Contract, renegotiation, and holdup: Results on the technology of trade and investment

KRISTY BUZARD
Department of Economics, University of California, San Diego

JOEL WATSON
Department of Economics, University of California, San Diego

This paper examines a class of contractual relationships with specific investment, a nondurable trading opportunity, and renegotiation. Trade actions are modeled as individual and trade-action-based option contracts ("nonforcing contracts") are explored. The paper introduces the distinction between *divided* and *unified* investment and trade actions, and it shows the key role this distinction plays in determining whether efficient investment and trade can be achieved. Under a nonforcing *dual-option contract*, the party without the trade action is made the residual claimant with regard to the investment action, which induces efficient investment in the divided case. The unified case is more problematic: here, efficiency is typically not attainable, but the dual-option contract is still optimal in a wide class of settings. More generally, the paper shows that, with ex post renegotiation, constraining parties to use "forcing contracts" implies a strict reduction in the set of implementable value functions.

KEYWORDS. Contract, renegotiation, holdup, forcing contracts, nonforcing contracts, specific investment, technology of trade, mechanism design.

JEL CLASSIFICATION. C70, D23.

The holdup problem arises in situations in which contracting parties can renegotiate their contract between the time they make unverifiable relation-specific investments and the time at which they can trade.[1] The severity of the holdup problem depends critically on the productive technology and on the timing of renegotiation opportunities. This paper contributes to the literature by examining how the nature of the "trade

action" in a contractual relationship influences the prospects for achieving an efficient outcome. We introduce a new distinction—whether the party who invests also is the one who consummates trade—that plays an important role in determining the outcome of the contractual relationship.

So that we can describe our modeling exercise more precisely, consider an example in which contracting parties Al and Zoe interact as follows. First Al and Zoe meet and write a contract that has an externally enforced element. Then one of them makes a private investment choice, which influences the *state* of the relationship. The state is commonly observed by the contracting parties, but is not verifiable to the external enforcer. Al and Zoe then send individual public messages to the external enforcer. After this, they have an opportunity to renegotiate their contract: this is called ex post renegotiation because it occurs after messages. Finally, the parties have a one-shot opportunity to trade and they also obtain external enforcement. Trade is verifiable to the external enforcer.

Because the investment is unverifiable, the investor cannot be directly rewarded for choosing the efficient investment level. Instead, incentives hinge on how the terms of trade can be made sensitive to the investment choice. Typically a conflict arises between the parties' joint interests prior to investment and their joint interests following investment and messages. In particular, investment incentives may be strengthened by specifying an inefficient trade action ex post in some off-equilibrium-path contingencies. But parties then would have the joint incentive to renegotiate and divide the surplus according to their bargaining power (holdup). Because parties rationally anticipate the renegotiation, the incentives to invest are distorted.

The description above obviously leaves the mechanics of trade and enforcement ambiguous. In reality, the parties have individual actions that determine whether and how trade is consummated. Let us suppose that Al selects the individual trade action, which we call a. This could be a choice of whether to deliver or to install an intermediate good, for example. We then have an *individual-action model*, whereby Al chooses a and the external enforcer compels a transfer t as a function of a and the messages that the parties sent earlier. In contrast, a *public-action model* (or *external-action model*) combines the trade action and the monetary transfer into a single public action (a, t) that is assumed to be taken by the external enforcer. With this modeling approach, the contract specifies how the public action is conditioned on the parties' messages.

Although the public-action model may typically be a bit unrealistic, it is simple and lends itself to elegant mechanism-design analysis (for example, as in Maskin and Moore 1999 and Segal and Whinston 2002). Alternatively, Watson (2007) demonstrates that analysis of the individual-action model can be straightforward as well. He also shows that the public-action model is equivalent to examining individual trade actions, but constraining attention to "forcing contracts" in which the external enforcer induces a particular trade action as a function of messages sent by the parties (so the trade action is constant in the state). Watson (2007) provides an example in which the restriction to forcing contracts has strictly negative efficiency consequences.

We deepen the examination of nonforcing contracts by investigating their efficacy in the context of different technologies of trade and investment. Specifically, we introduce the distinction between *divided* and *unified* investment and trade actions. In the

divided case, the investment and trade actions are chosen by different parties (Al takes the trade action and Zoe makes the investment). In the unified case, the investment and trade actions are selected by the same party (Al does both). We show that the prospects for inducing efficient investment and trade are very different in the divided and unified cases. In fact, the efficient outcome can always be achieved in the divided case (assuming investment has no immediate benefits), but typically cannot be achieved in the unified case.

Our analysis also highlights a simple contractual form that we call a *dual-option contract*. With the dual-option, Zoe sends a message that can be interpreted as a requested trade action or declaration of the state, and Al's subsequent trade action also serves as an option. We show that a dual-option contract is optimal in a large class of contractual relationships. For instance, it can be used to make Al's payoff constant in the state, gross of any investment costs, so that Zoe becomes the residual claimant with respect to the investment choice. This implies the efficiency result for the divided case. The dual option is also useful in the unified case, even though the efficient outcome typically cannot be achieved; specifically, we show that in a class of settings with a deterministic state, the dual-option contract is optimal.

Our analysis utilizes mechanism-design techniques. With both the individual-action and public-action modeling approaches, analysis of the contractual problem centers on calculating the set of implementable value functions from just after the state is realized (before messages are sent). Formally, an implementable value function is the state-contingent continuation value that results in equilibrium for a given contract. We provide simple tools to calculate the "punishment values" that determine the implementable sets for the class of relationships we analyze here. We use these tools to characterize optimal contracts and to find bounds on the set of implementable value functions.

In addition to the results on the divided and unified cases and dual-option contracts, we provide a general result on the comparison of forcing and nonforcing contracts, which shows that Watson's (2007) conclusions are robust over a large class of contractual relationships. In particular, in the important setting of ex post renegotiation described above, limiting attention to forcing contracts reduces the set of state-contingent continuation values. This does not mean that a more efficient outcome can always be achieved when actions are modeled as individual (because efficiency depends on what region of the implementable value set is relevant for giving appropriate investment incentives), but it underscores the importance of modeling trade actions as individual.

This is particularly salient for the setting of *cross/cooperative investment* (Che and Hausch 1999), where the investment by one party increases the benefit to the other party of subsequent trade. The literature has regarded cross-investment settings as especially prone to the holdup problem. Che and Hausch (1999) show that the optimal forcing contract is often "null" and leads to underinvestment. Our results establish that nonforcing contracts offer a significant improvement in efficiency, and our distinction between unified and divided investment and trade actions gives a basis for deeper analysis.

In the class of trade technologies that we study here, a single player (player 1, Al above) has the trade action. Examples of real settings with this property are contractual relationships in which the seller provides a service or good that does not require

the buyer's involvement (such as consulting, advertising, and some types of construction). In these settings, the seller has the trade action. Other settings with unilateral trade actions are ones in which the seller is the investor, production is inherent in the seller's investment, and trade is determined by whether the buyer installs or otherwise adopts the intermediate good; an example is specialized software. In these settings, the buyer has the trade action. We discuss the extension to multilateral trade actions in the Conclusion (Section 6).

The only assumption required for our first result, on making player 2 (Zoe) the residual claimant, is that investment does not confer a direct gain for some minimal trade action (an assumption satisfied by the most prominent models in the holdup literature). The key economic assumption behind our other results is that player 1's utility is supermodular as a function of the state and the trade action. That is, this player's marginal value of the trade action is monotone in the state. Our other assumptions are mainly weak technical conditions that guarantee well defined maxima, nontrivial settings, and the like. We argue that these conditions are likely to hold in a wide range of applications and that they are consistent with typical assumptions in the literature. Our result about the optimality of the dual-option contract in the unified case requires some additional assumptions on the technology of investment and trade.

The rest of the paper proceeds as follows. In the next section, we provide the details of the model. Section 2 presents an example that illustrates our main results. Section 3 contains our general results on optimal contracts and outcomes in the divided and unified cases. Readers interested in getting all of the basic ideas without the technical details can proceed from Section 3 straight to the Conclusion. Section 4 provides an overview of the basic tools for general analysis, which mostly restates material in Watson (2007). Section 5 contains our result on the difference in implementable sets based on variations regarding when renegotiation can occur and whether one restricts attention to forcing contracts. The Conclusion contains more discussion about the holdup problem and cross investment, as well as notes on the case of durable trading opportunities and multilateral trade actions. Most of the technical material and all of the proofs are contained in the Appendices.

1. The theoretical framework

We look at the same class of contracting problems and use the same notation as in Watson (2007), except that we add a bit of structure on the trade technology to focus our analysis. In particular, we examine the case in which a single player has a trade action. Throughout the paper, we use the convention of labeling the player with the trade action as player 1, and we call the other player 2. These two players are the parties engaged in a contractual relationship with a nondurable trading opportunity and external enforcement. Their relationship has the following payoff-relevant components, occurring in the order shown.

The *state of the relationship* θ. The state represents unverifiable events that are assumed to happen early in the relationship. The state may be determined by individual investment decisions and/or by random occurrences, depending on the

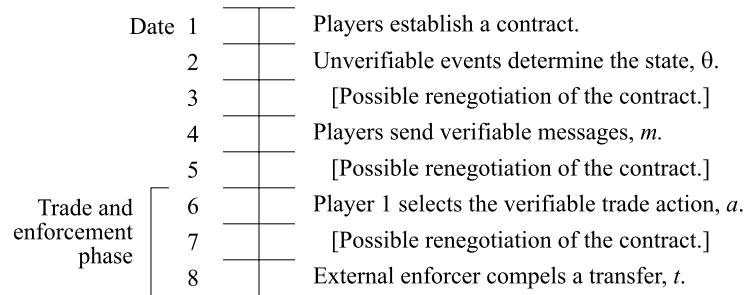|	Date	1		Players establish a contract.
		2		Unverifiable events determine the state, θ.
		3		[Possible renegotiation of the contract.]
		4		Players send verifiable messages, m.
		5		[Possible renegotiation of the contract.]
Trade and enforcement phase		6		Player 1 selects the verifiable trade action, a.
		7		[Possible renegotiation of the contract.]
		8		External enforcer compels a transfer, t.

FIGURE 1. Timeline of the contractual relationship.

setting. When the state is realized, it becomes commonly known by the players; however, it cannot be verified to the external enforcer. Let Θ denote the set of possible states.

The *trade action a*. This is an individual action chosen by player 1 that determines whether and how the relationship is consummated. The trade action is commonly observed by the players and is verifiable to the external enforcer. Let A be the set of feasible trade actions.

The *monetary transfers* $t = (t_1, t_2)$. Here t_i denotes the amount given to player i for $i = 1, 2$, where a negative value represents an amount taken from this player. Transfers are compelled by the external enforcer, who is not a strategic player, but, rather, who behaves as directed by the contract of players 1 and 2.[2] Assume $t_1 + t_2 \leq 0$.

We assume that the players' payoffs are additive in money and are thus defined by a function $u : A \times \Theta \rightarrow \mathbb{R}^2$. In state θ, with trade action a and transfer t, the payoff vector is $u(a, \theta) + t$. Define $U(a, \theta) \equiv u_1(a, \theta) + u_2(a, \theta)$, which is the joint value of the contractual relationship in state θ if trade action a is selected. We assume that, in each state θ, the joint value has a unique maximizer $a^*(\theta)$. We let $\gamma(\theta)$ denote the maximal joint payoff in state θ, so we have

$$\gamma(\theta) \equiv U(a^*(\theta), \theta) = \max_{a \in A} U(a, \theta). \tag{1}$$

In addition to the payoff-relevant components of their relationship, we assume that the players can communicate with the external enforcer using public, verifiable messages. Let $m = (m_1, m_2)$ denote the profile of messages that the players send and let M_1 and M_2 be the sets of feasible messages. The sets M_1 and M_2 are endogenous in the sense that they are specified by the players in their contract.

Figure 1 shows the timeline of the contractual relationship. At even-numbered dates through date 6, the players make joint observations and they make individual decisions—jointly observing the state at date 2, sending verifiable messages at date 4,

---

[2]That the external enforcer's role is limited to compelling transfers is consistent with what courts do in practice.

and selecting the trade actions at date 6. At date 8, the external enforcer compels transfers. At odd-numbered dates, the players make joint contracting decisions—establishing a contract at date 1 and possibly renegotiating it later.

The contract has an externally enforced component consisting of (i) feasible message spaces M_1 and M_2 and (ii) a *transfer function* $y : M \times A \to \mathbb{R}^2$ that specifies the transfer t as a function of the verifiable items m and a. That is, having seen m and a, the external enforcer compels transfer $t = y(m, a)$. The contract also has a self-enforced component, which specifies how the players coordinate their behavior for the times at which they take individual actions. Renegotiation of the contract amounts to replacing the original transfer function y with some new function y', in which case y' is the one submitted to the external enforcer at date 8.

We initially assume—and maintain throughout Sections 2 and 3—that the players can freely renegotiate at dates 3, 5, and 7. Renegotiation at date 5 is called *ex post renegotiation*. At date 3 it is called *interim renegotiation*.[3]

The players' individual actions at dates 2, 4, and 6 are assumed to be consistent with sequential rationality; that is, each player maximizes his expected payoff, conditional on what occurred earlier and on what the other player does, and anticipating rational behavior in the future. The joint decisions (initial contracting and renegotiation at odd-numbered periods) are assumed to be consistent with a cooperative bargaining solution in which the players divide surplus according to fixed bargaining weights π_1 and π_2 for players 1 and 2, respectively. The bargaining weights are nonnegative, sum to 1, and are written $\pi = (\pi_1, \pi_2)$. The negotiation surplus is the difference between $\gamma(\theta)$ and the joint value that would result if the players fail to reach an agreement, where the disagreement point is given by an equilibrium in the continuation in which the externally enforced component of the contract has not been altered.[4]

The effect of the renegotiation opportunity at date 7 is to constrain transfers to be "balanced," that is, satisfying

$$t \in \mathbb{R}_0^2 \equiv \{t' \in \mathbb{R}^2 \mid t_1' + t_2' = 0\}.$$

Thus, we simply assume that transfers are balanced and then otherwise ignore date 7. Also, as we explain later, the opportunity for ex post renegotiation implies that there is never any renegotiation surplus at date 3, so we can ignore interaction at date 3.

Much of our analysis does not depend on the details of date 2 interaction, but some of our key results concern the relation between the investment and the trade technologies, and for these we need to formally distinguish between different investment technologies. We assume that a single player makes an investment choice at date 2. This gives us two cases to consider:

---

[3]In Sections 4 and 5, we provide some analysis for the setting in which renegotiation is possible at date 3 but not at date 5.

[4]The generalized Nash bargaining solution has this representation. The rationality conditions identify a *contractual equilibrium*; see Watson (2006) for notes on the relation between "cooperative" and "noncooperative" approaches to modeling negotiation. The players obtain the joint value $\gamma(\theta)$ because, at the time of renegotiation, they know the state θ and can select a contract that forces the action $a^*(\theta)$, as described in the next subsection.

- *Unified case.* Player 1 has both the date 2 investment action and the date 6 trade action.

- *Divided case.* Player 2 has the date 2 investment action, whereas player 1 has the date 6 trade action.

