

FOUNDATIONS OF INCOMPLETE CONTRACTS*

by

Oliver Hart

Harvard University and London School of Economics

and

John Moore

London School of Economics and University of St Andrews

Contents:

Abstract

1. Introduction

2. A Model of Incomplete Contracts

3. Commitment

4. Property Rights

5. Interpreting Contractual Incompleteness

Appendix

References

Theoretical Economics Workshop
The Suntory Centre
Suntory and Toyota International Centres
for Economics and Related Disciplines
London School of Economics and Political
Science
Houghton Street
London WC2A 2AE
Tel.: 0171-405 7686

Discussion Paper
No. TE/98/358
September 1998

* We are grateful to Douglas Baird, Martin Hellwig, Eric Maskin, Andy Postlewaite, Ilya Segal, and Jean Tirole for helpful discussions on this topic. We also thank Philippe Aghion, Patrick Bolton, Malcolm Baker, V Bhaskar, Mathias Dewatripont, Bentley MacLeod, Leeat Yariv, and an anonymous referee for comments on an earlier draft. We acknowledge financial support from the US national Science Foundation, the LSE Financial Markets Group, and the Suntory and Toyota International Centres for Economics and Related Disciplines at the LSE.

Abstract

In the last few years, a new area has emerged in economic theory which goes under the heading of 'incomplete contracting'. However, almost since its inception the theory has been under attack for its lack of rigorous foundations. In this paper we evaluate some of the criticisms that have been made of the theory, in particular, those in Maskin and Tirole (1998a). In doing so, we develop a model that provides a rigorous foundation for the idea that contracts are incomplete.

Keywords: Incomplete contracting; rigorous foundations; economic theory.

JEL No.: D23

© the authors. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Contact address: Professor John Hardman Moore, Department of Economics, London School of Economics, Houghton Street, London WC2A 2AE.

1. Introduction

In the last ten to fifteen years, a new area has emerged in economic theory, which goes under the heading of "incomplete contracting".¹ This approach has been useful for understanding topics such as the meaning of ownership and the nature and financial structure of the firm. Yet almost since its inception, the theory has been under attack for its lack of rigorous foundations. In this paper, we evaluate some of the criticisms that have been made of incomplete contracting theory, in particular, those in the Maskin and Tirole (1998a) (henceforth MT). In doing so we develop a model that provides a rigorous foundation for the idea that contracts are incomplete.

Many papers in the incomplete contracts literature motivate the idea of contractual incompleteness as follows. Imagine a buyer, B, who requires a good (or service) from a seller, S. Suppose that the exact nature of the good is uncertain; more precisely, it depends on a state of nature which is yet to be realized. In an ideal world, the parties would write a contingent contract specifying exactly which good is to be delivered in each state. However, if the number of states is very large, such a contract would be prohibitively expensive. So instead the parties will write an incomplete contract. Then, when the state of nature is realized, they will renegotiate the contract, since at this stage they know what kind of good should be traded.

MT have criticized this informal story on the following grounds. They argue that there is a tension between the (standard) assumption made in the incomplete contracts literature that the parties are unboundedly rational, and the appeal to the transaction cost of describing the ex ante nature of trade. MT develop a number of irrelevance theorems showing that the parties can design clever contracts, involving the exchange of commonly held information via messages, that overcome the inability to describe trade in

¹The incomplete contracts literature can be seen as a development of the earlier transactions cost literature. See, e.g., Williamson (1975).

advance. They conclude that the informal justification of contractual incompleteness based on the ex ante indescribability of actions or trades is unconvincing.

In this paper we evaluate the MT critique. We argue that MT's irrelevance theorems are less damaging to the theory than might appear at first sight. Our argument has two parts. First, we develop a model, inspired by Segal (1995, 1998), in which it is costless to delineate the set of possible trades in advance ("trades are describable"), and yet where the "null contract"--the quintessentially incomplete contract--is optimal. Applied to our model, the MT irrelevance theorems say that the optimal contract without describability of trades cannot be worse than the optimal contract with describability, i.e., it must be the null contract too. However, this result in no way undermines the conclusion that the optimal contract is incomplete.

Second, moving beyond our model, one finds that what we consider to be MT's potentially most important result, their Theorem 4, requires quite restrictive assumptions. If these assumptions are relaxed, describability does matter. In fact, we provide an extension of our basic model in which the optimal contract with describability yields the first-best, whereas the optimal contract without describability is the null contract.

An important feature of our model, found also in Segal's work, is that there is no natural metric or ordering on the good to be traded; that is, it is not the case that a good can be represented by its quantity or quality (with higher quantity or quality being more valuable for the buyer). Instead, in each state of nature a particular type of good is uniquely appropriate.²

²In this respect our model differs from much of the literature on the foundations of incomplete contracts, where it is supposed that agents contract over the quantity of a homogeneous good to be traded. A paper that obtains similar results to ours, although using a different approach, is Che and Hausch (1996). In Che and Hausch's model, a good can be represented by its quality (with higher quality being better for the buyer), but there are externalities: the seller's investment affects the value of the good to the buyer, and the buyer's investment affects the seller's cost.

Although our model is closely related to that of Segal (1995, 1998), our formulation is somewhat different, and this greatly simplifies the analysis and proofs of the theorems.

We should emphasize that we are not wedded to the idea that contracts are totally incomplete. We also provide a simple extension of the basic model in which the optimal contract is partially incomplete. Moreover, we find in this extension that describability again matters: the degree of partial incompleteness depends on the parties' ability to describe the nature of trade.

Our conclusions rely heavily on the assumption that parties to a contract are unable to commit not to renegotiate their contract (and also to some extent on the assumption that they cannot commit not to collude with a third party). In contrast, MT take the point of view that, at least in an ideal world, commitment should be possible. Indeed, much of their analysis pertains to this case (see their Theorems 1 and 2). We examine their reasoning in some detail but do not find their arguments entirely convincing. Ultimately, however, our view is that the degree of commitment is something about which reasonable people can disagree: both cases--where there is and where there is not commitment--are worthy of study.

The paper is organized as follows. In Section 2 we discuss MT's notion of describability and their irrelevance theorems, using as a vehicle a simple buyer-seller model in which the buyer and seller cannot specify in advance the efficient good to trade. Section 3 is devoted to an examination of the commitment issue. In Section 4 we present a preliminary discussion of the role of property rights (asset ownership) when contracts are incomplete. Finally, Section 5 asks to what extent the optimal contracts we derive should be interpreted as incomplete, given that we employ the mechanism-design machinery of maximizing subject to incentive constraints, and do not impose exogenous restrictions on contracting.

2. A Model of Incomplete Contracts

Consider a buyer B and a seller S, who are involved in a two-date relationship, as illustrated in Figure 1. The parties meet and contract at

date 0, and trade at date 1. In between, at date 1/2, say, one or both of the parties may invest. For expositional purposes, in the text we assume that only S invests, but we deal with the more general case in the Appendix. Both parties are risk neutral (and are not wealth-constrained) and the rate of interest is zero. Also, both parties are sufficiently rational that they can compute and maximize expected payoffs.




Figure 1

We assume that the parties trade one unit of a good, which we call a widget. To capture the idea that it is hard to contract on this good in advance, we assume that there are N different widgets. In any state of nature, exactly one of these widgets should be traded.³ We call this the "special" (or specific) widget. It yields a (monetary) value v to B and costs S a (monetary) amount \tilde{c} to produce (this cost is incurred if and only if trade takes place). Here \tilde{c} is stochastic as of date 0. For simplicity, in the text we suppose that \tilde{c} takes on two values: $\tilde{c} = c_1$ with probability $\pi(\sigma)$ and $\tilde{c} = c_2$ with probability $1 - \pi(\sigma)$, where σ represents the cost of S's date 1/2 investment. We assume that $0 \leq c_1 < c_2 < v$; that $0 < \pi(\sigma) < 1$, $\pi'(\sigma) > 0$, $\pi''(\sigma) < 0$ for all $\sigma \geq 0$; and that $\pi'(0) = \infty$. Note that, among other things, there are always gains from trade at date 1.

The other $N-1$ widgets are "generic" (or general purpose) widgets. These generic widgets have cost g_n for S, where $g_n = c_1 + \frac{n}{N}(c_2 - c_1)$, $n = 1, \dots, N-1$. In other words, the costs of the generic widgets lie evenly

³We assume that it is technologically infeasible to trade more than one widget.

between c_1 and c_2 . It will become clear that the exact specification of these costs is not important. What matters is that as the number of widgets, N , increases, no large "gaps" remain between c_1 and c_2 . Also, it makes no difference if there are other generic widgets whose costs lie outside this range; see the analysis of the general case in the Appendix.

As will be seen, the value of a generic widget to B is unimportant, provided that such a widget generates strictly less surplus than the special widget (i.e., strictly less than $v - c_2$); we assume this in what follows.

We suppose that there is complete symmetry among the widgets at date 0, in the sense that each widget is equally likely to be the special widget or to be one of the $N-1$ generic widgets. That is, there are $2N!$ possible states of nature at date 1: for each of the two possible realizations of S's cost of producing the special widget (c_1 with probability $\pi(\sigma)$; c_2 with probability $1 - \pi(\sigma)$), there are $N!$ equally likely possible permutations of how costs are allocated across widgets.⁴

We assume that both parties observe the state of nature at date 1 (including the realization of \tilde{c}). However the state is not verifiable, and nor are the parties' final payoffs: that is, these things cannot be observed by outsiders, such as the courts. In the parlance of incomplete contract theory, the state and parties' payoffs are "observable, but not

⁴The assumption that S's investment affects only the cost of the special widget, and yet S does not know which widget this is when she invests, requires some justification. We have in mind a situation where S's investment is relationship-specific rather than widget-specific. For example, S's investment might be in learning how to do business with B more efficiently. Such an investment will pay off whatever widget B needs at date 1, i.e., whichever widget is special, but will not pay off with respect to the widgets B does not need (the generic ones).

Note that, if we had assumed that B invests rather than S, then it would be quite standard to suppose that B's investment pays off only if B receives the special widget from S (in spite of the fact that B does not know in advance which the special widget is).

verifiable."⁵

We follow MT in distinguishing between two cases. One is where the widgets can be described in advance, and the other is where they cannot.

Case D: The N widgets can be costlessly described at date 0.

Case ND: It is prohibitively costly to describe the N widgets at date 0, but it is costless to describe them at date 1.

