11 Norms and the theory of the firm

Oliver Hart

1 Introduction

Most standard models of incentives and/or organizations assume that economic agents are self-interested and must rely on formal contracts enforced by the courts to uphold their relationships. In reality, of course, many economic transactions are sustained by self-enforcing ("implicit") contracts, or norms of behavior, such as honesty or trust. An interesting question to ask is: does ignoring norms/self-enforcing contracts lead to misleading conclusions? That is, would a theory of incentives or organizations that incorporated norms look very different from the standard theory?

In this chapter, I will consider this question, focusing particularly on some of the attempts economists have made in the last ten years or so to integrate norms into the theory of the firm. I will argue that (a) although norms are undoubtedly very important both inside and between firms, incorporating them into the theory has been very difficult and is likely to continue to be so in the near future; (b) so far norms have not added a great deal to our understanding of such issues as the determinants of firm boundaries (the "make-or-buy" decision) – that is, at this point a norm-free theory of the firm and a norm-rich theory of the firm don't seem to have very different predictions.

2 Background

To begin with, it is worth mapping out some of the territory. I will follow Richard Posner in defining a norm as "a rule that is neither promulgated by an official source, such as a court or a legislature, nor enforced by the threat of legal sanctions, yet is regularly complied with" (see Posner 1997). I will focus on norms in and between organizations as opposed to societal norms, even though there is obviously an important connection between the two. For example, a society in which honesty is not taken very seriously is also one in which firms will have a lot of difficulty sustaining

trust. However, norms at the societal level are pretty slow to change, and, for many purposes, they can be taken as exogenous. In contrast, norms in and between organizations are capable of being designed.

A useful starting point is the idea that organizational norms matter when parties cannot write good contracts; more precisely, when transaction costs make contracts incomplete.[1] That is, in a world where parties can costlessly think and negotiate about the future, and judges are perfect, norms would not matter because parties' relationships could be governed by perfectly enforceable contracts. A leading source of contractual incompleteness stems from the fact that some economically significant variables are observable to the parties, but not to outsiders, such as a judge. (In the parlance of economics, these variables are "observable, but not verifiable.") For example, an ideal contract between an employer and an employee might specify that the employee would be given a bonus for good performance since this may encourage the employee to work hard. Both the employer and the employee may know after the fact whether the employee performed well or not, and therefore whether the bonus has been earned, but a judge may not have this information. As a result, the contract stating that the employer will pay the employee a bonus if the latter performs well is not legally enforceable. Here a norm of honesty would be very helpful. If the employer can be trusted to keep her word, the agreement that the employee will receive a bonus if he performs well can be sustained by informal means rather than by formal ones.

As another example, consider a company's promise to workers that it will not lay any of them off unless "things are really bad." Such a promise might serve an important role in providing risk averse workers with partial insurance about the future. However, enforcing such a promise in the courts is likely to be fraught with difficulty because of disagreement about the meaning of the phrase "things are really bad." (Without too much of a stretch of reality, it might be said that the event is observable but not verifiable.) Again, norms of honesty and decency can help here. If the firm can be trusted not to be opportunistic, then a flexible outcome can be achieved through an informal agreement: the company will reserve its right to shed workers if a disaster occurs, but will not abuse this right by laying off workers in events that are merely bad.

Given the link between norms and judicial imperfection, it is not surprising that much of the economic literature on norms in organizations goes under the heading of "self-enforcing contracts." However, it is important to realize that norms also matter when contracts are incomplete for other reasons, e.g. because the parties (themselves) are boundedly rational. For example, if the parties cannot think or negotiate ahead very well, then events will arise that their contract does not cover. A norm

of fairness can help to fill in the contractual gap in an appropriate manner. For reasons of tractability, most of the economic literature rules out bounded rationality among the contracting parties themselves, and so the role of fairness-type norms has not been much explored in an organizational context.[2] In my discussion, I will follow the literature in this regard; it should be emphasized, however, that a consequence of this is that much of interest may be left out.

3   Modeling difficulties

As I have already noted, theoretical progress on analyzing norms and organizations has been slow. The main reason is that economists do not have a very good way to formalize trust. Three main approaches have been tried, and each has significant drawbacks. In this section I will briefly describe them.