We assume that the investment influences the state. In the *deterministic* subcase, one of the players directly selects θ at date 2. More generally, the state may also depend on the outcome of a random variable.

Public-action modeling and forcing contracts

Because the trade action a is assumed to be taken by player 1, we have specified here an individual-action model. A public-action model, in contrast, abstracts by treating the trade action a as something that the external enforcer directly selects. Watson (2007) shows that specifying a public-action model is equivalent to examining the individual-action model but limiting attention to a particular class of contracts called *forcing contracts*, which, for any given message profile, prescribe that player 1 selects a particular trade action.

More precisely, a forcing contract specifies a large transfer from player 1 to player 2 in the event that player 1 does not take his contractually prescribed action. This transfer is sufficiently large to give player 1 the incentive to select the prescribed action in every state. Thus, the induced trade action is constant in the state, conditional on the messages sent earlier.

For example, holding the message profile fixed, the transfer function $\hat{y}$ defined as follows forces player 1 to select action $\hat{a}$ and imposes the transfer $\hat{t}$ (as though the external enforcer chose these in a public-action model):

Let L be such that $L > \sup_{a,\theta} u_1(a, \theta) - \inf_{a,\theta} u_1(a, \theta)$. Then define $\hat{y}(\hat{a}) \equiv \hat{t}$ and, for every $a \neq \hat{a}$, set $\hat{y}(a) \equiv \hat{t} + (-L, L)$.

We use the term *forcing* for any transfer function that, given the message profile, induces player 1 to select the same trade action over all of the states.[5] We use the term *nonforcing* for transfer functions that induce player 1 to select different actions in at least two different states.

Continuation value functions

A *(state-contingent) value function* is a function from Θ to $\mathbb{R}^2$ that gives the players' expected payoff vector from the start of a given date, as a function of the state. Such a value function represents the continuation values for a given outstanding contract and equilibrium behavior. We adopt the convention of not including any sunk investment

---

[5]One could add a public randomization device to the model for the purpose of achieving randomization over trade actions using forcing contracts. Allowing such randomization does not expand the set of implementable value functions here.

costs from date 2 in the function u or in the representation of continuation values from later dates.

The continuation values from the start of date 3 are important to calculate, because they determine the players' incentives to invest at date 2. Thus, our chief objective is to characterize the set of *implementable value functions* from the start of date 3. A value function v is implementable if there is a contract that, if formed at date 1, would lead to continuation value $v(\theta)$ in state θ from the start of date 3 for every $\theta \in \Theta$.

Related literature

Much of the recent contract-theory literature focuses on public-action mechanism-design models. For instance, Che and Hausch (1999), Hart and Moore (1999), Maskin and Moore (1999), Segal (1999), and Segal and Whinston (2002) have basically the same setup as we do except that their models treat trade actions as public (collapsing together the trade action and the enforcement phase), so they focus on forcing contracts.[6] In some related papers, the verbal description of the contracting environment identifies individuals who take the trade actions, but the actions are effectively modeled as public due to an implicit restriction to forcing contracts. In some cases, such as with the contribution of Edlin and Reichelstein (1996), simple forcing contracts (or breach remedies) are sufficient to achieve an efficient outcome and so the restriction does not have efficiency consequences.[7]

Examples of individual-action models in the literature, among others, are the articles of Hart and Moore (1988), MacLeod and Malcomson (1993), and Nöldeke and Schmidt (1995). Also relevant is the work of Myerson (1982, 1991), whose mechanism-design analysis nicely distinguishes between inalienable individual and public actions (he uses the term "collective choice problem" to describe public-action models).

Most closely related to our work is that of Evans (2006, 2008), who emphasizes how efficient outcomes can be achieved by conditioning external enforcement on costly individual actions. Evans (2006) examines general mechanism-design problems; Evans (2008), which we discuss more in the Conclusion, examines contracting problems with specific investment and durable trading opportunities. Related as well is the work of Lyon and Rasmusen (2004), which shares the theme of Watson (2007), and the recent work of Boeckem and Schiller (2008) and Ellman (2006).[8]

---

[6]Aghion et al. (1994) is another example. The more recent entries by Roider (2004) and Guriev (2003) have the same basic public-action structure. Demski and Sappington (1991), Nöldeke and Schmidt (1998), and Edlin and Hermalin (2000) examine models with sequential investments in a tradeable asset; in these models, as in Maskin and Tirole (1999), transferring the asset is essentially a public action.

[7]Stremitzer (forthcoming) elaborates on Edlin and Reichelstein (1996) by examining the informational requirements of standard breach remedies (specifically, partially verifiable investments).

[8]Also related are some studies of delegation in principal–agent settings with asymmetric information, where implementable outcomes depend on whether it is the principal or the agent who has the productive action. As Beaudry and Poitevin (1995) show, ex post renegotiation imposes less of a constraint in the case of "indirect revelation" (where the agent has the productive action). Thus, if it is possible to transfer "ownership" of the productive action to the agent, the threat of ex post renegotiation provides one reason for doing this.

In classifying the related literature, another major distinction to make is between models with cross investment and models with "own investment." In the latter case, investment enhances the investing party's benefit of trade. We discuss this distinction in more detail in the next two sections. Since the holdup problem is more problematic in the case of cross investment, and there the distinction between forcing and nonforcing contracts (public- versus individual-action modeling) is critical, we concentrate on settings with significant cross investment.

2. Example

In this section, we provide a simple example of specific investment and holdup to illustrate our results. We continue to call player 1 Al and player 2 Zoe. One of the players is the *investor* at date 2 (Al in the unified case, Zoe in the divided case). The investor selects $\theta \in [0, 9]$ at immediate cost $c(\theta) = 3\theta$. That is, the investor takes an action that determines the state. At date 6, Al selects a trade action $a \in [0, 9]$, which we interpret as a quantity of an intermediate good that he delivers to Zoe. Payoffs are given by

$$u_1(a, \theta) = 4\sqrt{a\theta} - 4a \quad \text{and} \quad u_2(a, \theta) = 4\sqrt{a\theta}.$$

As we assume throughout, the sunk investment cost is not included in these functions and in the value functions computed below. Assume that the players have equal bargaining weights.

The joint value of the relationship in state θ is $U(a, \theta) = 8\sqrt{a\theta} - 4a$, which is maximized at $a^*(\theta) = \theta$. Therefore, the maximal joint value in state θ is $\gamma(\theta) = U(a^*(\theta), \theta) = 4\theta$. Regardless of who makes the investment, we see that the efficient level of investment maximizes $4\theta - 3\theta$. Thus the optimal investment level is $\theta^* = 9$.

Note that, for any fixed trade action, Al's and Zoe's payoffs increase equally in θ. Thus, regardless of who is the investor, this example exhibits elements of both *cross investment* and *own investment*. Own investment refers to the investment boosting the investor's gain from trade, whereas cross investment refers to increasing the gain of the other party.[9] The cross-investment element is particularly problematic, as the literature shows, because there are contingencies (typically out of equilibrium) in which the non-investing party can extract surplus from the investor by threatening to hold up trade. This can distort the investor's incentives and lead to inefficient investment.[10]

In fact, Che and Hausch (1999) conclude that with significant cross investment, the "null contract"—forcing no trade, regardless of the messages—is best. These authors formulate a public-action model, which limits attention to forcing contracts. Indeed, it is straightforward to show that the null contract is the best forcing contract for our

---

[9]Che and Hausch (1999) use the term "cooperative investment" for cross investment.

[10]The literature demonstrates that forcing contracts can usually prevent the holdup problem in the own-investment case, where the investing party obtains a large share of the benefit created by the investment. See, for example, Chung (1991), Rogerson (1992), Aghion et al. (1994), Nöldeke and Schmidt (1995), and Edlin and Reichelstein (1996). An exception is the "complexity/ambivalence" setting studied by Segal (1999), Hart and Moore (1999), and Reiche (2006).

FIGURE 2. Value function and investment cost.

example, regardless of which player is the investor.[11] Unfortunately, the null contract leads to an inefficient level of investment. To see this, note that the players always renegotiate to take the ex post efficient trade action in each state. In our example this implies that in state θ, the renegotiation surplus equals the joint value 4θ. Since the investor receives half of the surplus (recall that the bargaining weights are $1/2$), the investor's value from date 3 is 2θ. This value and the investment cost are illustrated in part (a) of Figure 2. At date 2, the investor therefore chooses $\theta \in [0, 9]$ to maximize $2\theta - 3\theta$, and so the inefficiently low level $\theta = 0$ is chosen.

Implementation with a dual-option contract

We next demonstrate that by using nonforcing contracts, a more efficient outcome can be achieved (Watson's 2007 point) and that the unified and divided cases behave quite differently. Our analysis features a particular nonforcing contract that we call a *dual option*, which turns out to be an optimal contractual form in a wide range of settings. With the dual-option contract, Zoe sends a message at date 4 and then Al is forced to choose between two different trade actions at date 6, one of which depends on Zoe's message. Thus, Zoe's option is message-based, whereas Al's option later is his choice of the trade action.[12]

---

[11]The tools developed in Sections 4 and 5 can be used to show that if v is implemented by a forcing contract, then, for any θ, θ' with $\theta > \theta'$, $v_1(\theta) - v_1(\theta')$ and $v_2(\theta) - v_2(\theta')$ are bounded above by $2(\theta - \theta')$. The null contract achieves this bound. See Appendix B for more details.

[12]The literature emphasizes the importance of option contracts for aligning incentives. By laying out the details of the trade technology, we are able to differentiate between message-based and trade-action-based option components. Our dual-option contract is a novel addition to the theoretical literature because it has both of these components. By comparison, papers in the related literature examine either (i) one-sided or two-sided message-based contracts, or (ii) action-based options without messages. Che and Hausch (1999) were first to demonstrate the advantages of a sequential, two-sided message-based contract in a setting without renegotiation (where trade-action-based options offer no special advantage); Segal and Whinston (2002) examine general two-sided message contracts with renegotiation; Nöldeke and Schmidt (1995) look at action-based contracts without messages; Hart and Moore (1988) and MacLeod and Malcomson (1993) also have action-based contracts but with partial verifiability (which we further discuss in the Conclusion). Demski and Sappington (1991) briefly discuss a contract with both message- and trade-based option components, but they do not formally study this form.

Consider the following dual-option contract that is parameterized by a number $\alpha \in [0, 9]$. At date 4, Zoe must send a message $\hat{a} \geq 0$, which we interpret as a requested quantity for Al to deliver. Then Al is forced to choose between $a = \hat{a}$ and $a = \alpha$ at date 6, by having him pay Zoe a large amount if he selects any other trade action. That is, if the contract remains in place by date 8 and Al selects some $a \notin \{\alpha, \hat{a}\}$, then the enforcer compels a large transfer (say, 100) from Al to Zoe. If Al selects $a = \alpha$, then there is no transfer, whereas if Al selects $a = \hat{a}$, then the enforcer compels a transfer of $u_1(\alpha, \hat{a}) - u_1(\hat{a}, \hat{a}) = 4\sqrt{\alpha\hat{a}} - 4\alpha$ from Zoe to Al.

Let us construct a value function that this dual-option contract implements. Note that, given α and absent renegotiation at date 5, Al weakly prefers to choose $a = \alpha$ at date 6 if and only if

$$u_1(\alpha, \theta) \geq u_1(\hat{a}, \theta) + [u_1(\alpha, \hat{a}) - u_1(\hat{a}, \hat{a})].$$

Plugging in the values, this simplifies to

$$\sqrt{\theta}(\sqrt{\alpha} - \sqrt{\hat{a}}) \geq \sqrt{\hat{a}}(\sqrt{\alpha} - \sqrt{\hat{a}}).$$

Note that if Zoe requests $\hat{a} = \theta$ (the ex post optimal trade level for the realized state θ), then Al is indifferent between trade actions θ and α at date 6. Let us assume that Al selects $a = \theta$ in this contingency. Because $a = \theta$ is the efficient trade action in state θ, the players would not renegotiate at date 5, so the payoffs from date 4 are $u_1(\alpha, \theta)$ for Al and $\gamma(\theta) - u_1(\alpha, \theta)$ for Zoe.

Observe that Zoe can do no better by deviating from $\hat{a} = \theta$ at date 4. This is because Al can ensure himself a payoff of at least $u_1(\alpha, \theta)$ from date 6 by choosing $a = \alpha$. Renegotiation earlier can only add to Al's pocket, so his continuation payoff from date 4 is bounded below by $u_1(\alpha, \theta)$. Since the opportunity for renegotiation implies that Zoe's continuation value is $\gamma(\theta)$ minus Al's continuation value and since Zoe can hold Al down to $u_1(\alpha, \theta)$ by choosing $\hat{a} = \theta$, it is optimal for Zoe to send this message. Thus, in state θ, Zoe requests $\hat{a} = \theta$, there is no renegotiation, and Al delivers θ units. The contract implements the value function given by

$$v(\theta) = (u_1(\alpha, \theta), \gamma(\theta) - u_1(\alpha, \theta)) = (4\sqrt{\alpha\theta} - 4\alpha, 4\theta - 4\sqrt{\alpha\theta} + 4\alpha)$$

for all $\theta \in [0, 9]$. Incentives are the same if we include a constant transfer τ from Zoe to Al, in which case the implemented value function is given by

$$v(\theta) = (4\sqrt{\alpha\theta} - 4\alpha + \tau, 4\theta - 4\sqrt{\alpha\theta} + 4\alpha - \tau).$$

We next investigate the implications of this dual-option contract for the divided and unified cases.

Dual-option contract in the divided case

The implication of utilizing nonforcing contracts is dramatic in the divided case, where Zoe makes the investment. In fact, it is easy to see that the dual-option contract can be

used to make Zoe the residual claimant with respect to the post-investment joint value. In particular, set $\alpha = 0$ so that Al always has the option to deliver nothing and get no transfer. Also set $\tau = 0$. This dual-option contract implements

$$v(\theta) = (0, \gamma(\theta)) = (0, 4\theta)$$

for all θ. Zoe obtains exactly the joint value of her investment, so at date 2 she maximizes $\gamma(\theta) - c(\theta)$. Her optimal choice is $\theta^* = 9$ and thus efficient investment and trade are realized.[13]

Dual-option contract in the unified case

Next consider the unified case of the example, in which Al makes the investment. We see that this case is more problematic, but that positive investment can still be induced. Consider the dual-option contract with $\alpha = 9$ and $\tau = 27$. This contract implements

$$v_1(\theta) = u_1(9, \theta) + 27 = 12\sqrt{\theta} - 9$$

for all θ. The value function is shown in part (b) of Figure 2. At date 2, Al chooses θ to maximize $v_1(\theta) - c(\theta) = 12\sqrt{\theta} - 3\theta - 9$, and so his optimal investment choice is $\theta = 4$. Thus, the dual-option contract performs better than the null contract, but does not induce efficient investment.

