper. Ex post renegotiation generally imposes more constraints on the attainable outcomes than interim renegotiation. Segal and Whinston (2002) and Watson (2007) observe that this is the case in complete information environments, and Beaudry and Poitevin (1995) make a similar observation in the case of incomplete information environments.

²This result is known as the “renegotiation-proofness principle.” It was first introduced by Dewatripont (1989).

³There is a small literature that studies ex post renegotiation under multiple possible renegotiation procedures. In environments with complete information Sjöström (1999) and Amorós (2004) study the case of three or more

We consider quasi-linear environments with any number of agents under both complete and incomplete information. After presenting the basic setup in Section 2, in Section 3 we study ex post renegotiation in complete information environments. We show that ex post renegotiation-proof equilibria always result in ex post efficient outcomes (Lemma 1), and that for the case of three or more agents every ex post efficient social choice function can be implemented in a way that is ex post renegotiation-proof by a simple mechanism that requires two agents to report the state of the world and requires them to pay high penalties to other agents if they disagree (Proposition 2). When there are only two agents it is impossible to penalize both agents simultaneously when they disagree about the state of the world because such penalties would be undone by renegotiation. In this case, we show that the set of implementable social choice functions is strictly smaller than in the case of three or more agents: all ex post renegotiation-proof implementable social choice functions are outcome equivalent to budget balanced Groves mechanisms (Proposition 1). Many mechanism design environments do not admit the existence of budget balanced Groves mechanisms. In such cases, it is impossible to implement any social choice function in a way that is ex post renegotiation-proof. The implications and limitations of this result are discussed in Section 3.3.

Modeling ex post renegotiation in incomplete information environments is a more challenging task than in the case of complete information. In Section 4 we develop one possible approach under which the agents employ an incentive compatible renegotiation function through which they share their private information and decide collectively on how to renegotiate the outcome. If agents' values are interdependent then there may exist ex post inefficient renegotiation-proof equilibria (Example 5), but in the case of private values ex post renegotiation-proof equilibria are always ex post efficient (Lemma 2). When there are three or more agents and the agents' private information is correlated, we show that every ex post efficient social choice function can be implemented in a way that is ex post renegotiation-proof (Proposition 4). For the case of independent private values we show that every budget balanced "Groves in expectations" social choice function is ex post renegotiation-proof implementable.⁴ The opposite direction requires that no residual uncertainty remains after the mechanism is played (Proposition 3). As before, some mechanism design environments may not permit the implementation of any social choice function in a way that is ex post renegotiation-proof. The implications and limitations of this result in the context of incomplete information environments are discussed in Section 4.4.

All proofs are relegated to the Appendix unless otherwise stated.

2 Setup

A group of n agents must reach an agreement that involves the choice of a *social alternative* a from a set A and *monetary transfers* to the agents, $t = (t_1, \dots, t_n)$. An *outcome* (a, t) of the process of agents and Rubinstein and Wolinsky (1992) study a specific buyer and seller problem. Forges (1993, 1994) studies ex post renegotiation under incomplete information; Ausubel and Cramton (2004) allow for multiple resale scenarios following the Vickrey auction.

⁴"Groves in expectations" mechanisms are equivalent to Groves mechanisms from the agents' interim perspective. See, for example, Williams (1999).

negotiation among the agents is said to be *feasible* if $a \in A$ and $t \in \mathbb{T} = \{t \in \mathbb{R}^n : \sum_{i=1}^n t_i \leq 0\}$.⁵

The agents' preferences over outcomes depend on the state of the world θ that is chosen from a finite set Θ .⁶ Each agent i is an expected utility maximizer with a quasi-linear payoff function $v_i(a, \theta) + t_i$, where $v_i : A \times \Theta \rightarrow \mathbb{R}$ describes his preferences over social alternatives for different states of the world, and t_i denotes the monetary transfer he receives. We assume that for every state θ there is a single alternative $a^*(\theta)$ that maximizes the total surplus $\sum_{i=1}^n v_i(a, \theta)$; this alternative is called *ex post efficient in state θ* .⁷ A vector of transfers t that satisfies $\sum_{i=1}^n t_i = 0$ is called *budget balanced*. An outcome (a, t) is called *ex post efficient in state θ* if $a = a^*(\theta)$ and t is budget balanced.

A *social choice function* is a mapping $f : \Theta \rightarrow A \times \mathbb{T}$ from the set of states into feasible outcomes. A social choice function f is said to be *ex post efficient* if for every $\theta \in \Theta$ the outcome $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient in state θ .

We distinguish between complete and incomplete information environments as follows:

Complete information. Complete information environments are analyzed in Section 3. The state of the world θ is commonly known among the agents. It is convenient to denote the set of possible preferences of agent i by Θ_i . With this notation, the state of the world can be represented by a vector of the agents' payoff relevant types: $\theta = (\theta_1, \dots, \theta_n)$, where $\theta_i \in \Theta_i$, and, with a slight abuse of notation, $v_i(a, \theta) = v_i(a, \theta_i)$ for every i and a .

While in general $\Theta \subseteq \prod_{i=1}^n \Theta_i$, for some results we assume that the set of states has a structure of a Cartesian product.⁸

Condition 1 (Product domain) $\Theta = \prod_{i=1}^n \Theta_i$.

A *mechanism* $\langle S, m \rangle$ specifies a set of messages S_i for each agent i and an outcome rule $m : S \rightarrow \Delta(A \times \mathbb{T})$ from the set of message profiles $S = \prod_{i=1}^n S_i$ into the set of lotteries over feasible outcomes. A mechanism $\langle S, m \rangle$ together with a state of the world θ defines a complete information game, where a strategy of agent i is $\sigma_i(\theta) \in \Delta(S_i)$. We say that $\sigma = (\sigma_1, \dots, \sigma_n)$ is an *equilibrium* if for every $\theta \in \Theta$ the strategy profile $(\sigma_1(\theta), \dots, \sigma_n(\theta))$ is a Nash equilibrium of the complete information game in state θ . Equilibrium σ is said to be *ex post efficient* if the equilibrium outcome is ex post efficient in every state $\theta \in \Theta$.

A social choice function f is said to be *implementable* if there exists a mechanism $\langle S, m \rangle$ that has an equilibrium σ such that the equilibrium outcome coincides with $f(\theta)$ in every state θ .⁹

Incomplete information. Incomplete information environments are analyzed in Section 4. The set of states of the world is $\Theta = \prod_{i=1}^n \Theta_i$, there is a common prior distribution P over Θ ,

⁵In some environments the right-hand-side of the feasibility constraint on the sum of monetary transfers may depend on the chosen social alternative. Using a standard argument it is possible to redefine the agents' payoffs in a way that is consistent with our model. See, for example, Chapter 7 in Fudenberg and Tirole (1991).

⁶We assume that Θ is finite to simplify the exposition, but our results continue to hold if Θ is infinite.

⁷This assumption is generically satisfied if Θ is finite.

⁸This is a standard condition in implementation theory (sometimes called "independent domain"). See, for example, Moore (1992).

⁹Thus we employ a weak notion of implementation.

and each agent i privately observes his type $\theta_i \in \Theta_i$. We consider both private and interdependent values environments; by private values we mean the following.

Condition 2 (*Private values*) *The payoff of each agent depends only on his own type, i.e., with a slight abuse of notation:*

$$v_i(a, \theta_i, \theta_{-i}) = v_i(a, \theta_i) \quad \forall i \in \{1, \dots, n\}, a \in A, \theta_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}$$

Like in the case of complete information, a *mechanism* $\langle S, m \rangle$ specifies a set of messages S_i for each agent i and an outcome rule $m : S \rightarrow \Delta(A \times \mathbb{T})$ from the set of message profiles $S = \prod_{i=1}^n S_i$ into the set of lotteries over feasible outcomes. A mechanism $\langle S, m \rangle$ together with a common prior P defines a Bayesian game, where a strategy of agent i is $\sigma_i : \Theta_i \rightarrow \Delta(S_i)$. We say that $\sigma = (\sigma_1, \dots, \sigma_n)$ is an *equilibrium* if it is a Bayesian Nash equilibrium of the Bayesian game. Equilibrium σ is said to be *ex post efficient* if the equilibrium outcome is ex post efficient in every state $\theta \in \Theta$.

A social choice function f is said to be *implementable* if there exists a mechanism $\langle S, m \rangle$ that has an equilibrium σ such that the equilibrium outcome is payoff equivalent to f from the agents' interim perspective. That is, the equilibrium expected payoff of agent i of type $\theta_i \in \Theta_i$ is equal to

$$E_{\theta_{-i}|\theta_i} [v_i(a(\theta), \theta) + t_i(\theta)] = \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i} | \theta_i) (v_i(a(\theta), \theta) + t_i(\theta))$$

where $(a(\theta), t(\theta)) = f(\theta)$ for every $\theta \in \Theta$.

Although we do not discuss the subject of individual rationality in this paper (except for Section 4.4), it is straightforward to incorporate into the analysis an appropriate notion of individual rationality as an additional set of constraints on the social choice function and mechanism.

3 Ex post renegotiation-proofness under complete information

3.1 Preliminaries

In a complete information environment a state of the world $\theta \in \Theta$ is realized and becomes commonly known among the agents, but is not known to the mechanism designer. The agents then play the game induced by a given mechanism $\langle S, m \rangle$. After the outcome of playing the mechanism is determined, the agents have an opportunity to change it to some other outcome according to some renegotiation procedure.¹⁰ Our goal is to determine which equilibria of which mechanisms are robust against such renegotiation.

Rather than modeling various renegotiation procedures as games, we take the following shortcut. Suppose an outcome (a, t) was reached by the play of the mechanism in state θ . We say that the outcome (a, t) *can be renegotiated* if there exists another feasible outcome (a', t') that Pareto

¹⁰Alternatively one could allow renegotiation to take place before the lotteries from $\Delta(A \times \mathbb{T})$ prescribed by the outcome rule m are carried through. Since the agents have quasi-linear payoffs such an approach would not affect the results, but it would make the arguments more cumbersome.

dominates (a, t) in state θ . That is, in this state of the world all the agents weakly prefer, and at least one agent strictly prefers, the alternative outcome (a', t') to the original outcome (a, t) that was prescribed by the mechanism.¹¹

Consider an equilibrium of some mechanism. If we allow the possibility of ex post renegotiation, then an equilibrium outcome may be undermined in two ways. First, the equilibrium outcome reached by the mechanism may be renegotiated. Second, the original equilibrium strategy profile may cease to be an equilibrium once ex post renegotiation is allowed, since deviating when playing the mechanism may bring about an outcome which may itself be beneficially renegotiated. These two possibilities lead to the following definition.

Definition 1 *An equilibrium σ of a mechanism $\langle S, m \rangle$ is ex post renegotiation-proof (EPRP) if both of the following conditions hold:*

- (i) An outcome that is obtained under the equilibrium play of the mechanism cannot be renegotiated.*
- (ii) No agent can improve upon his equilibrium payoff in any state by a unilateral deviation from σ followed by renegotiation of the resulting outcomes.*

We are also interested in social choice functions that can be implemented in a renegotiation-proof way.

Definition 2 *A social choice function is ex post renegotiation-proof (EPRP) implementable if there exists a mechanism $\langle S, m \rangle$ that has an EPRP equilibrium σ such that the equilibrium outcome coincides with $f(\theta)$ in every state θ .*

3.2 Results

In this section we study what kind of equilibria of which mechanisms are EPRP, and characterize the social choice functions that are EPRP implementable. The first requirement for an equilibrium of some mechanism to be EPRP is that the equilibrium outcomes are not Pareto dominated in any state. This is a simple implication of part (i) of Definition 1 and is stated without proof.