The most commonly used approach is based on the framework of infinitely repeated games. Although this will be familiar to many, it is probably worth illustrating since I will use it later on. Suppose that a buyer B and a seller S want to trade a widget each period. S can deliver a high-quality widget or a low-quality widget; the former has value that exceeds its cost, while the latter has zero cost and zero value. The quality of the widget is observable (to B and S), but not verifiable (in a court of law). In a one-shot version of this game, trade will not occur if the parties are purely self-interested (and hence are not trustworthy). The reason is that if B promises to pay S as long as S supplies a high-quality widget, then it is always in B's interest to claim that the widget's quality was low, whether or not this is true, and, anticipating this, S has no incentive to supply high quality. (This example is isomorphic to the employer–employee example mentioned earlier.)

If this game is repeated infinitely often, however, trade at the high-quality level can be sustained. The way this works is (roughly) as follows. B promises to pay S a price P per period, where P lies between B's value and S's cost, as long as the widget quality is high in that period (recall that B observes widget quality). In return, S promises to supply a high-quality widget each period unless in some previous period B has broken her promise to pay, in which case S supplies low quality forever more.

It is easy to see that these promises are mutually self-enforcing, as long as the parties do not discount the future too much. The reason is that, while B can gain something each period by pretending that S's quality is low and withholding payment, this short-term gain is dwarfed by B's loss from never receiving a high-quality widget again.

Unfortunately, as is well known, this approach to explaining cooperation or trust runs into several difficulties.[3] First, it relies crucially on the assumption that there is no upper bound to the number of times the game is played. Suppose in contrast that it is known that the game will not be played more than τ times. Then, however large τ is, the parties will realize that in the last period B will break her promise to pay S (as in the one-shot game, there is no future to discourage her); anticipating this, S will supply a low-quality widget in the last period; hence B will have no incentive to pay in the previous period (she recognizes that this will have no effect on what happens in the last period), etc. In other words, the self-enforcing contract unravels. The conclusion is that, as in the one-period model, no trade will take place in any period, however big τ is.

Unfortunately, the assumption that there is no upper bound to the number of times the game will be played is hard to square with the fact that people have finite lives.

A second problem with the infinitely repeated game approach concerns the issue of renegotiation. Suppose B breaks her promise in some period. According to the equilibrium, S is meant to "punish" B by supplying a low-quality widget forever more (in effect, no trade occurs). However, by punishing B, S is also punishing himself since he won't get any payment. The question then is, why don't the parties let bygones be bygones and reinstate the cooperative outcome? After all, it is not as if S has learned anything adverse about B. B's characteristics are known, and the fact that B has broken her promise today tells S nothing about whether she will do so again.

The trouble is that, if B anticipates that cooperation will be restored after she breaks her promise, then this increases B's incentive to break her promise, and cooperation may not be sustainable. In other words, if the parties are rational enough to realize that they will renegotiate after a breach, then this may prevent cooperation occurring in the first place, i.e. the outcome may be as in the one-shot game.[4]

Partly because of these difficulties with the infinitely repeated game approach, another strand of the literature has instead supposed that the game is played finitely many times – t, say – but that the parties are not perfectly informed about each other: there is asymmetric information.[5] Suppose, for example, that there is a small probability that B is someone who always keeps her promises no matter what. (She is "irrational.") B knows whether she is the rational type or the irrational type, but S does not. Then in the early stages of the game, B has an incentive to pretend to be the irrational type even if she isn't, in order to encourage S to trade with her. In fact, it can be shown that, if t is large enough, then in every

equilibrium of the t-period game, cooperation will be sustained almost all of the time.

The asymmetric information approach has the advantage over the infinitely repeated game approach in that it does not require an infinite horizon and can deal with the problem of renegotiation. However, it faces another difficulty. It turns out that the approach is very sensitive to the precise characteristics of the irrational type, about which we as modelers know very little. One way to see this is the following. Suppose that in addition to the irrational honest type there is another "irrational type," who is totally dishonest but, with some probability, has an irresistible urge to propose an agreement to trade in any period. Then there is an equilibrium of the following form. The parties do not trade in any period. The irrational buyer who has an irresistible urge proposes to S that they should trade: S turns her down because he rationally sees that this type of buyer will never pay him. The other buyer types propose nothing because there is no point: they would be confused with the irresistible urge type and thought to be dishonest and not worth trading with. This way the no-trade equilibrium is sustained however large t is.

The conclusion is that the asymmetric information approach does not provide a very solid foundation for the idea that cooperation will necessarily occur when play is repeated many times.