Consider next a version of the dual-option contract with an additional parameter $\beta \in [0, 9]$, where (i) if Zoe's message satisfies $\hat{a} \geq \beta$, then Al is forced to deliver α units with transfer τ, and (ii) if $\hat{a} < \beta$, then the specifications are as described above. One can verify that Zoe optimally selects $\hat{a} = \theta$ as before, but now the parties renegotiate whenever $\theta \geq \beta$. Setting $\alpha = 9$ and $\tau = 27$ once again, this dual-option contract implements the value function given by

$$v_1^\beta(\theta) = \begin{cases} 12\sqrt{\theta} - 9 & \text{if } \theta < \beta \\ 2\theta + 9 & \text{if } \theta \geq \beta, \end{cases}$$

with $v_2^\beta(\theta) = \gamma(\theta) - v_1^\beta(\theta)$. The value function is pictured in part (c) of Figure 2.

The results we provide in the next section establish that this dual-option contract is, in fact, optimal for some β. The advantage of the cutoff β is that v_1 jumps up at this point (owing to the positive renegotiation surplus), giving player 1 an extra incentive to invest at level β. The highest investment level supported by this type of contract is $\theta = 6$, which is achieved by setting $\beta = 6$. An implication is that the efficient investment level cannot be achieved.

---

[13]One perspective on the related literature comes from considering how contractual elements act to effectively shift relative bargaining power away from the intrinsic level inherent in the exogenous bargaining protocol. The dual option here produces the same outcome that arises under a contract that forces no trade, if one were able to give Zoe all of the bargaining power in renegotiation. Thus, the action-based part of the dual option sets a reservation payoff level for Al that allows Zoe to extract the full surplus by sending the appropriate message. A similar type of intuition is at play in Aghion et al.'s (1994) use of a financial bond to effectively shift bargaining weights, and it is also present in MacLeod and Malcomson's (1993) use of an outside option in the context of ongoing renegotiation with a durable trading opportunity.

In summary, the example shows that by using nonforcing contracts, the parties can achieve a more efficient outcome than is possible with forcing contracts. Furthermore, the efficiency gain depends on the relation between the technology of trade and the technology of investment. In the divided case, the dual-option contract induces efficient investment and trade actions. In the unified case, the efficient outcome cannot be attained, but a nonforcing contract still is preferred.

3. INVESTMENT INCENTIVES AND RESIDUAL CLAIMANCY

In this section we provide general versions of the results shown in the example. We divide the analysis into two subsections, one dealing with the objective of giving investment incentives to player 2 (which is needed in the divided case) and one with the objective of giving such incentives to player 1 (for the unified case).

Before proceeding, it is useful to define some additional notation. Regardless of which player has the investment choice at date 2, let the investment be denoted $x \geq 0$. We normalize the investment variable so that the immediate cost of investment is exactly x for the investor. The state θ is then drawn from a distribution $G(x)$ that depends on the investment choice.[14] Recalling that $\gamma(\theta) = U(a^*(\theta), \theta)$ is the maximum joint value in state θ, we see that the efficient level of investment x^* solves

$$\max_{x \geq 0} \int \gamma(\theta) \, dG(x) - x.$$

In the *deterministic* case in which there is no random element, we suppose that $\theta \equiv x$ and so the investor selects θ directly; then we write θ^* as the efficient investment choice, which maximizes $\gamma(\theta) - \theta$.

Letting i denote the investing party, we want to implement a value function v so that $v_i(\theta)$ is increasing in θ to some particular extent. In this way, player i is rewarded for investing. Ideally, it would be possible to implement a value function that satisfies $v_i(\theta) = \gamma(\theta) - k$ for all $\theta \in \Theta$ and some fixed k, because this makes player i the residual claimant with respect to his/her investment decision. Player i's payoff from date 2 is then

$$\int v_i(\theta) \, dG(x) - x = \int \gamma(\theta) \, dG(x) - x - k,$$

and so player i optimally selects x^*, leading to efficient investment and trade. The players select k to divide the joint value at date 1.

It may also be possible to achieve an efficient outcome without having $v_i = \gamma - k$, but this is not always the case. More generally, in some settings we can characterize the optimal contract and the best (though inefficient) investment that can be induced. Our proofs utilize a dual-option contract and thus further demonstrate the utility of this simple contractual form.

---

[14]It is natural to assume that G is increasing in x in the sense of first-order stochastic dominance and that $U(a^*(\theta), \theta)$ is increasing in θ, so that higher investments increase the expected gains from trade, but these assumptions are not needed for Proposition 1 below.

Investment incentives for player 2

We start by showing that, assuming that investment conveys no instantaneous (direct) benefits, a dual-option contract can be used to make player 1's value from date 3 constant in the state. Player 2 then becomes the residual claimant. Thus, in the divided case in which player 2 is the investor, player 2 can be given the incentive to invest efficiently regardless of the distribution of the investment gains.

ASSUMPTION 1. *There exists a trade action $a^0 \in A$ such that $u_1(a^0, \theta) = u_2(a^0, \theta) = 0$ for every θ.*

This assumption is solely on the technology of trade. Think of a^0 as the "no trade" choice. The no-trade payoffs could be normalized to any level; we set them to zero here for simplicity. Note that the example satisfies Assumption 1 with $a^0 = 0$.

THEOREM 1. *Consider any contractual relationship that satisfies Assumption 1. Let k be any real number and define the value function v by $v_1(\theta) = k$ and $v_2(\theta) = \gamma(\theta) - k$ for all $\theta \in \Theta$. Then v is implementable.*

Appendix A contains the proof of this theorem, which is constructive and runs along the lines of the demonstration for the example. In particular, we show how to implement these value functions using a dual-option contract in which player 2 is required to declare a state $\hat{\theta}$ at date 4 and then player 1 is forced to tender either trade action $a^*(\hat{\theta})$ or trade action a^0.

Considering the implication of making the investor the residual claimant, Theorem 1 leads immediately to the following economic result.

PROPOSITION 1. *Under Assumption 1 and in the divided case in which player 2 is the investor and player 1 has the trade action, optimal contracting induces efficient investment and trade (the first-best outcome).*

Note that this result makes no restrictions on which party stands to gain from the investment. That is, the result holds for settings of cross investment, own investment, and any combination of the two. The key is simply that the investment action and the trade action are taken by different parties.

Investment incentives for player 1

Making player 1 the residual claimant with respect to the investment decision is considerably more difficult than is making player 2 the residual claimant. In fact, in the case of unified investment and trade actions, typically the efficient outcome cannot be induced. Identifying an optimal contract is also a challenge, but is possible with some additional assumptions and structure on the model. We present results along these lines below. The proofs require a full analysis of the conditions for implementation, which in turn requires the tools and results developed in Sections 4 and 5. Details of the analysis are provided in Appendix B, where we also present some general results and notes.

We begin with some intuition regarding the conditions for implementation. Suppose that γ is increasing. Our objective is to implement a value function v_1 that rises with γ, so that player 1's return on investment closely matches the joint return. The technical conditions for implementation imply upper bounds on the difference $v_1(\theta) - v_1(\theta')$ for $\theta > \theta'$. Observe that there are multiple such conditions involving each state. For example, for three states θ, θ', and θ'' with $\theta > \theta' > \theta''$, there are three conditions: $v_1(\theta) - v_1(\theta') \leq \rho$, $v_1(\theta') - v_1(\theta'') \leq \rho'$, and $v_1(\theta) - v_1(\theta'') \leq \rho''$ for some numbers ρ, ρ', and ρ''. We can call the first and second conditions *local*, or *inside*, conditions, whereas the last one is an *outside* condition. Note that by summing the inside conditions, we obtain a second bound on the difference $v_1(\theta) - v_1(\theta'')$; this bound is $\rho + \rho'$.

It turns out that, for a wide class of trade technologies, the outside condition is tighter than is the sum of the inside conditions; that is, $\rho'' < \rho + \rho'$. This means that implementability cannot be characterized by the local conditions alone, and some of these must hold with slack to ensure that the outside conditions are satisfied. As a result, it is not possible to implement a value function v_1 that rises smoothly and steeply. This is not such a big problem in the divided case, where we want v_1 to be constant, but recall that in the unified case we want v_1 to rise with γ. We find that the best way to give player 1 the incentive to invest is to implement a value function with some discrete jumps. We establish conditions under which a dual-option contract optimally performs in this way, as shown in the example.

We make several assumptions to structure the analysis. The first gives a set of mild technical restrictions that hold in most applications. We maintain this assumption throughout the rest of the paper.

ASSUMPTION 2. *(a) The sets A and Θ are compact subsets of $\mathbb{R}$ and contain at least two elements, and $u_1(\cdot, \theta)$ and $u_2(\cdot, \theta)$ are continuous functions of a for every $\theta \in \Theta$. Define $\underline{a} \equiv \min A$, and $\overline{a} \equiv \max A$, $\underline{\theta} \equiv \min \Theta$, and $\overline{\theta} \equiv \max \Theta$. (b) $U(\cdot, \theta)$ is strictly quasiconcave for every $\theta \in \Theta$. (c) Player 1's bargaining weight is positive: $\pi_1 > 0$.*

The next assumption is the main economic restriction that we impose hereinafter: that player 1's payoff is supermodular in the state and trade action.

ASSUMPTION 3. *The function u_1 is supermodular, meaning that $u_1(a, \theta) - u_1(a', \theta) \geq u_1(a, \theta') - u_1(a', \theta')$ whenever $a \geq a'$ and $\theta \geq \theta'$.*

With this assumption, player 1's marginal value of increasing his trade action rises weakly with the state. In other words, higher trade actions are weakly more attractive to him as the state increases. An implication is that, for any transfers specified as a function of the trade action, player 1's preferences satisfy the single-crossing property and he weakly prefers higher actions in higher states. This monotone structure helps us to characterize incentives at date 6.

Many interesting applications studied in the literature satisfy these assumptions. For instance, consider a buyer/seller relationship in which a is the number of units of an intermediate good to be transferred from the seller to the buyer. The buyer's benefit

of obtaining a units in state θ is $B(a, \theta)$. The seller's cost of production and delivery is $d(a, \theta)$, and we let $C(a, \theta) = -d(a, \theta)$. Suppose, as one typically does, that B is increasing and concave in a and that d is increasing and convex in a. If a is the buyer's action (he selects how many units to install, for example), then the buyer is player 1 and so we have $u_1 \equiv B$ and $u_2 \equiv C$. If the seller chooses a (she decides how many units to deliver, say), then the seller is player 1 and so we have $u_1 \equiv C$ and $u_2 \equiv B$. In either case, Assumption 2 is satisfied. Assumption 3 adds the weak supermodularity requirement on the payoff of the player who selects a. Our example from the previous section satisfies Assumptions 2 and 3.

Note that, in a given application, if u_1 is submodular, then one can redefine the trade action to be $-a$ and then Assumption 3 is satisfied. Also note that Assumption 3 is trivially satisfied in the case of *pure cross investment* in which u_1 does not depend on θ.

Our final assumption, which is needed only for the results in this subsection, pertains to the relative supermodularity and investment returns for u_1 and u_2.

ASSUMPTION 4. *(a) The expression $\pi_2 u_1 - \pi_1 u_2$ is supermodular ($\pi_2 u_1$ is relatively more supermodular than $\pi_1 u_2$). (b) For all θ, θ' with $\theta > \theta'$, $\pi_1[u_2(\overline{a}, \theta) - u_2(\overline{a}, \theta')] \geq \pi_2[u_1(\overline{a}, \theta) - u_1(\overline{a}, \theta')]$.*

Assumption 4 is clearly restrictive, limiting the class of trade technologies that we evaluate here, but it facilitates the identification of an optimal contract. Part (a) of the assumption contributes to a tight characterization of optimal punishments in the general mechanism-design exercise. A sufficient condition for Assumption 4(a) is that u_2 is submodular. Appendix B gives an alternative assumption—on the joint value at extreme trade actions—that can be used in place of Assumption 4(a). Assumption 4(b) requires that the cross-investment component is weakly larger than the own-investment component at the highest trade action. It is easy to check that our example satisfies Assumption 4.

These assumptions give us the following general result about implementation using the dual-option contract introduced in the example.

THEOREM 2. *Consider any contractual relationship that satisfies Assumptions 2–4. For any number β, define value function v^β by*

$$v_1^\beta(\theta) \equiv \begin{cases} u_1(\overline{a}, \theta) & \text{for } \theta < \beta \\ u_1(\overline{a}, \theta) + \pi_1 R(\overline{a}, \theta) & \text{for } \theta \geq \beta \end{cases}$$

and $v_2^\beta \equiv \gamma - v_1^\beta$ for all $\theta \in \Theta$. Then v^β is implemented by a dual-option contract in which (i) at date 4, player 2 sends a report $\hat{\theta}$ of the state; and (ii) if the report is at least β, then player 1 is forced to select $a = \overline{a}$ at date 6, and otherwise player 1 is forced to choose between $a^*(\hat{\theta})$ and $\overline{a}$.

The proof of Theorem 2 is provided in Appendix B, along with the proofs of Propositions 2 and 3 below. Note that we have not used Assumption 1 here.

We next show that value function v^β achieves the best possible investment incentives in the deterministic case where player 1's investment choice is to directly select θ. Thus, in this setting of unified investment and trade actions, the dual-option contract is optimal. Recall that we normalize so that the cost of investment is θ, and thus player 1 selects θ to maximize $v_1(\theta) - \theta$. The efficient choice θ^* maximizes $\gamma(\theta) - \theta$.

Let us say that *investment θ' is supported* by a contract if there is an implementable value function v' such that θ' solves $\max_{\theta \in \Theta} v_1'(\theta) - \theta$. Call θ^B the *best achievable* investment level if it maximizes $\gamma(\theta') - \theta'$ among all supportable θ'.

PROPOSITION 2. *Under Assumptions 2–4, and in the deterministic and unified case in which player 1 has both the investment and trade actions, the best achievable investment level θ^B is supported by value function v^{θ^B} as defined in Theorem 2 (that is, setting $\beta \equiv \theta^B$).*

Our next result gives conditions under which the efficient investment level can be supported; that is, when the best achievable investment level θ^B coincides with the efficient investment level θ^*.

PROPOSITION 3. *Suppose Assumptions 2–4 hold. The efficient level of investment θ^* is supported in the deterministic and unified case if and only if*

$$u_1(\overline{a}, \theta^*) + \pi_1 R(\overline{a}, \theta^*) - \theta^* \geq u_1(\overline{a}, \theta) - \theta \tag{2}$$

for all $\theta < \theta^*$.

The condition from Proposition 3 ensures that we can induce a large enough discontinuity in the value function at the efficient state so that player 1 maximizes his gain net of investment cost at θ^*. If this condition fails, efficiency cannot be attained in the unified case.

To summarize the results of this section, we have conditions under which player 2 can be made the residual claimant with respect to the investment action, which solves the contracting problem (yielding efficient investment and trade) in the divided case. We learn that it is generally not possible to make player 1 the residual claimant, given that player 1 has the trade action. As a result, the efficient outcome is typically not attainable in the case of unified investment and trade actions. However, for the class of trade technologies that satisfy Assumptions 2–4 and for the case of a deterministic state, we are able to characterize an optimal contract and provide conditions under which the efficient outcome can be achieved.

The results in this section are most pronounced when applied to settings of significant cross investment, where the optimal forcing contract is null and leads to an inefficient outcome. In the divided case, the nonforcing dual-option contract induces efficient investment and trade. In the unified case, a dual-option contract can outperform the null contract and sometimes induces the efficient outcome. We continue this theme in Section 5 by simply asking whether, in general, nonforcing contracts implement a wider range of value functions than do forcing contracts.