In Case D, for example, it is possible to write a specific performance contract at date 0 to the effect that S must supply a particular widget at date 1 -- although, by assumption, this widget will turn out to be the special widget only with probability 1/N. In Case ND, such a contract is infeasible.

The focus of MT's paper is on what differences, if any, there are between these polar cases, in terms of what contracts can achieve.

We begin our analysis of optimal contracting by considering the first-best. In the first-best, the special widget is always traded and the investment σ is chosen to maximize total expected surplus:

$$(2.1) \quad \underset{\sigma}{\text{maximize}} \quad \pi(\sigma)[v - c_1] + (1 - \pi(\sigma))[v - c_2] - \sigma.$$

⁵The assumption that the state of nature is unverifiable is important. If the state were verifiable, then the parties could achieve the first-best by writing a contract that specifies that S must supply the widget that is "special" in whatever state occurs, at a fixed price.

The assumption that the state of nature is observable is probably less important. We make this assumption for two reasons. First, it greatly simplifies the analysis. Second, the informal motivation of contractual incompleteness given in the introduction does not rely on an asymmetry of information between the parties.

One simple way to achieve the first-best is for the parties to write a contract that specifies the optimal value of σ ; and then rely on bargaining at date 1 to ensure that the special widget is traded. Following the literature (including MT), however, we assume that this is impossible: either σ is too complicated to describe in an enforceable way, or it is observed only by S.

2.1. Commitment

When the parties can commit not to renegotiate their contract, there is another way to achieve first-best: give S the right to make a take-it-or-leave-it offer to B at date 1. This way S captures all the date 1 surplus, and, anticipating this, makes the efficient investment decision at date 1/2. Moreover, such a contract does not require that the widgets are described at date 0, and so works in Case ND (even in Case ND, the parties can describe, and hence contract on, the special widget at date 1).

Proposition 1. Suppose Case ND holds. If the parties can commit not to renegotiate, then the first-best can be achieved.

Proof Under a take-it-or-leave-it contract, at date 1 S will ask B to pay v for the special widget, and B will agree. S's private incentive to invest is aligned with the social objective, (2.1), and the first-best is achieved.

Q.E.D.

If the first-best can be achieved when widgets cannot be described at date 0 (Case ND), then a fortiori it can also be achieved when widgets can be described (Case D).

Corollary. The conclusion of Proposition 1 also holds in Case D.

This corollary serves to illustrate MT's Theorem 1, albeit in a very particular setting. Their theorem states that, if the parties can commit not to renegotiate, then the inability to describe the nature of trade in advance does not matter. This is clearly true in Proposition 1, since there is an optimal contract that does not require S to describe the widget at date 0, but only at date 1 when she makes a take-it-or-leave-it offer.

The strength of MT's Theorem 1 lies in its general applicability to implementation problems. They have shown that in environments with complete information, where agents can commit not to renegotiate, implementation does not require the ex ante specification of a mechanism that maps messages into actions (widgets). Instead, the mechanism can be specified ex ante in terms of a mapping from messages into utilities (which are numbers). The description of actions (over which utilities are defined) can be postponed until ex post, as part of the message game. In other words, MT have shown that, with commitment, it does not matter if actions are impossible to describe ex ante, provided they can be described ex post -- as in Case ND. Unquestionably, this is a major contribution to implementation theory. However, as will become clear shortly, much depends on the assumption that agents are committed not to renegotiate the mechanism ex post.

2.2. No commitment

Proposition 1 relies heavily on the commitment assumption. The reason why B accepts S's take-it-or-leave-it offer is that if he rejects then the contract specifies "no trade", and, even though there are gains from trade still outstanding, the two parties are committed to abide by the contract.⁶

⁶The same applies were S to offer the wrong widget (one of the generic ones) and B were to accept: they precommit not to switch to the special widget.

Suppose instead that the parties cannot commit not to renegotiate.⁷ We will see that the positive conclusion of Proposition 1 is dramatically reversed.

For concreteness, let us assume that if the outcome of the contract at date 1 is inefficient and the parties renegotiate, then B has all the bargaining power. (The Appendix deals with the case of a general division of bargaining power.) Under this assumption, the take-it-or-leave-it contract of Proposition 1 works very badly, because B can reject S's offer and rely instead on bargaining to extract all the surplus, which gives S absolutely no incentive to invest at date 1/2 to reduce her expected costs. In fact, this take-it-or-leave-it contract works no better than no contract at all: under the "null contract", too, B extracts all the surplus in the date 1 bargain.

The question arises: can contracts achieve anything in these circumstances? For the moment we will assume that we are in Case D: the N widgets can be described at date 0. This puts us squarely in the world of mechanism design and classical implementation theory. That is, the parties can write into the date 0 contract a mechanism which is to be played at date 1, once they have learned the state; and, by design, the equilibrium of the mechanism will differ across states, so the contractual outcome is thereby indirectly conditioned on the state.⁸ The important difference here from most of the literature on implementation is that we are supposing that the parties are free to renegotiate the outcome once the mechanism has been played. For this class of problem, where there is renegotiation, Maskin and

⁷We assume that there is a (short) period of time after the formal provisions of the contract are carried out and before trade actually occurs; it is during this period that the parties have an incentive to renegotiate ex post (even though they would like to prevent such renegotiation ex ante). See Section 3 for a detailed discussion of commitment.

⁸See Moore (1992) for an introduction to the theory of mechanism design and implementation in environments with complete information. (Here, B and S both observe the state at date 1, and so have complete information.)

Moore (1998) provides a general characterization of the set of implementable choice rules. In what follows, although we implicitly rely on their characterization theorem, we will try to argue from first principles, so as to clarify the logic of the argument.

By "state" we mean actually two things: first, the realization of S's cost of producing the special widget; and, second, the realization of the permutation of the N widgets (i.e. the identity of the special widget, and the cost permutation of the generic widgets). Now, in terms of providing S with an incentive to reduce her expected costs, the only aspect of the state that matters is the realization of the cost of the special widget. That is, S is only concerned with the (expected) price p_1 she receives if her cost of producing the special widget is c_1 , as opposed to the (expected) price p_2 she receives if the cost is c_2 . The symmetry of the model -- the fact that all permutations are equally likely -- suggests that there can be no gain from having p_1 or p_2 depend on the realization of the permutation.⁹

S's choice of investment σ at date 1/2 is given by the solution to

$$(2.2) \quad \underset{\sigma}{\text{maximize}} \quad \pi(\sigma)[p_1 - c_1] + (1 - \pi(\sigma))[p_2 - c_2] - \sigma.$$

Notice that her private incentive to invest will be aligned with the first-best if p_1 and p_2 are equal: in that case, she enjoys the full benefit from any cost reduction. However, as the gap between the prices, $p_2 - p_1$,

⁹This assertion is justified within the formal proof of Proposition 2 below. Note that it would not be true if some permutations were more likely to occur than others. For example, suppose a particular widget were much more likely to be the special widget than any of the others. Then a specific performance contract -- the parties agree to trade that particular widget at a fixed price -- would work well. (Under the specific performance contract, p_1 would equal p_2 whenever the widget was the special widget, but not otherwise; and so the realization of the permutation would matter.)

grows, she shares any cost reduction with B, and her incentive to invest is diluted. Hence the parties' aim is to find a contractual mechanism that implements prices p_1 and p_2 for which the gap $p_2 - p_1$ is as small as possible.

As we have argued, the take-it-or-leave-it contract of Proposition 1 is useless, no better than having no contract at all. The point is that, because B has all the bargaining power at date 1, the difference between the prices, $p_2 - p_1$, ends up being equal to the difference between S's costs, $c_2 - c_1$; and from (2.2) this implies that S has no incentive to invest.

The surprising fact is that, for large N , no other contract can do significantly better!

Proposition 2. Suppose Case D holds. If the parties cannot commit not to renegotiate, then irrespective of the contract, as the number of widgets N tends to infinity, S's investment σ approaches zero. That is, in the limit, contracts cannot make any difference to expected total surplus, and the parties may as well use the null contract.

In fact, we will show that the most that can be gained from writing any contract is to reduce the price gap $p_2 - p_1$ from $c_2 - c_1$ to $\frac{N-1}{N}(c_2 - c_1)$: a reduction of the order of $O(\frac{1}{N})$.

Proposition 2 generalizes to any ex post division of bargaining power, and to a model in which both parties make investments, which may be multi-dimensional. See the Appendix.

Proposition 2 is closely related to Theorem 1 of Segal (1998), except that we dispense with his extreme widgets--what he terms "gold-plated" and "cheap imitation" widgets. Segal's construction relies on there being uncertainty over the number (and identity) of extreme widgets, and hence uncertainty over the ranking (in terms of cost) of the special widget. In our model, the generic widgets have costs similar to the special widget and, crucially, the ranking of the special widget depends only on S's investment: there is no other "aggregate uncertainty." This greatly simplifies the analysis and the proofs of the theorems.

The intuition behind Proposition 2 is as follows. Consider the problem the parties face in designing a contractual mechanism whose equilibrium depends on the realization of S's cost of producing the special widget. When this cost is c_i , $i = 1, 2$, let us say that "state i " has occurred.¹⁰ The parties are in effect playing a composite game: the contractual mechanism followed by renegotiation. Whichever widget, W say, is specified by the mechanism, even if it is one of the $N-1$ generic ones, renegotiation ensures that at the end of the day the special widget will be traded. Moreover, since B has all the bargaining power, S's payoff is minus the cost of producing W , $C(W)$ say. (This is gross of any transfer that the mechanism might specify. Note that such transfers can depend on the outcome of the mechanism, but not on c_i directly.) Now, given renegotiation, B's payoff and S's payoff sum to the gains from trade $v - c_i$. Hence B's payoff equals $C(W) + v - c_i$. Also, we must not forget that the mechanism might specify "no trade" as the outcome; this costs S nothing, and, following renegotiation, yields B a payoff $v - c_i$. Ignoring permutations, then, in either state there are $N+1$ possible (nonstochastic) outcomes to the composite game, each corresponding to a different point along the Pareto frontier. We list them in descending order of S's payoff (continuing to ignore transfers specified by the mechanism):

¹⁰This terminology is loose. Strictly speaking, both of these "states" comprise a subset of $N!$ states, each corresponding to a different permutation of the widgets.