Lemma 1 *In a complete information environment every EPRP equilibrium is ex post efficient.*

Next we present an example where an equilibrium fails to be EPRP in spite of being ex post efficient and in dominant strategies. In particular, the example shows that ex post renegotiation-proofness not only requires the transfers to be budget balanced in every state of the world, but also imposes restrictions on the relative size of the transfers across different states of the world.

¹¹A standard approach to modeling renegotiation, like in Maskin and Moore (1999), assumes that ex post renegotiation proceeds according to some given process (or a game). This process induces a mapping from outcomes and states into (possibly different) outcomes, and this mapping is commonly known by the agents and the mechanism designer. In contrast, we allow more flexibility in the ways renegotiation may proceed, and assume that the mechanism designer is ignorant of the exact way in which renegotiation will proceed. See Section 3.3 for further discussion.

Example 1 A buyer and a seller can trade a single good. The buyer values the good at V that can be either 0 or 2, the seller values the good at 1. The realization of V and the seller's valuation are commonly known between the agents. Consider a mechanism where the buyer is asked to report his value: after a report " $V = 2$ " the good is transferred from the seller to the buyer at a price p_2 , and after a report " $V = 0$ " there is no trade and the buyer pays p_0 to the seller. It is easy to see that the buyer has a dominant strategy to report his true valuation if $p_2 - p_0 \in (0, 2)$, and the resulting outcome is ex post efficient. However, as we show below, this equilibrium is not EPRP unless $p_2 - p_0 = 1$.

Suppose $p_2 - p_0 \in (1, 2)$. If the buyer with $V = 2$ reports " $V = 0$ " then the payoffs of the buyer and the seller (without renegotiation) would be $-p_0$ and p_0 , respectively. This outcome is Pareto dominated by a decision to trade at a new price \hat{p} that satisfies $\hat{p} - p_0 \in (1, 2)$. Hence, for any such $\hat{p} < p_2$, the buyer would prefer to misreport and then renegotiate the outcome to trade at the price \hat{p} rather than report his true valuation. Thus, the original equilibrium is not EPRP.¹²

This example deals with just one simple class of mechanisms. More generally, one could ask both the buyer and the seller about the state of the world, and then condition the outcomes on their reports. However, as we argue below, allowing for more complicated mechanisms does not expand the range of implementable outcomes in the environment considered in this example.

Note that in this example we can restate the EPRP restrictions on the transfers as a requirement that the seller's payoff (without renegotiation) be independent of the report of the buyer. If it is not the case, then the buyer could make a report that minimizes the seller's status quo payoff, and renegotiate the outcome while keeping all renegotiation surplus to himself. But if the seller's payoff is a constant independent of the state of the world, then budget balance implies that the buyer is a residual claimant of the total surplus. Hence, in the EPRP equilibrium of the considered mechanism the buyer's transfer equals the externality that his report imposes on the seller.

To generalize this idea we consider a class of *Groves mechanisms* from the mechanism design literature. Each agent i reports his own payoff relevant type $\theta_i \in \Theta_i$, an ex post efficient alternative given the reports is implemented, and each agent is given a transfer (as a function of the reports) that equals the externality he imposes on the other agents. There exists an equilibrium of this mechanism in dominant strategies where each agent tells the truth, and thus Groves mechanisms can be used to implement the following class of social choice functions.

Definition 3 Suppose there are two agents. A social choice function f is called Groves if $f(\theta) = (a^*(\theta), t(\theta))$ where a^* is ex post efficient and transfers t satisfy

$$t_i(\theta_i, \theta_j) = v_j(a^*(\theta_i, \theta_j), \theta_j) + H_i(\theta_j) \quad \forall i \in \{1, 2\}, \theta_i \in \Theta_i, \theta_j \in \Theta_j, \quad (1)$$

for some function $H_i : \Theta_j \rightarrow \mathbb{R}$.

¹²The argument for the case $p_2 - p_0 \in (0, 1)$ is similar. The buyer with $V = 0$ will find it profitable to report " $V = 2$ " and then renegotiate to "no trade" as long as a new payment \hat{p} is smaller than p_0 .

Next we describe which social choice functions for the case of two agents can be implemented in a renegotiation-proof way. If the *product domain* restriction (Condition 1) is satisfied, then only Groves social choice functions that are budget balanced are EPRP implementable. Otherwise, there may be other EPRP implementable social choice functions as well.¹³

Proposition 1 *Consider a complete information environment with two agents.*

- (i) *Every budget balanced Groves social choice function is EPRP implementable.*
- (ii) *Suppose the product domain condition is satisfied, and a social choice function f is EPRP implementable. Then f is budget balanced Groves.*

To prove part (i) of the result we consider a budget balanced Groves mechanism that corresponds to a given budget balanced Groves social choice function. This mechanism has an ex post efficient equilibrium where each agent truthfully reports his payoff relevant type. Moreover, like in Example 1, the truthful equilibrium of this mechanism has a property that each agent’s payoff (without renegotiation) is independent of the report of the opponent.¹⁴ Using the terminology of Segal and Whinston (2010) in this case there are “no harmful direct externalities” in the absence of renegotiation. Hence, any strategy of misreporting one’s own type and subsequent renegotiation is unprofitable: it can neither lower the opponent’s payoff, nor increase the total surplus (since the equilibrium outcome is ex post efficient).¹⁵

To illustrate the idea behind the proof of part (ii) of the result, first note that by Lemma 1 only ex post efficient social choice functions can potentially be EPRP implementable. Suppose we want to implement some ex post efficient social choice function that is not budget balanced Groves. A simple mechanism that works when there is no renegotiation is as follows. Each agent is asked to report the whole state of the world (and not just his payoff relevant type as in Groves mechanisms). If the agents’ reports agree, the outcome prescribed by the social choice function is implemented; if they disagree, both agents are penalized. However, penalties on both agents must involve ex post inefficient outcomes, which can be undone by renegotiation. Hence, if we want to sustain truth-telling in such a direct revelation mechanism, we need to ensure that no agent can profit from misreporting the state and renegotiating the resulting outcome.

Consider a simple case where agent 1 has two possible types, θ_1 and θ'_1 , and the preferences of agent 2 are fixed. Denote by $U_2(\alpha, \beta)$ the payoff of agent 2 in a direct revelation mechanism

¹³See Example 2 below.

¹⁴Suppose agent i has type θ_i and reports it truthfully, while agent j reports θ_j . Then i ’s payoff is

$$v_i(a^*(\theta_i, \theta_j), \theta_i) + t_i(\theta_i, \theta_j) = v_i(a^*(\theta_i, \theta_j), \theta_i) - t_j(\theta_i, \theta_j) = -H_j(\theta_i)$$

where the first equality is by budget balance, and the second by (1).

¹⁵Segal and Whinston (2010) consider an incomplete contracts model where the parties make ex ante investments that are noncontractible. In Proposition 5 they show that first best investments can be achieved if there are “no harmful direct externalities” between the players, and ex post bargaining is efficient. We thank a referee for bringing a connection with this result to our attention.

(without renegotiation) when agent 1's report is α and agent 2's report is β , where $\alpha, \beta \in \{\theta_1, \theta'_1\}$. The truth-telling conditions for agent 2 are

$$U_2(\theta_1, \theta_1) \geq U_2(\theta_1, \theta'_1) \quad \text{and} \quad U_2(\theta'_1, \theta'_1) \geq U_2(\theta'_1, \theta_1). \quad (2)$$

On the other hand, it can be shown that the renegotiation-proofness constraints for agent 1 boil down to the requirement of “no harmful direct externalities” of agent 1's report on agent 2's status quo payoff:

$$U_2(\theta'_1, \theta_1) \geq U_2(\theta_1, \theta_1) \quad \text{and} \quad U_2(\theta_1, \theta'_1) \geq U_2(\theta_1, \theta'_1). \quad (3)$$

Combining (2) and (3) we find that the payoff of agent 2 must be independent of the reports. Hence, agent 1 must be a residual claimant of the total surplus as in a Groves mechanism. This argument implies that in the environment considered in Example 1 only budget balanced Groves social choice functions are EPRP implementable.

If only one agent's payoff depends on the state of the world, then we can repeat the argument above for every pair of states, and show that every EPRP implementable social choice function must be budget balanced Groves.¹⁶ If the preferences of both agents depend on the state of the world, then the argument above can be used to link every pair of states where the preferences of only one agent vary. For product domain environments this generates enough restrictions on EPRP implementable social choice functions to imply that they must be budget balanced Groves. If the product domain condition is not satisfied, then, as we show in the next example, there may be other social choice functions that are renegotiation-proof.

Example 2 *A buyer and a seller can trade a single good. The buyer values the good at V , which can be either 0 or 2, and the seller values the good at C , which can be either 1 or 3. The realization of (V, C) is commonly known between the agents. Every budget balanced Groves social choice function prescribes trade if and only if the state is $(2, 1)$, the buyer's transfers satisfy $t_b(2, 1) - t_b(0, 1) = -1$ and $t_b(0, 3) - t_b(2, 3) = 0$, and the seller's transfers satisfy $t_s(2, 1) - t_s(2, 3) = 2$ and $t_s(0, 3) - t_s(0, 1) = 0$. In addition, budget balance requires that $t_b(V, C) + t_s(V, C) = 0$ for every (V, C) . This implies that*

$$0 = (t_b(2, 1) + t_s(2, 1)) + (t_b(0, 3) + t_s(0, 3)) - (t_b(0, 1) + t_s(0, 1)) - (t_b(2, 3) + t_s(2, 3)) = 1.$$

Hence, no budget balanced Groves social choice functions exist in this environment.

By part (ii) of Proposition 1 there are no EPRP implementable social choice functions if all combinations of V and C are possible. Suppose, however, that there are just two possible states: (V, C) is either equal $(0, 1)$ or $(2, 3)$. A social choice function that prescribes no trade together with a constant payment from one agent to another is EPRP implementable by a mechanism that simply ignores the agents' messages.

¹⁶Note that the product domain condition is trivially satisfied in this case.

Next, we briefly discuss the case of three or more agents. In this case, any ex post efficient social choice function f can be EPRP implemented using the following standard construction. Agents 1 and 2 are asked to report the state of the world; if their reports agree, the outcome prescribed by f is implemented; if their reports conflict, agents 1 and 2 pay large penalties to agent 3. There exists an equilibrium where agents 1 and 2 report the state truthfully if the penalties are large enough. This equilibrium is EPRP: (i) the outcome is ex post efficient if the agents tell the truth; (ii) the penalties can be chosen to be large enough so that any strategy of sending a wrong report followed by renegotiation would be unprofitable.¹⁷

Proposition 2 *Consider a complete information environment with $n \geq 3$ agents. A social choice function f is EPRP implementable if and only if $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient for every θ .*¹⁸

3.3 Discussion

Comparison with the standard approach. Ex post renegotiation is usually assumed to proceed according to some exogenously given process that is commonly known among the agents, and is also known to the mechanism designer who may take it into account when designing the mechanism.¹⁹ Since many papers on renegotiation under complete information belong to the hold-up and incomplete contracts literature, let us consider one such model that was studied by Edlin and Reichelstein (1996) and Che and Hausch (1999).²⁰ A buyer and a seller make relationship-specific investments that determine their payoffs from subsequent trade. Following the investment decisions the state of the world is realized and commonly observed, and then the agents may play a game that is induced by a given mechanism. In Edlin and Reichelstein (1996) the mechanism is simply a constant outcome determined by a fixed price contract, while Che and Hausch (1999) allow for nontrivial revelation mechanisms. The outcomes determined by the mechanism are renegotiated so that the post-renegotiation outcome is ex post efficient, and the renegotiation surplus is shared between the agents in fixed proportions. This process of ex post renegotiation uniquely determines a mapping from the outcomes reached by the mechanism and states into (possibly other) outcomes.

