

Munich Lectures

Lecture 1¹

In the last twenty years or so, a lot of my work has been on a topic that economists call the theory of the firm. I thought I would spend this first lecture saying a little about what I and others have been up to, and also mention a couple of applications of this work.

First, some background. Economists have written an enormous amount about markets. In fact, many people who don't know much about economics think that economists spend all their time analyzing perfect markets. While this is a great exaggeration, it's fair to say that one of the achievements of economics is to identify the conditions under which markets work well-- that is, to formalize Adam Smith's invisible hand. But economists also spend a lot of time thinking about market failure. And one of the most striking pieces of evidence of market failure is found in the capitalist system itself: the fact that firms exist. (For economists in the audience, I'm not talking about a market failure that requires government intervention, but one that the parties sort out themselves.)

Now the idea that the existence of firms proves that markets don't work well may be surprising to some people, so let me clarify what I mean. In a famous paper written in 1937, Ronald Coase, a British economist and future Nobel Laureate, explored the reason for firms. He started off by quoting another British economist, Sir Arthur Salter: "The normal economic system works itself. For its current operation, it is under no central control Over the whole range of human activity and human needs, supply is adjusted to demand, and production to consumption, by a process that is automatic, elastic and responsive." (Here Salter is, of course,

¹I am very grateful to Robin Greenwood for research assistance.

referring to the price mechanism.) Coase then points out that within a firm this description does not fit at all: the price mechanism hardly operates. If a worker moves from department X to department Y inside a firm, she does so not because she will be paid more in department Y, but because she is told to move from X to Y. In other words, outside the firm coordination is done by the price mechanism, while inside the firm it is done by direction by an entrepreneur or boss. One is the invisible hand, and the other is the visible hand, and they are clearly very different ways of coordinating production. The fact that we see the visible hand used a great deal in capitalist economies--some firms are huge and an enormous amount of economic activity takes place within firms--must mean that Adam Smith's invisible hand isn't perfect; because if it were everybody would be an independent entrepreneur.

But of course the market isn't that bad. We certainly don't see all economic activity being done inside firms. There is clearly a balance between the firm and the market. The interesting question is what determines this balance.

[The following quote by Dennis Robertson about the balance between the firm and the market is one of my favorites in economics, but I don't know whether it works in German. Robertson said that we find "islands of conscious power in this ocean of unconscious cooperation like lumps of butter coagulating in a pail of buttermilk."]

In the nearly seventy years since Coase's article, economists have come up with three main reasons why markets may not work well and why firms can sometimes, but will not always, do better. They are: the existence of relationship-specific investments; the difficulty of writing good contracts; and the fact that people can't be trusted. (In what follows, I will quote liberally from the work of Oliver Williamson; Ben Klein, R. Crawford and Armen Alchian; and from my

own work with Sandy Grossman and John Moore.)

Let me discuss each of these briefly. First, relationship-specific investments. If you read a standard economics textbook, you will be introduced to the idea of a competitive market. An example of a good traded in such a market might be wheat. The simplifying assumption here is that it doesn't matter whom you buy wheat from. All wheat is the same (for the sake of argument), and as a result I can buy wheat from one person today and from someone completely different tomorrow.

Many important economic relationships are not like this. They have a permanence. They are more like a marriage than a series of one-night stands. And the reason for this is that parties make relationship-specific investments. They commit or lock into each other.

For example, suppose I'm an electricity company and I plan to make electricity by burning coal. (This example is stimulated by the work of Paul Joskow.) One possibility is for me to locate a moderate distance from a number of different coal mines: in essence I buy coal on the open market. Another possibility is for me to locate right next to a particular coal mine, yours, say. In some circumstances, this may be more efficient because I save on the cost of transporting coal.

However, if I locate next to you, I am committing to a long-term relationship with you. This may be quite risky. My decision to locate next to you may be very profitable if I get cheap, high-quality coal for many years; but it might turn out to be a terrible mistake if our relationship breaks down after a year and you refuse to supply me any more coal or raise the price. At this point, it would be very costly for me to switch to another coal mine or buy coal on the open market; I might have to build another plant.

How do I guard against bad luck, or even worse opportunistic behavior by you, and at the same time ensure that we enjoy the efficiency benefits from my locating next to you? The natural way is for us to write a long-term contract. For example, the parties may agree that the coal mine will supply the electricity company with a certain quantity of coal at a certain price for the next thirty years. With such a contract in hand, the electricity company may be quite willing to commit to the coal mine and locate next to it. Now we come to the second reason for the existence of firms--the difficulty of writing good contracts. The future is uncertain and unpredictable, particularly over a 30-year period (and relationships between electric companies and coal mines do last this long). We are even more aware of this unpredictability since the events of September 11. Because the future is so uncertain, it is hard to write a long-term contract that takes into account all relevant contingencies: any contract will miss some potentially important future events--it will be incomplete.

