Internalization and incomplete contracts

This version: December 6, 2016

We have focused on the firm's location decision for most of this course, studying different motivations for the firm to operate outside its home country: market access (horizontal/export platform FDI), factor cost savings (vertical FDI), and tax minimization. Throughout these discussions, we have assumed that the firm will own the entity producing abroad. In this note, we ask the question: Why would a firm want to own a foreign affiliate, rather than contract with a separate firm for production in the foreign country? How does this decision depend on the characteristics of the goods being produced? How does it depend on the characteristics of the firm?

1 A model of make vs. buy

Consider a firm i that produces final goods from component parts. This *final-good firm* owns a production technology that creates the final good q, which can be sold at price p. This technology requires components m,

$$q = A_i m^{\alpha}. \tag{1}$$

The productivity parameter A_i is indexed by i. Different final-good firms have different productivity levels. We make two important assumptions about the components.

Relationship specificity. The components are relationship specific. The components are specially tailored to the final-good firm's application. The components have no value to anyone else but the final-good firm.

Difficult verification. It is difficult to verify the quality of the components. The final-good firm and the supplier can judge the quality of the components, but outside parties — like a court — cannot.

1.1 Component costs

There exists a *supplier firm* that can build components for the final-good firm. The supplier has an advantage in producing components. The supplier can produce components at a cost of p_m per unit.

The final-good firm could also build the components itself. If the final-good firm produces components, they cost the firm γp_m per unit, where $\gamma > 1$. In addition, the final-good firm would have to pay a fixed cost of operating the components production line, f^I.

1.2 Contracting

If the supplier produces the components, it saves the final-good firm the fixed cost of production and the components are made at a lower cost. So why not go with the supplier? The issue here is contractibility. Suppose it is not possible (or is very costly) to write a contract that specifies every aspect of the component to be delivered. One way this contracting problem can arise is if it is hard to verify the quality of the component.

If it is difficult for an outside arbiter to judge the quality of the component, then the final-good firm would have an incentive to claim that the components delivered did not meet the contract specifications and the firm would want to renegotiate the contract for a lower component price. Since the components are relationship specific, the supplier's outside option is zero. The components cannot be resold to another firm. The supplier foresees that the final-good firm will renegotiate the contract in the future, which will lead the supplier to produce fewer components than it would if the firms could write complete and enforceable contracts.

Without components, the final-good firm cannot produce. Since the value of the final-good firm's production technology is zero if it does not receive parts, the supplier has an incentive to renegotiate the contract for a higher component price.

The result is that both firms would like to renegotiate the contract, so no contract is written. The best the firms can do is bargain over the payment to the supplier after the components have been produced.

Hold up!

This situation is known as the *hold-up problem*. Each firm has an incentive to hold up the other firm for a better deal because the other firm has made a relationship-specific investment that is not valuable (or less valuable) to outsiders. The underlying problem is that the two firms cannot write a contract that is enforceable after production has occurred. Our model is a simple example of the hold-up problem.

The 2016 Prize in Economic Sciences in Memory of Alfred Nobel was awarded to Oliver Hart and Bengt Holmström for their work on hold-up problems and incomplete contracting.

1.3 Firm structure

The final-good firm weighs the inefficiency created by imperfect contracting against the inefficiency the final-good firm faces if it creates the component parts itself. A final-good firm i must choose between two ways to organize production.

1. **Integrate**. Both stages of production are done within the final-good firm.
2. **Outsource**. Contract with an arm's-length firm (the supplier) to produce the components. The final-good firm uses the components to produce the final good.

The final-good firm will choose the firm structure that delivers the highest profit.