Our approach to ex post renegotiation assumes that the mechanism designer is uncertain about how renegotiation works in this context. It may be unclear to him whether renegotiation always results in ex post efficient outcomes, and the sharing rule may be unknown to him and may vary depending on circumstances. In such a case, a mechanism designer who is interested in a particular social choice function may prefer to avoid uncertain renegotiation and try to achieve implementation through an equilibrium of a mechanism that is renegotiation-proof.

¹⁷Sjöström (1999) obtains a similar result in a somewhat different framework. He studies an exchange economy with agents that may be risk-neutral or risk-averse, and similarly to us assumes that the mechanism designer is ignorant of the exact way renegotiation will proceed. Theorem 1 in Sjöström (1999) shows that a social choice function is uniquely implementable in undominated Nash equilibrium if and only if it is Pareto efficient.

¹⁸The proof of Proposition 2 is available in the working paper version of the paper.

¹⁹See, for example, Maskin and Moore (1999). Some papers, like Chung (1991) and Aghion et al. (1994) assume that the mechanism designer has the power to affect the way the renegotiation will proceed.

²⁰Segal and Whinston (2002) discuss other papers that use a similar approach.

Existence of EPRP mechanisms and social choice functions. As explained above, EPRP implementable social choice functions may fail to exist when there are only two agents. If the preferences of one of the agents are independent of the state of the world, then it is easy to construct budget balanced Groves mechanisms using the same argument as in Example 1. However, as shown in Example 2, existence is not guaranteed when the preferences of both agents depend on the state of the world and the product domain condition is satisfied. To get a better sense of the scope of nonexistence suppose that each agent’s set of types is a connected open subset of an Euclidean space and that payoffs are continuously differentiable in types. Then any mechanism that implements ex post efficient alternatives in dominant strategies is outcome equivalent to a truthful equilibrium of a Groves mechanism.²¹ If we require in addition budget balance, then Groves mechanisms often fail to exist.²²

We realize that the EPRP requirements are strong, and so it is not surprising that some environments admit no EPRP mechanisms and implementable social choice functions. We hope that our theory of renegotiation-proof mechanisms can be viewed as a theory of “stable” institutions. If a mechanism retains its original form under a wide variety of circumstances, then such a mechanism must be renegotiation-proof in our strong sense. If, on the other hand, such a renegotiation-proof mechanism does not exist in a given environment, then one should not expect to observe an institution for governing transactions whose rules remain stable and not subject to perpetual renegotiation.

What next? There are several ways to weaken or circumvent the EPRP requirements.

1. The mechanism designer may be interested in implementing social choice correspondences, rather than social choice functions. For example, suppose the designer cares only about efficiency and not about how the surplus is distributed among the agents.²³ If renegotiation is uncertain, but we know that it leads to efficient outcomes, then there is no reason for the designer to restrict attention to renegotiation-proof mechanisms, and she should be content with any mechanism. If renegotiation is uncertain and imperfect, then exact efficiency may be unachievable, but the designer could try to make sure that outcomes are not “too inefficient.”

2. Renegotiation can be modeled as being less permissive, or being more predictable. For example, in certain situations renegotiation is costly. In other cases exogenous restrictions may be imposed on the kinds of renegotiations that are allowed.

3. For many specific applications the scope of renegotiation can be significantly reduced through careful contract design that pays close attention to the details of trade and renegotiation technology. For example, Watson (2007) shows that the effect of ex post renegotiation may be significantly

²¹See, for example, Holmström (1979).

²²See, for example, Green and Laffont (1979) and Walker (1980).

²³In the hold-up and incomplete contracts models it seems natural to study social choice functions, because the distribution of the surplus between the agents in different states of the world determines the incentives for ex ante investments. For example, in the case of “selfish” investments (when investment affects only an agent’s own payoff) Rogerson (1992) shows that mechanisms that are “Groves in expectations” (see Section 4.2) provide first-best incentives for ex ante investment. Hence, first-best investment incentives can be provided by EPRP mechanisms if there are three or more agents, and whenever budget balanced Groves mechanisms exist if there are only two agents.

reduced when the actions that consummate trade are modelled as inalienable actions of some agents and trade opportunities are nondurable. In related work, Evans (2009) shows that ex post renegotiation loses much of its power if sending a message in a mechanism entails an individual cost for the agent who sends it.

4. The mechanism designer may try to extract information about anticipated renegotiation from the agents through the course of play of the mechanism, and to condition the outcomes of the mechanism on this information. Such mechanisms are likely to be quite complex because they must provide the agents with the incentives to reveal all their private information – their payoff relevant private information and their beliefs about the expected renegotiation scenario.²⁴ In contrast, we study a simpler class of mechanisms that do not attempt to extract any information about expected renegotiation from the agents, but nonetheless implement the desired social choice function. Studying such mechanisms is a natural first step towards understanding of how to do mechanism design in the presence of multiple renegotiation scenarios.²⁵

5. In some two agent environments it may be possible to expand the set of EPRP implementable social choice functions by introducing a third party. This third party need not be able to observe the state of the world, its only role would be to collect the penalties from agents 1 and 2 when their reports about the state of the world disagree.²⁶

4 Ex post renegotiation-proofness under incomplete information

4.1 Preliminaries

In an incomplete information environment the state of the world $\theta = (\theta_1, \dots, \theta_n)$ is drawn according to a prior probability distribution P , each agent i learns his private information θ_i and employs Bayes’ rule to update his beliefs about the other agents’ types. Then, the agents play the Bayesian game that is induced by a given mechanism $\langle S, m \rangle$. After they “play the mechanism” the agents have an opportunity to change the outcome that is determined by the mechanism to some other outcome according to some renegotiation procedure. Our goal is to determine which equilibria of which mechanisms are robust against various renegotiation scenarios of this form.²⁷

²⁴We are aware of only two papers that allow such complex mechanisms. Green and Laffont (1994) consider a quasi-linear model with two players and costly renegotiation. They construct a renegotiation-proof mechanism for implementing the ex post efficient social choice correspondence. Amoros (2004) studies unique implementation in general environments with uncertain renegotiation. His characterization results require the mechanism designer to be able to destroy the agents’ endowments, which would lead to renegotiation in our model.

²⁵Conceptually, the mechanisms we study relate to the complex mechanisms outlined above in a way that is similar to the way that dominant strategy (and ex post) mechanisms relate to Bayesian mechanisms. Dominant strategy (and ex post) mechanisms are “belief-free” and rely only on the agents’ payoff relevant information, while Bayesian mechanisms may also rely on the agents’ higher-order beliefs (Bergemann and Morris, 2005; Chung and Ely, 2007). The simple mechanisms we study are “belief-free” with respect to anticipated renegotiation scenarios, while complex mechanisms may rely on the agents’ beliefs about future renegotiation.

²⁶While it is commonly argued that an introduction of third parties raises a possibility of collusion by a subset of agents, Sjöström (1999) and Baliga and Sjöström (2009) show that there are circumstances when this need not be the case.

²⁷As in the case of complete information, the standard approach to modeling renegotiation assumes that ex post renegotiation proceeds according to some given process (or a game). This process induces a mapping from outcomes,

Suppose that the play of the mechanism produced the outcome (a, t) . Under complete information we said that (a, t) can be renegotiated if there is another outcome (a', t') that Pareto dominates the original outcome given the state of the world, and thus each agent is willing to switch from (a, t) to (a', t') . There are several challenges with developing a natural counterpart of this definition for the case of incomplete information. First, the agents may still face some residual uncertainty about the state of the world even after the outcome of the mechanism is observed. The agents may have different preferences over (a, t) vs. (a', t') depending on their own private information, and depending on the private information of the other agents. Second, a model of renegotiation as a collective decision on whether to switch from the status quo outcome (a, t) to a single alternative outcome (a', t') seems overly restrictive. A more plausible model of renegotiation should allow the agents to communicate with each other, reveal some additional information, and then decide to which alternative outcome to switch conditional on the revealed information.

To provide a formal definition of our renegotiation procedure we first discuss the agents' beliefs at the point immediately after they have played the mechanism but before they have started renegotiating the outcome produced by the mechanism. Fix an equilibrium strategy profile σ of a given mechanism $\langle S, m \rangle$. For every outcome (a, t) that may be potentially reached we can derive the agents' beliefs conditional on this outcome being realized. First suppose agent i of type θ_i has played his equilibrium strategy σ_i . If outcome (a, t) is on the equilibrium path, then his beliefs are derived by Bayes' rule: he takes into account the reached outcome (a, t) , the equilibrium strategy profile σ , and his type θ_i . If outcome (a, t) was not supposed to happen under σ , then he may hold arbitrary beliefs.²⁸ Next suppose agent i of type θ_i has played σ'_i instead of his equilibrium strategy σ_i . For every outcome (a, t) the beliefs of this agent are derived by Bayes' rule: he takes into account the reached outcome (a, t) , the strategy profile (σ'_i, σ_{-i}) , and his type θ_i .

After a given outcome (a, t) has been reached by the mechanism, we allow it to be challenged by some *renegotiation mapping* $\rho : \Theta \rightarrow \Delta(A \times \mathbb{T})$. Note that ρ can be state contingent, because we want to allow the status quo outcome to be renegotiated to different alternative outcomes depending on the state of the world. Since the information about the state of the world is dispersed among the agents, appropriate incentives must be given to the agents for them to reveal it. Specifically, a mapping $\rho : \Theta \rightarrow \Delta(A \times \mathbb{T})$ is said to be *incentive compatible* for given agents' beliefs if under these beliefs each type of each agent weakly prefers the outcome according to ρ to any other outcome according to ρ where he lies about his type and the other agents tell the truth.²⁹ Finally, we want

states, and the agents' beliefs into outcomes and the agents' beliefs, and this mapping is commonly known among the agents and the mechanism designer. In contrast, we permit renegotiation to proceed in many different ways, and assume that the mechanism designer is ignorant of the exact way in which it will proceed. See Section 4.4 for further discussion.

²⁸There exist other reasonable approaches to specifying beliefs after the outcomes that are off the equilibrium path. For example, one can view these beliefs as being part of the description of the equilibrium of the mechanism, and, since we are using a weak notion of implementation, argue that these beliefs may be "chosen" by the mechanism designer. Since none of our results depend on the way these off-equilibrium beliefs are specified we have decided to adopt a more permissive approach and allow for arbitrary beliefs.

²⁹Incentive compatibility can equivalently be defined as follows. Consider a direct mechanism $\rho : \Theta \rightarrow \Delta(A \times \mathbb{T})$. Together with given agents' beliefs, ρ specifies an ex post renegotiation game of incomplete information. We say that ρ is incentive compatible for given agents' beliefs if ρ admits a truthful equilibrium.

the agents to be willing to switch from the outcome prescribed by the mechanism to whatever alternative outcome is generated by the renegotiation mapping ρ . Thus, ρ is called *posterior individually rational with respect to outcome* (a, t) if for every outcome (a', t') that is produced by ρ , no agent strictly prefers (a, t) to (a', t') given his beliefs and whatever has been revealed through the play of ρ .