4. Implementable value functions

This section summarizes how to calculate implementable value functions in general. Much of the analysis here repeats material in Watson (2007), so we keep this text brief and ask the reader to see Watson (2007) for more details. The culmination of the basic analysis here are some simple characterization results from Watson (2007), which we build on in the subsequent section.

In previous sections we assumed that the players can freely renegotiate at dates 3 and 5, but now we also consider the case in which renegotiation is possible at date 3 only (the interim phase). We let V^{EPF} be the set of implementable value functions from date 3 for the case of ex post renegotiation and with the restriction to forcing contracts. We let V^{EP} be the corresponding set for the case of ex post renegotiation and no contractual restrictions. Further, we let V^{I} be the set of implementable value functions for the case in which renegotiation can occur only at date 3.[15] We can characterize the implementable value functions by backward induction, starting with date 6 where player 1 selects the trade action.

State-contingent values from date 6

To calculate the value functions that are supported from date 6, we can ignore the payoff-irrelevant messages sent earlier (or equivalently, fix a message profile from date 4) and simply write the externally enforced transfer function as $\hat{y}\colon A \to \mathbb{R}^2$. That is, $\hat{y}$ gives the monetary transfer as a function of player 1's trade action.

Given the state θ, $\hat{y}$ defines a *trading game* in which player 1 selects an action $a \in A$ and the payoff vector is then $u(a, \theta) + \hat{y}(a)$. Focusing on pure strategies, we let $\hat{a}(\theta)$ denote the action chosen by player 1 in state θ. This specification is rational for player 1 if, for every $\theta \in \Theta$, $\hat{a}$ maximizes $u_1(a, \theta) + \hat{y}_1(a)$ by choice of a. The state-contingent payoff vector from date 6 is then given by the *outcome function* $w\colon \Theta \to \mathbb{R}^2$ defined by

$$w(\theta) \equiv u(\hat{a}(\theta), \theta) + \hat{y}(\hat{a}(\theta)). \tag{3}$$

Let W denote the set of supportable outcome functions. That is, $w \in W$ if and only if there are functions $\hat{y}$ and $\hat{a}$ such that $\hat{a}$ is rational for player 1 and, for every $\theta \in \Theta$, (3) holds. Furthermore, let W^{F} be the subset of outcomes that can be supported using forcing contracts. It is easy to see that $w \in W^{\text{F}}$ if and only if there is a trade action $\hat{a}$ and a transfer vector $\hat{t}$ such that $w(\theta) = u(\hat{a}, \theta) + \hat{t}$ for all $\theta \in \Theta$. We can compare individual-action and public-action models by determining whether the restriction to forcing contracts implies a significant constraint on the set of implementable value functions.

State-contingent values from date 5

We next step back to date 5. If there is no opportunity for ex post renegotiation, then nothing happens at date 5 and so W and W^{F} are the supported state-contingent value

---

[15]In the case of only interim renegotiation, a restriction to forcing contracts does not affect the implementable set.

sets from the start of date 5 as well. On the other hand, if ex post renegotiation is allowed, then at date 5 the players have an opportunity to discard their originally specified contract y and replace it with another, y'.

By picking a new contract y', the players are effectively choosing a new outcome function w' in place of the function w that would have resulted from the original contract y. The players can freely select w' from the set W or the set W^F, depending on whether they are restricted to forcing contracts. The players divide the renegotiation surplus according to the fixed bargaining weights π_1 and π_2. Dividing the surplus in this way is feasible because W and W^F are closed under constant transfers.

Clearly, we have $\gamma(\theta) = \max_{w \in W^F}[w_1(\theta) + w_2(\theta)]$ because the trade action that solves the maximization problem in (1) can be specified in a forcing contract to yield the desired outcome. The *renegotiation surplus* is

$$r(w, \theta) \equiv \gamma(\theta) - w_1(\theta) - w_2(\theta).$$

The bargaining solution implies that the players settle on a new outcome in which the payoff vector in state θ is $w(\theta) + \pi r(w, \theta)$.

We define an *ex post renegotiation outcome* to be the state-contingent payoff vector that results when, in every state, the players renegotiate from a fixed outcome in W. That is, a value function z is an ex post renegotiation outcome if and only if there is an outcome $w \in W$ such that $z(\theta) = w(\theta) + \pi r(w, \theta)$ for every $\theta \in \Theta$. Let Z denote the set of ex post renegotiation outcomes.[16] If trade actions are treated as public (and so attention is limited to forcing contracts), then the set of ex post renegotiation outcomes contains only the value functions of the form $z = w + \pi r(w, \cdot)$ with the constraint that $w \in W^F$. Let Z^F denote the set of ex post renegotiation outcomes under forcing contracts. Although the terminology is a bit loose, we refer to functions in Z and Z^F, in addition to functions in W and W^F, simply as "outcomes."

State-contingent values from dates 4 and 3

Analysis of contract selection and incentives at date 4 can be viewed as a standard mechanism-design problem. The players' contract is equivalent to a mechanism that maps messages sent at date 4 to outcomes induced in the trade and enforcement phase (possibly renegotiated at date 5). The revelation principle applies, so we can restrict attention to direct-revelation mechanisms defined by (i) a message space $M \equiv \Theta^2$ and (ii) a function that maps Θ^2 to the relevant outcome set that gives the state-contingent value functions from the start of date 5. The outcome set is either W, W^F, Z, or Z^F, depending on whether ex post renegotiation and/or nonforcing contracts are allowed. We concentrate on Nash equilibria of the mechanism in which the parties report truthfully in each state.[17]

Let us write $\psi^{\theta_1 \theta_2}$ for the outcome that the mechanism prescribes when player 1 reports the state to be θ_1 and player 2 reports the state to be θ_2. Note that, in any given

---

[16]All elements of Z are efficient in every state; also, Z and W are generally not ranked by inclusion.

[17]The revelation principle usually requires a public randomization device to create lotteries over outcomes (or that the outcome set is a mixture space), but it is not needed here.

state θ (the actual state that occurred), the mechanism implies a "message game" with strategy space Θ^2 and payoffs given by $\psi^{\theta_1 \theta_2}(\theta)$ for each strategy profile (θ_1, θ_2). For truthful reporting to be a Nash equilibrium of this game, it must be that $\psi_1^{\theta\theta}(\theta) \geq \psi_1^{\tilde{\theta}\theta}(\theta)$ and $\psi_2^{\theta\theta}(\theta) \geq \psi_2^{\theta\tilde{\theta}}(\theta)$ for all $\tilde{\theta} \in \Theta$.

We proceed using standard techniques for mechanism design with transfers, following Watson (2007). The key step is observing that, for any two states θ and θ', the outcome specified for the "off-diagonal" message profile (θ', θ) must be sufficient to simultaneously (i) dissuade player 1 from declaring the state to be θ' when the state is actually θ and (ii) discourage player 2 from declaring θ in state θ'. Thus, we require

$$\psi_1^{\theta\theta}(\theta) \geq \psi_1^{\theta'\theta}(\theta) \quad \text{and} \quad \psi_2^{\theta'\theta'}(\theta') \geq \psi_2^{\theta'\theta}(\theta').$$

Because the outcome sets are closed under constant transfers, we can choose the outcome to effectively raise or lower $\psi_1^{\theta'\theta}$ and $\psi_2^{\theta'\theta}$ while keeping the sum constant. Thus, a sufficient condition for these two inequalities is that the sum of the two holds. Letting $\psi \equiv \psi^{\theta\theta}$ and $\psi' \equiv \psi^{\theta'\theta'}$, we thus have the following necessary condition for implementing outcome ψ in state θ and outcome ψ' in state θ':

There exists an outcome $\hat{\psi}$ satisfying $\psi_1(\theta) + \psi'_2(\theta') \geq \hat{\psi}_1(\theta) + \hat{\psi}_2(\theta')$.

This condition, applied to all ordered pairs (θ, θ'), is necessary and sufficient for implementation. The sum $\hat{\psi}_1(\theta) + \hat{\psi}_2(\theta')$ is called the *punishment value* corresponding to the ordered pair (θ, θ'). The punishment value plays a central role in our analysis. Lower punishment values imply a greater set of implementable outcomes.

Interim renegotiation has the effect of requiring each "on-diagonal" outcome to be efficient in the relevant state; that is, for each θ, we need $\psi^{\theta\theta}$ to be efficient in this state. In the case of ex post renegotiation, allowing interim renegotiation entails no further constraint because every outcome in Z is efficient in every state. It is also the case that without ex post renegotiation, W and W^{F} yield the same set of implementable value functions from date 3. Therefore, we have three settings to compare: unrestricted contracts with ex post renegotiation, forcing contracts (public actions) with ex post renegotiation, and forcing contracts with interim (but not ex post) renegotiation.

Call a value function v *efficient* if $v_1(\theta) + v_2(\theta) = \gamma(\theta)$ for every $\theta \in \Theta$. The following results summarize the characterization of V^{EP}, V^{EPF}, and V^{I} and provide a general comparison.

RESULT 1 (Watson (2007)). Consider any value function $v : \Theta \to \mathbb{R}^2$.

- *Implementation with interim renegotiation.* Value function v is an element of V^{I} if and only if v is efficient and, for every pair of states θ and θ', there is an outcome $\hat{w} \in W^{\mathrm{F}}$ such that $v_1(\theta) + v_2(\theta') \geq \hat{w}_1(\theta) + \hat{w}_2(\theta')$.

- *Implementation with ex post renegotiation.* Value function v is an element of V^{EP} if and only if v is efficient and, for every pair of states θ and θ', there is an outcome $\hat{z} \in Z$ such that $v_1(\theta) + v_2(\theta') \geq \hat{z}_1(\theta) + \hat{z}_2(\theta')$.

- *Implementation with ex post renegotiation and forcing contracts.* Value function v is an element of V^{EPF} if and only if v is efficient and, for every pair of states θ and θ', there is an outcome $\hat{z} \in Z^{\mathrm{F}}$ such that $v_1(\theta) + v_2(\theta') \geq \hat{z}_1(\theta) + \hat{z}_2(\theta')$.

Furthermore, the sets V^{EP}, V^{EPF}, and V^{I} are closed under constant transfers.

RESULT 2 (Watson (2007)). *The implementable sets are weakly nested in that* $V^{\mathrm{EPF}} \subseteq V^{\mathrm{EP}} \subseteq V^{\mathrm{I}}$. *Furthermore,* $V^{\mathrm{EPF}} = V^{\mathrm{EP}}$ *if and only if, for every pair of states* $\theta, \theta' \in \Theta$ *and every* $\hat{z} \in Z$, *there is an ex post renegotiation outcome* $\tilde{z} \in Z^{\mathrm{F}}$ *such that* $\tilde{z}_1(\theta) + \tilde{z}_2(\theta') \leq \hat{z}_1(\theta) + \hat{z}_2(\theta')$. *Likewise,* $V^{\mathrm{EP}} = V^{\mathrm{I}}$ *if and only if, for all* $\theta, \theta' \in \Theta$ *and every* $\hat{w} \in W^{\mathrm{F}}$, *there is an ex post renegotiation outcome* $\hat{z} \in Z$ *such that* $\hat{z}_1(\theta) + \hat{z}_2(\theta') \leq \hat{w}_1(\theta) + \hat{w}_2(\theta')$.[18]

To summarize, we have thus far analyzed the players' behavior at the various dates in the contractual relationship, leading to a simple characterization of implementable value functions from date 3. The characterization is in terms of the minimum punishment values for each pair of states, which yields a way to relate the implementable sets for the cases of interim renegotiation, ex post renegotiation, and ex post renegotiation and forcing contracts. We next turn to investigate the relation more deeply.

5. A ROBUSTNESS RESULT FOR NON-FORCING CONTRACTS

The example from Watson (2007) and our example in Section 2 provide illustrations of $V^{\mathrm{EPF}} \neq V^{\mathrm{EP}} \neq V^{\mathrm{I}}$. Our main objective in this section is to examine the robustness of this conclusion. We consider the wide class of contractual relationships that satisfy Assumptions 2, 3, and 5 (which follows).

ASSUMPTION 5. *There exist states* θ^1, $\theta^2 \in \Theta$ *such that* $\theta^1 > \theta^2$, $a^*(\theta^1) > \underline{a}$, $a^*(\theta^2) < \overline{a}$ *and either* $U(\underline{a}, \theta^2) < U(\overline{a}, \theta^2)$ *or* $U(\underline{a}, \theta^1) > U(\overline{a}, \theta^1)$.

This is a weak assumption that removes a knife-edge case concerning the relative joint values of the extreme trade actions in the various states. For instance, if Θ has more than two elements and $U(\underline{a}, \theta) \neq U(\overline{a}, \theta)$ for some θ strictly between $\underline{\theta}$ and $\overline{\theta}$ with interior optimal actions, then Assumption 5 is satisfied.

We have the following robustness result.

THEOREM 3. *Consider any contractual relationship that satisfies Assumptions 2, 3, and 5. The sets of implementable value functions in the cases of unrestricted contracts with ex post renegotiation, forcing contracts with ex post renegotiation, and interim renegotiation are all distinct. That is,* $V^{\mathrm{EPF}} \neq V^{\mathrm{EP}} \neq V^{\mathrm{I}}$.

---

[18]Watson's (2005) Lemma 1 provides some of the supporting analysis (which was not explained fully in the relevant proof in Watson 2007). This lemma establishes that, for any given ordered pair of states θ and θ' and any supportable outcome ψ, there exists an implementable value function v for which $v_1(\theta) + v_2(\theta') = \psi_1(\theta) + \psi_2(\theta')$. Because the minimum punishment values exist, in each case we can let ψ equal the outcome that attains the minimum.

The analysis underlying Theorem 3 amounts to characterizing and comparing the minimum punishment values that can be supported for each of the settings of interest. Recall that the punishment value for the ordered pair (θ, θ') is the value $\psi_1(\theta) + \psi_2(\theta')$, where ψ is the outcome specified in the message game when player 1 reports the state to be θ' and player 2 reports the state to be θ. Lower punishment values serve to relax incentive conditions, so to characterize the sets of implementable value functions completely, we must find the minimum punishment values. We let P^{I}, P^{EP}, and P^{EPF} denote the minimum punishment values for the settings of interim renegotiation, ex post renegotiation, and ex post renegotiation and forcing contracts, respectively:

$$P^{\mathrm{I}}(\theta, \theta') \equiv \min_{w \in W^{\mathrm{F}}} w_1(\theta) + w_2(\theta')$$

$$P^{\mathrm{EP}}(\theta, \theta') \equiv \min_{\hat{z} \in Z} \hat{z}_1(\theta) + \hat{z}_2(\theta')$$

$$P^{\mathrm{EPF}}(\theta, \theta') \equiv \min_{\hat{z} \in Z^{\mathrm{F}}} \hat{z}_1(\theta) + \hat{z}_2(\theta').$$

Our assumptions on the trade technology guarantee that these minima exist.

