



Incomplete Contracts and Control

Prize Lecture, December 8, 2016

by Oliver Hart

Department of Economics, Harvard University, USA

1. INTRODUCTION

The work on incomplete contracts cited by the prize committee began in the summer of 1983, but it may be useful to say a bit about how I reached that point. As a graduate student, first at the University of Warwick and then at Princeton University, with a degree in mathematics behind me, I was drawn to general equilibrium theory, and my Ph.D. thesis was on general equilibrium theory with incomplete markets.¹ Although I ended up focusing on optimality and existence problems that could arise even in exchange economies, one of my primary interests was the theory of production. In a complete market Arrow-Debreu economy with perfect competition, it makes sense for a firm to maximize profit or net market value. But with incomplete markets, what is the generalization of this goal? More fundamentally, what happens if shareholders disagree about what the firm should do?²

I started to work on this topic after my thesis and as a result of a serendipitous assignment continued the work with Sanford Grossman in the summer of 1976.³ At some point we decided that, interesting though the disagreement between shareholders was, an empirically more important conflict was that between managers and shareholders.⁴ This led first to work on corporate takeovers as a mechanism for disciplining management and then to a paper on debt as a bonding device.⁵ At some stage we realized that since we were studying ways to incentivize management, maybe we should analyze directly the optimal incentive scheme between an owner and a manager. This led to Grossman and Hart (1983), a paper squarely in the principal-agent tradition.

This rather circuitous path helps to explain how my thinking evolved from markets to contracts as the unit of analysis, and provides the backdrop to the summer of 1983. Sitting in Grossman's University of Chicago office, the two of us were considering what to work on next. After some discussion we decided that a question that was ripe for analysis was: Why would one firm ever buy another firm rather than conduct business with that firm through a contract? In other words, what are the limits of contracts and why do we have firms?

Of course, this was hardly a new question: there is a literature on the boundaries of the firm that goes back to Coase (1937) and includes Oliver Williamson's many works (see, e.g., Williamson (1975)), and Klein et al. (1978). I think that it is fair to say that we were aware of this literature without being intimate with it. (Intimacy came later.) One thing that we knew for sure is that the literature was informal (or in the prize committee's felicitous language, "not formalized"). As economic theorists with a formal training we thought that we might be able to add something.

We worked on the difference between firms and contracts for ten very intense days. With apologies to John Reed these were ten days that shook my world.⁶ I recall that initially we thought that the difference had to do with authority. An employer can choose the task of an employee.⁷ But what is the conceptual difference between this and a requirements contract between two firms where one, the buyer say, can choose how many widgets q to buy from the other, the seller, with payment determined by a pre-agreed schedule: $p = p(q)$? One could say that the buyer has authority over q in this case. Perhaps more seriously, once we are in a world where a buyer obtains private information about a demand shock and a seller obtains private information about a supply shock, then, according to mechanism design theory, the quantity q should depend on *both* parties' reports about their shocks: neither party should have authority over q .

At some point it dawned on us that we were thinking about the issue the wrong way. We were viewing the problem in complete contracting terms. But what if the contract between the buyer and seller is incomplete?

2. INCOMPLETE CONTRACTS

The formal literature to that point was all about complete contracts. These are contracts where everything that can ever happen is written into the contract. There may be some incentive constraints arising from moral hazard or asymmetric information but there are no unanticipated contingencies.

Actual contracts are not like this, as lawyers have realized for a long time. They are poorly worded, ambiguous, and leave out important things. They are

incomplete. At some stage Grossman and I realized that a critical question that arises with an incomplete contract is, who has the right to decide about the missing things? We called this right the residual control or decision right. The question is, who has it?

Further thought led us to the idea that this is what ownership is. The owner of an asset has the right to decide on how the asset is used to the extent that its use is not contractually specified. This naturally leads to a theory of the difference between contracts and firms. Think of a firm as consisting of assets. If firm A and firm B sign an arms-length (incomplete) contract, then the owner of firm A has residual control rights over the A assets and the owner of firm B has residual control rights over the B assets. In contrast if, say, firm A buys firm B, then the owner of firm A has residual control rights over the A *and* B assets.

Why should it matter who has residual control rights? Residual control rights are like any other good: there is an optimal allocation of them. Sometimes it is more efficient for one owner to hold all the residual control rights, and sometimes it is more efficient for these control rights to be split between several owners. Which is the case will determine whether firms A and B should merge or stay as separate entities.

Grossman and I constructed a formal model along these lines (see Grossman and Hart (1986)), and I developed the ideas and model further in work with John Moore (see Hart and Moore (1990)). Collectively these papers are often referred to as “property rights theory” (PRT).

It is useful to illustrate the model with a real-world example. Consider a power plant that locates next to a coal mine with the purpose of burning coal to make electricity.⁸ One way to regulate the transaction is for the power plant to sign an arms-length long-term contract with the coal mine. Such a contract would specify the quantity, quality, and price of coal for many years to come. But any such contract will be incomplete. Events will occur that the parties could not foresee when they started out.

For example, suppose that the power plant needs the coal to be pure but that it is hard to specify in advance what purity means, given that there are many potential impurities. Imagine that ten years into the relationship, ash content is the relevant impurity and that high-ash-content coal is more expensive for the power plant to burn than low-ash-content coal but cheaper for the coal mine to produce. Given that the contract is incomplete, the coal mine may be within its rights under the contract to supply high-ash-content coal.

The power plant and coal mine can, of course, renegotiate the contract. However, the coal mine is in a strong bargaining position. It can demand a high price for switching to low-ash-content coal. The reason is that the power plant does not

have a good alternative: it may be very expensive for the power plant to transport coal from a different coal mine given that it is located next to this one.

Economists refer to this situation as the “hold-up” problem. The coal mine can hold up the power plant because the power plant, by locating next to the coal mine, has become dependent on it. The next point to realize is that although it may be impossible to write a contract that is complete enough to avoid hold-up, this does not mean that the parties will be unable to anticipate hold-up. Indeed, the assumption made in the theory is that the power plant *does* anticipate that it will be at the mercy of the coal mine, and that a substantial share of its future profit may be expropriated by the coal mine. Fearing such expropriation, the power plant may choose not to become so dependent on the coal mine in the first place. For example, it may locate at an equal distance between several coal mines rather than right next to this one, even though this may increase the cost of transporting coal.

At the risk of belaboring the point, it is worth pinpointing the source of the coal mine’s hold-up power. It arises because the owner of the coal mine has residual rights of control over the mine. In this case the key residual right of control is the decision about what kind of coal to mine: high-ash-content or low-ash-content.

Is there anything the power plant can do to avoid this situation? Short of writing a better contract, one thing it can do is to buy the coal mine in advance. That way the power plant as owner of the coal mine will have the key residual control right. The coal mine can no longer extract a high price by threatening to produce high-ash-content coal: the power plant can order the coal mine manager to mine low-ash-content coal. In an extreme case, if the coal mine manager threatens to disobey the order, the power plant can fire the manager and replace her with someone else.

One consequence is that the power plant may now be willing to become dependent on the mine. Given that it does not fear hold-up, it may locate next to it. Thus the theory identifies a benefit of integration, where integration in this case means the purchase of the coal mine by the power plant. The value of integration is that the power plant may undertake efficiency-enhancing relationship-specific investments—in this case locating next to the coal mine—that it would not carry out if it was protected only by an incomplete contract.

So far we have discussed the benefits of integration. But just as the transfer of residual control rights from the coal mine to the power plant empowers the owner of the power plant, it disempowers the owner of the coal mine, and this is likely to have costs in terms of *her* incentive to make relationship-specific investments. Assume that the coal mine was previously an owner-managed firm. After

the acquisition by the power plant the coal mine manager stays on but is now an employee of the power plant. Suppose that the coal mine manager has an idea about how to run the mine more efficiently. When the coal mine was separate, the manager had the authority (residual control rights) to implement, and benefit from, this idea. Now that the manager is an employee she has to get permission to implement the idea from her boss: the owner of the power plant has veto power. The owner of the power plant can use his veto right to extract some of the gains from the idea for himself. Knowing that she faces the risk of expropriation, the coal mine manager's incentive to innovate is reduced.