Final payoffs following renegotiation

(gross of transfers specified by mechanism)

	<u>Seller S</u>	<u>Buyer B</u>
<u>State 1</u>		
no trade	0	$v - c_1$
special widget	$- c_1$	v
generic widgets	$- c_1 - \frac{1}{N}(c_2 - c_1)$	$v + \frac{1}{N}(c_2 - c_1)$
	.	
	.	
	$- c_1 - \frac{N-1}{N}(c_2 - c_1)$	$v + \frac{N-1}{N}(c_2 - c_1)$
<u>State 2</u>		
no trade	0	$v - c_2$
generic widgets	$- c_1 - \frac{1}{N}(c_2 - c_1)$	$v - c_2 + c_1 + \frac{1}{N}(c_2 - c_1)$
	.	
	.	
	$- c_1 - \frac{N-1}{N}(c_2 - c_1)$	$v - c_2 + c_1 + \frac{N-1}{N}(c_2 - c_1)$
special widget	$- c_2$	v

Figure 2

Note that in state 1, the cheapest of the N widgets is the special one. And in state 2, the special widget is the most expensive. However this is only a matter of labeling: following renegotiation, all contractually specified outcomes, even "no trade", lead to the special widget being traded. In effect, we can ignore the labels ("no trade, "special widget", "generic widgets") that appear in the lists in Figure 2.

Unfortunately, this leaves very little to screen on. To see why, consider the limit $N \rightarrow \infty$, where the only difference between the lists is that a constant, $c_2 - c_1$, is added to B's payoff in state 1 relative to state 2. (This amount is the additional surplus from S having a lower cost of producing the special widget.) Clearly, for the purpose of mechanism design, adding a constant to one of the parties' payoffs does not help screen the states.

For finite N , we claim that S's payoff cannot differ by more than $\frac{1}{N}(c_2 - c_1)$ across states 1 and 2. To check this, consider a simple mechanism in which S is allowed to choose one of the N widgets. Then she will always select the cheapest widget -- viz., the special one (cost c_1) in state 1, and the cheapest generic one (cost $c_1 + \frac{1}{N}(c_2 - c_1)$) in state 2. Correspondingly, if B is allowed to choose, he will always select the most expensive widget -- viz., the most expensive generic one (cost $c_1 + \frac{N-1}{N}(c_2 - c_1)$) in state 1, and the special one (cost c_2) in state 2. Under either mechanism, S's payoff differs by only $\frac{1}{N}(c_2 - c_1)$ across states 1 and 2, as claimed.

Other, more sophisticated mechanisms might be considered. For example, one party might be entitled to veto a certain number of widgets prior to the other party choosing. It is easy to see that this doesn't succeed in widening the gap in S's payoff between states.

A specific performance contract does no better either. Consider the contract specifying that a certain widget should be traded at a fixed price. The payoffs associated with this widget are equally likely to be any one of those listed above (excepting "no trade"). Again, S's (expected) payoff differs by only $\frac{1}{N}(c_2 - c_1)$ across states 1 and 2.

An equivalent way of saying this is that the price difference $p_2 - p_1$ equals $\frac{N-1}{N}(c_2 - c_1)$. In terms of S's incentive to invest, this is only $O(\frac{1}{N})$

better than the null contract.

Now for the formal proof of the proposition.

Proof of Proposition 2

Take any abstract mechanism M . Consider a state, $(1, \tau)$ say, in which the special widget costs $S c_1$ to produce, and the N widgets are arranged according to some permutation τ .¹¹ Without loss of generality, suppose that τ is such that widget 1 is the special widget, and that widgets $2, \dots, N$ are generic widgets costing $c_1 + \frac{1}{N}(c_2 - c_1)$, \dots , $c_1 + \frac{N-1}{N}(c_2 - c_1)$ respectively.

In state $(1, \tau)$, let the equilibrium strategies of M for B and S be $\mu^B(1, \tau)$ and $\mu^S(1, \tau)$. And, following any renegotiation, let $p(1, \tau)$ denote the overall price that B pays S for delivery of the special widget: $p(1, \tau)$ equals any transfer specified in the mechanism plus any amount agreed by the parties during renegotiation. In other words, B 's and S 's final equilibrium payoffs are respectively $v - p(1, \tau)$ and $p(1, \tau) - c_1$.

Now consider another state, $(2, \tau^*)$ say, in which the special widget costs $S c_2$ to produce, and the N widgets are arranged according to some new permutation τ^* in which widgets $1, \dots, N-1$ are generic widgets costing $c_1 + \frac{1}{N}(c_2 - c_1)$, \dots , $c_1 + \frac{N-1}{N}(c_2 - c_1)$ respectively, and widget N is the special widget. Note that τ^* is a simple rotation of τ : widget $n = 1, \dots, N-1$ now has the characteristics that widget $n+1$ formerly had; and now widget N is the special widget rather than widget 1.

In state $(2, \tau^*)$, let the equilibrium strategies of M for B and S be $\mu^B(2, \tau^*)$ and $\mu^S(2, \tau^*)$. After renegotiation, let their respective payoffs be $v - p(2, \tau^*)$ and $p(2, \tau^*) - c_2$.

The question arises: What outcome does M specify if the strategy pair

¹¹State $(1, \tau)$ is thus one of the $N!$ constituent states of what in the text we loosely called "state 1".

$(\mu^b(2, \tau^*), \mu^s(1, \tau))$ is played? Suppose some general stochastic outcome is specified: B pays S an amount q ; widget $n = 1, \dots, N$ is traded with probability $\alpha_n \geq 0$; and there is no trade with probability $1 - (\alpha_1 + \dots + \alpha_N) \geq 0$.¹²

There are two incentive constraints that $q, \alpha_1, \dots, \alpha_N$ must satisfy. First, in state $(2, \tau^*)$, S must not have an incentive to deviate to $\mu^s(1, \tau)$. Second, in state $(1, \tau)$, B must not have an incentive to deviate to $\mu^b(2, \tau^*)$.

Suppose strategy pair $\mu^b(2, \tau^*), \mu^s(1, \tau)$ is played. The contractually-specified outcome is typically inefficient, and will be renegotiated: when the outcome of the lottery $(\alpha_1, \dots, \alpha_N)$ specifies either that a generic widget is traded, or that there is no trade, the parties must bargain in order to exploit the gains from trading the special widget (and, by assumption, B has all the bargaining power). The final payoffs depend on the state, as indicated in Figure 2. In state $(2, \tau^*)$, following the play of $\mu^b(2, \tau^*), \mu^s(1, \tau)$, S's final payoff is

$$q - \alpha_1 \left(c_1 + \frac{1}{N}(c_2 - c_1) \right) - \dots - \alpha_{N-1} \left(c_1 + \frac{N-1}{N}(c_2 - c_1) \right) - \alpha_N c_2,$$

which, according to the first incentive constraint, cannot be more than what she gets in equilibrium, $p(2, \tau^*) - c_2$. And in state $(1, \tau)$, following the play of $\mu^b(2, \tau^*), \mu^s(1, \tau)$, S's final payoff is

$$q - \alpha_1 c_1 - \alpha_2 \left(c_1 + \frac{1}{N}(c_2 - c_1) \right) - \dots - \alpha_N \left(c_1 + \frac{N-1}{N}(c_2 - c_1) \right),$$

which, according to the second incentive constraint, cannot be less than what she gets in equilibrium, $p(1, \tau) - c_1$, since if S were worse off B would be better off (all final payoffs lie along the Pareto frontier).

¹²As both parties are risk neutral, there is no gain from making the transfer q depend on which widget (if any) the mechanism specifies is traded.

Combining these two constraints, we have

$$\begin{aligned} p(2, \tau^*) - p(1, \tau) &\geq c_2 - c_1 - (\alpha_1 + \dots + \alpha_N) \frac{1}{N} (c_2 - c_1) \\ &\geq \frac{N-1}{N} (c_2 - c_1). \end{aligned}$$

Since this lower bound applies for any permutation τ (and associated rotation τ^*), we can take expectations across permutations, which are equally probable, to deduce that the difference between the expected price S receives if it costs her c_2 to produce the special widget and the expected price she receives if it costs her c_1 is at least $\frac{N-1}{N}(c_2 - c_1)$.

In other words, as the realization of S 's cost falls from c_2 to c_1 , her payoff rises by at most $\frac{1}{N}(c_2 - c_1)$. But this gives her only $O(\frac{1}{N})$ more incentive to reduce her expected costs than does the null contract (under which her payoff would be independent of the cost realization).

Q.E.D.

Given that, when there is no commitment, almost nothing can be achieved in Case D, it follows a fortiori that the same must be true in Case ND.

Corollary. The conclusion of Proposition 2 also holds in Case ND.

This corollary serves to illustrate MT's Theorem 4, albeit again in a very particular setting. Their theorem states that, subject to certain additional conditions, if the parties cannot commit not to renegotiate, then the inability to describe the nature of trade in advance does not matter. This is clearly true in Proposition 2, since even when the parties can describe widgets in advance, they achieve little more than under the null contract.

Before we move on, it is worth reviewing the role of our assumption

that the costs of the generic widgets fill the whole interval from c_1 to c_2 . Suppose this were not true, e.g., the costs of the generic widgets were spread evenly between $\frac{1}{2} c_1 + \frac{1}{2} c_2$ and c_2 . Then the following contract would be useful: S is allowed to choose which widget to supply at date 1 and receives a fixed price. The reason is that S will always choose the cheapest widget, which is the special widget if $\tilde{c} = c_1$ and the cheapest generic widget if $\tilde{c} = c_2$. Thus S's payoff differs by $(\frac{1}{2} c_1 + \frac{1}{2} c_2 - c_1)$ across the high cost and low cost states of nature, which gives S an incentive to invest in cost reduction.

2.3. Does describability matter?

The main message to emerge from the analysis of Sections 2.1 and 2.2 is that the inability to commit not to renegotiate makes a crucial difference. In contrast, the issue of whether or not actions (widgets) can be described ex ante -- whether Case D or ND prevails -- appears not to matter; this is consistent with MT's thesis.

These findings shed light on the informal story with which the paper began. It turns out that the usual "observable but not verifiable" assumption is enough to justify a high degree of contractual incompleteness (taking the null contract to be the quintessentially incomplete contract), provided (i) the parties cannot commit not to renegotiate, and (ii) the environment is rich enough (here, there are enough generic widgets). That is, the incentive constraints that emerge from dealing with Case D, and treating the problem as one of classical mechanism design (constrained by renegotiation), are enough to reduce massively the contractual possibilities. At a formal level, there is no need to invoke additional -- and less traditional -- assumptions like nondescribability in order to give the informal story solid theoretical foundations. This raises the question: Should contracts that are optimal subject to well-defined incentive constraints be thought of as "incomplete" at all? We address this question in Section 5.