Such an approach to ex post renegotiation appears to be quite general, and we would like to note that many alternative renegotiation procedures can be equivalently represented in this framework. For example, suppose that after a given outcome (a, t) has been reached, a single alternative outcome (a', t') is proposed, and the agents simultaneously vote whether to switch from (a, t) to (a', t') . If all agents vote in favor of switching, then (a', t') is implemented; otherwise (a, t) is implemented. The voting rules, together with given agents' beliefs, specify an ex post renegotiation game of incomplete information. For any equilibrium of this game we can construct an outcome equivalent incentive compatible renegotiation mapping ρ using a revelation principle-like argument. Since in this voting game each agent can veto the switch to (a', t') , the equilibrium conditions guarantee that ρ satisfies posterior individual rationality with respect to (a, t) .³⁰

We do not claim, however, that for any outcome of any reasonable renegotiation procedure there exists an equivalent representation in our framework. Here are a few considerations. First, one may argue that after the agents have played the mechanism agent i 's private information may contain other things beyond θ_i . For example, suppose the equilibrium strategy σ_i of agent i is mixed, and that the other agents are unable to infer the pure actions actually taken by agent i from observing the final outcome (a, t) . In this case the private history of actions taken by agent i becomes part of i 's private information. Second, in some settings the agents may receive additional signals about the state of the world after they have played the mechanism. This may relax the incentive compatibility constraints on the renegotiation mapping ρ that ensure truthful information revelation by the agents. As an extreme example, if the state of the world becomes commonly known at the ex post stage, then it is reasonable to drop the incentive compatibility constraints on ρ altogether.³¹ Third, one could replace the requirements of posterior individual rationality with weaker constraints of interim individual rationality.³² Fourth, there may exist alternative reasonable approaches to specifying the agents' beliefs at the point just before renegotiation.

We believe that our framework can be extended to incorporate the considerations above. At the moment, however, we have chosen to proceed with our current approach. Summarizing, here is our definition of renegotiation under incomplete information.

Definition 4 *Fix an equilibrium σ of a mechanism $\langle S, m \rangle$.*

³⁰We have studied this approach to renegotiation in the previous version of this paper.

³¹One of the approaches we considered in the previous version of the paper allowed for an external "oracle" device which could help the agents with renegotiation by revealing the state of the world.

³²We found such an approach to be too permissive. If agent j of type θ_j plays an equilibrium strategy of type θ'_j , then the other agents may end up with incorrect posterior beliefs about j . In such a case interim individual rationality (given incorrect beliefs) will not protect the other agents from accepting renegotiation proposals that are worse for them than the status-quo.

- (i) Suppose that an outcome (a, t) can be reached with a positive probability under σ . We say that (a, t) can be renegotiated with a positive probability on the equilibrium path if there exists an incentive compatible renegotiation mapping ρ that is posterior individually rational with respect to (a, t) , and at least one type of one agent that has a positive probability (given (a, t) and σ) strictly prefers ρ to (a, t) , where the beliefs of each agent i are derived by Bayes' rule (conditional on (a, t) , σ , and θ_i).
- (ii) Suppose that an outcome (a, t) can be reached with a positive probability following agent j 's unilateral deviation σ'_j from σ . We say that (a, t) can be renegotiated with a positive probability off the equilibrium path if there exists an incentive compatible renegotiation mapping ρ that is posterior individually rational with respect to (a, t) , and at least one type of one agent that has positive probability (given (a, t) and (σ'_j, σ_{-j})) strictly prefers ρ , where the beliefs of agent j are derived by Bayes' rule (conditional on (a, t) , (σ'_j, σ_{-j}) and θ_j), and the beliefs of each agent $i \neq j$ are derived by Bayes' rule (conditional on (a, t) , σ and θ_i) whenever possible and are arbitrary otherwise.

As in the case of complete information, there are two ways in which ex post renegotiation may undermine the equilibrium of a given mechanism. First, the equilibrium outcome that is reached by the mechanism may be renegotiated. Second, the original equilibrium strategy profile may cease to be an equilibrium once ex post renegotiation is allowed, since a deviation in the mechanism would bring about an outcome which may be potentially beneficially renegotiated. These two possibilities lead to the following incomplete information analog of Definition 1 (Section 3.1).³³

Definition 5 *An equilibrium σ of a mechanism $\langle S, m \rangle$ is ex post renegotiation-proof (EPRP) if both of the following conditions hold:*

- (i) *The outcomes obtained under the equilibrium play of the mechanism cannot be renegotiated with a positive probability on the equilibrium path.*
- (ii) *No agent can improve upon his interim equilibrium payoff by a unilateral deviation from σ followed by renegotiation of the resulting outcome with a positive probability off the equilibrium path.*

Here is an incomplete information analog of Definition 2 (Section 3.1).

Definition 6 *A social choice function is ex post renegotiation-proof (EPRP) implementable if there exists a mechanism $\langle S, m \rangle$ that has an EPRP equilibrium σ such that the equilibrium outcome is payoff equivalent to f from the agents' interim perspective.*

³³A closely related notion is that of "posterior efficiency," which is due to Forges (1993, 1994). An equilibrium of a given mechanism is posterior efficient if the outcomes obtained from equilibrium play cannot be renegotiated with a positive probability on the equilibrium path, where the approach to modeling renegotiation of an outcome is very similar to ours. Hence the definition of a posterior efficient equilibrium is similar to part (i) of our Definition 5. However, the concept of posterior efficiency assumes that ex post renegotiation is not anticipated by the agents when they play in the mechanism and thus cannot undermine the incentive compatibility of the original equilibrium. Thus the definition of posterior efficient equilibrium does not have an analog to part (ii) of our Definition 5.

4.2 Independent private values

In this section we study environments with private values. Note that each agent's preferences over any pair of outcomes, and in particular any original outcome produced by the mechanism and any alternative outcome, are independent of his beliefs over the other agents' types and depend only on his own type.

As in the case of complete information, in environments with private values there can be no ex post inefficient equilibria that are EPRP. In every state of the world each agent knows his own preferences, and thus everyone will consent to a Pareto improving change.³⁴ This implies the following lemma.

Lemma 2 *In an incomplete information environment with private values every EPRP equilibrium is ex post efficient.*

For the rest of this section we restrict attention to the case where the agents' private information is distributed independently. The case of correlated private information is discussed in Section 4.3.

Any mechanism that is ex post efficient under incomplete information has to induce the agents to reveal their types. For environments with independent private values the following class of mechanisms received special attention in the literature. Each agent i reports his own type $\theta_i \in \Theta_i$, an ex post efficient alternative given the reports is chosen, and each agent is given a transfer (as a function of the reports) that equals the expected externality that the agent imposes on the other agents. Such mechanisms may be called *Groves in expectation mechanisms*, because the interim expected equilibrium transfers of each agent under such a mechanism is identical to those of some Groves mechanism.³⁵ Each such mechanism has an equilibrium where the agents tell the truth, and thus Groves in expectation mechanisms can be used to implement the following class of social choice functions.

Definition 7 *A social choice function f is called Groves in expectation if $f(\theta) = (a^*(\theta), t(\theta))$ where a^* is ex post efficient and transfers t satisfy*

$$E_{\theta_{-i}}[t_i(\theta_i, \theta_{-i})] = E_{\theta_{-i}} \left[\sum_{j \neq i} v_j(a^*(\theta_i, \theta_{-i}), \theta_j) \right] + H_i \quad \forall i \in N, \theta_i \in \Theta_i, \quad (4)$$

for some $H_i \in \mathbb{R}$.

Next we describe the class of social choice functions that can be implemented in a renegotiation-proof way. As in the case of complete information with two agents, the requirements of ex post renegotiation-proofness restrict the relative size of the transfers across different states of the world. We show that Groves in expectations social choice functions that are budget balanced can always

³⁴There is a similar result for posterior efficient equilibria in Forges (1993, 1994).

³⁵The class of Groves in expectation mechanisms includes Groves mechanisms, AGV mechanisms (Arrow, 1979; d'Aspremont and Gerard-Varet, 1979), as well as other mechanisms.

be implemented in a renegotiation-proof way, and that no other social choice function can be EPRP implemented by an equilibrium of a mechanism that reveals “enough information”.

Proposition 3 *Consider an incomplete information environment with independent private values.*

- (i) *Every budget balanced Groves in expectation social choice function is EPRP implementable.*
- (ii) *Suppose that a social choice function f is EPRP implementable by an equilibrium σ in a mechanism $\langle S, m \rangle$ with the property that for every agent i with nontrivial private information ($\# |\Theta_i| > 1$), after $\langle S, m \rangle$ has been played, i 's belief about the other agents is concentrated on a single type profile θ_{-i} . Then f is budget balanced Groves in expectations.*

The idea of the proof of part (i) of the result is similar to that of part (i) of Proposition 1 (Section 3.2). Consider a budget balanced Groves in expectations mechanism that corresponds to a given budget balanced Groves in expectations social choice function. There is a truthful equilibrium of this mechanism that is ex post efficient, and that has the property that no agent's report imposes “harmful direct externalities” on the other agents (in expectation). Specifically, given the information of agent i , the expected sum of the other agents' payoffs (without renegotiation) is independent of i 's report.³⁶ Hence, an agent who misreports his type and consequently renegotiates the resulting outcomes cannot get more than his equilibrium payoff, because the other agents must be paid at least their equilibrium payoffs, and the total surplus is maximized in equilibrium.³⁷

Note that the property mentioned in part (ii) of the result is trivially satisfied in environments where only one agent has private information. Furthermore, in such environments budget balanced Groves in expectations social choice functions are simply budget balanced Groves. Hence, part (ii) of the result implies that in environments with one sided private information every EPRP implementable social choice function must be budget balanced Groves. To see the intuition, note that by Lemma 2 only ex post efficient social choice functions can be potentially EPRP implementable, and that to achieve efficiency one must rely on a report of the agent with private information. Then, by the same logic as in Example 1 (Section 3.2), renegotiation-proofness requires the reporting agent to be the residual claimant of the total surplus, or, equivalently, that the sum of the payoffs of the other agents be independent of his report. Otherwise he could report the type that minimizes the other agents' status quo payoffs, and renegotiate the outcome to an ex post efficient one while keeping all renegotiation surplus to himself.

³⁶Suppose agent i reports the type θ_i , and the other agents report their types truthfully. Then the expected sum of the other agents' payoffs is

$$E_{\theta_{-i}} \left[\sum_{j \neq i} v_j (a^* (\theta_i, \theta_{-i}), \theta_j) + \sum_{j \neq i} t_j (\theta_i, \theta_{-i}) \right] = E_{\theta_{-i}} \left[\sum_{j \neq i} v_j (a^* (\theta_i, \theta_{-i}), \theta_j) - t_i (\theta_i, \theta_{-i}) \right] = -H_i$$

where the first equality is due to budget balance, and the second follows from (4).

³⁷Ausubel and Cramton (2004) obtain a result similar to part (i) of Proposition 3 for multi-unit auction environments. They show that the truthful bidding remains an equilibrium of the Vickrey auction even in the presence of multiple resale opportunities.

Next, suppose that several agents have private information, and that after the mechanism is played each agent with private information can figure out the state of the world. Then, similarly to the case of one sided private information, renegotiation-proofness requires that each reporting agent fully internalizes the effect of his report on the social outcome. However, since the agents submit their messages before observing the types of the other agents, it is enough to assign them Groves in expectation transfers that would make them residual claimants in expectation, conditional on their private information at the interim stage.