In the case of the electricity company and the coal mine, the problem may be that the electricity company needs the coal it burns to be pure. However, it may be hard to specify in advance what "pure" coal is because there are many potential impurities. In the event, ash content may be the relevant impurity: it may turn out that it is cheaper for the coal mine to take coal out of a high ash-content deposit, but that this reduces the electricity company's profit a lot. Moreover, because the contract does not specify purity in sufficient detail, the coal mine may be within its rights to do this. As a result, even with a long-term contract, the electricity company is vulnerable and may lose money. Anticipating this, it may refuse to locate next to the coal mine.

So far I have argued that the electricity company may be unwilling to locate next to the coal mine when it is hard to write a good contract and when the parties can't trust each other.

Notice where trust comes in. If the electricity company could totally trust the coal mine to be honorable and not to save money on low-quality coal at great cost to the electricity company, then the absence of a good contract wouldn't matter. Unfortunately, even in the best of all possible worlds, there is a limit to how much business people can and do trust each other.

There is one final step to the argument. I have considered the case where the electricity company and coal mine are separate firms linked by an incomplete long-term contract. I want to interpret this as a market transaction given that the two companies are independent entities (even though it is not the kind of market transaction studied in standard textbooks). However, there is another possibility. The electricity company could buy the coal mine before locating next to it--this turns the transaction into one inside one large firm. How does this change things? In my work I have argued that the crucial change is that the electricity company manager, who is now the boss of the coal mine manager, acquires the right to intervene in the coal mine (he has residual control rights). So, for example, take the situation where the coal manager is thinking of mining high ash-content coal. When the coal mine is separate, the electricity company can do little about this except to bribe the coal mine to supply better coal (i.e., to pay the coal mine more). However, when the coal mine is part of the electricity company, the electricity company can intervene directly to forbid this action. The electricity company manager now owns or is in charge of the coal mine and can direct from which part of the mine the coal is taken. Moreover, if for some reason the coal company manager refuses to follow the electricity company's instructions, the electricity company can fire him and replace him with a better-behaved manager--something that was not possible when they were separate companies.

These interventive rights thus make the electricity company less vulnerable, and more

willing to locate next to the coal mine. However, there is a downside. The electricity company's ability to intervene in the coal mine's operations may make the coal manager more vulnerable. Suppose the coal manager has an idea about how to mine the coal more efficiently. When the coal mine was separate, this idea would lead to higher profits for the coal mine, some of which would end up in the coal manager's pocket given that he owns the mine or is in charge of it. Now that the electric company owns the coal mine, however, it is much more likely that the electricity company can implement the idea without rewarding the manager adequately. (The notion that an employee is not fully rewarded for his or her ideas is probably familiar to many of us.) Anticipating this, the coal manager has less incentive to have good ideas in the first place.

Let me summarize. If it is important to protect the electricity company from opportunistic behavior by the coal mine, in the form of supplying low-quality coal while sticking to the letter of the contract, then the electricity company should buy the coal mine--it is better to carry out the transaction inside the firm. On the other hand, if it is important to protect the coal manager from opportunistic behavior by the electricity company, in the form of not rewarding the managers for good ideas, then the electricity company and coal mine should remain independent--the transaction should be left in the marketplace.

[One should add that there is a third case where the coal mine buys the electricity generating plant and the coal manager acquires intervention rights, but I won't dwell on this.]

I want to spend the second half of this talk on two, brief applications of the above ideas. The first is to privatization or outsourcing.

Not so long ago it was taken for granted that certain services would be provided by the government even in a capitalist economy. But in the last twenty we have seen a huge

privatization movement--started by Margaret Thatcher in the U.K. and continued in many other countries--and now things look rather different. In fact, these days some people seem to hold the view that almost anything should be privatized.

While this is often presented as a political issue between left and right, really it's an economic issue--in fact one involving incomplete contracts. (One of the first people to recognize this is someone you know well, Klaus Schmidt. He published a paper about it in the Journal of Law, Economics and Organization in 1996. My discussion will follow more closely a paper I published in 1997 with Andrei Shleifer and Rob Vishny.)