The matter of describability should not be ignored altogether. It is clearly ridiculous to assume that all actions can be costlessly described in advance. And for this reason, MT's Theorem 4 is a potentially important result, because, stripped of its auxiliary conditions, the theorem appears to

conclude that nondescribability is irrelevant even when there is no commitment. We believe that such a broad-brush conclusion would be misleading, however. The inability to describe the widgets ex ante can make a big difference.

For example, consider a model where there is no uncertainty. Without loss of generality, suppose widget 1 is always the special widget, and that widgets 2, ..., N are always generic widgets costing $c_1 + \frac{1}{N}(c_2 - c_1)$, ..., $c_1 + \frac{N-1}{N}(c_2 - c_1)$ respectively. In this deterministic model, the first-best can obviously be achieved if widgets can be described at date 0 (Case D): the parties simply write a specific performance contract under which the parties agree to trade widget 1 at a fixed price, say v . Since this outcome is efficient, there is nothing to renegotiate, and S enjoys all of any cost saving: $p_1 = p_2 = v$. S has first-best incentives.

Things look very different, however, if the N widgets cannot be described at date 0 (Case ND). Now a specific performance contract is no longer feasible (widget 1 cannot be specified). The parties have to rely instead on a mechanism which reveals the identity of the special widget at date 1, while at the same time keeping $p_1 = p_2$. We assert that this implementation problem is essentially the same as that of Proposition 2, and that the conclusion is therefore the same: in approximate terms, no contract can do any better than the null contract.

A simple way to prove this assertion is to suppose that there are N "names" at date 0, each of which will describe a widget at date 1. However, it isn't known which name will attach to which widget: the meaning of the vocabulary (the list of N names) only becomes established at date 1. And there is no other way of describing widgets at date 0. In particular, even though B and S know at date 0 which widget is the special one, they have no words to describe it, other than the N names, any one of which may turn out to be appropriate at date 1. It is clear that this implementation problem is isomorphic to the problem in Section 2.2, where it wasn't known at date 0 which widget would have which costs/values. The conclusions of Proposition 2 thus carry over to the present setting, where widgets have fixed costs/values but cannot be described at date 0.

We introduce names here only as a device to simplify the argument. A

fortiori, our assertion holds when there are no names, or any other vocabulary for describing the widgets at date 0.

This example is not covered by MT's Theorem 4. Their result requires two conditions: first, that the set of states is "maximal"; and, second, that the final outcome is "renegotiation welfare neutral." Roughly, maximality means that every permutation of the N widgets has to be possible (which rules out our deterministic example); and renegotiation welfare neutrality means that final payoffs have to be the same across permutations. The first condition is not restrictive, but the second one is. To see this, notice that the example can be modified to include a small measure of fringe states so as to meet MT's maximality condition. Moreover, such a modification doesn't rob the example of its force: it is still the case that the specific performance contract is approximately first-best in Case D, whereas in Case ND the null contract is almost optimal. However, although this modified example satisfies maximality, it does not satisfy MT's renegotiation welfare neutrality condition, and so is not covered by their Theorem 4.

We may conclude this section as follows: nondescribability is generally an important constraint in the absence of a commitment not to renegotiate.

2.4. Partially incomplete contracts

Proposition 2 can be criticized for going too far: there is no point in writing any contract at all. What is needed is a theory of partial incompleteness.

The model can also be criticized on the grounds that the effects of S 's investment σ are too jagged: only the special widget's cost is affected by σ ; the costs/values of the other (generic) widgets are fixed. In principle, investment may reduce the cost of any widget.

Here we make a start in responding to these criticisms. Consider a variant of our model. There are N widgets, each of which can be described at date 0: we are in Case D. Of these N widgets, M are "defined". Defined widgets are simply a category of widget distinct from the others: e.g., they may all have a common shape, which is not shared by the other $N - M$ widgets. And each defined widget is distinct from the other defined widgets: e.g.,

each may have its own color. We take M to be large, and N large relative to M : $N \gg M \gg 1$.

There is one "special" widget, which is the widget that yields the greatest surplus at date 1. This is the widget that will be traded at date 1 (possibly following renegotiation). Let its value to B be v . Crucially, the special widget is always one of the M defined widgets: previously, we assumed that the special widget could be any of the N widgets.

The second important change we make to our model is to suppose that the cost of producing all of the M defined widgets is affected by σ . With probability $\pi(\sigma)$, the cost of producing a given defined widget, other than the special widget, is reduced by $\Delta > 0$.¹³ And with probability $\pi(\sigma)$, the cost of producing the special widget is reduced by $k\Delta$. We assume $k > 1$: in other words, we assume that σ has a greater impact on the expected cost of producing the special widget than on the other $M-1$ defined widgets. These cost reductions are perfectly correlated.¹⁴ That is, with probability $1 - \pi(\sigma)$, there are no cost reductions. Without cost reductions, we suppose that costs are evenly spread from \underline{c} to \bar{c} , with the cost of the special widget lying somewhere in the middle, say at c . Assume Δ is small enough that $\underline{c} < c - k\Delta$ and $c < \bar{c} - \Delta$; i.e., the special widget always has cost lying within the range of costs of the other defined widgets. All permutations of the M defined widgets (viz., the identity of the special widget, and the permutation of the costs of the other $M-1$ widgets) are equally probable at date 1.

The costs of the remaining $N-M$ widgets are unaffected by σ .¹⁵ We

¹³ $\pi(\sigma)$ is assumed to satisfy our earlier assumptions: $0 < \pi(\sigma) < 1$, $\pi'(\sigma) > 0$ and $\pi''(\sigma) < 0$ for all $\sigma \geq 0$; and $\pi'(0) = \infty$.

¹⁴ This is not an important assumption.

¹⁵ As in Section 2, we do not need to specify the values to B of the $N-1$ non-special widgets, other than to assume that their surplus (value minus cost) is always strictly less than the surplus of the special widget.

suppose that these costs are evenly spread from \underline{g} to \bar{g} , a range which encompasses the costs of the defined widgets: $\underline{g} < \underline{c} - \Delta$ and $\bar{c} < \bar{g}$. All cost permutations of these $N-M$ widgets are equally probable at date 1.

The expected total surplus is

$$\pi(\sigma)[v - c + k\Delta] + (1 - \pi(\sigma))[v - c] - \sigma.$$

And hence the first-best investment level σ^* satisfies

$$(2.4) \quad \pi'(\sigma^*) = 1/k\Delta.$$

Just as in our earlier model, the first-best cannot be attained, since the identity of the special widget is not known in advance. Moreover, as in Proposition 2, the incentive constraints impose severe restrictions on what can be achieved through contracting. However, unlike in Proposition 2, the parties can do appreciably better than under the null contract. With the null contract, S would have no incentive to invest since she gets no surplus from the date 1 bargain (the price B pays is perfectly correlated with the realised cost of producing the special widget). That is, σ would equal zero. The same would be true in a contract where, for example, either B or S were free to choose any of the N widgets (i.e., without limiting the choice to the defined widgets): S would always choose the cheapest widget (costing \underline{g}), B would always choose the most expensive (costing \bar{g}), and, either way, S's payoff would not depend on σ , which would give her no incentive to invest.

Instead, the parties can write a contract specifying that S supplies one of the defined widgets for a fixed price p . The choice of which defined widget may be left to S, in which case she will supply the cheapest, costing her either $\underline{c} - \Delta$ (with probability $\pi(\sigma)$) or \underline{c} (with probability $1 - \pi(\sigma)$).

This is not the special widget, and so the parties will renegotiate at date 1. But since S gets none of the surplus, her expected payoff, net of investment costs, equals

$$(2.5) \quad \pi(\sigma)[p - \underline{c} + \Delta] + (1 - \pi(\sigma))[p - \underline{c}] - \sigma.$$

And hence her choice of investment $\hat{\sigma}$ satisfies

$$(2.6) \quad \pi'(\hat{\sigma}) = 1/\Delta.$$

Comparing (2.6) with (2.4), we see that there is underinvestment: $\hat{\sigma} < \sigma^*$. However, there is more investment than under the null contract: $\hat{\sigma} > 0$. That is, a "partially incomplete" contract -- a contract defining the set of widgets from which S must supply one, and fixing the price -- is better than no contract; but it does not implement first-best.

Notice that if the contract allowed B to select from the set of defined widgets, then he would choose the most expensive, costing either $\bar{c} - \Delta$ (with probability $\pi(\sigma)$) or \bar{c} (with probability $1 - \pi(\sigma)$). And S's investment would again be given by $\hat{\sigma}$. In other words, aside from a transfer difference $\bar{c} - \underline{c}$, it does not matter if B or S has the right to select one of the defined widgets.

For large M, these partially incomplete contracts are almost as good as any contract can be. Consider a specific performance contract: at date 0 the parties agree that one specified defined widget should be supplied at date 1 for a fixed price. Since there is a small probability, $1/M$, that this widget will turn out to be the special one, S's incentives are slightly improved; but as $M \rightarrow \infty$ her investment drops to $\hat{\sigma}$.

We can appeal to the same logic of Proposition 2 to prove that this is as far as contracts can take us.

Proposition 3. Suppose Case D holds. If the parties cannot commit not to renegotiate, then as M , the number of defined widgets, tends to infinity, the optimal level of S 's investment converges to $\hat{\sigma}$, the solution to (2.6). In the limit, it is optimal to contract for the delivery of one of the defined widgets at a fixed price; the particular widget may be chosen by B or S at date 1, or may be specified in advance at date 0.

Proposition 3 goes some way towards meeting the criticism that we lack a theory of partial incompleteness. It should be recognized that this is not really a framework in which agents choose the degree of contractual incompleteness, because there are no "margins": in Proposition 3, the set of defined widgets is exogenously given.

Before leaving this model, we should point out that Proposition 3 does not hold in Case ND, where the widgets cannot be described at date 0. In fact, in this case, one can adapt the argument of Section 2.3 to show that as the total number of widgets N tends to infinity (and the ratio N/M also approaches infinity), S 's investment approaches zero, irrespective of the contract. Thus, here we have another example (like that in Section 2.3) where nondescribability matters.