From Result 2, we know that Theorem 3 is equivalent to saying that there exist states $\theta, \theta' \in \Theta$ such that $P^{\mathrm{I}}(\theta, \theta') < P^{\mathrm{EP}}(\theta, \theta')$ and there exist (possibly different) states $\theta, \theta' \in \Theta$ such that $P^{\mathrm{EP}}(\theta, \theta') < P^{\mathrm{EPF}}(\theta, \theta')$. Thus, to prove Theorem 3, we examine the punishment values achieved by various contractual specifications in the different settings. We develop some elements of the proof in the remainder of this section; Appendix C contains the rest of the analysis. We focus in this section on the relation between V^{EPF} and V^{EP}. The analysis of the relation between V^{EP} and V^{I} is considerably simpler and is wholly contained in Appendix C.

We establish $P^{\mathrm{EP}} < P^{\mathrm{EPF}}$ by comparing the punishment values implied by (i) the outcome in which player 1 is forced to take the trade action that yields the lowest punishment value among forcing contracts, and (ii) a related nonforcing specification in which player 1 is given the incentive to select some action a in state θ and a different action a' in state θ'. We derive conditions under which a and a' can be arranged to strictly lower the punishment value for (θ, θ'), relative to the best forcing case. We then find states θ^1 and θ^2 such that the conditions must hold for at least one of the ordered pairs (θ^1, θ^2) and (θ^2, θ^1).

To explore the possible outcomes in the cases of ex post renegotiation, consider player 1's incentives at date 6. For any given transfer function $\hat{y}$, necessary conditions for player 1 to select trade action a in state θ and action a' in state θ' are

$$u_1(a, \theta) + \hat{y}_1(a) \geq u_1(a', \theta) + \hat{y}_1(a')$$
$$u_1(a', \theta') + \hat{y}_1(a') \geq u_1(a, \theta') + \hat{y}_1(a). \tag{4}$$

Transfer function $\hat{y}$ can be specified so that player 1 is harshly punished for selecting any trade action other than a or a'. Then, in every state, either a or a' maximizes player 1's payoff from date 6. Thus, we can state the following fact.

FACT 1. Consider two states $\theta, \theta' \in \Theta$ and two trade actions $a, a' \in A$. Expression (4) is necessary and sufficient for the existence of a transfer function $\hat{y} \colon A \to \mathbb{R}_0^2$ (defined over all trade actions) such that player 1's optimal trade action in state θ is a and player 1's optimal trade action in state θ' is a'.

Summing the inequalities of expression (4), we see that there are values $\hat{y}(a), \hat{y}(a') \in \mathbb{R}_0^2$ that satisfy (4) if and only if

$$u_1(a, \theta) - u_1(a', \theta) \geq u_1(a, \theta') - u_1(a', \theta'). \tag{5}$$

Assumption 3 then implies the following fact.

FACT 2. If $\theta > \theta'$, then $a \geq a'$ implies inequality (5). If $\theta < \theta'$, then $a \leq a'$ implies inequality (5).

Note that Fact 2 gives sufficient conditions. In the case in which $u_1(\cdot, \cdot)$ is strictly supermodular (replacing weak inequalities in Assumption 3 with strict inequalities), player 1 can be given only the incentive to choose greater trade actions in higher states.
For any two states $\theta, \theta' \in \Theta$, define

$$E(\theta, \theta') \equiv \{(a, a') \in A \times A \mid \text{inequality (5) is satisfied}\}.$$

Also, for states $\theta, \theta' \in \Theta$ and trade actions $a, a' \in A$ with $(a, a') \in E(\theta, \theta')$, define

$$Y(a, a', \theta, \theta') \equiv \{\hat{y} \colon A \to \mathbb{R}_0^2 \mid \text{condition (4) is satisfied}\}.$$

Condition (4), combined with the identity $\hat{y}_1 = -\hat{y}_2$, implies the next fact.

FACT 3. For any $\theta, \theta' \in \Theta$ and $a, a' \in A$, with $(a, a') \in E(\theta, \theta')$, we have

$$\min_{\hat{y} \in Y(a, a', \theta, \theta')} \hat{y}_1(a) + \hat{y}_2(a') = u_1(a', \theta) - u_1(a, \theta).$$

Using the definition of the set W (recall expression (3)), any given $w \in W$ can be written in terms of the trade actions and transfers that support it. We have

$$w(\theta) = u(\hat{a}(\theta), \theta) + \hat{y}(\hat{a}(\theta))$$

and

$$w(\theta') = u(\hat{a}(\theta'), \theta') + \hat{y}(\hat{a}(\theta')),$$

where $\hat{a}$ gives player 1's choice of trade action as a function of the state and $\hat{y}$ is the transfer function that supports w.
For any state $\tilde{\theta}$ and trade action $\tilde{a}$, define $R(\tilde{a}, \tilde{\theta})$ to be the renegotiation surplus if, without renegotiation, player 1 would select $\tilde{a}$. That is, $R(\tilde{a}, \tilde{\theta}) = U(a^*(\tilde{\theta}), \tilde{\theta}) - U(\tilde{a}, \tilde{\theta})$. Combining the expressions for w in the previous paragraph with Fact 1 and the definition of ex post renegotiation outcomes, we obtain Fact 4.

FACT 4. Consider any two states $\theta, \theta' \in \Theta$ and let α be any number. There is an ex post renegotiation outcome $z \in Z$ that satisfies $z_1(\theta) + z_2(\theta') = \rho$ if and only if there are trade actions $a, a' \in A$ and a transfer function $\hat{y}$ such that $(a, a') \in E(\theta, \theta')$, $\hat{y} \in Y(a, a', \theta, \theta')$, and

$$\rho = u_1(a, \theta) + \hat{y}_1(a) + \pi_1 R(a, \theta) + u_2(a', \theta') + \hat{y}_2(a') + \pi_2 R(a', \theta'). \tag{6}$$

In the last line, the first three terms are $w_1(\theta)$ plus player 1's share of the renegotiation surplus in state θ, totaling $z_1(\theta)$. The last three terms are $w_2(\theta')$ plus player 2's share of the renegotiation surplus in state θ', totaling $z_2(\theta')$.

Finding the best (minimum) punishment value for states θ and θ' means minimizing $\hat{z}_1(\theta) + \hat{z}_2(\theta')$ by choice of $\hat{z} \in Z$. For now, holding fixed the trade actions a and a' that player 1 is induced to select in states θ and θ', let us minimize the punishment value by choice of $\hat{y} \in Y(a, a', \theta, \theta')$. To this end, we can use Fact 3 to substitute for $\hat{y}_1(a) + \hat{y}_2(a')$ in expression (6). This yields the punishment value for trade actions a and a' in states θ and θ', respectively, written

$$\lambda(a, a', \theta, \theta') \equiv u_1(a', \theta) + \pi_1 R(a, \theta) + u_2(a', \theta') + \pi_2 R(a', \theta'). \tag{7}$$

Next, we consider the step of minimizing the punishment value by choice of the trade actions a and a', which gives us a useful characterization of $P^{\mathrm{EP}}(\theta, \theta')$. Assumption 2(a) guarantees that $\lambda(a, a', \theta, \theta')$ has a minimum.

FACT 5. The minimum punishment value in the setting of ex post renegotiation is characterized as

$$P^{\mathrm{EP}}(\theta, \theta') = \min_{(a, a') \in E(\theta, \theta')} \lambda(a, a', \theta, \theta').$$

We obtain a similar characterization of the minimal punishment value for the setting in which attention is restricted to forcing contracts. The characterization is exactly as in Fact 5 except with the additional requirement that $a = a'$ because forcing contracts compel the same action in every state.

FACT 6. The minimum punishment value for the setting of forcing contracts and ex post renegotiation is characterized as

$$P^{\mathrm{EPF}}(\theta, \theta') \equiv \min_{a \in A} \lambda(a, a, \theta, \theta').$$

Recall that proving Theorem 3 requires us to establish that $P^{\mathrm{EP}}(\theta, \theta') < P^{\mathrm{EPF}}(\theta, \theta')$ for some pair of states $\theta, \theta' \in \Theta$. Appendix C finishes the analysis by exploring how one can depart from the optimal forcing specification in a way that strictly reduces the value $\lambda(a, a', \theta, \theta')$.

6. Conclusion

In this paper, we report the analysis of contractual relationships for a large class of trade technologies. We highlight the usefulness of nonforcing contracts (in particular, the dual-option contract) and the key distinction between the *divided* and *unified* cases of investment and trade actions. We show that the payoff of the party with the trade action can be neutralized so that the other party claims the full benefit of the investment, gross of investment costs, implying that the efficient outcome is achieved in the divided case. Holdup remains a problem in the unified case, although the dual-option contract is optimal in a class of settings and it can sometimes induce the efficient outcome. We also provide general results on the relation between individual-action and public-action models of contractual relationships, showing that limiting attention to forcing contracts has significant implications for implementability and hence inefficiency.

Our results reinforce the message of Watson (2007) on the usefulness of modeling trade actions as individual, particularly in settings of cross investment. The results suggest revisiting some of the conclusions of public-action models in the existing literature. In particular, settings with cross investment are generally not as problematic as previous modeling exercises (Che and Hausch 1999, Edlin and Hermalin 2000, and others) find. Efficient outcomes can be achieved in the case of divided investment and trade actions. Our results show the importance, for applied work, of differentiating between the cases of divided and unified investment and trade actions. This distinction may be just as important as the distinction between own and cross investment (on which the literature has focused until now).

In our model, the trading opportunity is nondurable in that there is a single moment in time when trade can occur. One might wonder if the results differ substantially in settings with durable trading opportunities (where if trade does not occur at one time, then it can still be done at a later date). This issue is explored by Evans (2008) and Watson and Wignall (2009), both of whom examine individual-action models. Evans' (2008) elegant model is very general in terms of the available times at which the players can trade and renegotiate. He constructs equilibria in which, by having the players coordinate in different states on different equilibria in the infinite-horizon trade/negotiation game, the holdup problem is partly or completely alleviated. Evans' strongest result (in which the efficient outcome is reached) requires the ability of the players to commit to a joint financial hostage; that is, money is deposited with a third party until trade occurs, if ever.

Watson and Wignall (2009) examine a cross investment setting without the possibility of joint financial hostages, and their model is more modest in other dimensions. They show that the set of implementable post-investment payoff vectors in the setting of a durable trading opportunity is essentially the same as in the setting of a nondurable trading opportunity. This suggests that, in general, the results from the current paper carry over to the durability setting. Watson and Wignall also show that, in the divided case, there are nonstationary contracts that uniquely support the efficient outcome.

Our modeling exercise, combined with the recent literature, suggests some broad conclusions about the prospect of efficient investment and trade in contractual relationships. First, the holdup problem is not necessarily severe, and efficient outcomes

can often be achieved. Durability of the trading opportunity does not worsen the holdup problem and may soften it in some cases, but it depends on the investment and trade technologies. Inefficiency may be unavoidable in the following problematic cases.

- When there is cross investment and unified investment and trade actions, as identified herein.

- When trade involves "complexity/ambivalence" as described by Segal (1999), Hart and Moore (1999), and Reiche (2006).

- When multiple parties make cross/cooperative investments.

- When the investment conveys a significant direct benefit (not requiring trade) on the noninvesting party, in addition to any benefit contingent on trade.

On the last point, Ellman's (2006) model provides intuition in terms of the notion of specificity. Settings in which multiple parties make cross investments are similar in nature to settings of team production (studied by Holmstrom 1982 and others).

In each of the cases above, the holdup problem would be reduced if the parties had some way to create joint financial hostages, as explored by Evans (2008), Boeckem and Schiller (2008), and Baliga and Sjöström (2009). Bull (2009) provides a cautionary note on the inability of such financial arrangements to withstand side contracting.

Regarding extensions of our analysis here, future research would be useful on different classes of trade technologies, in particular those in which both parties take trade actions (either simultaneously or sequentially). For instance, consider the setting with trade action profile $a = (a_1, a_2)$, where $a_1 \in [0, \overline{a}]$ is the verifiable quantity of an intermediate good that player 1 produces and delivers to player 2, and $a_2 \in \{\text{accept, reject}\}$ is player 2's verifiable choice of whether to accept or reject delivery. For simplicity, suppose that the players choose their trade actions simultaneously; the case of sequential choices works out similarly. Suppose that, for every state, "accept" is part of an ex post efficient trade-action profile.[19] Then a version of our results can be proved by utilizing contracts that force player 2 to accept delivery and are dual options with respect to a_1.

More precisely, if (0, accept) satisfies Assumption 1, then the conclusions of Theorem 1 and Proposition 1 go through, so an efficient outcome can be achieved in the divided case. Suppose further that, fixing $a_2 = $ accept and considering u as a function of a_1, Assumptions 2–4 hold. Then the conclusion of Theorem 2 is valid and the conclusions of Propositions 2 and 3 hold within the class of contracts that force player 2 to accept delivery, and are otherwise arbitrary (nonforcing and message-based).

It is not always possible to make player 1 the residual claimant by building a dual-option contract on player 2's acceptance/rejection choice, because there is typically not a single quantity a_1 that figures in the ex post efficient outcome in every state. Thus, we generally are not able to apply the argument for Theorem 1 to make player 1 the residual claimant. However, in some cases, a nonforcing specification for a_2 yields an improvement on the contracts that force acceptance of delivery.

---

[19]For instance, if positive trade is inefficient in some state, then (0, accept) is an ex post efficient outcome where player 1 delivers nothing and player 2 accepts.

There are technologies of trade for which it is possible to make either player the residual claimant with respect to the investment choice. Consider a setting with trade-action profile $a = (a_1, a_2)$ and, for any player i let j denote the other player. Here are assumptions that imply that player i can be made the residual claimant using a dual-option contract: Assume that there is a trade action $\hat{a}_i$ for player i and a trade action a_j^0 for player j such that, for every state θ, (i) $\hat{a}_i$ is an efficient action with an appropriate choice of a_j and (ii) $u_i(a_j^0, \hat{a}_i, \theta) = u_j(a_j^0, \hat{a}_i, \theta) = 0$. In words, the first condition says there is a trade action that player i can be forced to take, such that efficiency can be achieved in each state for some selection of player j's trade action. This condition is trivially satisfied in the setting of one-sided trade actions that we study. The second condition is Assumption 1 for j's trade action. Under these assumptions, player i can be made the residual claimant using the arguments behind Theorem 1. Clearly, it is possible for the assumptions to hold for both $i = 1$ and $i = 2$, although this seems to be a rather special case.

We have not begun an analysis of settings with more complicated multilateral trade actions, but we expect it to be a fruitful line of future research. Evans' (2008) model has a dynamic version of the trade technology described above, where one player makes a delivery choice and the other chooses whether to accept or reject delivery. It would also be interesting to look at a wide range of settings with partially verifiable trade actions. For example, a court may observe whether a particular trade was made but have trouble identifying which party disrupted trade (in the event that trade did not occur).[20]

Finally, recall that in the modeling exercise here, we assume that each party's productive actions are exogenously given. However, in some settings it may be possible to arbitrarily assign a particular task (such as delivering an object from one place to another) to an individual player. Our model indicates that the parties would have preferences over task assignment. Thus, it would be useful to determine whether physical trade actions are assignable in some real settings and to develop a model of optimal assignment. One might imagine a theory of firm boundaries that is based on the optimal assignment of different types of tasks over time.