Finally, suppose that several agents have private information, and that some residual uncertainty about the state of the world remains after the mechanism is played. The presence of asymmetric information at the renegotiation stage may reduce the profitability of misreporting followed by renegotiation, because even in the most favorable renegotiation scenario a deviating agent may fail to capture the entire difference between the efficient total surplus and the status quo payoffs of the other agents. In such cases the renegotiation-proofness constraints are weaker, and it may be possible that social choice functions that are different from budget balanced Groves in expectation functions are EPRP implementable.

For illustration we consider two buyer-seller examples with two sided private information. First we revisit the setting of Example 2 (Section 3.2) under incomplete information.

Example 3 *A buyer and a seller can trade a single good. The buyer values the good at V which can be either 0 or 2 with equal probability, and the seller values the good at C which can be either 1 or 3 with equal probability, and V and C are independently distributed. Unlike in Example 2 each agent privately observes the realization of his value. Every budget balanced Groves in expectation social choice function prescribes trade if and only if $(V, C) = (2, 1)$, the buyer's transfers satisfy $t_b(2, 1) - t_b(0, 3) = -\frac{3}{2}$ and $t_b(2, 3) - t_b(0, 1) = \frac{1}{2}$, and the seller's transfers can be found from the budget balance conditions, namely: $t_s(V, C) = -t_b(V, C)$ for every (V, C) . By part (i) of Proposition 3 such social choice functions are EPRP implementable. It is possible to show that in this environment there are no other EPRP implementable social choice functions.³⁸*

The next example exhibits EPRP implementable social choice functions that are not budget balanced Groves in expectations.

Example 4 *Consider the same setting as in Example 3, but suppose V is either 0 or 3, and C is either 1 or 2. Then every budget balanced Groves in expectation social choice function prescribes trade if and only if $V = 3$, the buyer's transfers satisfy $t_b(3, 1) - t_b(0, 2) = t_b(3, 2) - t_b(0, 1) = -\frac{3}{2}$, and the seller's transfers are the opposite of the buyer's transfers. By part (i) of Proposition 3 such social choice functions are EPRP implementable, but, as we show below, there are others as well.*

³⁸If such a social choice function existed, then by part (ii) of Proposition 3 an equilibrium of a mechanism that implements it must preserve some asymmetry of information between the agents after the play of the mechanism. However, it can be shown that even if the other agents' private information is not automatically revealed, it may nonetheless be possible for a deviating agent to extract the other agents' private information at no cost, renegotiate the trade decision to an efficient one, and capture the whole renegotiation surplus. For details see Lemma 3 in the working paper version of this paper.

Suppose the buyer is asked to report his value: after a report “ $V = 3$ ” the good is transferred from the seller to the buyer at a price $p \in (1, 2)$, and after a report “ $V = 0$ ” there is no trade and no payment. The buyer has a dominant strategy to tell the truth, and there will be no renegotiation on the equilibrium path because the outcome is *ex post* efficient. Suppose the buyer with $V = 3$ reports “ $V = 0$ ” which results in no trade and no payment. If the seller’s cost was known, then the most profitable renegotiation scenario for the buyer would be to renegotiate to trade and pay the seller no more than his cost, which would result in an expected payoff of $\frac{1}{2}(3 - 1) + \frac{1}{2}(3 - 2) = \frac{3}{2}$. However, the seller’s cost is his private information, and it is not revealed in the course of the play of the mechanism. Hence, to induce both types of the seller to trade, the buyer has to pay at least 2, which results in an expected payoff of at most $3 - 2 = 1$, and this is less than the equilibrium payoff $3 - p$. To induce only the seller with $C = 1$ to trade, the buyer has to pay at least 1, which results in an expected payoff of at most $\frac{1}{2}(3 - 1) = 1 < 3 - p$. The argument for the buyer with $V = 0$ is similar.

4.3 Correlated incomplete information

In this section we first discuss what happens if the agents have interdependent valuations. Then we present an analog of Proposition 2 (Section 3.2) for the case of three or more agents with correlated beliefs.

In environments with interdependent valuations *ex post* renegotiation-proofness no longer implies *ex post* efficiency. Suppose that every agent prefers the outcome (a', t') to the outcome (a, t) when the state of the world is $\theta = (\theta_1, \dots, \theta_n)$. Nevertheless, type θ_i of agent i may be reluctant to switch to (a', t') if there exists another state $(\theta_i, \tilde{\theta}_{-i})$ in which he prefers (a, t) to (a', t') . The next example illustrates this possibility.

Example 5 *A buyer and a seller can trade a single good. There are two equally likely states: in state L the buyer’s value is 0 and the seller’s value is 1, and in state H the buyer’s value is 2 and the seller’s value is $\frac{3}{2}$. The seller privately observes the realization of the state, while the buyer is uninformed. Consider a mechanism that prescribes no trade and no payment (independently of the agents’ messages). The trivial equilibrium of this mechanism is EPRP despite being *ex post* inefficient in state H.*

Suppose the agents are offered to renegotiate this outcome to trade at a price \hat{p} . Note that \hat{p} must be at least $\frac{3}{2}$ to give the seller incentives to trade in state H. But then the seller would like to trade in state L as well. So, no trade will take place since the buyer is not willing to pay more than his expected value, which is equal to $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2 = 1$.

In mechanism design theory the case of incomplete information with correlated agents’ beliefs is very similar to the complete information case. In the latter case the mechanism designer can ask the agents to report the state of the world and verify their reports against each other. In the former case the mechanism designer can ask the agents to report their types and “stochastically

compare” them against each other. For our next result we introduce the following condition from Kosenok and Severinov (2008).

Condition 3 *The prior probability distribution P over Θ satisfies:*

(i) *For every agent i and every type θ'_i :*

$$P(\cdot | \theta'_i) \neq \sum_{\theta_i \in \Theta_i} c_{\theta_i, \theta'_i} P(\cdot | \theta_i) \quad \forall c_{\theta_i, \theta'_i} \in \mathbb{R}_+, \theta_i, \theta'_i \in \Theta_i$$

(ii) *For every probability distribution $P' \neq P$ there exists agent i of type θ'_i (with a positive P' -probability) such that:*

$$P'(\cdot | \theta'_i) \neq \sum_{\theta_i \in \Theta_i} c_{\theta_i, \theta'_i} P(\cdot | \theta_i) \quad \forall c_{\theta_i, \theta'_i} \in \mathbb{R}_+, \theta_i, \theta'_i \in \Theta_i$$

Part (i) of Condition 3 requires the conditional beliefs of each type of each agent to be linearly independent (in a restricted sense). It allows to design personalized transfers for each agent (as a function of reports) that ensure truth-telling.³⁹ If we want the system of personalized transfers for the agents to satisfy budget balance, then this may distort the agents’ incentives to tell the truth. Part (ii) of Condition 3 takes care of this issue.

This condition allows us to obtain the following incomplete information analog of the “if” part of Proposition 2.

Proposition 4 *Consider an incomplete information environment with $n \geq 3$ agents with a prior distribution P that satisfies Condition 3. A social choice function f is EPRP implementable if $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient for every θ .*

4.4 Discussion

Comparison with the standard approach. As in the case of complete information, ex post renegotiation under incomplete information is usually assumed to proceed according to some process that is commonly known among the agents, and is also known to the mechanism designer. For example, Beaudry and Poitevin (1995) study a model with two agents who first play a game induced by a given mechanism, and then have a chance to renegotiate the realized outcomes as follows. One of the agents proposes a new mechanism; if the other agent agrees, then they play the new mechanism which determines the final outcome; otherwise the outcome prescribed by the original mechanism is implemented. The proposer of the new mechanism takes into account the prescribed outcome, his beliefs about the state of the world, beliefs about the opponent’s beliefs, etc. The equilibrium of the subsequent game (consisting of the ratification procedure and the new mechanism) determines new outcomes and revised beliefs for every state of the world. This

³⁹See, for example, Crémer and McLean (1988).

procedure uniquely determines the renegotiation mapping from the outcomes, states and the agents' beliefs into outcomes and the agents' beliefs.

In contrast we assume that there are multiple possible renegotiation possibilities, and that the mechanism designer is uncertain about which one the agents will follow. Such a mechanism designer would presumably be interested in a mechanism that implements the desired social choice function that is robust against all reasonable renegotiation possibilities, if it exists.

Existence of EPRP mechanisms and social choice functions. In incomplete information environments with independent private values, budget balanced Groves in expectations mechanisms are known to exist.⁴⁰ Hence, by Proposition 3, in such environments EPRP implementable social choice functions exist. It should be noted, however, that budget balanced Groves in expectations social choice functions do not always satisfy interim individual rationality.⁴¹ To get a sense of the scope of the problem suppose that each agent's set of types is a connected open subset of an Euclidean space and that payoffs are continuously differentiable in types. Then any mechanism that is both ex post efficient and interim incentive compatible is payoff equivalent to a Groves in expectations mechanism.⁴² It is well known that in several economically important environments with incomplete information there are no individually rational ex post efficient mechanisms.⁴³ Whenever this is the case, the existence of interim individually rational EPRP implementable social choice functions and mechanisms is not guaranteed.

In the cases of interdependent values and correlated beliefs considered in Section 4.3 EPRP implementable social choice functions were shown to exist. However, we do not know whether existence can be ensured in general.

What next? If the EPRP requirements appear to be too stringent for a particular problem, one could use alternative approaches 1 through 4 from Section 3.3.

Another approach that could be useful is to restrict the allowed renegotiation mappings to be "credible", so that the resulting post-renegotiation state contingent allocations could not be renegotiated again in some states of the world.⁴⁴ While it is possible to show that a restriction to credible renegotiation proposals does not alter our results in the case of complete information,

⁴⁰This follows from the existence of AGV mechanisms mentioned in footnote 35.

⁴¹In Example 3 from Section 4.2 no budget balanced Groves in expectation social choice function (and thus no EPRP implementable social choice function) satisfies interim individual rationality. To see this note that the interim expected payoffs of the buyer with $V = 0$ and of the seller with $C = 3$ are $\frac{1}{2}t_b(0, 1) + \frac{1}{2}t_b(0, 3)$ and $\frac{1}{2}t_s(0, 3) + \frac{1}{2}t_s(2, 3)$, respectively. Summing up these payoffs and using budget balance and the restrictions on the transfers, it follows that $\frac{1}{2}t_b(0, 1) - \frac{1}{2}t_b(2, 3) = -\frac{1}{4}$. Hence, the payoffs of the buyer with $V = 0$ and of the seller with $C = 3$ cannot both be nonnegative.

⁴²See, for example, Williams (1999).

⁴³For example, bilateral trade (Myerson and Satterthwaite, 1983), public good provision (Mailath and Postlewaite, 1990), litigation and settlement (Spier, 1994; Klement and Neeman, 2005).

⁴⁴Such an approach is considered, for example, in Krasa (1999) who studied a different notion of ex post renegotiation-proofness called *unimprovability*. Unimprovability is a property of a state contingent allocation rather than a property of an equilibrium of some mechanism, and the process by which the agents arrive at a given allocation is not modeled. The allocation is called unimprovable if the agents do not want to change it by re-trading goods or revealing further information. In this respect the definition of unimprovable allocation is related to the requirement on EPRP equilibrium outcomes given in part (i) of our Definition 5.

this is not the case under incomplete information. We have chosen not to restrict attention to “credible” renegotiation proposals, because we think that a more permissive notion of renegotiation goes well with our goal of determining which equilibria of which mechanisms are robust against various renegotiation scenarios. For example, in some cases the agents might have time for just one round of renegotiation, and then some “non-credible” renegotiation agreements may be realized.