1.3.1 Best-case scenario

We begin by computing a benchmark best case scenario. Suppose the two firms could write complete and enforceable contracts. If so, the parties would like to maximize the joint profit of the venture. How the profit is split between the firms depends on the details of the contract, but we do not need to know the profit allocation for our purposes. What is important for us is to determine the quantity of components m that maximizes joint profits. This quantity is the solution to

$$\max_{m} \quad \pi_F + \pi_S = pA_i m^{\alpha} - p_m m. \tag{2}$$

In this problem, the components are made at the low cost p_m and the final-good firm does not incur the fixed cost of integration. To solve this problem we find the first-order condition. The first-order condition says that the derivative of the profit function (with respect to m) should be equal to zero. The first-order condition with respect to the choice of m is

$$\alpha p A_i m^{\alpha} - p_m = 0. \tag{3}$$

Solving this equation for m tells us that the supplier firm will make

$$\begin{aligned} \alpha p A_i m^{\alpha-1} - p_m &= 0 \\ \alpha p A_i m^{\alpha-1} &= p_m \\ m^* &= \left(\frac{\alpha p A_i}{p_m}\right)^{\frac{1}{1-\alpha}} \end{aligned} \tag{4}$$

units of components to use in producing the final good. We denote this quantity m^*.

Numerical example

Let $\alpha = 0.75$, $A = 2$, $p_m = 1.1$, $p = 1.5$. How many components are produced, m^*? What are the associated quantity of final goods q^*, revenues R^*, and joint profits $\pi_S^* + \pi_F^*$?

We find that components produced are $m^* = 17.5$*, final goods produced are* $q^* = 17.1$*, revenues are* $R^* = 25.7$*, and final-good firm profits are* $\pi_F^* + \pi_S^* = 6.4$.

This scenario sets the efficient benchmark, which we will use to measure the inefficiencies created by the incomplete contracts problem. In this scenario, the quantity of components is chosen so that the marginal revenue equals the marginal cost, which can be seen in (3) — this is the best we can do. The characteristics of the components, however, make this benchmark unattainable, as this contract cannot be written. We now turn to the firm structures that are available to the final-good firm.

1.3.2 Outsourcing with ex post bargaining

In this case, the final-good firm i and the supplier cannot write an enforceable contract. The timing of events:

1. The supplier chooses how much m to produce.
2. The final-good firm and supplier bargain over the revenue the components will generate.
3. The final good is made and sold at price p.
4. The revenue from selling the final good is split between the two firms according to the deal struck in step 2.

Since firms cannot write enforceable contracts, the bargaining over revenue happens after the production decision has been made.[1] How much revenue does each party get from the bargain?

[1] The Latin phrase *ex post facto* translates to "after the fact." The arrangement we are studying is called ex post bargaining because the bargaining happens after production has occurred.

You can take entire courses on the study of bargaining problems, but we will keep things simple. The share of the revenue that goes to firm i is $1-\beta$ and the share that goes to the supplier firm is β.[2] We can think of β as the bargaining power of the supplier. We do not model β, but take it as a model primitive. Notice that we are assuming that the bargain reached in stage 2 is enforceable. Presumably, it is easy enough to observe the revenue generated from selling the goods.

The supplier's problem is to choose m to maximize its profit. The supplier's profit is its share of the revenues βpq minus the costs it paid to create the components, $p_m m$. Notice the difference between this maximization problem and the one in the efficient benchmark (2). When choosing m, the supplier only considers the effect of its choice on its own profit. The supplier chooses m to solve

$$\max_{m} \quad \pi_S = \beta p A_i m^{\alpha} - p_m m. \tag{5}$$

We take the first-order condition and solve it for m,

$$\begin{aligned} \alpha\beta p A_i m^{\alpha-1} - p_m &= 0 \\ \alpha\beta p A_i m^{\alpha-1} &= p_m \\ m^B &= \left(\frac{\alpha\beta p A_i}{p_m}\right)^{\frac{1}{1-\alpha}}. \end{aligned} \tag{6}$$

We denote this outcome as m^B. We can pull the β out of the expression in (6) and see that

$$m^B = \beta^{\frac{1}{1-\alpha}}\left(\frac{\alpha p A_i}{p_m}\right)^{\frac{1}{1-\alpha}} = \beta^{\frac{1}{1-\alpha}} m^*. \tag{7}$$

Since $\beta^{\frac{1}{1-\alpha}} < 1$ the amount of m that the supplier produces is less than it produces when the two firms can commit to a level of production.