So integration has costs as well as benefits. Whether it makes sense for the power plant to purchase the coal mine will depend on whether the distortion in the power plant's investment is more important than the distortion in the coal mine manager's investment. It is also worth noting that a further possibility is for the coal mine to purchase the power plant. This is not the same as the power plant purchasing the coal mine since now the residual control rights are concentrated in the hands of the coal mine manager rather than the power plant manager. Finally, the theory can be generalized beyond the case of two managers and two assets to many assets and many workers, and to more general ownership structures, such as joint and shared ownership; see Hart and Moore (1990). That paper also shows that synergistic assets should be owned together and that assets should be owned by indispensable human capital.

Let us pause to make some observations. First, the formal models in Grossman and Hart (1986) and Hart and Moore (1990) take the view that *ex post* renegotiation of an incomplete contract occurs under conditions of symmetric information—both parties can see what has been left out of the contract—and that, given that there are no wealth constraints, the bargaining proceeds efficiently, as in Coase (1960). Inefficiencies arise solely because *ex ante* relationship-specific investments are distorted.

Second, the distortion in *ex ante* investments can be overcome if these investments are contractible. In this case the parties can write a contract that specifies, say, that the power plant must locate next to the coal mine in return for an upfront payment: in effect the mine compensates the power plant for its later hold-up power. For the theory to work, one has to suppose that some aspects of the investment are not contractible (or are costly to contract on): for example, even if the location decision is contractible, whether the power plant installs a boiler that burns this particular mine's coal efficiently might not be. Similarly, one must assume that the coal mine manager's investment in innovation is non-contractible (a highly plausible assumption).

Third, the focus on distortions in noncontractible investments distinguishes

the approach from those of Williamson and Klein et al. Most of Williamson's work is concerned with *ex post* bargaining inefficiencies and how integration can reduce these. Klein et al. do discuss *ex ante* inefficiencies but do not distinguish between contractible and noncontractible investments.⁹

Fourth, the prior work of Coase and Williamson emphasizes authority over human capital as the defining feature of the firm: An employer can tell an employee what to do. In contrast, PRT emphasizes control over physical (more generally non-human) assets. When the power plant purchases the coal mine it acquires residual control rights over the mine. To see the difference, note that, according to PRT, purchasing the mine would not be worth much if the coal mine manager is indispensable. In that case the manager would retain her hold-up power even as an employee. If the power plant wants a shift from high-ash-content coal to low-ash-content coal, the coal mine manager could demand a huge increase in salary for doing this. It is because typically the coal mine manager is somewhat replaceable that the power plant is in a stronger bargaining position after it has acquired the mine than before.

3. APPLICATION TO FINANCIAL CONTRACTING

As well as helping us to understand asset ownership and firm boundaries, PRT has a number of applications. One is to financial contracting.¹⁰

Understanding the financial structure of a firm has been challenging since Modigliani and Miller (1958) showed that under some plausible assumptions, a firm's financial structure has no effect on its total value. One strand of the literature (notably, Jensen and Meckling (1976)) argues that the Modigliani-Miller irrelevance result no longer holds if managers cannot be relied on to act on behalf of shareholders. However, a problem with this approach is that it supposes that financial structure is used to solve an incentive problem. Once incentive schemes are allowed to fulfill this task the irrelevance result is restored.

PRT offers a different perspective in which financial structure is thought of in control terms.¹¹ To understand the approach, replace the power plant by a financial investor. Specifically, suppose that the coal mine needs money to expand/modernize, and approaches an investor with deep pockets. How does it persuade this person to invest?

One possibility is to offer him a share of the future profit from the coal mine. But this may not be enough. The reason is that the financial contract between the investor and the coal mine is likely to be incomplete. There are many actions or decisions that will be taken during the course of the relationship that the contract will not (cannot) specify.

For example, the investor may worry that the manager of the coal mine will divert earnings: she could pay herself a large salary or reinvest profits rather than paying dividends. Another possibility is that the manager could adopt a strategy for running the coal mine that the investor does not approve of. Or the manager might hold on to her position as CEO even when a different CEO might be better.

Opportunistic behavior is similar to hold-up in the previous analysis. One way to protect the investor against such behavior is to give him residual control rights or votes. For example, the investor could become the owner of the coal mine, rather than having an arms-length contract with the coal mine. This would allow him to intervene to stop opportunistic behavior, e.g., he could control the manager's salary or replace her.

But as we have seen, there can be a downside to taking away control from the manager. According to the previous analysis one cost is that the manager's incentive to have ideas may be reduced. Thus there will be an optimal balancing of control between the investor and the manager.

This optimal balancing of control has been analyzed in an important paper by Aghion and Bolton (1992). Aghion and Bolton dispense with noncontractible investments by the manager (the incentive to have ideas) and focus instead on her private benefits.¹² These private benefits include the psychic satisfaction from pursuing a vision for the company, the job satisfaction from being CEO, and the remuneration associated with being in a position of power. The private benefits are something that only the manager enjoys; they cannot be transferred to the investor. In contrast, monetary returns are verifiable and can be transferred to the investor.

In the Aghion-Bolton model the cost of allocating control to the investor is that the investor may pursue a ruthless profit-maximizing strategy that destroys managerial private benefits. The manager could try to offer a side-payment to the investor to persuade the investor not to choose such a strategy, but the problem is that the manager is wealth constrained and so her ability to renegotiate with the investor is limited. The Coase theorem fails because one party is wealth-constrained. Thus there is a fundamental trade-off. On the one hand, allocating all the control rights to the manager means that the manager may pursue private benefits at the expense of profit; as a result, the investor may not be compensated adequately and may not invest in the project. On the other hand, allocating all the control rights to the investor means that ex post decisions may not be first-best efficient.

Aghion and Bolton show that under some assumptions, the solution to this trade-off is to make control state-contingent. Specifically, the manager will have control in states of the world where private benefits are important relative to

profits and the investor in states of the world where profits are important relative to private benefits. To the extent that private benefits do not vary much with the state, but profits do, this suggests that the investor should have control in bad states of the world. In a bad state the manager may want to keep the company going to preserve her private benefits even if the assets could have much greater value if deployed elsewhere.

Significant support for the Aghion-Bolton model can be found in Kaplan and Stromberg's work on venture capital contracts (see Kaplan and Stromberg (2003)). Kaplan and Stromberg study start-up deals in the information technology, software, and telecommunication sectors. They find that the allocation of voting and control rights is frequently made contingent on verifiable measures of financial performance. For instance, the venture capitalist may obtain voting or board control if the firm's earnings before interest and taxes fall below a pre-specified level or if the firm's net worth falls below a threshold. If the firm performs poorly, the venture capitalist obtains full control. As firm performance improves, the entrepreneur retains/obtains more control rights, and, if the firm performs very well, the venture capitalist relinquishes most of his control rights.

One interesting feature of both the Aghion-Bolton model and the Kaplan-Stromberg study is that control does not shift to the investor as a result of the manager's failure to make a promised payment. Rather control shifts because a particular state of the world occurs. In other words, the financial contract does not correspond to a classic debt contract. In the venture capital context one reason for this may be that start-ups do not generate significant cash flows for a while. Still debt contracts are ubiquitous in other settings and it is desirable to explain them.