APPENDIX A: PROOF OF THEOREM 1

This appendix provides a proof of the first theorem. For any fixed k, consider the following contract. In the message phase (date 4), player 2 must declare the state. Let $\hat{\theta}$ denote player 2's announcement. If player 1 subsequently selects action $a^*(\hat{\theta})$, then the enforcer compels a transfer of $\hat{t} = (k - u_1(a^*(\hat{\theta}), \hat{\theta}), u_1(a^*(\hat{\theta}), \hat{\theta}) - k)$. If player 1 selects action a^0, then the transfer is $\underline{t} = (k, -k)$. If player 1 chooses any other trade action, then the enforcer compels transfer $(-\tau, \tau)$, where τ is set large enough so that player 1 is

---

[20]Hart and Moore's (1988) model has this feature. It is straightforward to incorporate partial verifiability into the modeling framework developed here. One can represent the external enforcer's information about the trading game as a partition of the space of action profiles. One can then simply assume that the contracted transfers y must be measurable with respect to this partition. Note that MacLeod and Malcomson (1993) and De Fraja (1999) examine settings with partially verifiable trade actions (along the lines of Hart and Moore 1988), although they make assumptions about the renegotiation protocol and the timing of outside options that weaken the affect of renegotiation compared to the rest of the literature.

forced to choose between $a^*(\hat{\theta})$ and a^0. That is, regardless of $\hat{\theta}$, in no state does player 1 have the incentive to choose $a \notin \{a^*(\hat{\theta}), a^0\}$.

Suppose that date 6 is reached without renegotiation and that the state is θ. Note that, by Assumption 1, player 1 gets a payoff of k if he chooses a^0. Alternatively, his payoff is

$$u_1(a^*(\hat{\theta}), \theta) + k - u_1(a^*(\hat{\theta}), \hat{\theta})$$

if he chooses $a^*(\hat{\theta})$. Thus, it is rational for player 1 to choose $a^*(\hat{\theta})$ if $u_1(a^*(\hat{\theta}), \theta) \geq u_1(a^*(\hat{\theta}), \hat{\theta})$ and to select a^0 otherwise, which we suppose is how player 1 behaves.

Consider next how player 2's payoff from date 4 depends on $\hat{\theta}$. Let θ be the actual state and divide the analysis into three cases. First, if player 2 declares $\hat{\theta} = \theta$, then, under the original contract, player 1 chooses $a^*(\hat{\theta})$ at date 6 and there is nothing to be jointly gained by renegotiating at date 5. In this case, the payoffs from date 4 are k for player 1 and

$$u_1(a^*(\theta), \theta) + u_2(a^*(\theta), \theta) - k = \gamma(\theta) - k$$

for player 2.

Second, if player 2 instead were to declare the state to be some $\hat{\theta} \neq \theta$ such that $u_1(a^*(\hat{\theta}), \theta) < u_1(a^*(\hat{\theta}), \hat{\theta})$, then the players anticipate that player 1 will select a^0 at date 6 under the original contract. Incorporating the impact of renegotiation at date 5, player 1's payoff from date 4 then is $k + \pi_1 R(a^0, \theta)$, where $R(a^0, \theta)$ is the renegotiation surplus in state θ if, without renegotiation, the players anticipate that a^0 will be the chosen trade action. Since $R(a^0, \theta) \geq 0$, player 1's payoff from date 4 weakly exceeds k and we conclude that player 2's payoff is weakly less than $\gamma(\theta) - k$.

Finally, suppose that player 2 declares the state to be $\hat{\theta} \neq \theta$ such that $u_1(a^*(\hat{\theta}), \theta) > u_1(a^*(\hat{\theta}), \hat{\theta})$. In this case, the players anticipate that player 1 will select $a^*(\hat{\theta})$ at date 6 under the original contract. Incorporating renegotiation at date 5, player 1's payoff from date 4 then is

$$u_1(a^*(\hat{\theta}), \theta) + k - u_1(a^*(\hat{\theta}), \hat{\theta}) + \pi_1 R(a^*(\hat{\theta}), \theta),$$

where $R(a^*(\hat{\theta}), \theta)$ is the renegotiation surplus in state θ if, without renegotiation, the players anticipate that $a^*(\hat{\theta})$ will be the chosen trade action. The first and third terms sum to weakly more than zero, so the entire expression weakly exceeds k. This implies that player 2's payoff is weakly less than $\gamma(\theta) - k$.

The foregoing analysis demonstrates that player 2 optimally tells the truth at date 4; that is, she declares $\hat{\theta} = \theta$. The payoffs from date 3 are thus k for player 1 and $\gamma(\theta) - k$ for player 2, which means that the contract implements the desired value function. □

Appendix B: Additional analysis and proofs for Section 3

In this appendix, we first provide necessary and sufficient conditions for making player 1 the residual claimant. We follow this with a note on when the conditions fail. We then provide a characterization of the punishment values defined in Sections 4 and 5. We use this characterization to examine the relation between the "inside conditions" and the

"outside conditions" discussed in Section 3, which shows the difficulties of determining the second-best contract in the general unified case. Analysis of the inside and outside conditions for implementability motivates Theorem 2 and Propositions 2 and 3. This appendix concludes with their proofs.

Making player 1 the residual claimant

To make player 1 the residual claimant, we need to implement a value function v that satisfies, for some constant k, $v_2(\theta) = k$ and $v_1(\theta) = \gamma(\theta) - k$ for all $\theta \in \Theta$. Consider two states θ and θ', and order them so that $\theta > \theta'$. The conditions for implementation associated with these two states (for (θ, θ') and (θ', θ)) are

$$v_1(\theta) + v_2(\theta') \geq P^{\mathrm{EP}}(\theta, \theta') \tag{8}$$

and

$$v_1(\theta') + v_2(\theta) \geq P^{\mathrm{EP}}(\theta', \theta). \tag{9}$$

Using Fact 5 from Section 5, these conditions are equivalent to the existence of trade actions a, a', b, and b' such that $(a, a') \in E(\theta, \theta')$, $(b', b) \in E(\theta', \theta)$,

$$v_1(\theta) + v_2(\theta') \geq \lambda(a, a', \theta, \theta')$$

and

$$v_1(\theta') + v_2(\theta) \geq \lambda(b', b, \theta', \theta).$$

Substituting for v_1 and v_2 using the identities $v_2(\theta) = k$ and $v_1(\theta) = \gamma(\theta) - k$, these two inequalities become

$$\lambda(a, a', \theta, \theta') \leq \gamma(\theta) \tag{10}$$

and

$$\lambda(b', b, \theta', \theta) \leq \gamma(\theta'). \tag{11}$$

The following lemma summarizes.

LEMMA 1. *Consider any contractual relationship that satisfies Assumptions 2(a) and 3. Let k be any real number and define value function v by $v_2(\theta) = k$ and $v_1(\theta) = \gamma(\theta) - k$ for all $\theta \in \Theta$. Then $v \in V^{\mathrm{EP}}$ if and only if, for all pairs of states θ, θ' with $\theta > \theta'$, there are trade actions a, a', b, and b' such that $(a, a') \in E(\theta, \theta')$, $(b', b) \in E(\theta', \theta)$, and inequalities (10) and (11) hold.*

One can use these conditions to establish whether efficient investment can be obtained in specific examples with unified investment and trade actions, but sufficient conditions are stronger than are the assumptions we make in this paper.

For an illustration of cases where the conditions of Lemma 1 fail, suppose that the strict version of Assumption 3 is satisfied, meaning u_1 is strictly supermodular. Further suppose that Assumption 2 holds. Also suppose that U is strictly increasing in θ and

that $U(\overline{a}, \overline{\theta}) > \gamma(\underline{\theta})$. That is, the joint value of the highest trade action in the highest state exceeds the maximal joint value in the lowest state (gross of investment cost).

Using (7), $U = u_1 + u_2$, and some algebra, we can rewrite inequality (11) as

$$\pi_1[U(b, \theta) - U(b', \theta')] \le \pi_2[\gamma(\theta') - \gamma(\theta)] - [u_1(b, \theta') - u_1(b, \theta)].$$

Examining the case of $\theta = \overline{\theta}$ and $\theta' = \underline{\theta}$, this becomes

$$\pi_1[U(b, \overline{\theta}) - U(b', \underline{\theta})] \le \pi_2[\gamma(\underline{\theta}) - \gamma(\overline{\theta})] - [u_1(b, \underline{\theta}) - u_1(b, \overline{\theta})]. \tag{12}$$

Because u_1 is strictly supermodular, $b \ge b'$ is required. From Assumption 2(b), that $U(\overline{a}, \overline{\theta}) > \gamma(\underline{\theta})$, and that U is strictly increasing in θ, we conclude that the left side of inequality (12) is strictly positive and bounded away from zero.[21] We also have that the first bracketed term on the right side is strictly negative.

Thus, if $|u_1(b, \underline{\theta}) - u_1(b, \overline{\theta})|$ is small relative to $\pi_2|\gamma(\underline{\theta}) - \gamma(\overline{\theta})|$, then inequality (12) fails to hold and there is no way to implement value functions that make player 2's payoff constant in the state. In other words, in the case of unified investment and trade actions, the first-best level of investment generally cannot be induced.

Optimal punishments and investment incentives

To construct the second-best contract in general, we first need to determine how to optimally punish deviations from truth-telling in the message phase (date 4). That is, we must calculate the punishment values. Although there are two implementation conditions for each pair of states, as shown above, we can focus on the one that bounds the rise in v_1 in the state. This condition corresponds to expression (11). Another way to look at this condition is to start with inequality (9), substitute for $v_2(\theta) = \gamma(\theta) - v_1(\theta)$, and rearrange terms to obtain

$$v_1(\theta) - v_1(\theta') \le \gamma(\theta) - P^{\mathrm{EP}}(\theta', \theta) \tag{13}$$

for $\theta > \theta'$. So lowering (improving) the punishment value for states (θ', θ) relaxes the constraint on how v_1 may rise with the state. Incidentally, from inequality (8), we get the corresponding bound for player 2:

$$v_2(\theta) - v_2(\theta') \le \gamma(\theta) - P^{\mathrm{EP}}(\theta, \theta').$$

These two inequalities are usefully employed to determine the limits of contracting in applications. For instance, we used the forcing-contract version (with P^{EPF}) to verify that the null contract is the optimal forcing contract in our example.

Our next step is to calculate $P^{\mathrm{EP}}(\theta', \theta)$, the minimal punishment value, for any pair of states (θ', θ) with $\theta' < \theta$. This turns out to be straightforward under our assumptions. An alternative to Assumption 4(a) follows.

---

[21]To see this, consider two cases. If $U(\underline{a}, \overline{\theta}) \ge U(\overline{a}, \overline{\theta})$, because U is strictly quasiconcave in a, every point on the graph of $U(\cdot, \overline{\theta})$ is above every point on the graph of $U(\cdot, \underline{\theta})$ and so the result is immediate. If $U(\underline{a}, \overline{\theta}) < U(\overline{a}, \overline{\theta})$, U strictly increasing in θ implies that the result holds over the range $[\underline{a}, a^*(\overline{\theta})]$. Over the range $[a^*(\overline{\theta}), \overline{a}]$, the problem reduces to the first case.

ASSUMPTION 4. *(a′) For all $\theta \in (\underline{\theta}, \overline{\theta}]$, $U(\underline{a}, \theta) \geq U(\overline{a}, \theta)$.*

LEMMA 2. *Under Assumptions 2, 3, and either Assumption 4(a) or 4(a′), for any pair of states (θ', θ) with $\theta' < \theta$, the optimal punishment involves inducing player 1 to select $a^*(\theta')$ in state θ' and $\overline{a}$ in state θ. That is, $P^{\mathrm{EP}}(\theta', \theta) = \lambda(a^*(\theta'), \overline{a}, \theta', \theta)$.*

PROOF. From (7), the punishment value for (θ', θ), $\lambda(a', a, \theta', \theta)$, is given by

$$u_1(a, \theta') + \pi_1 R(a', \theta') + u_2(a, \theta) + \pi_2 R(a, \theta),$$

which can be rewritten as

$$u_1(a, \theta') + \pi_1 U(a^*(\theta'), \theta') - \pi_1 U(a', \theta') + \pi_1 u_2(a, \theta) + \pi_2 U(a^*(\theta), \theta) - \pi_2 u_1(a, \theta).$$

The optimal punishment value is obtained by choosing a and a' to minimize this objective function under the constraint that $a \geq a'$ (because of supermodularity of u_1). Ignoring the constant terms that do not contain a or a', and substituting $\pi_2 = (1 - \pi_1)$ and $\pi_1 u_1(a, \theta) + \pi_1 u_2(a, \theta) = \pi_1 U(a, \theta)$, the objective function becomes

$$-[u_1(a, \theta) - u_1(a, \theta')] + \pi_1 U(a, \theta) - \pi_1 U(a', \theta'). \tag{14}$$

The number a' affects only the last term; to minimize it (that is, maximize $U(a', \theta')$) without consideration of the constraint $a \geq a'$, it is optimal to set $a' = a^*(\theta')$. From supermodularity of u_1, the negative bracketed term is minimized by choosing $a = \overline{a}$. By Assumption 4(a′), the final term is also minimized at $\overline{a}$. Thus, $\lambda(a', a, \theta', \theta)$ attains its lowest value when $a = \overline{a}$ and $a' = a^*(\theta')$.

To see that the same result holds with Assumption 4(a) in place of Assumption 4(a′), observe that if the minimizing values a and a' satisfy $a > a'$ then it must be that $a = \overline{a}$ and $a' = a^*(\theta')$. That $a' = a^*(\theta')$ is an implication of strict quasiconcavity of $U(\cdot, \theta')$, for if $a' < a$ and $a' \neq a^*(\theta')$, then it must be that $a^*(\theta') > a$, but then raising a' to a strictly increases $U(a', \theta')$. The conclusion that $a = \overline{a}$ follows from strict quasiconcavity of $U(\cdot, \theta)$ and from supermodularity of u_1, which imply that it is not optimal to set $a \in (a', \overline{a})$.

So we know that either it is optimal to have $a = \overline{a}$ and $a' = a^*(\theta')$, or $a = a'$ is optimal. In the latter case, Assumption 4(a) implies that $a = a' = \overline{a}$ is best. This is apparent by rearranging terms to show that, by substituting $a = a'$ into expression (14), the expression becomes

$$\pi_1[u_2(a, \theta) - u_2(a, \theta')] - \pi_2[u_1(a, \theta) - u_1(a, \theta')],$$

which is decreasing in a. This yields a contradiction because we can lower a' to $a^*(\theta')$ to strictly decrease the objective function.   □

A note about inside and outside constraints on value functions

Lemma 2 allows us to easily calculate the lowest possible punishment values for any unilateral deviation from truth-telling, and with this in hand we can begin to evaluate how

the conditions for implementability come together to constrain the value function. For the unified case, we want to know whether we can implement a value function so that v_1 increases at the same rate as does γ. One way to get at this is to examine constraints on $v_1(\theta) - v_1(\theta')$ at the margin where θ and θ' are very close, and then chain together these inside conditions to characterize the optimal implementable value function.