3. Commitment

In Section 2 we saw that, although nondescribability matters in incomplete contracting models, the crucial assumption is the lack of commitment. If the parties can commit not to renegotiate their contract, then they can achieve the first-best (Proposition 1).

In this section we consider how reasonable it is to assume that the parties cannot commit not to renegotiate.¹⁶ We will also discuss the role of third parties to a contract. It will be convenient from an expositional point of view to gear our discussion fairly closely to that of MT.

¹⁶For an interesting discussion of this question, see Tirole (1998).

One obvious way for B and S to commit not to renegotiate is for them to write in their contract an irrevocability clause; or, equivalently, a clause that says that B must pay S a huge sum of money if renegotiation occurs. The problem with this is that, under the current legal system, there is nothing to stop B and S from writing a new contract that cancels the irrevocability clause or waives the penalty. The point is that the courts will enforce the new contract rather than the original one. Anticipating that the irrevocability clause will not stand, B will decline S's take-it-or-leave-it offer in the model of Section 2, and renegotiate.

MT argue that this justification for lack of commitment is unsatisfactory because in an ideal world B and S could register their first contract with the court, and could instruct the court to enforce the first contract and ignore any revised contract. However, a registration system like this does not exist anywhere in the world as far as we know, and would require a system-wide institutional change.¹⁷

Even if such a registration system were put in place, it might not prevent renegotiation. B and S might be able to renegotiate indirectly by writing side-contracts with third parties. For example, B and S could agree ex post to operate through a middleman: S will supply the widget to the middleman, who will supply it to B; in return B pays the middleman, who pays S. These side-contracts do not violate the first contract (which states that no renegotiation will occur) because they do not involve B and S directly.

Of course, the original contract could state that not only can it not be renegotiated, but also no equivalent set of contracts with third parties should be enforced. The question, though, is: What is an "equivalent set of

¹⁷Such a change might be quite costly and the benefits may not be all that great: the majority of contracting parties may choose not to register their contracts, since they recognize that they will think of new things to include as time passes (they are "boundedly rational"). Thus, to the extent that a registration system has a fixed cost, it might not be worth introducing for the minority of people like B and S, who are unboundedly rational and for whom contract renegotiation is an impediment.

contracts with third parties"? There may be quite legitimate sequences of trade linking S to B through various middlemen, and it may be hard for a judge to distinguish between the legitimate ones and the ones that are designed to circumvent a "no renegotiation" provision.

Note that side-contracting also interferes with the use of third parties as a commitment device. Suppose that B and S sign a contract with a third party T stating that B and S will each pay T a huge sum of money if renegotiation occurs. Then ex post B and S can avoid the penalty by contracting indirectly through a middleman. Moreover, if B and S try to prevent this ex ante by promising to pay a penalty in the event that a middleman is used, then the same problem arises as above: it might be hard to distinguish between cases where the middleman is used for legitimate business purposes and cases where he is used to circumvent renegotiation.

Now that we have raised the issue of third parties, it is worth asking whether they can be used in other ways than just to prevent renegotiation. The answer is yes. Third parties can drive a wedge between B's and S's payoffs: B can be penalized without rewarding S, and vice versa. This may improve incentives even if the parties cannot commit not to renegotiate their contract.

To see how a third party can improve matters, consider the following contract in the model of Section 2.2:

At date 1, B chooses between the following possibilities:
(1) B pays $c_2 - c_1$ to S and no trade occurs; or (2) B pays nothing to S. If B chooses (2), S has a choice of "accepting" B's offer or "rejecting" B's offer. If S "accepts" B's offer, no trade occurs. If S "rejects" B's offer, then S supplies a widget of her choice to B, and B pays $c_1 + \frac{1}{2N}(c_2 - c_1)$ to S and a fine $F \geq c_2 - c_1$ to a third party T.

An important part of this contract is that it does not prohibit

renegotiation. Thus, if a "no trade" outcome occurs, or if S "rejects" B's offer and supplies an inefficient widget, then this is not the end of the matter: since there are gains from trade the parties will always renegotiate and trade the efficient widget (the special widget). As in Section 2, we assume that B has all the bargaining power in the renegotiation process.

Consider first the case where S's cost of producing the special widget at date 1 equals c_1 : state 1 occurs. If B chooses (1), then S's payoff is $c_2 - c_1$, and B's payoff is $(v - c_1) - (c_2 - c_1) = v - c_2$, since B gets all the gains from renegotiation. On the other hand, if B chooses (2), then it is easy to see that S will prefer to "reject" (she supplies the special widget, costing c_1 , and her payoff is $-c_1 + c_1 + \frac{1}{2N}(c_2 - c_1) = \frac{1}{2N}(c_2 - c_1) > 0$) than to "accept" (her payoff is zero). Hence B has to pay the fine, which reduces his payoff below $v - c_2$. The conclusion is that in state 1, B chooses (1), and S's payoff is $c_2 - c_1$.

Consider next the case where S's cost of producing the special widget at date 1 equals c_2 : state 2 occurs. If B chooses (1), then S's payoff is $c_2 - c_1$, and B's is $(v - c_2) - (c_2 - c_1) = v + c_1 - 2c_2$. If B chooses (2), then it is easy to see that S prefers to "accept" (her payoff is zero) than to "reject" (she supplies the cheapest generic widget, costing $c_1 + \frac{1}{N}(c_2 - c_1)$, and her payoff is $-c_1 - \frac{1}{N}(c_2 - c_1) + c_1 + \frac{1}{2N}(c_2 - c_1) = -\frac{1}{2N}(c_2 - c_1) < 0$). Hence, under (2) B avoids the fine and his payoff is $v - c_2 > v + c_1 - 2c_2$ (he gets all the gains from renegotiation). The conclusion is that in state 2, B chooses (2), and S's payoff is zero.

We see that S's payoff decreases from $c_2 - c_1$ in state 1 to zero in state 2. In other words, her payoff falls by the exact amount that her costs rise. But this means that S has first-best investment incentives: at date 1/2, S will solve

$$\underset{\sigma}{\text{maximize}} \quad \pi(\sigma)[c_2 - c_1] - \sigma,$$

which is equivalent to (2.1). Since renegotiation ensures that the efficient widget is supplied, the three-party contract yields the first-best outcome.

Let us now discuss some potential problems with the above contract. First, the mechanism is very fragile. It relies on the fact that there are discrete differences between the costs of the various widgets; the mechanism is designed so that when state 1 is the true state "rejection" by S gives her a small positive payoff, whereas when state 2 is the true state "rejection" by S gives her a small negative payoff. There is reason to think that if there is a continuum of widgets, it may be impossibly difficult to screen state 1 from state 2, even if third parties are allowed; certainly, contracts like that above will not work. Moreover, general implementation theorems (in particular, Theorem 3 in Maskin and Moore (1998)), which employ devices, such as getting one agent to announce the state of nature (which can then be challenged by the other agent), are unlikely to be operational because the description of a state is so rich: the entire vector of costs has to be announced.¹⁸ And, since the distinction between states 1 and 2 is infinitesimally small--a difference in the cost of one widget taken from a continuum--it does not suffice for agents to announce anything less than the entire state. We conjecture that in some properly articulated model with a continuum of widgets, the conclusions of Proposition 2 will hold, even allowing for third parties; but this awaits further research.

Even if we stick to the case of a finite number of generic widgets, the above contract is vulnerable to collusion. In particular, either B and T or S and T can gain by writing a (secret) side-contract in which it is agreed that any fine received by T is handed over to the other party.

Consider a side-deal between B and T. Suppose state 1 occurs. Then if B chooses (2) in the above game, S will "reject" as before, but B's payoff will be $v - c_1 - \frac{1}{2N}(c_2 - c_1) > v - c_2$, since B pays the fine to himself. Thus B will choose (2), and S's payoff is $\frac{1}{2N}(c_2 - c_1)$.

On the other hand, if state 2 occurs, collusion makes no difference: B chooses (2), S "accepts" and S's payoff is zero.

¹⁸Note that these mechanisms are designed to rule out unwanted equilibria. If uniqueness is not required, then the mechanisms can be much simpler.

We see that collusion has a devastating effect on S's investment incentives. Recognizing that B and T will collude after she has made her investment decision, S will choose σ to solve:

$$\underset{\sigma}{\text{maximize}} \quad \pi(\sigma) \left[\frac{1}{2N} (c_2 - c_1) \right] - \sigma,$$

which leads to a very low value of σ . In fact, as $N \rightarrow \infty$, $\sigma \rightarrow 0$, which is the same outcome as when there is no contract at all (see Proposition 2).¹⁹

Can collusion be avoided? An obvious approach is to prohibit collusion in the original three-party contract, i.e., to instruct the courts not to enforce any side-deals between a subset of the parties. The difficulty with this, however, is that B and T (or S and T) can disguise their side-deal by using a middleman.

A second approach to avoiding collusion, suggested by MT, is to replace the single third party T by a collection of third parties. For example, MT propose that the contract between B and S could state that any fines should be paid to the "community of citizens", i.e., to the general public. The idea is that it is hard -- if not impossible -- for B or S to collude with a whole community.

However, such an arrangement raises new problems. First, as a matter of contract law, there appears to be nothing to stop B and S from canceling

¹⁹It is not difficult to show that S and T also have an incentive to collude (assuming B and T do not). Collusion between S and T makes no difference in state 1, when B chooses (1) and S's payoff is zero. However, in state 2, S will "reject" if B chooses (2), since S now receives the fine. Thus, if B understands the collusion between S and T, B will choose (1) in order to avoid paying the fine. Hence S's payoff is zero in state 2 as well as in state 1. The conclusion is that S will choose $\sigma = 0$, which is the same outcome as when there is no contract.

the fine, after S has "rejected" an offer from B, but before the fine has been paid, i.e., they could simply change their mind. The point is that the community is not a true party to the contract -- the taxpayers never signed anything -- and so they would have no grounds to complain or sue. Of course, B and S could pick a representative of the community to be a signatory, but this would raise the possibility of collusion between B (or S) and the representative.

Second, even if we put this issue aside, it is unclear who would collect the fine on behalf of the community. For example, suppose the contract states that S must reject B's offer by placing an advertisement in a designated newspaper, and this advertisement must include an offer from B to pay F dollars to the first person who responds to it, e.g., by sending an e-mail to a particular address. Then S could always tip off a friend, thereby ensuring that the friend is the first to respond, i.e., in effect S receives the fine herself.