A third approach is to move away from thinking about the trustworthy type as a fringe, irrational agent and instead to recognize that all agents are trustworthy to some extent. One way to do this is to suppose that each agent incurs a psychic cost \$$C$ if she breaks a promise, where C is distributed in the population according to a known probability distribution and a person's C, although known to her, may or may not be known to others. This approach, like the asymmetric information approach, can explain cooperation in a finite horizon model.[6] However, not surprisingly, as with the asymmetric information approach, its conclusions are very sensitive to assumptions made about the distribution of C in the population and also about the nature of C – matters that again the modeler knows little about. For example, suppose B pays S slightly less than what she promised. Does she incur the whole psychic cost C or just part of it? Or suppose B promises n different sellers that she will pay them if they perform well (they are workers, say) and then simultaneously breaks her promise to them all. Does B incur a total psychic cost of \$$C$ or \$$nC$? The nature of the optimal self-enforcing contract is likely to be very dependent on these features of the model.

Not only are the asymmetric information and psychic cost approaches quite sensitive to the precise modeling assumptions made, but also it turns out that these approaches are not that easy to work with in a contractual

or organizational setting. For these reasons, most researchers have used the infinitely repeated game approach, in spite of its shortcomings. In what follows, I will do the same. In the next section, I use the approach to illustrate the effects of self-enforcing contracts on the determinants of firm boundaries.

4 Norms and firm boundaries

A good application of norms in the organizational context is to the issue of the determinants of firm boundaries (the "make-or-buy" decision). Trust helps to sustain agreements both inside the firm and between firms. An interesting question to ask is: does trust favor one type of transaction relative to the other?

In the last fifteen years or so a formal literature – the property-rights approach – has developed that tries to explain firm boundaries in terms of the optimal allocation of asset ownership (see Grossman and Hart 1986 and Hart and Moore 1990).[7] This literature shares with the earlier transaction cost literature of Williamson (1975, 1985) and Klein, Crawford and Alchian (1978) the view that firms are important when contracts are incomplete. It departs from the transaction-cost literature in being more explicit about the role of decision rights and the link between decision rights and asset ownership. According to the property-rights view, the owner of a non-human asset has residual rights of control over the asset, i.e. the right to make all decisions concerning that asset that have not been specified in a contract or that are not inconsistent with some law. (When there are multiple owners of an asset or firm, they will typically delegate some of the residual control rights to a board of directors.) Residual control or decision rights are like any other good: there will be an optimal allocation of them. For example, suppose that individuals 1 and 2 are involved in an economic relationship. If it is important to encourage 1 to make an asset- or relationship-specific investment, it may be efficient to allocate ownership of some key non-human assets to 1. This way individual 1 is protected to some extent against "hold-up" by 2 since, if the economic relationship with 2 doesn't work out, 1 always has the option to take her assets away and trade with someone else. However, while allocating assets to 1 protects 1 from hold-up by 2, it has the opposite effect on 2: since 2 has fewer assets to take elsewhere, 2 is now more vulnerable to hold-up and so will be less willing to make an asset- or relationship-specific investment himself. Typically it will be optimal to divide the assets between the parties so that each party has some. If we view each set of assets with a common owner as a firm, this yields a theory of firm boundaries.

The property-rights theory has in the main been applied to static or one-shot situations where parties are self-interested and not trustworthy. However, it is natural to ask how the optimal allocation of assets or firm boundaries changes when norms and trust operate. Some papers that study this issue include Baker, Gibbons and Murphy (2002) and Halonen (2000). In what follows I will discuss some of the ideas behind these papers, using as a vehicle the paper on trucking by Baker and Hubbard (2000) (the Baker–Hubbard paper is not itself about norms or trust).

Consider a shipper S who at date 0 wants goods shipped from A to B. The shipper hires a trucker T to do this. The trucker may come with his own truck, in which case he is an independent contractor, or the shipper may provide the truck, in which case the trucker is an employee. We will assume that the shipper and trucker can contract on the shipment from A to B (known as the front-haul), but that they cannot contract on several other things. First, the shipper may sometimes want the trucker to engage in a back-haul, i.e. transport a second shipment of goods from B to another destination C. However, whether there will be a back-haul and its nature – how valuable the second shipment is, whether it is easy to transport, and the identity of the destination C – are variables that are hard to forecast and become known only when the trucker arrives at B, at date 1, say. So contracting about the back-haul must wait until then.[8] Second, the parties cannot contract on maintenance: how well the trucker drives the truck. The trucker may have an incentive to drive fast, take time off to visit a friend, and then speed again to reach B; this may be pleasant for the driver, but is bad for the truck. To make things simple, we will assume – at some cost of realism – that maintenance is observable to the trucker and shipper but is not verifiable.