One attempt to do so is contained in Hart and Moore (1998) and Hart and Moore (1994).¹³ In Hart and Moore (1998), the assumption that monetary returns are verifiable and transferable is dropped. Instead it is supposed that the manager can walk off with them. What can persuade the manager to pay over some of these cash flows to the investor is that the investor has a threat: he can seize the assets underlying the project and liquidate them. Here liquidation means using the assets in some second-best manner, perhaps for some other activity or with a different manager. (It could refer to a sale of the assets.)

In this setting, Hart and Moore (1998) show that a debt contract works well. With a debt contract the manager of the coal mine promises to make a fixed stream of payments to the investor. As long as these payments are made the manager remains in charge, that is, she retains (residual rights of) control over the coal mine. If a payment is not made control shifts to the investor, who can decide whether to liquidate the mine. At this stage renegotiation is possible.

The motivation for the manager to make a debt payment is very simple: she wants to retain control of the assets. Why is control valuable? Because the manager can use the assets to produce future monetary returns that she can then pocket. To put it another way, there are two reasons why the manager may default on a debt payment. One is if she cannot make the payment: revenue is too small as a result of an adverse shock, say. This corresponds to an involuntary default. The other reason is that the manager does not want to make the payment. In turn there can be two explanations for this. The first is that future revenues, which the manager can pocket, are worth less than what she is being asked to pay. For example, suppose that the assets will last for one more period and will generate \$100, but the current debt payment is \$120. (Ignore discounting.) It is not worth it for the manager to pay \$120 to be able to earn \$100 in the future; it is better to default and pocket the \$120 now. The second explanation is that, even though the debt payment is less than the future revenues (say the debt payment is \$80), the manager may be able to default and renegotiate the payment down to close to the liquidation value of the assets (which might be \$60).

These last cases, where the manager can pay but won't pay, correspond to a voluntary or strategic default.

This model has several interesting features. First, it shows how important collateral is. An investor will be less concerned about strategic default if the liquidation value of the assets is high, since the manager cannot renegotiate the debt below this level. Thus the manager will be able to borrow more in this case and more good projects can go forward. Similarly, if the assets are durable—their liquidation value remains high over time—the maturity of the debt can be longer: the investor will not be vulnerable to strategic debt renegotiation late in the project's life. Empirical support for both these predictions has been obtained by Benmelech (2009) and Benmelech and Bergman (2008).¹⁴

Second, the model has the feature that inefficient liquidation can occur. Go back to the numerical example where the assets generate \$100 next period, the current debt payment is \$80, and the liquidation value is \$60. Suppose that current revenue is \$40. Clearly the manager will default since her \$40 does not cover the debt payment. The investor can liquidate for \$60 but the assets are worth more than this—\$100—if they are left in place. In an ideal world, a Coasian renegotiation would ensure that the assets are indeed left in place. In such a renegotiation the manager would compensate the investor for the \$60 liquidation value that he gives up by promising part of next period's \$100. However, the parties are not operating in an ideal world. The promise to pay part of the \$100 next period is not credible. Since this is the end of the project, and the assets will

have no further value, the investor knows that he will have no leverage then: the manager can pocket all the \$100 with impunity. Thus the only way for the investor to be paid is to liquidate now.

In Hart and Moore (1994), the assumption that the manager can pocket the monetary returns is replaced by the assumption that the manager can withdraw her human capital. Suppose that a project costs \$100 at date 0 and yields \$120 at date 2. The manager borrows the \$100 and promises to repay this amount at date 2. At date 1 the manager could threaten to withdraw her human capital unless the debt is reduced. If the parties have equal bargaining power, and the project has zero value without the manager, then the debt can be renegotiated down to \$60, and an investor who foresees this will not lend money. Collateral can again help here. If the assets have an alternative use at date 1, then the investor is at least partially protected against strategic default.

The Hart-Moore (1994) model reminds us again of the distinction between human and non-human assets. A project that consists mainly of human capital is difficult to finance because an investor is subject to hold-up by the human capital. Conversely, a project that has significant non-human assets can be financed without the fear of hold-up.¹⁵

In some interesting recent work, Kaplan et al. (2009) have investigated the importance of human vs. non-human assets in the context of start-ups. Their paper, whose title suggestively begins, “Should Investors Bet on the Jockey or the Horse?,” finds evidence that non-human assets in the form of a business plan are an important and durable source of value. However, Bernstein et al. (2017) find that human assets are also very important in the early stages of a start-up. Indeed, as Rajan (2012) suggests, the balance may shift over time: part of the transformation of a start-up into a successful, mature firm may be a standardization process that ensures that no individual’s human capital is that important.

This latest research is notable because it connects to one of the original motivations of Grossman and Hart (1986) and Hart and Moore (1990). Both papers start off with the same question: What is a firm? The answer given is that non-human assets are an important part of any firm; they are the glue that holds the firm together (see also Hart (1995)). The work by the above authors is helping to clarify this.¹⁶

4. APPLICATION TO PUBLIC VS. PRIVATE OWNERSHIP

Economists generally agree that there are some goods and services that will not be provided at an adequate level through private markets, and that therefore need to be financed by the government. Clear examples are national defense, the

police, foreign policy, and prisons. Examples about which there might be greater debate are health care and schools.

Government financing does not necessarily imply government provision. The government has a choice about whether to provide the good or service itself or to contract with a private provider. The incomplete contracting approach can be helpful in elucidating the trade-off.

Andrei Shleifer, Robert Vishny, and I (Hart et al. (1997)) explore the public-private choice, focusing particularly on prisons.¹⁷ Suppose that the government, acting on behalf of society, wants to incarcerate prisoners. The government may have several goals: to prevent the prisoners from escaping, to treat them humanely, and to maximize the chance that they can return to society as well-functioning citizens. The government can own a prison and use government employees to run it; or it can contract with a private company to run the prison. Which is better?

The first point to note is that in an ideal world, where everything can be anticipated and written into a contract, the choice does not matter since a complete contract will be written in both cases. In a complete contracting world, ownership and residual rights of control are irrelevant since all decisions are specified in the contract. The presence of asymmetric information and moral hazard does not change this conclusion: these factors simply lead to the addition of various incentive-compatibility constraints in the solution for the optimal contract.

When contracts are incomplete, residual control rights become important. The contracts that the federal or state governments write with private prison companies are in fact quite elaborate and cover a number of aspects of prisoner treatment including food, hygiene, health care, work, education, recreation, etc. However, Hart et al. (1997) argue that the contracts are significantly incomplete with respect to two important factors: the use of force by guards and the quality of personnel. As a result of this incompleteness a private contractor can use its residual rights of control to save money by hiring cheap, unqualified guards. These guards may not have the skill to respond to violent situations effectively.

The prison company hiring cheap unskilled guards is analogous to the coal mine choosing to mine high-ash-content rather than low-ash-content coal. In each case the supplier is choosing an action permitted by the contract that saves money at the expense of quality: in one case the quality effect is borne by the power plant, in the other case by the government or society. Of course, if the quality reduction has greater value than the saving in costs, *ex post* renegotiation of the contract should occur. Indeed, the model in Hart et al. (1997) supposes that it does. However, there is still a distortion: the private provider will have an excessive incentive to develop cost-saving, quality-reducing ideas.

If the government owns the prison the same problem does not arise. Just as the power plant can require the provision of low-ash-content coal if it owns the coal mine, the government can forbid the prison warden from hiring cheap, unskilled guards.

Hart et al. (1997) cite some evidence that the level of violence is indeed higher in private prisons.¹⁸

Of course, private provision can yield some benefits too. In the case of the coal mine, we argued that the manager will have a greater incentive to innovate when the mine is independent. The same is true of prisons. The warden of a private prison will have a greater incentive to find socially efficient ways of saving money or to develop socially valuable rehabilitation programs.¹⁹ This might tip the balance in favor of private provision in situations where innovation matters and violence is a relatively small problem, e.g., half-way houses or youth correctional facilities. However, in maximum security prisons, where the prevention of violence by prisoners against guards and other prisoners is paramount, Hart et al. (1997) conclude that the case for private provision is weak.