The Use of Lotteries

MT have proposed an even more ingenious way to solve the problem of third-party collusion: eliminate third parties altogether and use lotteries to introduce a wedge between what S receives and what B pays. Specifically, suppose that B is (at least slightly) risk averse, rather than risk neutral.²⁰ Then it is possible to find a random variable \tilde{p} whose mean, $E\tilde{p}$, equals $c_1 + \frac{1}{2N}(c_2 - c_1)$, but the certainty equivalent of $-\tilde{p}$ is very low to B. (Simply raise the variance of \tilde{p} .) Now replace the previous three-party contract with a two-party one, where there are no fines, but if S "rejects" B's offer, B pays the random amount \tilde{p} to S. This is equivalent to the previous contract since the lottery has the effect of a penalty on B.²¹

²⁰We continue to suppose that S is risk neutral. The argument can easily be modified if S is also risk averse.

²¹We are assuming that B is not wealth-constrained. Wealth constraints may limit the maximum penalty that can be imposed through a lottery.

This approach has its own difficulties, however. The simplest way to introduce randomness in \tilde{p} is to make \tilde{p} contingent on an objective, nonmanipulable event, e.g., the change in a stock market index over a short interval following S's announcement.²² However, the problem is that, if the event is objective, B can insure against it in advance, i.e., B could go to a (competitive) insurance company and agree, conditional on S's announcement and a particular realization of the stock market index, to exchange the actual value of \tilde{p} for its expected value, $E\tilde{p}$. If S "rejects", this makes B's combined payment to S and the insurance company equal to $E\tilde{p} = c_1 + \frac{1}{2N}(c_2 - c_1)$, and the effect of the penalty is removed.

In a private communication, Eric Maskin has suggested that B and S could avoid the possibility of insurance by making \tilde{p} depend on the realization of a subjective event, or an event whose probability distribution is private information to B and S. For example, B and S could construct a randomization device, e.g., a machine, whose structure is known only to B and S. However, this would seem to open the door to manipulation of the device by B or S. That is, there appears to be a trade-off: the more objective a lottery is, the less it can be manipulated, but the more it can be insured against; the less objective a lottery is, the less it can be insured against, but the more it can be manipulated.²³

²²The reason the interval must be short is that, if it were not, then this would give time for B and S to renegotiate the contract after S's announcement, to avoid the unwanted randomness.

²³Yet another possibility, suggested to us by Andy Postlewaite, is that, instead of constructing a machine, the parties can induce endogenous (i.e., subjective) randomness by agreeing to play a game, with a publicly observed outcome, that has a mixed strategy equilibrium. The idea is that, although the players' strategies are not observable to outsiders, they are self-enforcing. See Barany (1992) for an analysis of this kind of idea. The advantage of a game over a machine is that a game is not manipulable by one party. However, the disadvantage of a game is that it may be difficult to arrange that the parties play the game simultaneously with S's "rejection" of

At this point, perhaps we ought to bring to a halt this rather protracted tennis match between the believers in commitment and the believers in no commitment. Let the match be declared an honorable draw. To repeat what we said in the Introduction: the degree of commitment is something about which reasonable people can disagree.

4. Property Rights

We saw in Section 2 that the inability to specify trade in advance and/or the inability to commit not to renegotiate can lead to an inefficient outcome, in which S underinvests. In this section we consider whether some form of vertical integration can alleviate the situation. In particular, following the property rights literature (see Grossman and Hart (1986), Hart and Moore (1990) and Hart (1995)), we ask whether S would invest more if she owned B's (nonhuman) assets.

We will assume that, as owner, S has residual rights of control over B's assets, in the sense that S has access to B's downstream technology. What this means is that S can turn the special widget into final output that can be sold on the downstream market at price $\hat{v} < v$. Here $v - \hat{v}$ represents the contribution of B's (specialized) human capital; even though S owns B's nonhuman assets, S cannot capture this amount.

To see how downstream integration changes S's incentives, let us stick to the setup of Section 2.2 where B has all the bargaining power in renegotiation and where, in the limit $N = \infty$, it is optimal to have no contract at date 0. Consider first the case where $\hat{v} \geq c_2 > c_1$. Then, at date 1, S can always threaten to produce without B and obtain $\hat{v} - c_1$ in state 1 or $\hat{v} - c_2$ in state 2. Of course, renegotiation will occur, since there is an extra $v - \hat{v}$ to be gained if B participates, but since B has all the

B's offer. If there is even a short interval of time between S's "rejection" and the playing of the game, then the parties can (and will) renegotiate the contract during this interval, to avoid the randomness induced by the mixed strategies (see also footnote 21).

bargaining power this does not affect S's payoff. Thus S chooses σ to maximize $\pi(\sigma)[\hat{v} - c_1] + (1 - \pi(\sigma))[\hat{v} - c_2] - \sigma$, which yields the same solution as (2.1), i.e., the first-best.

On the other hand, suppose $c_2 > \hat{v} > c_1$.²⁴ Then, S will obtain $\hat{v} - c_1$ in state 1 but nothing in state 2 (it is not worth her while producing without B's participation in this state). Hence S solves:

$$(4.1) \quad \underset{\sigma}{\text{maximize}} \quad \pi(\sigma)[\hat{v} - c_1] - \sigma,$$

which yields the first order condition $\pi'(\sigma) = 1/(\hat{v} - c_1)$. Comparing this with the first order condition for (2.1), $\pi'(\sigma) = 1/(c_2 - c_1)$, we see that S will underinvest relative to the first-best, but will generally set $\sigma > 0$, i.e., downstream integration has a positive effect.

The conclusion so far is that a reallocation of property rights can help when contracts are incomplete. In fact, this is also the conclusion obtained in Maskin and Tirole (1998b). However, MT make two further observations. First, they point out that the parties can obtain an even better outcome by including a third party in their contract. Second, they note that there may be several property rights allocations which yield the same outcome, i.e., the theory lacks predictive power.

Although we have given arguments against the use of third parties in Section 3, it is worthwhile to consider MT's first point. Since the first-best can be achieved without a third party if $\hat{v} \geq c_2$, the interesting case to study is where $c_1 < \hat{v} < c_2$. The contract MT propose is the following. B and S agree at date 0 that they will jointly own B's assets, i.e., that neither party has the right to use the assets without the other's permission. However, B has the option to sell his share in the joint venture

²⁴If $\hat{v} \leq c_1$, downstream integration obviously has no effect, since S will not produce in either state without B's participation.

to S at date 1 at price $0 < P < \hat{v} - c_1$; moreover, if B exercises his option, S must not only pay P to B but also a fine F to a third party.

To see how this works, suppose first that S's cost of producing the special widget at date 1 equals c_1 : state 1 occurs. Then, if B does not exercise his option to sell, S obtains a zero payoff in the absence of renegotiation since S cannot use B's assets without B's permission. Since B gets all the gains from renegotiation, B's post-renegotiation payoff is $v - c_1$, and S's is zero. On the other hand, if B exercises his option to sell, S's payoff equals $\hat{v} - c_1 - P - F$, and B's payoff equals $v - \hat{v} + P$, since B obtains the full amount $v - \hat{v}$ in renegotiation. It follows that, since $P < \hat{v} - c_1$, B will choose not to exercise his option, and so S's payoff in state 1 is zero.

Consider next the case where S's cost of producing the special widget at date 1 equals c_2 : state 2 occurs. Then, whether or not B exercises his option to sell, S will not use the assets without B's participation since $\hat{v} < c_2$. Hence B will obtain the full surplus $v - c_2$ in renegotiation and his payoff will equal $v - c_2$ if he does not exercise his option and $v - c_2 + P$ if he does. Since $P > 0$, B prefers to exercise his option, and so S's payoff in state 2 equals $-(P + F)$.

Notice that as S's cost of producing the special widget at date 1 rises from c_1 to c_2 , the fall in her payoff can be made equal to $c_2 - c_1$ if the parties choose $F = c_2 - c_1 - P$. In other words, the appropriate choice of F induces S to make the first-best level of investment.

Given the discussion of Section 3, it is not altogether surprising that third parties can be used to enhance a simple property rights outcome (or for that matter to achieve the first-best).²⁵ However, we argued in Section 3 that third parties may be problematic because they are vulnerable to collusion. These considerations apply with equal force to the property rights model presented here. Thus, in practice B and S may find it difficult

²⁵MT show that their scheme generalizes to the case where B invests as well as S.

to use a third party in the way that MT suggest.

Even in the absence of third parties, MT's joint ownership scheme has interesting properties. In particular, set $F = 0$ (i.e., eliminate the third party). Then the difference between S's gross payoffs in states 1 and 2 equals P , which can be set as high as $\hat{v} - c_1$. Thus, the same outcome can be achieved with joint ownership, plus an option to sell, as by letting S have 100% ownership of the asset: in both cases, S will solve (4.1).

This is in fact MT's second point: there may be more than one allocation of property rights that sustains the second-best optimum.²⁶ However, this observation does not seem terribly damaging to the property rights approach. First, the theory is still capable of ruling out most allocations of property rights (for example, in the case where only S invests, ownership structures in which S (resp., B) owns B's assets with probability ρ (resp., $(1 - \rho)$) are suboptimal for all $0 \leq \rho < 1$). Second, MT's joint ownership contract seems very fragile. P must be chosen so that B has an incentive to exercise his option in state 2 but not in state 1. However, if, say, \hat{v} is stochastic with support (c_1, c_2) , then B will exercise his option to sell with positive probability in state 1, which will reduce S's incentives to invest. In contrast, the contract where S owns B's assets with probability 1 is robust to the introduction of uncertainty in \hat{v} . In fact, we conjecture that the indeterminacy in optimal ownership structure will be much reduced, and may disappear, in a world of uncertainty.

In concluding this section, we should point out that the above analysis falls some way short of providing a fully satisfactory foundation for a theory of ownership. The property rights approach takes the view that an

²⁶The reader may wonder whether the solution to (4.1) does indeed represent the second-best optimum in the two-party case where only S invests; or whether the parties could do better by having ownership of S's assets be a function of verifiable messages sent by the parties at date 1. The answer is that they cannot do any better (under the assumption that S's gross payoff in the absence of renegotiation with B, \hat{v} , is nonverifiable). This can be demonstrated using the results of Maskin and Moore (1998).

owner has residual control rights. However, the above model does not distinguish between specific and residual rights, and in fact equates residual rights with complete control rights (in particular, the right to have access to B's downstream technology). A more satisfactory model would proceed by assuming that certain decisions need to be taken with respect to assets, some of which can be described in advance (the specific control rights), but others of which cannot. (These nonspecifiable decisions are similar to the characteristics of the widget that cannot be described in advance in Section 2, Case ND.) Suppose that there is a large number of hard-to-describe uses of the assets, only one of which, say, will be relevant in a particular state of nature; and the states of nature are equally likely. Then it seems probable that an analysis along the lines of Section 2 will lead to the conclusion that the best the parties can do is to allocate residual control rights over assets, i.e., ownership rights. However, a formal demonstration of this must await further work.