Third, the trucker can spend time searching for alternative customers as he drives from A to B. (He has a mobile phone/access to the internet, etc.) For those searches to pay off the trucker must be able to drive the truck away at date 1. Some such searches are productive – they pay off in the absence of a profitable back-haul from B to C – but others are carried out to improve the bargaining power of the trucker when he negotiates over the terms of the back-haul at date 1. To simplify we will follow Baker–Hubbard in assuming that all search activities are on average *unproductive*, i.e. their return is less than their (effort) cost.

Finally, we will assume that the owner of the truck bears all the increases or decreases in the value of the truck; he is the residual income claimant. This may seem like a rather traditional view of ownership, and it is extreme (it rules out value-sharing agreements between the shipper and the trucker), but it is consistent with the residual control rights approach in

the following sense: the owner has the (residual) right to decide to whom to sell the truck, when, and at what price. To the extent that the owner can always sell the truck for 1¢ (the verifiable price) and at the same time agree to supply another service to the buyer for an exorbitant price, he can ensure that he never has to share the sales revenue with anyone else.

The key question is: who should own the truck? In the static or one-shot version of the model, the trade-off is the following. If the trucker owns the truck he will maintain it (he bears the value consequences), but he will engage in search or rent-seeking activities (as owner of the truck, he can exploit these activities since he has the right to drive away the truck at date 1). On the other hand, if the shipper owns the truck, the trucker will not maintain it at all (he does not bear the value consequences), but neither will he engage in rent-seeking activities (these do not pay off given that the trucker does not have the right to drive the truck away).

To simplify matters, I will assume that in the one-shot model encouraging maintenance is more important than discouraging rent-seeking and so it is best for T to own the truck, i.e. T should be an independent contractor rather than an employee. To the extent that S owns other assets than the truck and T doesn't, I will refer to this arrangement as "non-integration," and to the arrangement where S owns the truck (and therefore has all the assets) as "integration."

So far we have analyzed asset ownership or firm boundaries in a trust-less environment. I now want to ask the following question: how does trust affect the boundaries of the firm? To the extent that there is a conventional wisdom on this matter, I suspect that it is that an increase in trust will make it more likely that the parties will "use the market," i.e., choose to be independent (non-integration) – and to be linked by a relational contract – rather than to become one firm (integration). This conventional wisdom can probably be traced to the fact that transaction-cost economics (TCE) tends to see the market as the first choice if it is feasible, and in a high-trust environment it is likely to be feasible.

To analyze this choice more formally, let us proceed as in section 3 and suppose that the relationship between S and T is repeated infinitely often and that both parties discount the future at the common discount factor Δ, where $0 < \Delta < 1$. We are led to consider the following self-enforcing contract: T promises to maintain the truck well and to engage in minimal rent-seeking activity (search). In return, S promises a fixed payment P per period. The self-enforcing contract is sustained as follows: if either party breaches, we revert to the equilibrium of the one-shot game described above forever more. (In contrast to section 3, this equilibrium involves some trade rather than no trade.) We will also suppose (following Baker, Gibbons and Murphy 2002, but in contrast to Halonen 2000) that

ownership of the truck can be transferred at this point, i.e. if S owns the truck T will buy it. (Recall that, given our assumptions, it is efficient for T to own the truck in the one-shot game.)

Note that $\Delta = 0$ corresponds to the one-shot game, since, if the future does not matter at all, no cooperation can be sustained. At the other extreme $\Delta = 1$ corresponds to the case where trust can easily be sustained since the future overwhelms the present in importance. Thus an increase in Δ can be interpreted as a move to a higher-trust environment.

Thus the question: how does trust affect asset ownership or firm boundaries? can be rephrased as: how does an increase in Δ affect asset ownership or firm boundaries?