Hart et al. (1997) use the same logic to argue that private provision makes sense for garbage collection, does not make sense for the army, police or foreign policy, and may or may not make sense for schools and health care. Competition strengthens the case for privatization since actions that reduce quality will elicit a negative market response. Competition may work fairly well in the case of schools and hospitals, but it is hard to imagine it operating in the case of prisons.

Perhaps the most valuable lesson of the analysis is that it suggests that the public-private choice should be seen as a matter of efficiency, not ideology.

5. FOUNDATIONS

The property rights theory described in Section 2 is based on the idea that ownership of non-human assets is a source of bargaining power when contracts are incomplete. The formal models of Grossman and Hart (1986) and Hart and Moore (1990) justified contractual incompleteness by appealing to the idea that it may be difficult to describe in advance what kind of good a buyer wants from a seller; this may depend on a future state of the world and there may be many such states. In contrast, once the state is realized, it is easy to describe the good and hence a perfect spot contract can be written. Unfortunately, when bargaining takes place *ex post*, *ex ante* investments will already have been sunk, hold-up is possible, and, anticipating this, the parties will choose these investments inefficiently.

The challenge is to turn this informal story into a formal one.

It turns out that this is not easy. Consider first the issue of bargaining power. Return to the power plant-coal mine example. Suppose that it is desirable for the power plant to be in a strong bargaining position *ex post* in order to give it an incentive to locate next to the coal mine *ex ante*. Why not allocate bargaining power directly rather than through asset ownership? For example, the initial contract could specify that whenever renegotiation occurs it will proceed by the power plant making an offer to the coal mine that the coal mine can only accept or reject. Any attempt by the coal mine to make a counter-offer will be heavily penalized.

Indeed imagine that this is all the contract says; it does not specify the quantity, quality or price of coal at all!

Consider the situation where the power plant wants the coal mine to produce low-ash-content coal. Then the power plant will make an offer to the coal mine for low-ash-content coal in return for a payment that slightly exceeds the cost c of mining low-ash-content coal; say the price $p = c + \varepsilon$. Assume that there is a deadline by which the deal must be done. The coal mine can always threaten to reject such a low offer, but the power plant can ignore such threats, confident that at the last moment the coal mine will accept the offer since ε is better than nothing. This is the power of subgame perfection. As a result the power plant obtains all the *ex post* surplus from the transaction, and will locate next to the coal mine.

Indeed, as Hart and Moore (1988) show, it is not even necessary to stipulate in the contract that the coal mine will be penalized for making counter-offers. The contract could instead state that both sides can make offers and that neither party needs to accept an offer before trade occurs; instead they can trade and then produce a signed acceptance later. Suppose that absent any signed document, trade is assumed to be a gift from the coal mine to the power plant. Imagine now that the power plant makes the proposal described above while the coal mine proposes a much higher price, p' , say. Then trade will occur at price p . Why? On the one hand the coal mine can ensure itself p by trading and then signing and disclosing the power plant's offer in the event of a dispute. On the other hand the power plant will never sign and disclose the coal mine's proposal in the event of a dispute since it would prefer to reveal nothing and claim that the coal is a gift.

Of course, allocating all the bargaining power to the power plant gives it an incentive to locate next to the coal mine, but it does not give the manager of the coal mine an incentive to innovate. Suppose that the manager finds a way to reduce the cost of coal from c to c' . Then the power plant will change its offer from $c + \varepsilon$ to $c' + \varepsilon$. The coal mine's innovation will be completely expropriated by the power plant. This suggests that asset ownership may be useful after all.

Allocate bargaining power to the power plant through the contract but assign ownership of the coal mine to the mine manager. Ownership of the coal mine may provide the manager with a good outside option—perhaps selling the coal elsewhere—and since this outside option will increase if the manager becomes more efficient, some of the gains from increased efficiency will now accrue to the manager.

Unfortunately, this is not the end of the story. Suppose that the manager's outside option if she owns the mine $= \bar{p} - \lambda c$, where c is the cost of supplying coal to the power plant and $0 < \lambda < 1$ reflects the fact an increase in efficiency from supplying the power plant translates partially but not completely into gains from supplying coal to others. Obviously, asset ownership would be irrelevant if the coal mine manager could be offered an incentive scheme directly of the form $\bar{p} + (1 - \lambda)c$. (Since the manager incurs the cost of supplying coal this makes her net payment $\bar{p} - \lambda c$.) This would indeed be possible if c were verifiable, but at first sight seems impossible if c is merely observable.

But there is a way to do it, as Eric Maskin and Jean Tirole (1999) have pointed out, drawing on work of Maskin (1999) and Moore and Repullo (1988). The following game, to be played *ex post*, is written into the contract. The coal mine manager announces her cost of supplying (low-ash-content) coal to the power plant, c^* , say. The power plant can accept this and pay the manager $\bar{p} + (1 - \lambda)c^*$ or he can challenge, claiming the cost is c^{**} (presumably lower than c^* , so that the manager is paid less). If the power plant challenges, the coal mine manager pays a large fine F to a third party. At this point the challenge is tested: the mine manager is asked whether she wants to supply coal at a price $= 1/2c^* + 1/2c^{**}$. If she does supply, this establishes that the coal mine manager lied since she would be losing money by supplying if her true cost were c^* . In this case the power plant receives F from the third party. If the mine manager does not supply, that is, the challenge is proved incorrect, the power plant pays F to the third party.

The unique subgame perfect equilibrium of this mechanism is for the coal mine manager to tell the truth about her cost and for the manager to receive a net payment of $\bar{p} - \lambda c$. The need for asset ownership has been avoided.

Are there ways of ruling out Maskin-Tirole mechanisms? One objection is that collusion could take place between the buyer or seller and the third party.²⁰ However, it is not clear how such a collusive agreement would be enforced since the parties could make clear in the contract that it is illicit. Also Maskin and Tirole have shown that if at least one of the parties is risk-averse, a cleverly-designed lottery can substitute for the third party. Nonetheless, if one is prepared to assume that (a) the buyer and the seller are both risk neutral, (b) third parties cannot be used, (c) the parties can always renegotiate the contract after any

procedures for revising it have been completed, e.g., because there is no clear deadline; then the incomplete contracting story, and the role for asset ownership, can be resurrected. This is shown in Segal (1999) and Hart and Moore (1999).

But these are strong assumptions and I, for one, do not feel very comfortable with them. If the model accurately captures reality, one would expect to see some attempts to use Maskin-Tirole mechanisms, as well as to allocate bargaining power contractually. I know of no cases of the first and not many of the second. Also the model as it stands cannot explain *ex post* inefficiency (except if parties are wealth-constrained). This seems a significant limitation since the earlier work of Coase and Williamson argues convincingly that reducing *ex post* inefficiency is at least one of the rationales for the existence of firms.

Of course, one could try to incorporate *ex post* inefficiency by supposing that the parties are asymmetrically informed.²¹ However, as long as there is symmetric information at the contracting date, a further set of mechanisms (also not seen in practice) can be used to overcome this.²² For these reasons in recent work I have turned to a different approach.

6. DROPPING RATIONALITY

If parties are fully rational I do not see why they would not include in their contracts mechanisms of the type suggested by Maskin and Tirole. As I have said, as far as I know, there are no examples of this in practice. One can, of course, always put the blame on judges: they would not understand and/or would not enforce such mechanisms.²³ But this just pushes the question one step further: there are many smart judges and if mechanism design is the solution to incomplete contracting problems one would expect judges eventually to understand, embrace, and enforce contracts based on the mechanism design approach. It has been eighteen years since the publication of Maskin and Tirole's paper and I do not see any move in this direction.