5. Interpreting Contractual Incompleteness

A question that is sometimes asked about incomplete contracting models is: are the optimal contracts really incomplete? To put it another way, what is an incomplete contract?²⁷

Some contracts are manifestly incomplete in the sense that they leave something out or are ambiguous.²⁸ For example, consider a contract that says that S must supply B with a widget on February 29, 1998, even though no such date exists. Or, to give a deeper example, consider a specific performance contract which says that S must supply B with a particular widget, but which does not indicate the damages if performance turns out to be impossible. Incompleteness like this is very common in reality, but unfortunately it is very hard to model. It would seem necessary to assume that the parties are

²⁷For illuminating discussions of this question, see Ayres and Gertner (1992) and Tirole (1998).

²⁸Ayres and Gertner (1992) call such contracts "obligationally incomplete."

boundedly rational in the sense that they do not foresee even relatively obvious events. In contrast, in this paper, we have assumed that the parties are constrained in contracting only by the fact that complicated states of nature cannot be verified.²⁹

The contracts that we have derived in this paper are therefore not incomplete in the above sense. In particular, the parties' obligations are fully specified in all circumstances. This is true even of the null contract that was (approximately) optimal in Proposition 2. The null contract is complete in that it is absolutely clear what everybody's obligations are: nobody has any!

However, there is another sense in which one can say that a contract is incomplete: it is incomplete if the parties would like to add contingent clauses, but are prevented from doing so by the fact that the state of nature cannot be verified (or because states are too expensive to describe ex ante).³⁰ For example, a contract that says that S must supply X widgets to B at a fixed price (and pay huge damages if she does not supply) is incomplete if the parties would really have liked to make the number of widgets contingent on the state.

Viewed in this way, the optimal contract in Proposition 2 (under the noncommitment assumption) -- that is, no contract -- is incomplete. It is true that the parties' obligations are fully specified and that renegotiation at date 1 always "completes" the contract (i.e., makes it contingent). However, the way the contract is completed is not optimal from an ex ante

²⁹Actually, it is not entirely clear that parties who write obligatorily incomplete contracts are boundedly rational. Douglas Baird has suggested that it may be rational for parties to write contracts with missing provisions or ambiguities, to the extent that they anticipate that the courts will fill in the gaps or remove the ambiguities. Viewed this way, obligatorily incomplete contracts and the "insufficiently state contingent" contracts described below are not fundamentally different.

³⁰Ayres and Gertner (1992) refer to such contracts as "insufficiently state contingent."

perspective. The parties would like to ensure that the price of the special widget is independent of S's cost, but, as we have seen, this may not be compatible with their ex post incentive constraints.

Of course, at some level this is all a matter of semantics -- one could just as well call the contracts in Section 2 "optimal complete contracts subject to commitment and incentive constraints (and possibly also describability constraints)." However, we believe that there is a qualitative difference between the contracts of Section 2 and the complete or comprehensive contracts studied in the traditional mechanism design literature (including that based on asymmetric information or moral hazard) -- not least because the contracts of this paper provide the beginnings of a foundation for a theory of ownership or property rights -- and thus it is reasonable to have a different term for them.

In this Appendix we generalise Proposition 2 to the case where both parties make investments at date 1/2: B invests β , costing $C^b(\beta)$; and S invests σ , costing $C^s(\sigma)$. Now β and σ may be multi-dimensional. Also, we allow for any division of surplus in the date 1 bargaining. Specifically, we suppose that, if the outcome of the contractually-specified mechanism is inefficient, then with probability λ S makes a take-it-or-leave-it offer to B, and with probability $1-\lambda$ B makes a take-it-or-leave-it offer to S.

As in the text, there are N widgets, numbered $1, \dots, N$, which can be described at date 0 (we are in Case D).³¹ In each state of nature at date 1, one of these widgets is special, yielding value v to B and costing S c to produce. The stochastic mapping from (β, σ) to (v, c) is arbitrary, except that we suppose v and c are always ranked and bounded: there exist $\bar{v} < \infty$ and $\underline{c} \geq 0$ such that $\bar{v} \geq v \geq c \geq \underline{c}$ with probability 1.³² The fact that the mapping from (β, σ) to (v, c) is arbitrary means that we can allow for any degree of correlation between v and c , and also for externalities where B's investment β affects S's cost c , or where S's investment σ affects B's value v .

The remaining $N-1$ widgets are generic, and, for $n = 1, \dots, N-1$, have cost $g_n = \underline{c} + \frac{n}{N}(\bar{v} - \underline{c})$. To simplify matters, we now assume that the value of each of these special widgets equals its cost.³³

³¹As in the Corollary to Proposition 2, our result holds a fortiori if the N widgets cannot be described in advance (Case ND).

³²The restriction $v \geq c$ is made so that, at worst, the special widget is like a generic widget for which value equals cost (see below). Our results still hold if the value of a widget (special or generic) to B is strictly less than the cost to S. In such a case, the parties will typically renegotiate in order to avoid inefficient trade.

³³As in Section 2, the exact specification of these costs/values is not important. What matters is that as the number of widgets, N , increases, no

We assume that there is complete symmetry among the widgets at date 0, in the sense that each widget is equally likely to be the special widget or to be one of the $N-1$ generic widgets. (This uniform distribution over permutations is unaffected by β and σ .)

At date 1, both B and S observe the state: the realised permutation of the N widgets, and the cost c to S and the value v to B of the special widget. However, no-one else observes the state. In other words, the state is observable (to the two parties) but not verifiable (to outsiders, such as the courts).

As a preliminary exercise, suppose a contractual mechanism specifies that some widget W is traded. If the state of nature is such that, first, W is a generic widget with cost/value g_n , and, second, the special widget has cost c to S and value v to B, then, following renegotiation, S's final payoff will be $-g_n + \lambda(v - c)$, and B's final payoff will be $g_n + (1 - \lambda)(v - c)$. (These payoffs are gross of any transfer that the mechanism might specify.) Equally, if the mechanism specifies that there is no trade, then, following renegotiation, S's final payoff will be $\lambda(v - c)$, and B's final payoff will be $(1 - \lambda)(v - c)$.

To provide a benchmark for Proposition 2*, let β^0 and σ^0 denote the equilibrium investment levels at date 1/2 if no contract were written at date 0 and the terms of trade were bargained from scratch at date 1 -- in other words, if the "null contract" were in place. Under the null contract, for a given realization of c and v , the trade price equals $(1-\lambda)c + \lambda v$, S's payoff equals $\lambda(v - c)$, and B's payoff equals $(1 - \lambda)(v - c)$. So β^0 and σ^0 jointly solve

large "gaps" in cost or value remain between \underline{c} and \bar{v} . (In fact, if one knows the value of λ , this range can be reduced; but, by "covering" the full range from \underline{c} to \bar{v} , we ensure that Proposition 2* below holds for all λ .) Also, it makes no difference if there are other generic widgets whose costs/values lie outside this range.

$$(A.1) \quad \beta^0 = \operatorname{argmax}_{\beta} E[(1 - \lambda)(v - c) - C^b(\beta) | \beta, \sigma^0],$$

$$\sigma^0 = \operatorname{argmax}_{\sigma} E[\lambda(v - c) - C^s(\sigma) | \beta^0, \sigma],$$

where $E[\cdot | \beta, \sigma]$ denotes the expectation operator with respect to the date 1 joint distribution of c and v conditional on investments β and σ having been made at date 1/2. We assume that the solution (β^0, σ^0) to (A.1) is unique.

Proposition 2*. Suppose Case D holds. If the parties cannot commit not to renegotiate, then, irrespective of the contract, as $N \rightarrow \infty$ their investments β and σ approach the values β^0 and σ^0 given by (A.1). That is, in the limit, contracts cannot make any difference to expected total surplus, and the parties may as well use the null contract.

Proof

Take any abstract mechanism M . A state is defined by the permutation τ , say, of the widgets, and the cost c to S and value v to B of the special widget. In state (c, v, τ) , let the equilibrium strategies of M for B and S be $\mu^b(c, v, \tau)$ and $\mu^s(c, v, \tau)$. And, following any renegotiation, let $p(c, v, \tau)$ denote the overall price that B pays S for delivery of the special widget: $p(c, v, \tau)$ equals any transfer specified in the mechanism plus any amount agreed by the parties during renegotiation. In other words, B 's and S 's final equilibrium payoffs are respectively $v - p(c, v, \tau)$ and $p(c, v, \tau) - c$.

Let $P(c, v)$ denote the expected value of $p(c, v, \tau)$ taken over all permutations τ (which are equally probable).

Without loss of generality, we may suppose that τ is such that widget 1 is the special widget, and that widgets 2, ..., N are generic widgets costing $\underline{c} + \frac{1}{N}(\bar{v} - \underline{c})$, ..., $\underline{c} + \frac{N-1}{N}(\bar{v} - \underline{c})$ respectively.

Now consider another state (c^*, v^*, τ^*) , where τ^* is a rotation of τ :

widgets 1, ..., N-1 are generic widgets costing $\underline{c} + \frac{1}{N}(\bar{v} - \underline{c})$, ..., $\underline{c} + \frac{N-1}{N}(\bar{v} - \underline{c})$ respectively, and widget N is the special widget. Without loss of generality, we suppose that

$$(A.2) \quad (1 - \lambda)(c^* - c) + \lambda(v^* - v) \geq 0.$$

The proof of Proposition 2* proceeds via two lemmas, which together pin down the size of the gap between the expected prices, $P(c^*, v^*) - P(c, v)$. The first lemma provides a lower bound.