The answer is that it all depends: an increase in Δ does not have a clear-cut effect on the choice between integration and non-integration (see Baker, Gibbons and Murphy 2002). To see why, note that an increase in Δ improves all organizational forms. If Δ is close to 1, the first-best – where T maintains the truck well and does not engage in rent-seeking – can be sustained under a self-enforcing contract whether S owns the truck or T does. The reason is that no one wants to breach a self-enforcing contract since the future gains from cooperation are so large relative to the short-run gain from breaching. On the other hand, if Δ is close to zero, then non-integration is best (given our assumptions). This suggests that there is no simple monotonic relation between optimal organizational form and the discount factor Δ.

Specifically, it is easy to construct cases where integration is superior to non-integration when Δ is fairly close to 1, even though non-integration is superior to integration when Δ is close to zero. (Such cases turn the conventional wisdom on its head – a higher-trust environment favors large firms.) To see why, suppose that the cost of maintenance is very low but the value is very high. In the static model (one-shot game), there will be no maintenance under integration, which is highly inefficient. But in the dynamic model it is easy to get maintenance by offering T a small bonus if he looks after the truck. Since the bonus covers his (small) cost, T will maintain the truck as long as he expects to receive the bonus; and S will pay the bonus since, given that it is small, there is little gain from not doing so. Finally, there is no incentive for T to engage in rent-seeking under integration since he can't drive away with the truck. So in this case the first-best can be achieved under integration in the repeated game even for moderate discount factors Δ.

In contrast, under non-integration, while T will maintain the truck (as in the static model), he may need quite a large bonus from S to be deterred from engaging in rent-seeking behavior; but the promise of a large bonus gives S a strong incentive to breach. Hence it may be impossible to sustain

the first-best under a self-enforcing contract for moderate levels of Δ when T owns the truck.

Note that, in spite of what I earlier called the conventional wisdom, there is some evidence that trust does indeed favor large firms rather than small ones; on this, see La Porta *et al.* (1997) and Kumar, Rajan and Zingales (1999).

It should be emphasized that, while in this example non-integration is optimal when Δ is small and integration is optimal when Δ is large, it is easy to construct another example based on the same model that yields the opposite conclusion.

I think that the correct conclusion to draw from this discussion is the following. The boundaries of the firm will be drawn to elicit appropriate actions from the parties – in this case, truck maintenance and (absence of) rent-seeking. In broad terms the choice between the two organizational forms will depend on the importance of these goals and the ease with which they can be achieved. It is easier to encourage maintenance if T owns the truck and to discourage rent-seeking if S does. This is true both in the static model and the repeated game. Thus in qualitative terms trust does not change things that much.[9]

5   The role of formal contracts

So far I have discussed the role of norms in situations where the opportunities for writing formal contracts have been quite limited. In section 3 formal contracts were impossible and in section 4 the only formal contracts concerned the allocation of asset ownership and spot (one-period) deals between S and T.

In this section I will make some brief remarks about the general impact of formal contracts on the sustainability of self-enforcing contracts, and mention one implication for judicial attitudes toward firms. Formal contracts have at least two effects on self-enforcing contracts. First, the better formal contracts are, the smaller is the surplus remaining for the parties to try to exploit via a self-enforcing contract. This reduces the incentive of parties to breach a self-enforcing contract, since, given that there is less at stake, the gains from opportunistic behavior are lower. Second, however, if a self-enforcing contract is breached, the penalty is also lower since the parties can always rely on formal contracts in the post-breach, no-trust environment; as a result, as argued by Baker, Gibbons and Murphy (1994), the incentive to breach may *rise*.

Because these two effects are opposing, it is hard to draw clear-cut conclusions about whether formal contracts will make it easier to sustain self-enforcing contracts (i.e. formal and informal contracts are complements),

or more difficult (i.e. formal and informal contracts are substitutes). Which way it goes would seem to depend on the circumstances.

In their interesting recent paper, Rock and Wachter (2001) take the position that one would expect to see few formal contracts inside the firm given the concentration of residual control rights in the hands of one party (the board of directors): rather the firm is a place where informal agreements will flourish.[10] My interpretation of (one part of) their argument is that it is hard to imagine two divisions of a firm being bound by a formal contract. The reason is that either party can be prevented from fulfilling the contract by the board of directors, who can always *ex post* deny the members of the divisions (including the division heads) access to key non-human assets or key decision-making authority. Division members are unlikely to be prepared to enter into formal agreements which require them to pay damages in the event of breach, given that they have so little power to ensure that these agreements are implemented.