My view is that the reason we do not see these mechanisms is that parties are not fully rational. This is in many ways an unfortunate conclusion since while there is one way to model rationality there are many ways, perhaps infinitely many, to model irrationality. The lack of discipline makes many economists uncomfortable, and they are willing to hold on to the rationality approach at all cost. However, I think that there is no alternative but to abandon it.

I began to do this in a paper with John Moore (Hart and Moore (2008)). Looking back, one can say that our motivation was three-fold. In no particular order of importance: First, we wanted to develop a theory immune to the Maskin-Tirole critique. Second, we wanted to explain why parties do not allocate

bargaining power contractually to solve the hold-up problem. Third, we wanted to develop a model that allows for ex post inefficiency. It is this multiplicity of motives that perhaps explains why we did not try to introduce cognitive limitations for the parties but instead focused on ideas of fairness and reasonable behavior. In this respect we were much influenced by the large behavioral literature on the latter topic.²⁴

To understand our approach, consider a simple situation of a buyer B and a seller S, who meet at date 0. At that time there is a competitive market for buyers and sellers, but after date 0 B and S will pair off and will be isolated from the market. At date 1 there are gains from trade. S can supply one widget at cost c and B obtains value $v > c$ from it. All returns are measured in money (but these returns are not verifiable).

For simplicity, suppose that the seller's reservation utility determined in the date 0 market for buyers and sellers is zero. One contract that B could offer to S that will give B all the gains from trade is the following: The contract states that at date 1, B will make an offer to S that S can accept or reject; S cannot make any offers to B. As we have seen, under standard rationality assumptions, B will offer S just above c at date 1, S will supply the widget, and B will receive the full surplus $v - c$.

Now the evidence from the famous ultimatum game experiments suggests that things may not work this way in practice. These experiments find that in situations like this B will end up offering considerably more than c , and that moreover if B does not do so S will reject the offer.²⁵ Note, however, that the parallel between the ultimatum game and our case is not exact since there is no prior contract in the ultimatum game.

Moore and I could have constructed a model based directly on the idea that S will turn down ungenerous offers. We did not do so for two reasons. First, we were concerned about the possible criticism that the ultimatum game evidence concerns relatively small payoffs.²⁶ Second, we wanted our model to apply to more general situations than just ones where a seller can choose not to trade. For example, what is the analogy of ultimatum game behavior when parties are asked to play a Maskin-Tirole mechanism?

We therefore proceeded as follows. We assumed that even ex post perfect contracts cannot be written and so it is possible for both the buyer and seller to provide less than ideal performance while staying within the terms of the contract: we refer to less than ideal performance as "shading." In the buyer-seller example, the seller might shade by supplying a widget of deficient quality, while the buyer might shade by not providing information that would make the seller's task easier. A critical assumption is that a party will shade if and only if he does not feel well treated. So in the case of the seller who receives a low ball, but

nonetheless profitable, offer, the seller will accept the offer but then punish the buyer by shading.

We made the further crucial assumption that the initial contract circumscribes what parties feel is fair. The date 0 competitive market for buyers and sellers is an important element here. The idea is that the broad terms of the contract are regarded by both parties as reasonable since they are negotiated at arms-length and neither party blames the other for the equilibrium terms of trade. As a result, neither B nor S feels entitled to an outcome outside the contract. In contrast any discretionary decision made by one of the parties at date 1—when the competitive market is no longer there to provide an objective benchmark—may be found unreasonable by the other party and may lead to shading. To make things as stark as possible, Hart and Moore (2008) suppose that each party is subject to an extreme self-serving bias and feels that a reasonable outcome is one that maximizes that party's payoff over all outcomes permitted by the contract that are individually rational for the other party.

Return to the widget example and the contract that specifies that B will make a take-it-or-leave-it offer to S. Let B offer a price just above c . S will consider this unreasonable given that B could have been more generous. Indeed, the best outcome for S under the contract would be for B to offer v (anything more than v would involve B's making a loss and so would not be individually rational). How much does S shade given the actual offer c ? Hart and Moore (2008) assume that shading is a fraction of how much S is shortchanged or aggrieved, where the latter is the difference between the payoff S feels entitled to—here $v - c$ —and what she gets—zero. Specifically, S reduces B's payoff by $\theta(v - c)$, where $0 < \theta < 1$. Shading does not affect the payoff of the party doing the shading.

In sum, under the contract that gives B the right to make a take-it-or-leave-it offer to S, there will be a deadweight loss of $\theta(v - c)$. Note that there is no way of negotiating around this. Coasian bargaining fails because shading is noncontractible. B could, of course, offer more than c to reduce S's aggrievement, but it is not in his interest to do this: offering a dollar more increases B's cost by a dollar but reduces shading by only θ .

There is, however, a solution to this problem in this simple example. B and S could fix the price in advance: they could write a contract at date 0 that specifies the date 1 price of the widget to be c . In this case neither party has any discretion at date 1. B and S both regard the price c as fair since it is negotiated at arms-length in a competitive market at date 0. There will be no shading or deadweight losses at date 1 and the full surplus $v - c$ will be earned. The first-best is achieved.

This simple framework suggests an alternative rationale for the existence of a contract. A contract negotiated before parties are locked into each other gets

them onto the same page and avoids bad feeling later on. This is different from (but complementary to) the traditional view that contracts are useful to encourage noncontractible investments. (In the above example there are no noncontractible investments.)

Once we depart from the case of certainty, it will typically not be possible to achieve the first-best. To illustrate this, consider a simplified version of the Hart-Moore model in which the further assumption is made that S has zero wealth; this version will also be helpful for describing some experimental work. Suppose that $v = 20$ for sure but $c = 16$ with probability π and 10 with probability $1 - \pi$. The uncertainty about c will be resolved shortly before date 1 and the realization of c is then observable to both parties. However, c is not verifiable. The probability distribution of c is common knowledge *ex ante*. Assume further that *ex post* trade is voluntary: either party can refuse to trade and not be penalized, perhaps because a third party cannot verify who is responsible for the absence of trade. B and S are both risk neutral. There are many more buyers than sellers in the date 0 market and so the reservation utility level for S is zero. Finally, ignore renegotiation for the moment.

What is an optimal contract for B to offer in this setting? There are only two possibilities. Either B wants to ensure trade in both states or only in the low cost state. In the first case the optimal contract will specify a price range $[10, 16]$ and allow B to pick from this range at date 1. That way B can guarantee trade whether c is high or low, given that trade is voluntary. Moreover, this is the smallest price range that will do the job, which minimizes aggrievement and shading.

With such a contract B will choose $p = 10$ when $c = 10$ and $p = 16$ when $c = 16$. In the low cost state S will be aggrieved since B could have been more generous and have chosen the best outcome for S , $p = 16$. S 's level of aggrievement is 6. S punishes B by shading by 6θ , and so B 's net payoff = $10 - 6\theta$. In the high cost state, S is not aggrieved since she receives the highest price permitted by the contract. B 's payoff = 4.

The expected payoffs for the two parties are, respectively,

$$U_B = (10 - 6\theta)(1 - \pi) + 4\pi, \quad (1)$$

$$U_S = 0. \quad (2)$$

Call this flexible contract, contract 1.

On the other hand, B can choose a contract that permits trade only in the low cost state. The best such contract fixes the price at 10. The expected payoffs of the two parties are, respectively,

$$U_B = 10(1 - \pi), \quad (3)$$

$$U_S = 0. \quad (4)$$

Call this rigid contract, contract 2.

Obviously, contract 2 is better than contract 1 if and only if

$$10(1 - \pi) > (10 - 6\theta)(1 - \pi) + 4\pi. \quad (5)$$

This will be true if π is small.

In other words, B will offer S a fixed price contract that precludes trade in the high cost state if that state is unlikely to occur. The intuition is simple. It is not worth expanding the price range from 10 to $[10, 16]$ just to realize trade in the high cost state if it has low probability, given that this causes a large deadweight loss from shading in the low cost state that has high probability.