Unfortunately, there are also outside conditions to examine; they give constraints on $v_1(\theta) - v_1(\theta')$ for θ and θ' that are far apart. We demonstrate that the outside conditions are typically tighter than the sum of the inside conditions, so a triangle inequality fails. Thus, one cannot rely on marginal analysis to calculate bounds on implementable value functions (a cautionary note relative to Segal and Whinston 2002).

Consider any three states satisfying $\theta^{\mathrm{L}} < \theta^{\mathrm{M}} < \theta^{\mathrm{H}}$, such that the optimal trade action in state θ^{M} is interior so that $a^*(\theta^{\mathrm{M}}) \in (\underline{a}, \overline{a})$. Also assume that the assumptions for Lemma 2 hold. Using inequality (13), we have three necessary conditions for implementation:

$$v_1(\theta^{\mathrm{M}}) - v_1(\theta^{\mathrm{L}}) \leq \gamma(\theta^{\mathrm{M}}) - P^{\mathrm{EP}}(\theta^{\mathrm{L}}, \theta^{\mathrm{M}})$$

$$v_1(\theta^{\mathrm{H}}) - v_1(\theta^{\mathrm{M}}) \leq \gamma(\theta^{\mathrm{H}}) - P^{\mathrm{EP}}(\theta^{\mathrm{M}}, \theta^{\mathrm{H}})$$

$$v_1(\theta^{\mathrm{H}}) - v_1(\theta^{\mathrm{L}}) \leq \gamma(\theta^{\mathrm{H}}) - P^{\mathrm{EP}}(\theta^{\mathrm{L}}, \theta^{\mathrm{H}}). \tag{15}$$

Summing the first two yields

$$v_1(\theta^{\mathrm{H}}) - v_1(\theta^{\mathrm{L}}) \leq \gamma(\theta^{\mathrm{H}}) - P^{\mathrm{EP}}(\theta^{\mathrm{M}}, \theta^{\mathrm{H}}) + \gamma(\theta^{\mathrm{M}}) - P^{\mathrm{EP}}(\theta^{\mathrm{L}}, \theta^{\mathrm{M}}). \tag{16}$$

We want to know whether (16) is a weakly tighter bound than is (15), which would mean that the inside conditions imply the outside conditions and allow implementability to be characterized by marginal analysis. So we must establish whether the following triangle inequality holds:

$$P^{\mathrm{EP}}(\theta^{\mathrm{L}}, \theta^{\mathrm{H}}) + \gamma(\theta^{\mathrm{M}}) \leq P^{\mathrm{EP}}(\theta^{\mathrm{L}}, \theta^{\mathrm{M}}) + P^{\mathrm{EP}}(\theta^{\mathrm{M}}, \theta^{\mathrm{H}}).$$

Expanding terms using expression (7) and Lemma 2, we get

$$u_1(\overline{a}, \theta^{\mathrm{L}}) + u_2(\overline{a}, \theta^{\mathrm{H}}) + \pi_2 R(\overline{a}, \theta^{\mathrm{H}}) + \gamma(\theta^{\mathrm{M}})$$

$$\leq u_1(\overline{a}, \theta^{\mathrm{L}}) + u_2(\overline{a}, \theta^{\mathrm{M}}) + \pi_2 R(\overline{a}, \theta^{\mathrm{M}}) + u_1(\overline{a}, \theta^{\mathrm{M}}) + u_2(\overline{a}, \theta^{\mathrm{H}}) + \pi_2 R(\overline{a}, \theta^{\mathrm{H}}),$$

which simplifies to

$$\gamma(\theta^{\mathrm{M}}) \leq u_2(\overline{a}, \theta^{\mathrm{M}}) + \pi_2 R(\overline{a}, \theta^{\mathrm{M}}) + u_1(\overline{a}, \theta^{\mathrm{M}}).$$

This is equivalent to

$$\gamma(\theta^{\mathrm{M}}) \leq U(\overline{a}, \theta^{\mathrm{M}}) + \pi_2 \gamma(\theta^{\mathrm{M}}) - \pi_2 U(\overline{a}, \theta^{\mathrm{M}}),$$

which simplifies to $\gamma(\theta^{\mathrm{M}}) \leq U(\overline{a}, \theta^{\mathrm{M}})$. This inequality, coupled with Assumption 2(b), requires that $a^*(\theta^{\mathrm{M}}) = \overline{a}$, which contradicts what we assumed earlier.

Proof of Theorem 2

Consider the following contract: In the message phase (date 4), player 2 must declare the state. Let $\hat{\theta}$ denote player 2's announcement. If $\hat{\theta} \geq \beta$, then player 1 is forced to select $\overline{a}$ at date 6. Otherwise, player 1 is forced to choose between $a^*(\hat{\theta})$ and $\overline{a}$. In this case, if player 1 selects action $a^*(\hat{\theta})$, then the enforcer compels a transfer of

$$\hat{t} = \big(u_1(\overline{a}, \hat{\theta}) - u_1(a^*(\hat{\theta}), \hat{\theta}), u_1(a^*(\hat{\theta}), \hat{\theta}) - u_1(\overline{a}, \hat{\theta})\big), \tag{17}$$

and if player 1 selects action $\overline{a}$, then the transfer is $\overline{t} = (0, 0)$. The forcing arrangement is achieved by specifying a transfer of $(-\tau, \tau)$ if player 1 picks any other trade action, where τ is set large enough to keep him from doing so.

We show that this contract implements the value function v^β defined in the text. First note that, in any state θ, if at date 4 player 2 declares the state to be $\hat{\theta} \in [\underline{\theta}, \beta)$ and the players do not renegotiate at date 5, then player 1 obtains at least $u_1(\overline{a}, \theta)$ because he has the option of choosing $\overline{a}$ with no transfer. Further, if player 1 selects $a^*(\hat{\theta})$, then (from expression (17)) he gets

$$u_1(a^*(\hat{\theta}), \theta) + u_1(\overline{a}, \hat{\theta}) - u_1(a^*(\hat{\theta}), \hat{\theta}),$$

whereas he gets $u_1(\overline{a}, \theta)$ by choosing $\overline{a}$. The latter payoff weakly exceeds the former if and only if

$$u_1(\overline{a}, \theta) - u_1(a^*(\hat{\theta}), \theta) \geq u_1(\overline{a}, \hat{\theta}) - u_1(a^*(\hat{\theta}), \hat{\theta}).$$

From the supermodularity of u_1 and given that $\overline{a} \geq a^*(\hat{\theta})$, we know that it is rational for player 1 to choose $\overline{a}$ in the case of $\hat{\theta} < \theta$ and it is rational for player 1 to choose $a^*(\hat{\theta})$ in the case of $\hat{\theta} > \theta$. Player 1 is indifferent if $\hat{\theta} = \theta$.

We can therefore prescribe the following behavior, for any state θ.

- If player 2 declares $\hat{\theta} \in [\theta, \beta)$, then, absent renegotiation, player 1 chooses $a^*(\hat{\theta})$ at date 6.

- For any other message (either $\hat{\theta} \geq \beta$ or $\hat{\theta} < \theta$), absent renegotiation, player 1 chooses $\overline{a}$.

It is clear that, given player 1's behavior just specified, it is optimal for player 2 to report truthfully at date 4. For instance, in a state $\theta < \beta$, if player 2 reports honestly, then there is no renegotiation and she gets $\gamma(\theta) - u_1(\overline{a}, \theta)$. If she reports a different state, then player 1 is expected to take an ex post inefficient trade action that gives him at least $u_1(\overline{a}, \theta)$, so player 2 fares less well.

With the specified behavior for the players, in any state $\theta < \beta$ there is no renegotiation and player 1 obtains the payoff $u_1(\overline{a}, \theta)$. In any state $\theta \geq \beta$, the players renegotiate away from the anticipated action of $\overline{a}$ and player 1 gets $u_1(\overline{a}, \theta) + \pi_1 R(\overline{a}, \theta)$. Thus, value function v^β is implemented. $\qquad\square$

Proof of Proposition 2

Consider any two states θ' and θ with $\theta' < \theta$. Using Lemma 2 for the pairing (θ', θ), which shows that the minimum punishment value for such a pair involves inducing player 1 to select $a^*(\theta')$ in state θ' and $\overline{a}$ in state θ, we have that

$$P^{\mathrm{EP}}(\theta', \theta) = u_1(\overline{a}, \theta') + u_2(\overline{a}, \theta) + \pi_2 R(\overline{a}, \theta).$$

We use this equality to substitute for P^{EP} in the necessary condition (13). Rearranging terms yields the following upper bound on the value difference between the two states for player 1:

$$v_1(\theta) - v_1(\theta') \le \pi_1 R(\overline{a}, \theta) + u_1(\overline{a}, \theta) - u_1(\overline{a}, \theta'). \tag{18}$$

Consider any contract that supports the best achievable investment level θ^{B} and let v^{B} be the implemented value function. By definition of θ^{B}, we know that θ^{B} solves player 1's investment problem of maximizing $v_1^{\mathrm{B}}(\theta) - \theta$.

Define $\beta \equiv \theta^{\mathrm{B}}$ and consider the value function v^{β} defined in Theorem 2. To see that v^{β} supports θ^{B} (that is, θ^{B} maximizes $v_1^{\beta}(\theta) - \theta$), first observe that for any state $\theta < \theta^{\mathrm{B}}$, we have

$$v_1^{\beta}(\theta^{\mathrm{B}}) - v_1^{\beta}(\theta) = u_1(\overline{a}, \theta^{\mathrm{B}}) + \pi_1 R(\overline{a}, \theta^{\mathrm{B}}) - u_1(\overline{a}, \theta),$$

so v^{β} meets the upper bound on player 1's payoff difference between θ and θ^{B}, as identified in inequality (18). Since v^{B} also must satisfy the bound (18), we conclude that $v_1^{\beta}(\theta^{\mathrm{B}}) - v_1^{\beta}(\theta) \ge v_1^{\mathrm{B}}(\theta^{\mathrm{B}}) - v_1^{\mathrm{B}}(\theta)$ for all $\theta < \theta^{\mathrm{B}}$. This implies that the maximizer of $v_1^{\beta}(\theta) - \theta$ must be no less than θ^{B}.

The final step is to consider the implications of θ^{B} not maximizing $v_1^{\beta}(\theta) - \theta$. In this case, let $\tilde{\theta} > \theta^{\mathrm{B}}$ denote an investment that player 1 strictly prefers. We then have $v_1^{\beta}(\tilde{\theta}) - \tilde{\theta} > v_1^{\beta}(\theta^{\mathrm{B}}) - \theta^{\mathrm{B}}$. Plugging in the implemented values of v_1^{β}, this is equivalent to

$$u_1(\overline{a}, \tilde{\theta}) + \pi_1 R(\overline{a}, \tilde{\theta}) - \tilde{\theta} > u_1(\overline{a}, \theta^{\mathrm{B}}) + \pi_1 R(\overline{a}, \theta^{\mathrm{B}}) - \theta^{\mathrm{B}}.$$

Rearranging terms, we see that this is equivalent to

$$\pi_1 \gamma(\tilde{\theta}) - \tilde{\theta} - [\pi_1 \gamma(\theta^{\mathrm{B}}) - \theta^{\mathrm{B}}] > \pi_2 U(\overline{a}, \theta^{\mathrm{B}}) - \pi_2 U(\overline{a}, \tilde{\theta}) + u_2(\overline{a}, \tilde{\theta}) - u_2(\overline{a}, \theta^{\mathrm{B}}).$$

It is not difficult to verify that Assumption 4(b) implies that the expression on the right side is weakly positive and thus the expression on the left side is strictly positive. This further implies that $\gamma(\tilde{\theta}) > \gamma(\theta^{\mathrm{B}})$, which means that

$$\gamma(\tilde{\theta}) - \tilde{\theta} - [\gamma(\theta^{\mathrm{B}}) - \theta^{\mathrm{B}}] > 0,$$

contradicting that θ^{B} is the best achievable investment level. Thus, we know that θ^{B} maximizes $v_1^{\beta}(\theta) - \theta$.   □

Proof of Proposition 3

If θ^* is supported, then Proposition 2 implies that it is supported by the value function v^{θ^*} from Theorem 2. We then know that θ^* maximizes $v_1^{\theta^*}(\theta) - \theta$ by choice of θ, which implies that inequality (2) holds for all $\theta < \theta^*$. Thus, the condition of the proposition is necessary. Sufficiency requires not only that inequality (2) hold for all $\theta < \theta^*$, but also that player 1 prefer not to invest $\tilde{\theta} > \theta^*$ when v^{θ^*} is the implemented value function. This follows from the argument in the final paragraph of the proof of Proposition 2, replacing θ^{B} with θ^* and "best achievable" with "efficient." □

APPENDIX C: PROOF OF THEOREM 3

In this appendix, we complete the proof of Theorem 3. We start with the comparison of V^{EPF} and V^{EP} and then provide the analysis for the comparison of V^{EP} and V^{I}.

Completion of the proof that $V^{\mathrm{EPF}} \neq V^{\mathrm{EP}}$

We pick up from the analysis at the end of Section 5. Consider a pair of states θ^1 and θ^2 that satisfies Assumption 5. That is, we have $\theta^1 > \theta^2$ and either $U(\underline{a}, \theta^2) < U(\overline{a}, \theta^2)$ or $U(\underline{a}, \theta^1) > U(\overline{a}, \theta^1)$. Let b^1 denote a solution to the forcing-contract problem

$$\min_{a \in A} \lambda(a, a, \theta^1, \theta^2)$$

and let b^2 denote a solution to the forcing-contract problem

$$\min_{a \in A} \lambda(a, a, \theta^2, \theta^1).$$

It is easy to show that $b^1 \geq a^*(\theta^1) > \underline{a}$ and $b^2 \leq a^*(\theta^2) < \overline{a}$ follow from Assumptions 2(b) and 5. We use these facts below. We demonstrate that either $P^{\mathrm{EP}}(\theta^1, \theta^2) < P^{\mathrm{EPF}}(\theta^1, \theta^2)$ or $P^{\mathrm{EP}}(\theta^2, \theta^1) < P^{\mathrm{EPF}}(\theta^2, \theta^1)$ or both, which implies that $V^{\mathrm{EPF}} \neq V^{\mathrm{EP}}$.

Let us evaluate the minimum punishment value that corresponds to the ordered pair of states (θ^1, θ^2). Specifically, compare the optimal forcing-contract punishment (forcing player 1 to select b^1 in both states) with a nonforcing specification in which player 1 is induced to select b^1 in state θ^1 and $\underline{a}$ in state θ^2. This is a valid nonforcing contractual specification because, by Fact 2, $\theta^1 > \theta^2$ and $b^1 > \underline{a}$ imply $(b^1, \underline{a}) \in E(\theta^1, \theta^2)$.