Not only do Rock and Wachter provide a persuasive argument as to why formal contracts may be difficult to sustain inside the firm, but also the discussion of this section suggests a reason why formal contracts may be undesirable even if they *are* feasible: they may in some cases make it harder to sustain self-enforcing contracts (the case of substitutes described above). This may provide some justification for the view that the courts should be hesitant to intervene in the firm's informal business; that is, they should take a hands-off attitude even in cases where they have the ability or expertise to intervene.

6   Summary and open questions

In this chapter, I have argued that it has been difficult to incorporate norms into the theory of organizations; and also that, although there has been some interesting recent work on this topic, this work has not to date greatly changed our views about the determinants of organizational form.

I want to conclude by making a further qualification about the material discussed above. The infinitely repeated game models of sections 3 and 4 are really models of individual reputation or trustworthiness. That is, while it is tempting to think of the buyer and the seller in section 3, and the shipper in section 4, as representing firms, an extra step is really required for the argument to work. This step involves explaining why a particular set of norms or reputation is associated with a firm or organization rather than with an individual or set of individuals who work there.

To put it in stark terms: what ensures that, when the CEO of a company that is known for high trust leaves, the high-trust reputation doesn't go with her? Somehow there has to be some stickiness in the firm or system,

so that a firm's reputation can be separated from that of key personnel. To put it another way, a firm's reputation has to have some of the characteristics of a non-human asset. However, exactly how this comes about is far from obvious.

One attempt to explain how a reputation can be embodied in a firm rather than a set of individuals is contained in a paper by Tadelis (1999).[11] Tadelis considers the relationship between a firm and its consumers. Think of the way a firm treats its customers, e.g. the way it services its product, as a norm. Tadelis assumes that every consumer observes this norm, i.e. they know how past customers have been treated, but that consumers do not know who owns (or manages) the firm. If ownership changes, customers do not see this and so assume that the firm will continue to treat its customers in the same way. As a result a firm that has treated its customers well in the past will have a valuable reputation: moreover, outside buyers may be prepared to pay a lot for this reputation since at least in the short run – until and unless they show that they cannot maintain the reputation – they can charge more for their product than if they started from scratch (without a reputation).

The Tadelis model provides a useful starting point in helping to understand why a firm's intangible assets can be valuable. However, the idea that a firm's reputation matters only when (a significant fraction of) consumers cannot observe a change in ownership is not that plausible. It is to be hoped that in the future it will be possible to relax the informational assumptions of the model. For the moment the creation of a theory of norms attached to a firm or organization seems an even more challenging goal than the development of such a theory for the case of an individual.

NOTES

1. For a discussion of the implications of transaction costs for contractual relationships, see, for example, Williamson (1975).
2. But see Fehr and Gächter (2000) for a discussion.
3. For a discussion, see Fudenberg and Tirole (1991, chapter 5).
4. To be a bit more precise, suppose that the gains from renegotiation are split in a fixed (exogenous) way. Then if B gets most of the gains she has a large incentive not to pay S in any period; while if S gets most of the gains he has an incentive to renounce the self-enforcing agreement at the beginning of a period (i.e. refuse to supply) and negotiate a better deal.
5. See Kreps *et al.* (1982).
6. See, e.g., Hart and Holmström (1987).
7. For a summary of this literature, see Hart (1995).

8. For a formal justification of the idea that, when the future is uncertain, many aspects of a contract will be negotiated *ex post* rather than *ex ante*, see Hart and Moore (1999a).

9. A possible qualification should be noted. In the static models of Grossman and Hart (1986) and Hart and Moore (1990), joint ownership of an asset is never optimal. In contrast, the repeated game model described in this section can explain joint ownership of an asset if it is supposed that ownership of the asset cannot be transferred after the breach of a self-enforcing contract (see Halonen 2000). The reason is that, since joint ownership is sub-optimal in the static model, the threat of it can support cooperative behavior in the dynamic model. Note, however, that joint ownership can be optimal in more complicated versions of the static model, where it is important to discourage rent-seeking behavior of *both* parties (see, e.g., Rajan and Zingales 1998). (If neither party can walk away with the asset, then each party's incentive to search for alternative trading partners is reduced.) Thus in fact joint ownership (or joint ventures) can be explained both in the static (no-trust) model and in the dynamic (trust) model.

10. A related, but distinct, idea is that firms will arise in situations where it is important to suppress individual incentives and foster cooperative behavior. (See Holmström 1999.)

11. For earlier work, see Kreps (1990) and Tirole (1996).