Some background. Economists generally agree that some goods or services in society must be paid for by the government. Clearly in this category are military expenditure, the police, prisons, and various emergency services. Possibly in this category are hospitals, schools and garbage collection. Why can't we rely on the market for these goods? The answer is that they are public goods--their benefits are enjoyed by society as a whole. (Would you pay individually for a missile or for a policeman?)

But it doesn't follow that the government must make these goods just because the government pays for them, i.e., that the government must own the facilities or employ the workers that produce them. In principle it is quite possible for the government to contract with a private provider of police or prison or military services on behalf of society.

What determines whether the government should make these goods or contract for them? The answer is that the trade-off is similar to that facing the electricity company in deciding whether to buy the coal mine. To be concrete, let me focus on the case of prisons (as I did in my paper with Shleifer and Vishny). (As a piece of background, all German prisons are state-

owned; but in the U.S. there are many private prisons.)

The government is like the electricity company and the provider of services is like the coal mine. Just as the electricity company wants the coal mine to supply coal, the government wants a prison provider to house and feed prisoners (and stop them from escaping!). In a world where perfect contracts can be written the government can rely on a contract to get high quality prison services at a reasonable price--in the same way as the electricity company can rely on such a contract to get coal. [For people who know the literature, in a world of perfect contracts, it also doesn't matter if the government owns the prison, but I will not dwell on this.]

But as I've argued, we live in a world where contracts are incomplete. In this case, the government may want to own the prison in order to prevent the provider from supplying low-quality prison services while sticking to the letter of the contract. The downside of government ownership is that its ability to intervene in the prison may reduce the prison provider's incentives to be efficient and innovative.

The two most important areas of incompleteness in the case of prisons seem to involve the use of force and the quality of personnel. It turns out to be difficult to specify (and enforce) clearly in a contract how skilled prison workers should be or the kind of training they should receive. Unskilled and untrained workers tend to be bad at dealing with emergencies--they sometimes use too little force to control a situation and sometimes too much. Force is another thing that it is difficult to deal with contractually: the exact conditions under which it should be used are hard to specify.

In our paper we expressed the concern that private providers would have an incentive to save money by hiring low-quality prison guards and giving them poor training (in the same way

that an independent coal mine would have an incentive to mine low-quality coal). And in fact, there is quite a lot of evidence--some of it casual--that private prisons in the U.S. do cut costs at the expense of quality. A recent survey found that private prisons have 50% more inmate-on-staff assaults and 2/3 more inmate-on-inmate assaults than public prisons. In the last few years all the major private prison companies have been involved in lawsuits concerning the way they treat prisoners. There have been some notorious incidents. Six very violent prisoners escaped with ease from a facility operated by CCA, one of the biggest private providers. At the same prison, two prisoners were murdered. At the facility of another major provider, Wackenhut, which housed junior offenders, some girls were raped and one of them committed suicide.

Excessive force is not the only problem. Sometimes force can be insufficient. In our 1997 paper, we discussed the case of ESMOR, which ran a private detention facility in New Jersey that houses foreign nationals caught attempting to enter the U.S. illegally (not on average a hard-core criminal crowd). When a riot broke out in 1995, the under-trained guards ran away and called the police from a public telephone!

[This is not to say that bad things don't happen in public prisons--they do.]

Of course, as we have noted, private providers may have certain advantages: they have stronger incentives to be efficient and to be innovative, e.g., they might develop interesting new education, training or rehabilitation programs for prisoners (for which they would be reimbursed by the government). However, there does not seem to be much evidence of such innovation in the case of prisons, and the documented cost savings can be attributed mainly to the fact that private prisons hire lower-quality, non-union labor rather than to efficiency gains.

Efficiency and innovation effects are likely to be particularly small in the case of high

security prisons that house the most violent criminals; here making sure that guards carry out routine tasks correctly and react to emergencies appropriately is probably paramount. Shleifer, Vishny and I therefore concluded that there is a strong case for public ownership of prisons generally and for high-security ones in particular.

[In this discussion I have ignored several issues. One is competition--which is not very relevant in the case of prisons. Another is corruption and patronage--which is (see also Glaeser's recent paper). As a sign of the influence of our paper, note that the private prison population increased from 78,000 in 1996 to 276,000 in 2001!]