If $V^{\mathrm{EP}} = V^{\mathrm{EPF}}$, then it must be that $\lambda(b^1, b^1, \theta^1, \theta^2) \leq \lambda(b^1, \underline{a}, \theta^1, \theta^2)$. Applying the definition of λ, this is

$$u_1(b^1, \theta^1) + \pi_1 R(b^1, \theta^1) + u_2(b^1, \theta^2) + \pi_2 R(b^1, \theta^2)$$
$$\leq u_1(\underline{a}, \theta^1) + \pi_1 R(b^1, \theta^1) + u_2(\underline{a}, \theta^2) + \pi_2 R(\underline{a}, \theta^2).$$

Canceling the second term on each side and using the definition of R, we arrive at

$$u_1(b^1, \theta^1) + u_2(b^1, \theta^2) - \pi_2 U(b^1, \theta^2) \leq u_1(\underline{a}, \theta^1) + u_2(\underline{a}, \theta^2) - \pi_2 U(\underline{a}, \theta^2).$$

Substituting $u_2(\cdot, \theta^2) = U(\cdot, \theta^2) - u_1(\cdot, \theta^2)$ on both sides, we have

$$u_1(b^1, \theta^1) + U(b^1, \theta^2) - u_1(b^1, \theta^2) - \pi_2 U(b^1, \theta^2)$$
$$\leq u_1(\underline{a}, \theta^1) + U(\underline{a}, \theta^2) - u_1(\underline{a}, \theta^2) - \pi_2 U(\underline{a}, \theta^2).$$

Finally, rearranging this expression a bit and using $\pi_1 + \pi_2 = 1$, we conclude that $\lambda(b^1, b^1, \theta^1, \theta^2) \leq \lambda(b^1, \underline{a}, \theta^1, \theta^2)$ is equivalent to

$$u_1(b^1, \theta^1) - u_1(\underline{a}, \theta^1) - [u_1(b^1, \theta^2) - u_1(\underline{a}, \theta^2)] \leq \pi_1[U(\underline{a}, \theta^2) - U(b^1, \theta^2)]. \qquad (23)$$

Similarly, ordering states θ^1 and θ^2 in the opposite way, we compare the optimal forcing-contract punishment (forcing player 1 to select b^2 in both states) with a non-forcing specification in which player 1 is induced to select b^2 in state θ^2 and $\overline{a}$ in state θ^1. Note that $\theta^2 < \theta^1$ and $b^2 < \overline{a}$ imply $(b^2, \overline{a}) \in E(\theta^2, \theta^1)$. If $V^{\mathrm{EP}} = V^{\mathrm{EPF}}$, then it must be that $\lambda(b^2, b^2, \theta^2, \theta^1) \leq \lambda(b^2, \overline{a}, \theta^2, \theta^1)$, which similar algebraic manipulation reveals to be equivalent to

$$u_1(\overline{a}, \theta^1) - u_1(b^2, \theta^1) - [u_1(\overline{a}, \theta^2) - u_1(b^2, \theta^2)] \leq \pi_1[U(\overline{a}, \theta^1) - U(b^2, \theta^1)]. \qquad (24)$$

The foregoing analysis shows that if $V^{\mathrm{EPF}} = V^{\mathrm{EP}}$, then expressions (23) and (24) hold. Assumption 3 then implies that the left sides of these inequalities are nonnegative, which implies

$$U(\underline{a}, \theta^2) \geq U(b^1, \theta^2) \quad \text{and} \quad U(\overline{a}, \theta^1) \geq U(b^2, \theta^1).$$

Using Assumption 2(b), and that $b^1 > \underline{a}$ and $b^2 < \overline{a}$, we obtain the following fact.

FACT 7. *If $V^{\mathrm{EPF}} = V^{\mathrm{EP}}$, then $U(\underline{a}, \theta^2) \geq U(\overline{a}, \theta^2)$ and $U(\overline{a}, \theta^1) \geq U(\underline{a}, \theta^1)$.*

Assumption 5 and the contrapositive of Fact 7 provide the contradiction that proves $V^{\mathrm{EPF}} \neq V^{\mathrm{EP}}$.

Proof that $V^{\mathrm{EP}} \neq V^{\mathrm{I}}$

We next prove the claim about the relation between V^{I} and V^{EP}. Since forcing contracts are sufficient to construct V^{I}, we can state the following fact.

FACT 8. *The minimum punishment value in the setting of interim renegotiation is characterized as*

$$P^{\mathrm{I}}(\theta, \theta') = \min_{a'' \in A} u_1(a'', \theta) + u_2(a'', \theta').$$

Remember that, by Result 2, $V^{\mathrm{I}} = V^{\mathrm{EP}}$ if and only if $P^{\mathrm{EP}}(\theta, \theta') = P^{\mathrm{I}}(\theta, \theta')$ for all $\theta, \theta' \in \Theta$. We can again compare the minimization problems to determine if this is the case.

Take θ^1 and θ^2 as satisfying Assumption 5. Consider any solution to the minimization problem that defines $P^{\mathrm{EP}}(\theta^1, \theta^2)$ and denote it (b, b'). That is, (b, b') solves

$$\min_{(a,a')\in E(\theta^1,\theta^2)} u_1(a', \theta^1) + \pi_1 R(a, \theta^1) + u_2(a', \theta^2) + \pi_2 R(a', \theta^2).$$

Then $P^{\mathrm{EP}}(\theta^1, \theta^2) = P^{\mathrm{I}}(\theta^1, \theta^2)$ is equivalent to

$$u_1(b', \theta^1) + \pi_1 R(b, \theta^1) + u_2(b', \theta^2) + \pi_2 R(b', \theta^2) = \min_{a''\in A} u_1(a'', \theta^1) + u_2(a'', \theta^2).$$

Because $R(\cdot, \cdot) \geq 0$, we see that $P^{\mathrm{EP}}(\theta^1, \theta^2) = P^{\mathrm{I}}(\theta^1, \theta^2)$ only if b' solves the minimization problem on the right side of the above equation and also $R(b, \theta^1) = R(b', \theta^2) = 0$.

By Assumption 2(b), $R(b', \theta^2) = 0$ if and only if $b' = a^*(\theta^2)$. Combining this with the requirement that b' must minimize $u_1(\cdot, \theta^1) + u_2(\cdot, \theta^2)$, we derive that

$$u_1(a^*(\theta^2), \theta^1) + u_2(a^*(\theta^2), \theta^2) \leq u_1(a'', \theta^1) + u_2(a'', \theta^2)$$

for all a''. In particular, the following inequality must hold:

$$u_1(a^*(\theta^2), \theta^1) + u_2(a^*(\theta^2), \theta^2) \leq u_1(\underline{a}, \theta^1) + u_2(\underline{a}, \theta^2).$$

Using the identity $u_2 = U - u_1$ and rearranging terms, we see that this is equivalent to

$$\begin{aligned} u_1(a^*(\theta^2), \theta^1) - u_1(\underline{a}, \theta^1) &- \left[u_1(a^*(\theta^2), \theta^2) - u_1(\underline{a}, \theta^2) \right] \\ &\leq U(\underline{a}, \theta^2) - U(a^*(\theta^2), \theta^2). \end{aligned} \tag{25}$$

Similarly, ordering states θ^1 and θ^2 in the opposite way, it is necessary that $a^*(\theta^1)$ must solve $P^{\mathrm{I}}(\theta^2, \theta^1)$ so that $P^{\mathrm{EP}}(\theta^2, \theta^1) = P^{\mathrm{I}}(\theta^2, \theta^1)$. In particular, we must have

$$u_1(a^*(\theta^1), \theta^2) + u_2(a^*(\theta^1), \theta^1) \leq u_1(\overline{a}, \theta^2) + u_2(\overline{a}, \theta^1).$$

This inequality is equivalent to

$$\begin{aligned} u_1(\overline{a}, \theta^1) - u_1(a^*(\theta^1), \theta^1) &- \left[u_1(\overline{a}, \theta^2) - u_1(a^*(\theta^1), \theta^2) \right] \\ &\leq U(\overline{a}, \theta^1) - U(a^*(\theta^1), \theta^1). \end{aligned} \tag{26}$$

By Assumption 3, the left sides of expressions (25) and (26) must be nonnegative, which implies both $U(\underline{a}, \theta^2) \geq U(a^*(\theta^2), \theta^2)$ and $U(\overline{a}, \theta^1) \geq U(a^*(\theta^1), \theta^1)$. From Assumption 2(b), we see that this is only possible if $\underline{a} = a^*(\theta^2)$ and $\overline{a} = a^*(\theta^1)$. If this is the case, Assumption 2(b) also implies that $U(\underline{a}, \theta^2) \geq U(\overline{a}, \theta^2)$ and $U(\overline{a}, \theta^1) \geq U(\underline{a}, \theta^1)$. Thus we obtain our last fact.

FACT 9. If $V^{\mathrm{I}} = V^{\mathrm{EP}}$, then $U(\underline{a}, \theta^2) \geq U(\overline{a}, \theta^2)$ and $U(\overline{a}, \theta^1) \geq U(\underline{a}, \theta^1)$.

The contrapositive of Fact 9 combined with Assumption 5 provides the contradiction that proves $V^{\mathrm{I}} \neq V^{\mathrm{EP}}$. $\qquad\square$

References

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey (1994), "Renegotiation design with unverifiable information." *Econometrica*, 62, 257–282. [290, 291, 294]

Baliga, Sandeep and Tomas Sjöström (2009), "Contracting with third parties." *American Economic Journal: Microeconomics*, 1, 75–100. [308]

Beaudry, Paul and Michel Poitevin (1995), "Contract renegotiation: A simple framework and implications for organization theory." *Canadian Journal of Economics*, 28, 302–335. [290]

Boeckem, Sabine and Ulf Schiller (2008), "Option contracts in supply chains." *Journal of Economics and Management Strategy*, 17, 219–245. [290, 308]

Bull, Jesse (2009), "Third-party budget breakers and side contracting in team production." Unpublished paper, Florida International University. [308]

Che, Yeon Koo and Donald B. Hausch (1999), "Cooperative investments and the value of contracting." *American Economic Review*, 89, 125–147. [285, 290, 291, 292, 307]

Che, Yeon Koo and József Sákovics (2008), "Hold-up problem." In *The New Palgrave Dictionary of Economics* (Lawrence E. Bloom and Steven N. Durlauf, eds.), Palgrave Macmillan, Hampshire, UK. [283]

Chung, Tai Yeong (1991), "Incomplete contracts, specific investments, and risk sharing." *Review of Economic Studies*, 58, 1031–1042. [291]

De Fraja, Gianni (1999), "After you sir. Hold-up, direct externalities, and sequential investment." *Games and Economic Behavior*, 26, 22–39. [309]

Demski, Joel S. and David E. M. Sappington (1991), "Resolving double moral hazard problems with buyout agreements." *RAND Journal of Economics*, 22, 232–240. [290, 292]

Edlin, Aaron S. and Benjamin E. Hermalin (2000), "Contract renegotiation and options in agency problems." *Journal of Law, Economics, and Organization*, 16, 395–423. [290, 307]

Edlin, Aaron S. and Stefan Reichelstein (1996), "Holdups, standard breach remedies, and optimal investment." *American Economic Review*, 86, 478–501. [290, 291]

Ellman, Matthew (2006), "Specificity revisited: The role of cross-investments." *Journal of Law, Economics, and Organization*, 22, 234–257. [290, 308]

Evans, Robert (2006), "Mechanism design with renegotiation and costly messages." Unpublished paper, University of Cambridge. [290]

Evans, Robert (2008), "Simple efficient contracts in complex environments." *Econometrica*, 76, 459–491. [290, 307, 308, 309]

Grossman, Sanford J. and Oliver D. Hart (1986), "The costs and benefits of ownership: A theory of vertical and lateral integration." *Journal of Political Economy*, 94, 691–719. [283]

Grout, Paul A. (1984), "Investment and wages in the absence of binding contracts: A Nash bargaining approach." *Econometrica*, 52, 449–460. [283]

Guriev, Sergei (2003), "Incomplete contracts with cross-investments." *Contributions in Theoretical Economics*, 3. [290]

Hart, Oliver D. and John H. Moore (1988), "Incomplete contracts and renegotiation." *Econometrica*, 56, 755–785. [283, 290, 292, 309]

Hart, Oliver D. and John H. Moore (1999), "Foundations of incomplete contracts." *Review of Economic Studies*, 66, 115–138. [290, 291, 308]

Holmstrom, Bengt (1982), "Moral hazard in teams." *Bell Journal of Economics*, 13, 324–340. [308]

Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian (1978), "Vertical integration, appropriable rents, and the competitive contracting process." *Journal of Law and Economics*, 21, 297–326. [283]

Lyon, Thomas P. and Eric Rasmusen (2004), "Buyer-option contracts restored: Renegotiation, inefficient threats, and the hold-up problem." *Journal of Law, Economics, and Organization*, 20, 148–169. [290]

MacLeod, W. Bentley and James M. Malcomson (1993), "Investments, holdup, and the form of market contracts." *American Economic Review*, 83, 811–837. [290, 292, 294, 309]

Maskin, Eric S. and John H. Moore (1999), "Implementation and renegotiation." *Review of Economic Studies*, 66, 39–56. [284, 290]

Maskin, Eric S. and Jean Tirole (1999), "Two remarks on the property-rights literature." *Review of Economic Studies*, 66, 139–149. [290]

Myerson, Roger B. (1982), "Optimal coordination mechanisms in generalized principal–agent problems." *Journal of Mathematical Economics*, 10, 67–81. [290]

Myerson, Roger B. (1991), *Game Theory*. Harvard University Press, Cambridge, Massachusetts. [290]

Nöldeke, Georg and Klaus M. Schmidt (1995), "Option contracts and renegotiation: A solution to the hold-up problem." *RAND Journal of Economics*, 26, 163–179. [290, 291, 292]

Nöldeke, Georg and Klaus M. Schmidt (1998), "Sequential investments and options to own." *RAND Journal of Economics*, 29, 633–653. [290]

Reiche, Sönje (2006), "Ambivalent investment and the hold-up problem." *Journal of the European Economic Association*, 4, 1148–1164. [291, 308]

Rogerson, William P. (1992), "Contractual solutions to the hold-up problem." *Review of Economic Studies*, 59, 777–793. [291]

Roider, Andreas (2004), "Asset ownership and contractibility of interaction." *RAND Journal of Economics*, 35, 787–802. [290]

Segal, Ilya (1999), "Complexity and renegotiation: A foundation for incomplete contracts." *Review of Economic Studies*, 66, 57–82. [290, 291, 308]

Segal, Ilya and Michael D. Whinston (2002), "The Mirrlees approach to mechanism design with renegotiation (with applications to hold-up and risk sharing)." *Econometrica*, 70, 1–45. [284, 290, 292, 314]

Stremitzer, Alexander (forthcoming), "Standard breach remedies, quality thresholds, and cooperative investments." *Journal of Law, Economics, and Organization*. [290]

Watson, Joel (2005), "Contract and mechanism design in settings with multi-period trade." Unpublished paper, University of California, San Diego. [303]

Watson, Joel (2006), "Contract and game theory: Basic concepts for settings with finite horizons." Unpublished paper, University of California, San Diego. [288]

Watson, Joel (2007), "Contract, mechanism design, and technological detail." *Econometrica*, 75, 55–81. [284, 285, 286, 289, 290, 292, 300, 302, 303, 307]

Watson, Joel and Chris Wignall (2009), "Hold-up and durable trading opportunities." Unpublished paper, University of California, San Diego. [307]

Williamson, Oliver E. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press, New York. [283]

Williamson, Oliver E. (1979), "Transaction cost economics: The governance of contractual relations." *Journal of Law and Economics*, 22, 233–261. [283]