Throughout the paper we have assumed that after a mechanism is played the realized outcome is revealed to the players. In principle there may be other possible disclosure policies. Some mechanisms may in addition reveal some or all of the actions taken by the players in the course of the play of the mechanism. Other mechanisms may reveal to each player only part of the information about the realized outcome, say, only his own transfer and some minimal information about the chosen social alternative. We conjecture that Propositions 3 and 4 would continue to hold true under any disclosure policy.

5 Appendix

Proof of Proposition 1. (i) Suppose $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient in every $\theta \in \Theta$, i.e. $a(\theta) = a^*(\theta)$ and $t(\theta)$ is budget balanced, and transfers t satisfy (1). Consider a mechanism where each agent i reports his own payoff relevant type $\theta_i \in \Theta_i$, and the outcome rule as a function of reports is given by f . In this mechanism truthtelling is a dominant strategy, since for each agent i of type θ_i

$$\begin{aligned} v_i(a^*(\theta_i, \theta_j), \theta_i) + t_i(\theta_i, \theta_j) &= v_i(a^*(\theta_i, \theta_j), \theta_i) + v_j(a^*(\theta_i, \theta_j), \theta_j) + H_i(\theta_j) \\ &\geq v_i(a^*(\theta'_i, \theta_j), \theta_i) + v_j(a^*(\theta'_i, \theta_j), \theta_j) + H_i(\theta_j) \\ &= v_i(a^*(\theta'_i, \theta_j), \theta_i) + t_i(\theta'_i, \theta_j) \quad \forall \theta'_i \neq \theta_i \end{aligned}$$

where the equalities use (1), and the inequality uses the fact that $a^*(\theta_i, \theta_j)$ is ex post efficient in state (θ_i, θ_j) . Hence, the truthful equilibrium of this mechanism implements f .

Suppose the above equilibrium is not EPRP. First, note that there will be no renegotiation on the equilibrium path since the outcome $(a^*(\theta), t(\theta))$ is ex post efficient for every θ . Suppose agent 1 in state $\theta = (\theta_1, \theta_2)$ can improve upon his truthtelling payoff by sending a wrong report $\theta'_1 \neq \theta_1$ and consequent renegotiation of the outcome from $(a^*(\theta'_1, \theta_2), t(\theta'_1, \theta_2))$ to another outcome (\hat{a}, \hat{t}) . This strategy is beneficial for agent 1 if

$$v_1(\hat{a}, \theta_1) + \hat{t}_1 > v_1(a^*(\theta_1, \theta_2), \theta_1) + t_1(\theta_1, \theta_2) \tag{5}$$

agent 2 agrees to such renegotiation if

$$v_2(\hat{a}, \theta_2) + \hat{t}_2 \geq v_2(a^*(\theta'_1, \theta_2), \theta_2) + t_2(\theta'_1, \theta_2) \tag{6}$$

Sum up (5) and (6):

$$\begin{aligned}
v_1(\widehat{a}, \theta_1) + v_2(\widehat{a}, \theta_2) + \widehat{t}_1 + \widehat{t}_2 &> v_1(a^*(\theta_1, \theta_2), \theta_1) + v_2(a^*(\theta'_1, \theta_2), \theta_2) + t_1(\theta_1, \theta_2) + t_2(\theta'_1, \theta_2) \\
&= v_1(a^*(\theta_1, \theta_2), \theta_1) + v_2(a^*(\theta'_1, \theta_2), \theta_2) + t_1(\theta_1, \theta_2) - t_1(\theta'_1, \theta_2) \\
&= v_1(a^*(\theta_1, \theta_2), \theta_1) + v_2(a^*(\theta_1, \theta_2), \theta_2)
\end{aligned}$$

where the first equality uses budget balance ($t_2(\theta'_1, \theta_2) = -t_1(\theta'_1, \theta_2)$), and the second equality uses (1). But $a^*(\theta_1, \theta_2)$ is ex post efficient in state (θ_1, θ_2) and $\widehat{t}_1 + \widehat{t}_2 \leq 0$:

$$v_1(a^*(\theta_1, \theta_2), \theta_1) + v_2(a^*(\theta_1, \theta_2), \theta_2) \geq v_1(\widehat{a}, \theta_1) + v_2(\widehat{a}, \theta_2) + \widehat{t}_1 + \widehat{t}_2$$

which gives a contradiction. Hence, the truthful equilibrium of this mechanism is EPRP.

(ii) Suppose a mechanism $\langle S, m \rangle$ has an EPRP equilibrium σ that implements the social choice function f . For every state $\theta = (\theta_1, \theta_2) \in \Theta$ define a maximized total surplus: $S(\theta_1, \theta_2) = \sum_{i=1}^2 v_i(a^*(\theta_1, \theta_2), \theta_i)$.

First, by Lemma 1 note that $f(\theta) = (a(\theta), t(\theta))$ must be ex post efficient: $a(\theta) = a^*(\theta)$ and $t_1(\theta) + t_2(\theta) = 0$ for every θ . Next, suppose agent 1 in state θ uses strategy $\sigma_1(\theta')$ where $\theta' = (\theta'_1, \theta'_2) \in \Theta$. Since σ is EPRP, agent 1 cannot improve upon his equilibrium payoff even if the resulting outcomes are renegotiated in the most profitable way. That is, every outcome $(\widetilde{a}, \widetilde{t})$ is renegotiated to an ex post efficient outcome $(a^*(\theta), \widehat{t})$ such that agent 2 receives a payoff that leaves him indifferent between the two outcomes:

$$v_2(a^*(\theta), \theta_2) + \widehat{t}_2 = v_2(\widetilde{a}, \theta_2) + \widetilde{t}_2$$

The payoff of agent 1 in this case is

$$v_1(a^*(\theta), \theta_1) - \widehat{t}_2 = v_1(a^*(\theta), \theta_1) + v_2(a^*(\theta), \theta_2) - (v_2(\widetilde{a}, \theta_2) + \widetilde{t}_2)$$

The renegotiation-proofness constraint to prevent such deviations by agent 1 is

$$v_1(a^*(\theta), \theta_1) + t_1(\theta) \geq v_1(a^*(\theta), \theta_1) + v_2(a^*(\theta), \theta_2) - E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_2(\widetilde{a}, \theta_2) + \widetilde{t}_2] \quad (7)$$

where the expectation for the last term is taken with respect to a probability measure over outcomes $(\widetilde{a}, \widetilde{t})$ induced by the outcome rule m and the strategy profile $(\sigma_1(\theta'), \sigma_2(\theta))$. Rewrite (7):

$$t_1(\theta) + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\widetilde{t}_2] \geq v_2(a^*(\theta), \theta_2) - E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_2(\widetilde{a}, \theta_2)] \quad \forall \theta' \neq \theta \quad (8)$$

Repeat the same argument for agent 2:

$$t_2(\theta) + E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [\widetilde{t}_1] \geq v_1(a^*(\theta), \theta_1) - E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [v_1(\widetilde{a}, \theta_1)] \quad \forall \theta' \neq \theta \quad (9)$$

Sum up (8) and (9), and use $t_1(\theta) + t_2(\theta) = 0$:

$$\begin{aligned} & E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [\tilde{t}_1] + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_2] \\ \geq & S(\theta_1, \theta_2) - E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [v_1(\tilde{a}, \theta_1)] - E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_2(\tilde{a}, \theta_2)] \end{aligned} \quad (10)$$

Switch the roles of θ and θ' in (10):

$$\begin{aligned} & E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_1] + E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [\tilde{t}_2] \\ \geq & S(\theta'_1, \theta'_2) - E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_1(\tilde{a}, \theta'_1)] - E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [v_2(\tilde{a}, \theta'_2)] \end{aligned} \quad (11)$$

Sum up (10) and (11):

$$\begin{aligned} & E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [\tilde{t}_1 + \tilde{t}_2] + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_1 + \tilde{t}_2] \\ \geq & S(\theta_1, \theta_2) + S(\theta'_1, \theta'_2) - E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [v_1(\tilde{a}, \theta_1) + v_2(\tilde{a}, \theta'_2)] \\ & - E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_1(\tilde{a}, \theta'_1) + v_2(\tilde{a}, \theta_2)] \end{aligned} \quad (12)$$

Note that the left hand side of (12) is nonpositive since transfers \tilde{t} are feasible. Using this fact and rearranging we get

$$\begin{aligned} & E_{m \circ (\sigma_1(\theta), \sigma_2(\theta'))} [v_1(\tilde{a}, \theta_1) + v_2(\tilde{a}, \theta'_2)] + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [v_1(\tilde{a}, \theta'_1) + v_2(\tilde{a}, \theta_2)] \\ \geq & S(\theta_1, \theta_2) + S(\theta'_1, \theta'_2) \end{aligned} \quad (13)$$

Thus a necessary condition for ex post renegotiation-proofness is

$$S(\theta_1, \theta'_2) + S(\theta'_1, \theta_2) \geq S(\theta_1, \theta_2) + S(\theta'_1, \theta'_2). \quad (14)$$

A reverse inequality in (14) is obtained if we repeat the same argument for states (θ_1, θ'_2) and (θ'_1, θ_2) instead of (θ_1, θ_2) and (θ'_1, θ'_2) . Hence, a necessary condition for ex post renegotiation-proofness is

$$S(\theta_1, \theta'_2) + S(\theta'_1, \theta_2) = S(\theta'_1, \theta'_2) + S(\theta_1, \theta_2) \quad (15)$$

and thus all inequalities used to arrive at (15) must hold as equalities. In particular, the left side of (13) must be equal to the left side of (14), which implies that $m \circ (\sigma_1(\theta), \sigma_2(\theta'))$ is equal to $a^*(\theta_1, \theta'_2)$ with certainty, and $m \circ (\sigma_1(\theta'), \sigma_2(\theta))$ is equal to $a^*(\theta'_1, \theta_2)$ with certainty. Thus (8) becomes

$$t_1(\theta) + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_2] = v_2(a^*(\theta), \theta_2) - v_2(a^*(\theta'_1, \theta_2), \theta_2) \quad \forall \theta' \neq \theta \quad (16)$$

and if we switch the roles of θ and θ' in (9):

$$t_2(\theta') + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_1] = v_1(a^*(\theta'), \theta'_1) - v_1(a^*(\theta'_1, \theta_2), \theta'_1) \quad \forall \theta \neq \theta' \quad (17)$$

Take $\theta = (\theta_1, \theta_2)$ and $\theta' = (\theta'_1, \theta_2)$, and add up (16) and (17):

$$t_1(\theta_1, \theta_2) + t_2(\theta'_1, \theta_2) + E_{m \circ (\sigma_1(\theta'), \sigma_2(\theta))} [\tilde{t}_1 + \tilde{t}_2] = v_2(a^*(\theta_1, \theta_2), \theta_2) - v_2(a^*(\theta'_1, \theta_2), \theta_2) \quad (18)$$

Another implication of the argument above is that all inequalities used to arrive at (15) must hold as equalities is $\tilde{t}_1 + \tilde{t}_2 = 0$. Rewrite (18) using budget balance ($t_2(\theta'_1, \theta_2) = -t_1(\theta'_1, \theta_2)$):

$$t_1(\theta_1, \theta_2) - t_1(\theta'_1, \theta_2) = v_2(a^*(\theta_1, \theta_2), \theta_2) - v_2(a^*(\theta'_1, \theta_2), \theta_2)$$

Hence, the transfer rule satisfies (1). The argument for agent 2 is identical. ■

Proof of Lemma 2. Let σ be an EPRP equilibrium of a mechanism $\langle S, m \rangle$. Suppose σ is not ex post efficient in some state $\theta = (\theta_1, \dots, \theta_n)$. That is, there exists some outcome (a, t) reached with positive probability by σ when the profile of types is θ , such that (a, t) is not ex post efficient in state θ .