Note also the importance of S's wealth constraint. In the absence of such a constraint, B could offer a contract that specifies $p = 16$, leading to trade in both states. B could charge S upfront $6(1 - \pi)$ for this contract, thus recouping all of S's expected profit.

This model achieves the main goals described above. First, it is immune to the Maskin-Tirole critique. Mechanisms or take-it-or-leave-it offers do not achieve the first-best. Indeed, contract 1 contains such a mechanism and leads to shading. Second, there can be ex post inefficiency. If (5) holds, B will deliberately choose a contract that causes trade not to occur with some probability.²⁷

The Hart-Moore model relies on a number of non-standard assumptions. While several of these are similar to behavioral assumptions that have been validated, there are some significant differences. It therefore seemed desirable to test the model directly, and Ernst Fehr, Christian Zehnder, and I have done this in the lab.²⁸ The following is a rough description of the lab experiment in Fehr et al. (2011); see also Fehr et al. (2009). (Some simplifications and liberties have been taken in describing it.) We divide the student participants into buyers and sellers; their roles stay the same during the experiment. Each buyer meets with two sellers, who can bid for the buyer's contract. (The purpose of this is to achieve ex ante competition.) The buyer can choose between two types of contract: one, a flexible contract of the form $[p, 16]$, the other, a fixed price contract p . Once the contract type has been selected the sellers compete for the contract through a clock auction. The auction determines the level of p : p starts off at 10 and rises a small amount every second until one of the sellers accepts. As in many experiments of this type the resulting p is close to 10, and we shall treat it as 10 in what follows.

Thus at the end of the auction a seller has agreed to a contract that is either $[10, 16]$ if B chose a flexible contract at the earlier stage, or $p = 10$ if B chose a rigid contract. The buyer and the winning seller then move to the next stage, date

1. A randomization device determines c and both parties see the outcome. Under the flexible contract B then chooses p from the range $[10, 16]$; p is required to be at least c . Trade then occurs. At this point S can choose whether to shade. In the experiment shading is a discrete action that imposes a small cost on the seller and a large cost on the buyer.

In contrast, under the rigid contract, trade takes place only if $c = 10$. Again after trade the seller can choose whether to shade.

The probability π of the high cost state is chosen to be relatively small. The experiment is repeated several times with buyers and sellers being re-matched randomly each time.

If buyers and sellers are fully rational, we would predict a very simple outcome. Start at the end. Since shading is costly a rational seller will never engage in it. (Notice that this is why we modified the assumption from the theoretical model that shading is costless.) Hence a buyer can safely ignore shading and will choose the flexible contract since this guarantees trade in both states. B will choose $p = 10$ when $c = 10$ and $p = 16$ when $c = 16$. The first-best is achieved.

This is not what happens in the experiment. Buyers choose rigid contracts a significant fraction of the time, and these contracts are more profitable than flexible contracts. With the flexible contract buyers offer more than 10, and significant shading occurs, in the low cost state. Shading is rare in the rigid contract.

These results are broadly consistent with the Hart-Moore model. It is particularly striking that there is little shading in the rigid contract. This is in spite of the fact that this contract allocates all of the surplus to the buyer. It appears that the seller accepts and is not angry about her low (zero) payoff given that this was determined competitively.²⁹

Fehr et al. (2015) extend the experiment to permit communication and renegotiation. Allowing the buyer to send a message to the seller at the contract formation stage, explaining how he plans to choose price in each state, improves the efficiency of the flexible contract but the trade-off between rigid and flexible contracts remains.³⁰ Allowing for renegotiation improves the efficiency of the rigid contract relative to the flexible contract. Under the rigid contract trade occurs if $c = 10$. On the other hand, if $c = 16$, the buyer can offer a new (renegotiated) contract. Some shading occurs in this case, since the seller is aggrieved that she does not receive 20. Still trade now takes place in both states. Interestingly, we find that the seller does not expect renegotiation to occur when $c = 10$, and is not aggrieved and does not shade when this fails to happen.

The Hart-Moore model can be used to re-examine some of the issues discussed in property rights theory. First, given that there may be some ex post inefficiency, asset ownership will matter, but for different reasons from before.

Since an inefficient allocation may occur (as in the rigid contract without renegotiation), outside options, determined by asset ownership, may be exercised even if they are first-best inefficient. Hence *ex ante* asset allocation will affect *ex post* surplus. Also asset ownership will determine the size of the gains from trade, and therefore entitlements and shading. Hence, even if the final outcome is independent of asset ownership, the deadweight losses from shading will generally depend on the asset allocation. Hart (2009) uses these ideas to develop a theory of asset ownership based on payoff uncertainty rather than relationship-specific investments.

Second, the model can throw light on the employment relationship. Suppose that, instead of uncertainty about the seller's cost, there is uncertainty about what type of good a buyer wants from a seller. Refer to the type of good as the seller's task. In Hart and Moore (2008), it is shown that under some assumptions the optimal contract will take the following form: the price paid to the seller is fixed and one of the parties is given the right to choose the task. If the buyer is allocated the right, this can be interpreted as an employment contract. If the seller is allocated the right, it can be interpreted as independent contracting. The reason for fixing the price is that there will be less shading if *ex post* disagreement over entitlements is limited to the choice of task rather than to the choice of the price *and* the task.

This version of the model relates to earlier ideas about the employment relationship (see Coase (1937), Simon (1951), and Alchian and Demsetz (1972)). But in the earlier work of Coase and Simon it was assumed that the price or wage is fixed. Here there is an explanation.

This model of employment is used in Hart and Holmström (2010) to develop a theory of firm scope. Consider two firms that want to coordinate on a technological platform. They can do this as separate firms and write a contract. Or they can merge. In the first case since it is hard to specify all the details of coordination there will be *ex post* inefficiency, either because renegotiation fails or because of shading. In the second case, a boss of the merged firm can require coordination (the boss chooses the task), but in doing so will not fully internalize the workers' costs of shifting to the new technology. Hart and Holmström (2010) show that the relative importance of these effects will determine which organizational form is better.

Finally, in some recent work with Maija Halonen-Akatwijuka, I have explored dynamic issues that arise when contracts are reference points. If parties write a sequence of contracts, then the first contract will become a reference point for the second contract, and so on. This interdependence can make it easy for the parties to write a new contract if not much changes over time, since the parties

can agree that the terms of the contract should be roughly the same as before. However, if outside conditions change then notions of fairness and good faith emanating from earlier contracts may make it difficult for the parties to adjust to the new environment, and they may fail to trade even though this is efficient. For details, see Halonen-Akatwijuka and Hart (2016).

7. LOOKING AHEAD

When Sanford Grossman and I sat in his office in the summer of 1983, our goal was to develop a formal model of the limits of contracts and the boundaries of the firm. I can report only partial success in this endeavor. There is as yet no tractable, widely agreed upon, theory of incomplete contracts. Indeed, to the extent that, as I have argued, one must depart from rationality to make progress, there may never be.

At the same time I believe that the incomplete contracts approach yields some valuable insights. I have tried to describe some of them here, but there are many other applications, including in the areas of law and political science.

Economists are drawn to areas with simple, elegant, and uncontroversial models. The area of incomplete contracts is not like that; it is messy. But contracts are incomplete in reality and contractual incompleteness underlies numerous significant phenomena, some of which have great policy relevance. I hope that economists, particularly young ones, will, in spite of the messiness, continue to work on this challenging topic.