Similar arguments can be applied to understand the provision of other public goods. For example, in the case of garbage collection or weapons production, contracts can quite effectively prevent the damage to quality from cost reduction; and given the importance of innovation (particularly in the case of weapon production), one would expect private supply to have a strong edge over public supply. In contrast, in the very topical (at least in the U.S.) case of airline security, it is plausible that innovation is less important than simply not cutting corners--airline security is more like the case of high security prisons--and so the case for public provision is quite strong, contrary to what many U.S. politicians seems to think.

My second application is rather different--it concerns bankruptcy reform. Here I will be extremely brief. The discussion so far has concerned whether certain activities should be carried out inside a firm or outside a firm. However, similar ideas--concerning incomplete contracts and intervention rights--can be used to understand how a firm should finance itself--whether it should borrow or issue shares--and also what should happen if it cannot pay its debts. (My discussion of the financing choice follows Aghion-Bolton, 1992).

Suppose you are an entrepreneur and you need capital for a new project. One thing you can do is to issue shares to one or more outside investors. Of course, the investors may be worried about simply handing over the money to you and hoping that one day you will pay them back. They may want some protection. And one way for them to get protection is for them to acquire voting control, that is, intervention rights. If they have the majority of votes, they can replace the board of directors--and you--if they think you are not doing a good job.

However, giving voting control may have a significant downside. Just as with the coal manager, your incentive to innovate may fall. Also, you may worry that if things start to go badly the investors will terminate the project far before you wish to do so.

Is there anything that can be done to weaken the investors' intervention rights and yet ensure that they are protected? A good compromise may be debt finance. Suppose the entrepreneur borrows from a bank or financier. The entrepreneur promises to make a sequence of payments at specified dates. If the entrepreneur makes the payments, he retains control of the project. If he fails to make the payments, he loses control: the financier (bank) can foreclose on the project and liquidate it.

With a debt contract, the entrepreneur has an incentive to generate high earnings in order to repay his debt and faces no threat of intervention as long as he does so. At the same time the financier is protected: the entrepreneur will be encouraged to make debt repayments rather than reinvest earnings in the project since otherwise he will lose control.

In the above situation I talked about borrowing from a single bank or financier. In reality, however, people often have many creditors/lenders. Now things get more complicated if the entrepreneur cannot pay his debts. (This is a clear example of incomplete contracts. The

contract typically does not specify what happens if the entrepreneur defaults to many creditors.) Previously I suggested that, if the entrepreneur defaults on his debt, the bank should have the right to liquidate the project. But now there are many creditors/lenders. If each one has the right to liquidate part of the project, to get his or her money back, then the problem is that there may be a race to be first, and, in the rush to seize assets, a potentially valuable project may be dismantled (torn apart). That is, while it may be in the collective interest of the creditors, that the project continues (and the debts are restructured), individual action may not lead to this.

There is an obvious solution to this problem. Instead of allowing each creditor to liquidate part of the project, turn the creditors into owners. That is, carry out an automatic debt-equity swap whereby the former creditors become the new owners of the company (with votes), i.e., if 100 creditors/lenders are each owed 1 DM, then each owns 1% of the project after the debt-equity swap. Now let the new owners decide collectively--through a vote--whether to liquidate the project or keep it going. If the project is worth more alive than dead, the owners will vote to keep it going; if not, they will vote to liquidate it. The case where a valuable project is torn apart is no longer a problem.

This is the essence of the bankruptcy proposal suggested in Aghion, Hart and Moore (1992). Of course, reality is much more complicated than the above example, not least because there are different types of debt with different securities. The proposal can be adapted to this, but I will not discuss this here.

Several countries in the world have expressed an interest in the AHM proposal. Germany is not one of them! (The German bankruptcy procedure seems to be a combination of U.S. Chapter 11 and U.K. administration.) A few countries including Colombia, the Philippines, and

Switzerland (for banks), have adopted or may adopt some version of AHM. (Colombia apparently has already adopted a version.)

It is a pity, in my opinion, that Switzerland doesn't currently have such a proposal for airlines. I believe that many of the recent problems facing Swissair are the result of creditors' collective action problems--in particular, the fact that senior creditors like UBS and Credit Suisse have no real interest in saving Swissair since they can be repaid through asset liquidation rather than going concern value. These problems of creditor conflict would be removed if an automatic debt-equity swap had occurred along the lines of AHM soon after the debt default. Then UBS and Credit Suisse, as owners, would have wanted to maximize the long-run value of Swissair. While I am not making any prediction about the efficient outcome for Swissair--there is plenty of evidence that it is highly inefficient and should be significantly slimmed down--it probably should not be closed down and I doubt that its planes would have been grounded if a debt-equity swap had occurred.

To conclude . . .