Suppose the following indirect renegotiation mechanism is offered to the agents after outcome (a, t) is reached. The agents simultaneously announce Y or N . If at least one agent have announced N , then outcome (a, t) is realized. If all agents have announced Y , then outcome $(a^*(\theta), \hat{t})$ is realized, where $a^*(\theta)$ maximizes the total surplus in state θ , and transfers \hat{t} are defined as

$$\hat{t}_i = v_i(a, \theta_i) - v_i(a^*(\theta), \theta_i) + t_i + \frac{1}{n} \left(\sum_{j=1}^n v_j(a^*(\theta), \theta_j) - \sum_{j=1}^n v_j(a, \theta_j) \right).$$

Transfers \hat{t} are feasible since transfers t are feasible: $\sum_{i=1}^n \hat{t}_i = \sum_{i=1}^n t_i \leq 0$.

Note that agent i of type θ_i strictly prefers $(a^*(\theta), \hat{t})$ to (a, t) for any beliefs about the opponents' types:

$$v_i(a^*(\theta), \theta_i) + \hat{t}_i = v_i(a, \theta_i) + t_i + \frac{1}{n} \left(\sum_{j=1}^n v_j(a^*(\theta), \theta_j) - \sum_{j=1}^n v_j(a, \theta_j) \right) > v_i(a, \theta_i) + t_i$$

Hence, there exists an equilibrium of this renegotiation mechanism such that in state θ all agents choose Y , and thus the outcome is $(a^*(\theta), \hat{t})$. Using a standard argument it is straightforward to construct an outcome equivalent incentive compatible direct mechanism ρ . By construction, this mechanism ρ is posterior individually rational with respect to (a, t) , and it satisfies the requirements of Definition 3(i).

This gives a contradiction, and hence an EPRP equilibrium σ must be ex post efficient. ■

Proof of Proposition 3. (i) Suppose $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient in every $\theta \in \Theta$, i.e. $a(\theta) = a^*(\theta)$ and $t(\theta)$ is budget balanced, and transfers t satisfy (4). Consider a direct revelation mechanism where each agent i reports his type $\theta_i \in \Theta_i$, and the outcome rule as a function of reports is given by f . There is a truthtelling equilibrium of this mechanism, since for each agent i

of type θ_i

$$\begin{aligned}
E_{\theta_{-i}} [v_i(a^*(\theta_i, \theta_{-i}), \theta_i) + t_i(\theta_i, \theta_{-i})] &= E_{\theta_{-i}} \left[v_i(a^*(\theta_i, \theta_{-i}), \theta_i) + \sum_{j \neq i} v_j(a^*(\theta_i, \theta_{-i}), \theta_j) \right] + H_i \\
&\geq E_{\theta_{-i}} \left[v_i(a^*(\theta'_i, \theta_{-i}), \theta_i) + \sum_{j \neq i} v_j(a^*(\theta'_i, \theta_{-i}), \theta_j) \right] + H_i \\
&= E_{\theta_{-i}} [v_i(a^*(\theta'_i, \theta_{-i}), \theta_i) + t_i(\theta'_i, \theta_{-i})] \quad \forall \theta'_i \neq \theta_i
\end{aligned}$$

where the equalities use (4), and the inequality uses the fact that $a^*(\theta_i, \theta_{-i})$ is ex post efficient in state (θ_i, θ_{-i}) . Hence, the truthful equilibrium of this mechanism implements f .

Suppose the truthtelling equilibrium σ is not EPRP. First, suppose that some outcome (a, t) is reached with positive probability by the equilibrium play, and it can be renegotiated with positive probability. Note that the beliefs of each agent i over Θ_{-i} are derived by Bayes' rule from P conditional on (a, t) , σ , and θ_i . Thus we can say that at this point the agents share a common "prior" over Θ , which is derived by Bayes' rule from P conditional on (a, t) and σ . Also note that it is commonly known among the agents that (a, t) is ex post efficient. Hence, by the "no-trade theorem" (Milgrom and Stokey, 1982) there does not exist a feasible incentive compatible mapping ρ that improves upon outcome (a, t) as required by Definition 3(i).

Next, suppose agent 1 of type θ_1 can improve upon his interim payoff from truthtelling by sending a wrong report $\theta'_1 \neq \theta_1$ and consequent renegotiation of the resulting outcomes. For example, if the true state is (θ_1, θ_{-1}) , then outcome $(a^*(\theta'_1, \theta_{-1}), t(\theta'_1, \theta_{-1}))$ will be reached by the play of the mechanism. By Definition 3(ii) every viable renegotiation mechanism $(\hat{a}, \hat{t}) : \Theta \rightarrow \Delta(A \times \mathbb{T})$ must satisfy posterior individual rationality with respect to the reached outcome for each agent $i \neq 1$

$$v_i(\hat{a}(\theta_1, \theta_{-1}), \theta_i) + \hat{t}_i(\theta_1, \theta_{-1}) \geq v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_i(\theta'_1, \theta_{-1}) \quad (19)$$

Sum up (19) over $i \neq 1$ and take expectation over Θ_{-1} :

$$E_{\theta_{-1}} \left[\sum_{i \neq 1} v_i(\hat{a}(\theta_1, \theta_{-1}), \theta_i) + \sum_{i \neq 1} \hat{t}_i(\theta_1, \theta_{-1}) \right] \geq E_{\theta_{-1}} \left[\sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + \sum_{i \neq 1} t_i(\theta'_1, \theta_{-1}) \right] \quad (20)$$

Let us derive an upper bound on the interim expected payoff of agent 1 from such a strategy:

$$\begin{aligned}
& E_{\theta_{-1}} [v_1(\widehat{a}(\theta_1, \theta_{-1}), \theta_1) + \widehat{t}_1(\theta_1, \theta_{-1})] \\
\leq & E_{\theta_{-1}} \left[v_1(\widehat{a}(\theta_1, \theta_{-1}), \theta_1) - \sum_{i \neq 1} \widehat{t}_i(\theta_1, \theta_{-1}) \right] \\
\leq & E_{\theta_{-1}} \left[v_1(\widehat{a}(\theta_1, \theta_{-1}), \theta_1) + \sum_{i \neq 1} v_i(\widehat{a}(\theta_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} t_i(\theta'_1, \theta_{-1}) \right] \\
= & E_{\theta_{-1}} \left[\sum_{i=1}^n v_i(\widehat{a}(\theta_1, \theta_{-1}), \theta_i) \right] + H_i + E_{\theta_{-1}} \left[\sum_{i=1}^n t_i(\theta'_1, \theta_{-1}) \right]
\end{aligned}$$

where the first inequality uses feasibility condition for transfers ($\widehat{t}_1(\theta_1, \theta_{-1}) \leq -\sum_{i \neq 1} \widehat{t}_i(\theta_1, \theta_{-1})$), the second inequality uses (20), and the equality uses definition of Groves in expectations transfers (4). Using budget balance ($\sum_{i=1}^n t_i(\theta'_1, \theta_{-1}) = 0$) and the fact that $a^*(\theta_1, \theta_{-1})$ is ex post efficient in state (θ_1, θ_{-1}) we have

$$\begin{aligned}
& E_{\theta_{-1}} [v_1(\widehat{a}(\theta_1, \theta_{-1}), \theta_1) + \widehat{t}_1(\theta_1, \theta_{-1})] \leq E_{\theta_{-1}} \left[\sum_{i=1}^n v_i(a^*(\theta_1, \theta_{-1}), \theta_i) \right] + H_i \\
= & E_{\theta_{-1}} [v_1(a^*(\theta_1, \theta_{-1}), \theta_1) + t_1(\theta_1, \theta_{-1})]
\end{aligned}$$

where the equality is by definition of Groves in expectations transfers (4). This gives a contradiction. Hence, the truthful equilibrium of this mechanism is EPRP.

(ii) Suppose a mechanism $\langle S, m \rangle$ has an EPRP equilibrium σ that implements the social choice function f , and satisfies the property described in the statement of the result.

First, by Lemma 2 note that $f(\theta) = (a(\theta), t(\theta))$ must be ex post efficient: $a(\theta) = a^*(\theta)$ and $\sum_{i=1}^n t_i(\theta) = 0$ for every θ . Next, since σ is EPRP, no agent should be able to improve upon his interim equilibrium payoff by deviating from his equilibrium strategy even if the resulting outcomes are renegotiated in the most profitable way. In order to write down these renegotiation-proofness constraints, we first need to find the highest payoff that a given agent can get from renegotiation.

Specifically, suppose the state is (θ_1, θ_{-1}) , and outcome (a, t) was reached by the mechanism when every agent $i \neq 1$ has played his equilibrium strategy, and agent 1 has played strategy $\sigma_1(\theta'_1)$ where $\theta'_1 \neq \theta_1$.^{45,46} By ex post efficiency a must be equal to $a^*(\theta'_1, \theta_{-1})$ with probability one, and t must be budget balanced with probability one. We will assume that $a = a^*(\theta'_1, \theta_{-1})$ and $\sum_{i=1}^n t_i = 0$ hold everywhere, and it is straightforward to adapt the argument to allow for inefficiencies to happen on zero probability events. By Definition 3(ii) every viable renegotiation mechanism must satisfy posterior individual rationality with respect to the reached outcome for

⁴⁵If $\sigma_1(\theta'_1)$ is a mixed strategy, then suppose agent 1 has played a particular pure strategy in its support.

⁴⁶If $\#\Theta_1 = 1$, then pick any other agent j such that $\#\Theta_j > 1$.

each agent $i \neq 1$. Thus every outcome (\hat{a}, \hat{t}) to which renegotiation may lead to must satisfy

$$v_i(\hat{a}, \theta_i) + \hat{t}_i \geq v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_i \quad (21)$$

Sum up (21) over $i \neq 1$ and rearrange

$$-\sum_{i \neq 1} \hat{t}_i \leq \sum_{i \neq 1} v_i(\hat{a}, \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} t_i \quad (22)$$

Let us derive an upper bound on the payoff of agent 1 in this case

$$\begin{aligned} v_1(\hat{a}, \theta_1) + \hat{t}_1 &\leq v_1(\hat{a}, \theta_1) - \sum_{i \neq 1} \hat{t}_i \\ &\leq \sum_{i=1}^n v_i(\hat{a}, \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_1 \\ &\leq \sum_{i=1}^n v_i(a^*(\theta_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_1 \end{aligned}$$

where the first inequality uses feasibility condition for transfers ($\hat{t}_1 \leq -\sum_{i \neq 1} \hat{t}_i$), the second inequality uses (22) and budget balance ($t_1 = -\sum_{i \neq 1} t_i$), and the last inequality is by definition of $a^*(\theta_1, \theta_{-1})$.

Next we show that for every θ_{-1} it is possible to achieve the above upper bound on the renegotiation payoff of agent 1 of type θ_1 who has mimicked type θ'_1 . The key condition which allows to show this is the assumption that the knowledge of the outcome is sufficient to compute θ_{-1} (assuming that every agent $i \neq 1$ has played his equilibrium strategy).