REFERENCES

- Aghion, P. and P. Bolton (1992), An Incomplete Contracts Approach to Financial Contracting, *Review of Economic Studies* **59**: 473–494.
- Aghion, P., D. Fudenberg, R. Holden, T. Kunitomo, and O. Tercieux (2012), Subgame-Perfect Implementation under Information Perturbations, *Quarterly Journal of Economics* **127**: 1843–1881.
- Alchian, A.A. and H. Demsetz (1972), Production, Information Costs, and Economic Organization, *American Economic Review* **62**: 777–795.
- Andersen, S., U. Gneezy, M. Hoffman, S. Ertac, and J.A. List (2011), Stakes Matter in Ultimatum Games, *American Economic Review* **101**: 3427–3439.
- Antras, P. (2003), Firms, Contacts, and Trade Structure, *Quarterly Journal of Economics* **118**: 1375–1418.
- Antras, P. and E. Helpman (2004), Global Sourcing, *Journal of Political Economy* **112**: 552–580.
- Baker, G. and T.N. Hubbard (2003), Make versus Buy in Trucking: Asset Ownership, Job Design, and Information, *American Economic Review* **93**: 551–572.

- Baker, G. and T.N. Hubbard (2004), Contractibility and Asset Ownership: On-Board Computers and Governance in U.S. Trucking, *Quarterly Journal of Economics* **119**: 1443–1479.
- Baliga, S. and T. Sjöström (2016), A Theory of the Firm Based on Haggling, Coordination and Rent Seeking, Northwestern University working paper.
- Bartling, B. and K. Schmidt (2015), Reference Points, Social Norms, and Fairness in Contract Renegotiations, *Journal of the European Economic Association* **13**: 98–129
- Benmelech, E. (2009), Asset Salability and Debt Maturity: Evidence from Nineteenth-Century Railroads, *Review of Financial Studies* **22**: 1545–1584.
- Benmelech, E. and N. Bergman (2008), Liquidation Values and the Credibility of Financial Contract Renegotiation: Evidence from U.S. Airlines, *Quarterly Journal of Economics* **123**: 1635–1677.
- Bernstein, S., A. Korteweg, and K. Laws (2017), Attracting Early Stage Investors: Evidence from a Randomized Field Experiment, forthcoming in *Journal of Finance*.
- Bolton, P. and D. Scharfstein (1990), A Theory of Predation Based on Agency Problems in Financial Contracting, *American Economic Review* **80**: 93–106.
- Bolton, P. and D. Scharfstein (1996), Optimal Debt Structure and the Number of Creditors, *Journal of Political Economy* **104**: 1–25.
- Borek, C.T., A. Frattarelli, and O. Hart (2014), Tax Shelters or Efficient Tax Planning? A Theory of the Firm Perspective on the Economic Substance Doctrine, *Journal of Law and Economics* **57**: 975–1000.
- Brandts, J., G. Charness, and M. Ellman (2016), Let's Talk: How Communication Affects Contract Design, *Journal of the European Economic Association* **14**: 943–974.
- Camerer, C. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction*, Princeton University Press.
- Coase, R.H. (1937), The Nature of the Firm, *Economica* **4**: 386–405.
- Coase, R.H. (1960), The Problem of Social Cost, *Journal of Law and Economics* **3**: 1–44.
- d'Aspremont, C. and L.A. Gérard-Varet (1979), Incentives and incomplete information, *Journal of Public Economics* **11**: 25–45.
- Drèze, J. (1974), Investment Under Private Ownership: Optimality, Equilibrium and Stability, In J. Drèze ed. *Allocation Under Uncertainty: Equilibrium and Optimality*, Chp. 9, 129–166, Macmillan, London.
- Fehr, E. and K.M. Schmidt (2003), Theories of Fairness and Reciprocity: Evidence and Economic Applications. In Mathias Dewatripont, Lars Peter Hansen and Stephen J Turnovsky eds. *Advances in Economics and Econometrics, Econometric Society Monographs*, Eighth World Congress, Volume 1, 208–257.
- Fehr, E., O. Hart, and C. Zehnder (2009), Contracts, Reference Points, and Competition—Behavioral Effects of the Fundamental Transformation, *Journal of the European Economic Association* **7**: 561–572.
- Fehr, E., O. Hart, and C. Zehnder (2011), Contracts as Reference Points—Experimental Evidence, *American Economic Review* **101**: 493–525.
- Fehr, E., O. Hart, and C. Zehnder (2015), How do Informal Agreements and Revision Shape Contractual Reference Points?, *Journal of the European Economic Association* **13**: 1–28.
- Gale, D., and M. Hellwig (1985), Incentive-Compatible Debt Contracts: The One-Period Problem, *Review of Economic Studies* **52**: 647–663.

- Grossman, S. and O. Hart (1979), A Theory of Competitive Equilibrium in Stock Market Economies, *Econometrica* **47**: 293–330.
- Grossman, S. and O. Hart (1980) Takeover Bids, the Free-Rider Problem, and the Theory of the Corporation, *Bell Journal of Economics* **11**: 42–64.
- Grossman, S. and O. Hart (1982), Corporate Financial Structure and Managerial Incentives. In J. McCall ed. *The Economics of Information and Uncertainty*, University of Chicago Press: Chicago IL.
- Grossman, S. and O. Hart (1983), An Analysis of the Principal-Agent Problem, *Econometrica* **51**: 7–45.
- Grossman, S. and O. Hart (1986), The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration, *Journal of Political Economy* **94**: 691–719.
- Güth, W., R. Schmittberger, and B. Schwarze (1982), An Experimental Analysis of Ultimatum Bargaining, *Journal of Economic Behavior and Organization* **3**: 367–388.
- Hadfield, G.K. and I. Bozovic (2016), Scaffolding: Using Formal Contracts to Support Informal Relations in Support of Innovation, *Wisconsin Law Review* **5**: 981–1032.
- Halonen-Akatwijuka, M. and O. Hart (2016), Continuing Contracts, Harvard University working paper.
- Hart, O. (1975), On the Optimality of Equilibrium when the Market Structure is Incomplete, *Journal of Economic Theory* **11**: 418–443.
- Hart, O. (1979), On Shareholder Unanimity in Large Stock Market Economies, *Econometrica* **47**: 1057–1083.
- Hart, O. (1995), Firms, Contracts, and Financial Structure, Clarendon Press: Oxford UK.
- Hart, O. (2009), Hold-Up, Asset Ownership, and Reference Points, *Quarterly Journal of Economics* **124**: 267–300.
- Hart, O. and B. Holmström (2010), A Theory of Firm Scope, *Quarterly Journal of Economics* **125**: 483–513.
- Hart, O. and J. Moore (1988), Incomplete Contracts and Renegotiation, *Econometrica* **56**: 755–785.
- Hart, O. and J. Moore (1990), Property Rights and the Nature of the Firm, *Journal of Political Economy* **98**: 1119–1158.
- Hart, O. and J. Moore (1994), A Theory of Debt Based on the Inalienability of Human Capital, *Quarterly Journal of Economics* **109**: 841–879.
- Hart, O. and J. Moore (1998), Default and Renegotiation: A Dynamic Model of Debt, *Quarterly Journal of Economics* **113**: 1–41.
- Hart, O. and J. Moore (1999), Foundations of Incomplete Contracts, *Review of Economics Studies* **66**: 115–138.
- Hart, O. and J. Moore (2008), Contracts as Reference Points, *Quarterly Journal of Economics* **73**: 1–48.
- Hart, O., A. Shleifer, and R. Vishny (1997), The Proper Scope of Government: Theory and an Application to Prisons, *Quarterly Journal of Economics* **112**: 1127–1161.
- Herweg, F. and K.M. Schmidt (2015), Loss Aversion and Inefficient Renegotiation, *Review of Economic Studies* **82**: 297–332.
- Holmström, B. (1979), Moral Hazard and Observability, *Bell Journal of Economics* **10**: 74–91.
- Holmström, B. and P. Milgrom (1991), Multi-Task Principal Agent Analyses, *Journal of Law, Economics, and Organization* **7**: 24–52.