Lemma 1. $P(c^*, v^*) - P(c, v) \geq (1 - \lambda)(c^* - c) + \lambda(v^* - v) - \frac{1}{N}(\bar{v} - \underline{c}).$

Proof of Lemma 1

Consider the outcome of the mechanism M if B plays strategy $\mu^b(c^*, v^*, \tau^*)$ and S plays strategy $\mu^s(c, v, \tau)$. Suppose M specifies: B pays S an amount q ; widget $n = 1, \dots, N$ is traded with probability $\alpha_n \geq 0$; and there is no trade with probability $1 - (\alpha_1 + \dots + \alpha_N) \geq 0$. (As both parties are risk neutral, there is no loss of generality in not having the transfer q depend on which widget -- if any -- the mechanism specifies is traded.)

There are two incentive constraints that $q, \alpha_1, \dots, \alpha_N$ must satisfy. First, in state (c^*, v^*, τ^*) , S must not have an incentive to deviate to $\mu^s(c, v, \tau)$. Second, in state (c, v, τ) , B must not have an incentive to deviate to $\mu^b(c^*, v^*, \tau^*)$.

Suppose strategy pair $\mu^b(c^*, v^*, \tau^*), \mu^s(c, v, \tau)$ is played. The contractually-specified outcome is typically inefficient, and will be renegotiated: when the outcome of the lottery $(\alpha_1, \dots, \alpha_N)$ specifies either that a generic widget is traded, or that there is no trade, the parties must bargain in order to exploit the gains from trading the special widget. The final payoffs depend on the state. In state (c^*, v^*, τ^*) , following the play of strategy pair $\mu^b(c^*, v^*, \tau^*), \mu^s(c, v, \tau)$, S's final payoff is

$$\begin{aligned}
q &+ \alpha_1 \left(-\underline{c} - \frac{1}{N}(\bar{v} - \underline{c}) + \lambda(v^* - c^*) \right) + \dots \\
&+ \alpha_{N-1} \left(-\underline{c} - \frac{N-1}{N}(\bar{v} - \underline{c}) + \lambda(v^* - c^*) \right) + \alpha_N \left(-c^* \right) \\
&+ (1 - \alpha_1 - \dots - \alpha_N) \left(\lambda(v^* - c^*) \right),
\end{aligned}$$

which, according to the first incentive constraint, cannot be more than what she gets in equilibrium, $p(c^*, v^*, \tau^*) - c^*$. And in state (c, v, τ) , following the play of strategy pair $\mu^b(c^*, v^*, \tau^*)$, $\mu^s(c, v, \tau)$, S's final payoff is

$$\begin{aligned}
q &+ \alpha_1 \left(-c \right) + \alpha_2 \left(-\underline{c} - \frac{1}{N}(\bar{v} - \underline{c}) + \lambda(v - c) \right) + \dots \\
&+ \alpha_N \left(-\underline{c} - \frac{N-1}{N}(\bar{v} - \underline{c}) + \lambda(v - c) \right) \\
&+ (1 - \alpha_1 - \dots - \alpha_N) \left(\lambda(v - c) \right),
\end{aligned}$$

which, according to the second incentive constraint, cannot be less than what she gets in equilibrium, $p(c, v, \tau) - c$, since if S were worse off B would be better off (all final payoffs lie along the Pareto frontier).

Combining these two constraints, we have

$$\begin{aligned}
p(c^*, v^*, \tau^*) - p(c, v, \tau) &= (1 - \lambda)(c^* - c) - \lambda(v^* - v) + \frac{1}{N}(\bar{v} - \underline{c}) \\
&\geq \alpha_1 \left(c - \underline{c} + \lambda(v - c) \right) + \alpha_N \left(\bar{v} - v^* + (1 - \lambda)(v^* - c^*) \right) \\
&+ (1 - \alpha_1 - \dots - \alpha_N) \left(\frac{1}{N}(\bar{v} - \underline{c}) \right),
\end{aligned}$$

which is nonnegative, given that $\bar{v} \geq v^* \geq c^* \geq \underline{c}$ and $\bar{v} \geq v \geq c \geq \underline{c}$.

Since this lower bound applies for any permutation τ (and associated rotation τ^*), we can take expectations across permutations to deduce that

$$P(c^*, v^*) - P(c, v) \geq (1 - \lambda)(c^* - c) + \lambda(v^* - v) - \frac{1}{N}(\bar{v} - \underline{c}).$$

This proves Lemma 1.

The second lemma provides an upper bound for the gap between the expected prices, $P(c^*, v^*) - P(c, v)$.

Lemma 2. $P(c^*, v^*) - P(c, v) \leq (1 - \lambda)(c^* - c) + \lambda(v^* - v)$.

Proof of Lemma 2

Consider the outcome of the mechanism M if B plays strategy $\mu^b(c, v, \tau)$ and S plays strategy $\mu^s(c^*, v^*, \tau)$. (Note that, in contrast to the proof of Lemma 1, the selected states here differ only in the costs and values of the special widget; in both states, the permutation of the widgets is τ . And recall that the choice of τ is without loss of generality.) Suppose M specifies: B pays S an amount \hat{q} ; widget $n = 1, \dots, N$ is traded with probability $\hat{\alpha}_n \geq 0$; and there is no trade with probability $1 - (\hat{\alpha}_1 + \dots + \hat{\alpha}_N) \geq 0$.

There are two incentive constraints that $\hat{q}, \hat{\alpha}_1, \dots, \hat{\alpha}_N$ must satisfy. First, in state (c, v, τ) , S must not have an incentive to deviate to $\mu^s(c^*, v^*, \tau)$. Second, in state (c^*, v^*, τ) , B must not have an incentive to deviate to $\mu^b(c, v, \tau)$.

Suppose strategy pair $\mu^b(c, v, \tau), \mu^s(c^*, v^*, \tau)$ is played. The contractually-specified outcome is typically inefficient, and will be renegotiated: when the outcome of the lottery $(\hat{\alpha}_1, \dots, \hat{\alpha}_N)$ specifies either that a generic widget is traded, or that there is no trade, the parties must bargain in order to exploit the gains from trading the special widget. The

final payoffs depend on the state. In state (c, v, τ) , following the play of strategy pair $\mu^b(c, v, \tau), \mu^s(c^*, v^*, \tau)$, S's final payoff is

$$\begin{aligned} \hat{q} + \hat{\alpha}_1(-c) + \hat{\alpha}_2\left(-c - \frac{1}{N}(\bar{v} - c) + \lambda(v - c)\right) + \dots \\ + \hat{\alpha}_N\left(-c - \frac{N-1}{N}(\bar{v} - c) + \lambda(v - c)\right) \\ + (1 - \hat{\alpha}_1 - \dots - \hat{\alpha}_N)\left(\lambda(v - c)\right), \end{aligned}$$

which, according to the first incentive constraint, cannot be more than what she gets in equilibrium, $p(c, v, \tau) - c$. And in state (c^*, v^*, τ) , following the play of strategy pair $\mu^b(c, v, \tau), \mu^s(c^*, v^*, \tau)$, S's final payoff is

$$\begin{aligned} \hat{q} + \hat{\alpha}_1(-c^*) + \hat{\alpha}_2\left(-c - \frac{1}{N}(\bar{v} - c) + \lambda(v^* - c^*)\right) + \dots \\ + \hat{\alpha}_N\left(-c - \frac{N-1}{N}(\bar{v} - c) + \lambda(v^* - c^*)\right) \\ + (1 - \hat{\alpha}_1 - \dots - \hat{\alpha}_N)\left(\lambda(v^* - c^*)\right), \end{aligned}$$

which, according to the second incentive constraint, cannot be less than what she gets in equilibrium, $p(c^*, v^*, \tau) - c^*$, since if S were worse off B would be better off.

Combining these two constraints, we have

$$\begin{aligned} p(c^*, v^*, \tau) - p(c, v, \tau) - (1 - \lambda)(c^* - c) - \lambda(v^* - v) \\ \leq -\hat{\alpha}_1\left((1 - \lambda)(c^* - c) + \lambda(v^* - v)\right), \end{aligned}$$

which is nonpositive by (A.2).

Since this upper bound applies for any permutation τ , we can take expectations across permutations to deduce that

$$P(c^*, v^*) - P(c, v) \leq (1 - \lambda)(c^* - c) + \lambda(v^* - v).$$

This proves Lemma 2.

Together, Lemma 1 and Lemma 2 show that as $N \rightarrow \infty$, the gap between $P(c^*, v^*)$ and $P(c, v)$ approaches $(1 - \lambda)(c^* - c) + \lambda(v^* - v)$, whatever the contractually-specified mechanism. But this means that, up to an additive constant, the expected price paid by B to S is the same as if the parties bargained from scratch at date 1 -- as if the null contract were in place. This in turn means that B and S will make equilibrium investments β^0 and σ^0 at date 1/2. Proposition 2* is proved.

Q.E.D.

REFERENCES

- AYRES, I. and GERTNER, R. (1992), "Strategic Contractual Inefficiency and the Optimal Choice of Legal Rules," *The Yale Law Journal*, 101, 729-773.
- BARANY, I. (1992), "Fair Distribution Protocols or How the Players Replace Fortune," *Mathematics of Operations Research*, 17, 327-340.
- CHE, Y.-K. and HAUSCH, D. (1998), "Cooperative Investments and the Value of Contracting," *American Economic Review*, forthcoming.
- GROSSMAN, S. and HART, O. (1986), "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94, 691-719.
- HART, O. (1995) *Firms, Contracts, and Financial Structure*, (Oxford University Press).
- HART, O. and MOORE, J. (1990), "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119-1158.
- MASKIN, E. and MOORE, J. (1998), "Implementation and Renegotiation," *Review of Economic Studies*, forthcoming.
- MASKIN, E. and TIROLE, J. (1998a), "Unforeseen Contingencies and Incomplete Contracts," *Review of Economic Studies*, forthcoming.
- MASKIN, E. and TIROLE, J. (1998b), "Two Remarks on the Property-Rights Literature," *Review of Economic Studies*, forthcoming.
- MOORE, J. (1992), "Implementation in Environments with Complete Information," in Laffont, J. J. (ed.), *Advances in Economic Theory* (Cambridge University Press).
- SEGAL, I. (1995), "Essays on Commitment, Renegotiation, and Incompleteness of Contracts" (Ph.D. dissertation, Harvard University).

SEGAL, I. (1998), "Complexity and Renegotiation: A Foundation for Incomplete Contracts," *Review of Economic Studies*, forthcoming.

TIROLE, J. (1998), "Incomplete Contracts: Where Do We Stand?" (mimeo, Institut d'Economie Industrielle, Toulouse).

WILLIAMSON, O. (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, (New York: Free Press).