Specifically, suppose outcome (a, t) was reached. If the implied θ_{-1} is such that $a^*(\theta_1, \theta_{-1}) = a^*(\theta'_1, \theta_{-1})$, then renegotiation cannot improve the payoff of agent 1, and the upper bound is reached without renegotiation. Otherwise the following indirect renegotiation mechanism is offered to the agents. The agents simultaneously announce Y or N . If at least one agent have announced N , then outcome (a, t) is realized. If all agents have announced Y , then outcome $(a^*(\theta_1, \theta_{-1}), \hat{t})$ is realized, where $a^*(\theta_1, \theta_{-1})$ maximizes the total surplus in state (θ_1, θ_{-1}) , and transfers \hat{t} are defined as

$$\begin{aligned} \hat{t}_i &= v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) - v_i(a^*(\theta_1, \theta_{-1}), \theta_i) + t_i \quad \forall i \neq 1 \\ \hat{t}_1 &= \sum_{i \neq 1} v_i(a^*(\theta_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_1 \end{aligned}$$

Note that transfers \hat{t} are budget balanced since transfers t are budget balanced: $\sum_{i=1}^n \hat{t}_i = \sum_{i=1}^n t_i = 0$. By construction, every agent $i \neq 1$ of type θ_i is indifferent between $(a^*(\theta_1, \theta_{-1}), \hat{t})$ and (a, t) for any beliefs about the opponents' types, while agent 1 of type θ_1 strictly prefers $(a^*(\theta_1, \theta_{-1}), \hat{t})$ to (a, t) .

Hence, there exists an equilibrium of this renegotiation mechanism such that in state (θ_1, θ_{-1}) all agents choose Y , and thus the outcome is $(a^*(\theta_1, \theta_{-1}), \hat{t})$. Using a standard argument it is straightforward to construct an outcome equivalent incentive compatible direct mechanism ρ . By construction, this mechanism ρ is posterior individually rational with respect to (a, t) , and it satisfies the requirements of Definition 3(ii).

Take an expectation of the upper bound on the renegotiation payoff of agent 1 of type θ_1 over possible outcomes (a, t) and the opponents' type profiles θ_{-1}

$$\begin{aligned} & E_{\theta_{-1}} \left[E_{m \circ (\sigma_1(\theta'_1), \sigma_{-1}(\theta_{-1}))} \left[\sum_{i=1}^n v_i(a^*(\theta_1, \theta_{-1}), \theta_i) - \sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) + t_1 \right] \right] \\ = & E_{\theta_{-1}} \left[\sum_{i=1}^n v_i(a^*(\theta_1, \theta_{-1}), \theta_i) - \sum_{i=1}^n v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) \right] + E_{\theta_{-1}} [v_1(a^*(\theta'_1, \theta_{-1}), \theta'_1) + t(\theta'_1, \theta_{-1})] \end{aligned}$$

where the equality follows from the fact that σ implements f , and thus the resulting outcome is interim payoff equivalent to $(a^*(\theta'_1, \theta_{-1}), t(\theta'_1, \theta_{-1}))$. Hence, the renegotiation-proofness constraint to prevent deviations of this sort by agent 1 of type θ_1 is

$$\begin{aligned} & E_{\theta_{-1}} [v_1(a^*(\theta_1, \theta_{-1}), \theta_1) + t_1(\theta_1, \theta_{-1})] \tag{23} \\ \geq & E_{\theta_{-1}} \left[\sum_{i=1}^n v_i(a^*(\theta_1, \theta_{-1}), \theta_i) - \sum_{i=1}^n v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) \right] + E_{\theta_{-1}} [v_1(a^*(\theta'_1, \theta_{-1}), \theta'_1) + t(\theta'_1, \theta_{-1})] \end{aligned}$$

Rewrite (23):

$$E_{\theta_{-1}} [t_1(\theta_1, \theta_{-1})] - E_{\theta_{-1}} [t_1(\theta'_1, \theta_{-1})] \geq E_{\theta_{-1}} \left[\sum_{i \neq 1} v_i(a^*(\theta_1, \theta_{-1}), \theta_i) \right] - E_{\theta_{-1}} \left[\sum_{i \neq 1} v_i(a^*(\theta'_1, \theta_{-1}), \theta_i) \right] \tag{24}$$

Switching the roles of θ_1 and θ'_1 we find that (24) must hold with equality. Thus the transfer rule satisfies (4). The argument for other agents i such that $\#|\Theta_i| > 1$ is identical. ■

Proof of Proposition 4. Suppose $f(\theta) = (a(\theta), t(\theta))$ is ex post efficient in every $\theta \in \Theta$, i.e. $a(\theta) = a^*(\theta)$ and $t(\theta)$ is budget balanced. Consider a mechanism where each agent i reports his type $\theta_i \in \Theta_i$, and the outcome rule as a function of reports $\theta = (\theta_1, \dots, \theta_n)$ specifies $\alpha(\theta) = a^*(\theta)$ and budget balanced transfers $\tau(\theta)$ such that

$$E_{\theta_{-i}|\theta_i} [\tau_i(\theta_i, \theta_{-i})] = E_{\theta_{-i}|\theta_i} [t_i(\theta_i, \theta_{-i})] \quad \forall i, \theta_i \tag{25}$$

and

$$E_{\theta_{-i}|\theta_i} [\tau_i(\theta'_i, \theta_{-i})] < -L \quad \forall i, \theta'_i \neq \theta_i \tag{26}$$

where $L \in \mathbb{R}$ is chosen to be sufficiently large. The existence of such transfer rule τ follows from

the results in Kosenok and Severinov (2008).⁴⁷

There exists a truthful equilibrium since the expected penalty L is chosen to be large enough. To see that this equilibrium is EPRP first note that there will be no renegotiation on the equilibrium path since the outcome $(\alpha(\theta), \tau(\theta))$ is ex post efficient for every θ . If agent i misreports, then his expected payment is more than L , which is chosen to be large enough to make such a deviation unprofitable, even if the deviator consequently renegotiates the outcome and captures the implied increase in the total surplus. ■

References

- [1] Aghion, P., Dewatripont, M., and P. Rey (1994) “Renegotiation design with unverifiable information,” *Econometrica* 62, 257-282.
- [2] Amoros, P. (2004) “Nash implementation and uncertain renegotiation,” *Games and Economic Behavior* 49, 424-434.
- [3] Arrow, K. (1979) “The property rights doctrine and demand revelation under incomplete information,” in *Economics and Human Welfare*, ed. M. Boskin. Academic Press.
- [4] d’Aspremont, C. and L.-A. Gérard-Varet (1979) “Incentives and incomplete information,” *Journal of Public Economics* 11, 25-45.
- [5] Ausubel, L. and P. Cramton (2004) “Vickrey auctions with reserve pricing,” *Economic Theory* 23, 493–505.
- [6] Baliga, S. and T. Sjöström (2009) “Contracting with third parties,” *American Economic Journal: Microeconomics* 1, 75–100.
- [7] Beaudry, P. and M. Poitevin (1995) “Contract renegotiation: a simple framework and implications for organization theory,” *Canadian Journal of Economics* 28, 302-335.
- [8] Bergemann, D. and S. Morris (2005) “Robust mechanism design,” *Econometrica* 73, 1771-1813.
- [9] Che, Y. K. and D. B. Hausch (1999) “Cooperative investment and the value of contracting,” *American Economic Review* 89, 125-147.
- [10] Chung, T. (1992) “Incomplete contracts, specific investment, and risk sharing,” *Review of Economic Studies* 58, 1031-1042.
- [11] Chung, K. S., and J. Ely (2007) “Foundations of dominant strategy mechanisms,” *Review of Economic Studies* 74, 447-476.

⁴⁷The result follows from their Corollary 1 (p.134) to Theorem 1 (p.132-133) if we note that: (i) the uniqueness of the ex post efficient social alternative for every state implies that the agents’ equilibrium interim transfers (given by (25)) are uniquely determined; (ii) the agents’ interim transfers from misreporting (given by the left side of (26)) can be made arbitrarily large (p.149).

- [12] Crémer, J. and R. McLean (1988) “Full extraction of the surplus in bayesian and dominant strategy auctions,” *Econometrica* 56, 1247-1257.
- [13] Dewatripont, M. (1989) “Renegotiation and information revelation over time: the case of optimal labor contracts,” *Quarterly Journal of Economics* 104, 589-619.
- [14] Edlin A. and S. Reichelstein (1996) “Holdups, standard breach remedies, and optimal investment,” *American Economic Review* 86, 478-501.
- [15] Evans, R. (2009) “Mechanism design with renegotiation and costly messages,” forthcoming in *Econometrica*.
- [16] Forges, F. (1993) “Some thoughts on efficiency and information,” in *Frontiers of Game Theory*, ed. K. Binmore, A. Kirman, and P. Tani, MIT Press.
- [17] Forges, F. (1994) “Posterior efficiency,” *Games and Economic Behavior* 6, 238-261.
- [18] Fudenberg, D. and J. Tirole (1991) *Game Theory*, MIT Press.
- [19] Garratt, R. and T. Tröger (2006) “Speculation in standard auctions with resale,” *Econometrica* 74, 753–769.
- [20] Green, J. R. and J.-J. Laffont (1979) *Incentives in Public Decision Making*, North-Holland.
- [21] Green, J. and J.-J. Laffont (1994) “Non verifiability, costly renegotiation and efficiency,” *Annales D’Economie et De Statistique* 36, 81-95.
- [22] Groves, T. (1973) “Incentives in teams,” *Econometrica* 41, 617-631.
- [23] Holmström, B. (1979) “Groves’ scheme on restricted domains,” *Econometrica* 47, 1137-1144.
- [24] Klement, A. and Z. Neeman (2005) “Against compromise: a mechanism design approach,” *Journal of Law, Economics, and Organization* 21, 285-314.
- [25] Kosenok, G. and S. Severinov (2008) “Individually rational, budget-balanced mechanisms and allocation of surplus,” *Journal of Economic Theory* 140, 126-161.
- [26] Krasa, S. (1999) “Unimprovable allocations in economies with incomplete information,” *Journal of Economic Theory* 87, 144-168.
- [27] Mailath and Postlewaite (1990) “Asymmetric information bargaining problems with many agents,” *Review of Economic Studies* 57, 351-367.
- [28] Maskin, E. and J. Moore (1999) “Implementation and renegotiation,” *Review of Economic Studies* 66, 39-56.

- [29] Moore, J. (1992) "Implementation, contracts, and renegotiation in environments with complete information," in *Advances in Economic Theory, Sixth World Congress*, ed. by J-J. Laffont. Cambridge University Press, Cambridge.
- [30] Myerson, R. B. and M. A. Satterthwaite (1983) "Efficient mechanisms for bilateral trading," *Journal of Economic Theory* 29, 265-281.
- [31] Rogerson, W. (1992) "Contractual solutions to the hold-up problem," *Review of Economic Studies* 59, 777-793.
- [32] Rubinstein A., and A. Wolinsky (1992) "Renegotiation-proof implementation and time preferences," *American Economic Review* 82, 600-614.
- [33] 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.
- [34] Segal, I., and M. Whinston (2010) "Property rights," in *Handbook of Organizational Economics*, ed. by R. Gibbons and J. Roberts, forthcoming.
- [35] Sjöström, T. (1999) "Undominated Nash implementation with collusion and renegotiation," *Games and Economic Behavior* 26, 337-352.
- [36] Spier, K. E. (1994) "Pretrial bargaining and the design of fee shifting rules," *RAND Journal of Economics* 25, 197-214.
- [37] Walker, M. (1980) "On the nonexistence of a dominant strategy mechanism for making optimal public decisions," *Econometrica* 48, 1521-1540.
- [38] Watson, J. (2007) "Contract, mechanism design, and technological detail," *Econometrica* 75, 55-81.
- [39] Williams, S. R. (1999) "A characterization of efficient, Bayesian incentive compatible mechanisms," *Economic Theory* 14, 155-180.