- Holmström, B. and J. Tirole (1998), Private and Public Supply of Liquidity, *Journal of Political Economy* **106**: 1–40.
- Innes R. (1990), Limited Liability and Incentive Contracting with Ex Ante Action Choices, *Journal of Economic Theory* **52**: 45–67.
- Jensen, M. and W. Meckling (1976), Theory of the Firm: Managerial Behavior, Agency Costs, and Ownership Structure, *Journal of Financial Economics* **3**: 305–360.
- Joskow, P. (1987), Contract Duration and Relationship Specific Investments: Empirical Evidence from Coal Markets, *American Economic Review* **77**: 168–185.
- Kaplan, S. and P. Strömberg (2003), Financial Contracting Theory Meets the Real World, *Review of Economic Studies* **70**: 281–315.
- Kaplan, S., B. Sensoy, and P. Strömberg (2009), Should Investors Bet on the Jockey or the Horse? Evidence from the Evolution of Firms from Early Business Plans to Public Companies, *Journal of Finance* **64**: 75–115.
- Klein, B., R. Crawford, and A. Alchian (1978), Vertical Integration, Appropriable Rents and the Competitive Contracting Process, *Journal of Law and Economics* **21**: 297–326.
- Maskin, E. (1999), Nash Equilibrium and Welfare Optimality, *Review of Economic Studies* **66**: 23–38.
- Maskin, E. and J. Tirole (1999), Unforeseen Contingencies and Incomplete Contracts, *Review of Economic Studies* **66**: 83–114.
- Matouschek, N. (2004), Ex Post Inefficiencies in a Property Rights Theory of the Firm, *The Journal of Law, Economics, and Organization* **20**: 125–147.
- Mirrlees, J.A. (1999), The Theory of Moral Hazard and Unobservable Behavior, *Review of Economic Studies* **66**: 3–21.
- Modigliani, F. and M. Miller (1958), The Cost of Capital, Corporation Finance and the Theory of Investment, *American Economic Review* **68**: 261–297.
- Moore, J. and R. Repullo (1988), Subgame Perfect Implementation, *Econometrica* **56**: 1191–1220.
- Rajan, R. (2012), Presidential Address: The Corporation in Finance, *Journal of Finance* **67**: 1173–1217.
- Schmidt, K. (1996), The Costs and Benefits of Privatization: An Incomplete Contracts Approach, *Journal of Law, Economics, and Organization* **12**: 1–24.
- Segal, I. (1999), Complexity and Renegotiation: A Foundation for Incomplete Contracts, *Review of Economic Studies* **66**: 57–82.
- Shleifer, A. and R. Vishny (1997), A Survey of Corporate Governance, *Journal of Finance* **52**: 737–783.
- Simon, H.A. (1951), A Formal Theory of the Employment Relationship, *Econometrica* **19**: 293–305.
- Townsend, R. (1979), Optimal Contracts and Competitive Markets with Costly State Verification, *Journal of Economic Theory* **21**: 265–293.
- U.S. Department of Justice (2016), Review of the Federal Bureau of Prisons' Monitoring of Contract Prisons, <https://oig.justice.gov/reports/2016/e1606.pdf#page=1>
- Williamson, O. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*, The Free Press: New York.
- Woodruff, C. (2002), Non-contractible Investments and Vertical Integration in the Mexican Footwear Industry, *International Journal of Industrial Organization* **20**: 1197–1224.

NOTES

1. See Hart (1975).
2. See, e.g., Dreze (1974).
3. This led to Grossman and Hart (1979); see also Hart (1979).
4. There was already active work on this question. See, e.g., Mirrlees (1999) and Holmström (1979) on principal-agent theory and Jensen and Meckling (1976) on corporate finance.
5. See Grossman and Hart (1980) and Grossman and Hart (1982).
6. I refer here to the American journalist who died in 1920, not the former CEO of Citigroup.
7. This is very much along the lines of Coase (1937) and Simon (1951).
8. Paul Joskow has studied such situations and the relationships can last for decades (see Joskow (1987)).
9. Empirical support for the idea that ownership affects noncontractible investments can be found in Baker and Hubbard (2003, 2004) and Woodruff (2002). An interesting application of PRT has been to explain the outsourcing decisions of multinational firms. See Antras (2003) and Antras and Helpman (2004).
10. The control rights approach is only one strand of a very active literature on financial contracting that has developed over the last thirty or so years. Important work has been done on asymmetric information by Townsend (1979) and Gale and Hellwig (1985), and on moral hazard by Innes (1990) and Holmström and Tirole (1998). The control rights approach is complementary to these contributions.
11. In what follows we consider the relationship between one party who needs funds and one party who has funds. For a more general discussion of corporate governance and control, see Shleifer and Vishny (1997).
12. They also assume that the financial investment is contractible. This is no longer an analytical issue since the manager, being wealth constrained, cannot compensate the investor for making it.
13. The first version of the Hart-Moore (1998) paper appeared in 1989 and circulated for several years before being published. A simple exposition of it is contained in Hart (1995, Chapter 5). Related and companion papers are Bolton and Scharfstein (1990, 1996).
14. It should be noted that while a debt contract is optimal when project returns are perfectly certain, contracts that shift control in more elaborate ways can perform better under uncertainty. For conditions for a debt contract to be optimal, see Hart and Moore (1998).
15. In reality the distinction between human and non-human assets is not quite as stark as we have suggested. There are ways that people can reduce their hold-up power, e.g., by signing non-compete contracts. However, there are limits to the control of human assets that do not exist in the case of non-human assets.
16. I have been involved in two legal cases where the question of what is a firm was important. In *Black and Decker v. U.S.A.*, Black and Decker argued that it created a new entity to manage employee and retiree health care benefits for efficiency rather than tax reasons. I argued for the U.S. government that the new entity was equivalent

to a division of Black and Decker since Black and Decker retained control. The case was settled. In a second case, WFC Holdings (Wells Fargo) v. U.S.A., WFC argued that there was a business purpose from moving its real estate lease operations into a separate subsidiary. I argued, again on behalf of the U.S. government, that the subsidiary was equivalent to a division of WFC given that WFC had total control. The case went to trial and the U.S. government won. For a discussion of these cases, see Borek et al. (2014).

17. Related ideas are developed in Schmidt (1996).
18. Further evidence can be found in the recent report of the U.S. Department of Justice (2016). Another way to discourage the choice of cost-saving, quality-reducing actions is to set up the private prison as a non-profit. However, to the extent that the non-profit can use cost savings to increase salaries this is unlikely to fully resolve the problem.
19. The idea that the prison manager can engage in some cost-reducing activities that are socially desirable and other quality-reducing activities that are socially undesirable is in the spirit of the multi-tasking model of Holmström and Milgrom (1991).
20. Another objection is that Maskin-Tirole mechanisms are not robust to small departures from common knowledge. See Aghion et al. (2012).
21. See, e.g., Matouschek (2004).
22. See, e.g., d'Aspremont and Gerard-Varet (1979).
23. For recent work along these lines, see Baliga and Sjöström (2016).
24. See, e.g., Camerer (2003) and Fehr and Schmidt (2003).
25. See, e.g., Güth et al. (1982).
26. See Andersen et al. (2011).
27. Herweg and Schmidt (2015) develop an alternative and complementary theory of inefficient outcomes. In their model parties write a contract that has to be renegotiated after the state of the world is realized. The parties take the initial contract as a reference point to which they compare the gains and losses of the renegotiated transaction. As a result of loss aversion some efficient renegotiations will not occur.
28. For some non-laboratory evidence consistent with the idea that contracts are reference points, see Hadfield and Bozovic's recent study of how companies involved in innovation manage their relationships (Hadfield and Bozovic (2016)).
29. There is some evidence that terms implied by a contract are regarded as fair even in the absence of competition. See Bartling and Schmidt (2015).
30. Brandts et al. (2016), in a somewhat different setting, find that under free-form (ex post) communication flexible contracts dominate rigid contracts.