Numerical example

Continuing from the earlier example, what are m^B, q^B, R^B, π_F^B, and π_S^B when $\beta = 0.7$

We find that $m^B = 4.2$, $q^B = 5.9$, $R^B = 8.8$, $\pi_F^B = 2.6$, *and* $\pi_S^B = 1.5$. *Notice the strong under-provision of* m: $m^B/m^* = 0.24$. *Joint profit falls from* 6.4 *under the best-case scenario to* $2.6 + 1.5 = 4.2$.

When the two firms cannot write an enforceable contract, the incentives for the supplier to produce become distorted. The supplier understands that it will not receive the full value (the marginal revenue) of the components, so the supplier provides fewer components than in the best-case scenario. The severity of the under-provision depends on two parameters. The larger is the supplier's bargaining power β, the smaller is the inefficiency: If $\beta = 1$ there is no under-provision of components. The smaller is α, the smaller are deviations from the optimal component quantity.

[2] This is the outcome under an assumption about bargaining called *Nash Bargaining*. The underlying assumption that drives this revenue-splitting rule is that the outside option of each firm is zero. If the two firms cannot come to an agreement, then both get nothing: The component parts are not useful to other firms, and the final-good firm cannot operate its technology without these specific component parts. The solution would change a little if the goods could be resold to an outside party for less than they are worth to the final-good firm, but our main results would remain unchanged.

It is worth stressing the inefficiency created by bargaining. The issue is not the bargaining over revenues per se, but that the bargaining leads to an under-provision of components. This under-provision shrinks the total profit created: The firms are splitting a smaller pie. Notice in the numerical example above, that if firms could write a contract that supports m^* being produced, they could increase both firms' profits compared to what they earn when bargaining. Better contacting technology makes both firms better off.

1.3.3 Integration

If the final-good firm chooses to create the component parts in-house, it avoids the contracting and bargaining problems that led to under-provision of components from the supplier, but the final-good firm incurs extra costs to run its own component production line. These extra costs are modeled as a proportional increase in the cost of producing components relative to the supplier firm and a fixed cost of operating the production line. Components produced by the final-good firm cost γp_m per unit, where $\gamma > 1$. The fixed cost of component production is f^I.

If the final-good firm integrates, it chooses m to solve

$$\max_{m} \quad \pi_F = pA_i m^\alpha - \gamma p_m m - f^I. \tag{8}$$

The first-order condition with respect to the choice of m is

$$\alpha p A_i m^\alpha - \gamma p_m = 0. \tag{9}$$

Solving this equation for m yields the optimal component choice,

$$m^I = \left(\frac{\alpha p A_i}{\gamma p_m}\right)^{\frac{1}{1-\alpha}} = \left(\frac{1}{\gamma}\right)^{\frac{1}{1-\alpha}} m^*. \tag{10}$$

Again, the component choice is smaller than in the best-case scenario. In this case, though, the smaller m is not the result of the misaligned incentives of the final-good firm and the supplier, but because the final-good firm cannot produce components as cheaply as the supplier can. Notice that α plays the same role in the integrated firm as it does in the outsourcing case.

1.3.4 Optimal firm structure

The final-good firm compares the profits from integrating to outsourcing, and chooses the firm structure that delivers the largest profit. It chooses to integrate and produce parts in-house when

$$\pi_F^I \geq \pi_F^B. \tag{11}$$

Numerical example

Continuing from the earlier example, what are m^I, q^I, R^I, and π_F^I when $\gamma = 1.3$ and $f^I = 0.25$? Should the final-good firm integrate, or purchase from a supplier, despite the ex post bargaining?

We find that $m^I = 6.1$, $q^B = 7.8$, $R^B = 11.7$, $\pi_F^B = 2.7$. In this example, the firm should integrate. This firm structure maximizes profits.

1.4 When is outsourcing optimal?

In this section, we will study how profits — and the optimal firm structure — depend on the underlying parameters of the model. Do more efficient final-good firms outsource more often? How does firm structure depend on production costs? How does firm structure depend on bargaining power?

Bargaining power

The supplier firm's bargaining power (β) is not a choice variable. The firms bargaining power comes from outside the model. We can learn quite a bit about how the model works, however, by studying the ways the model outcomes change when β changes.

Figure 1: Component production.

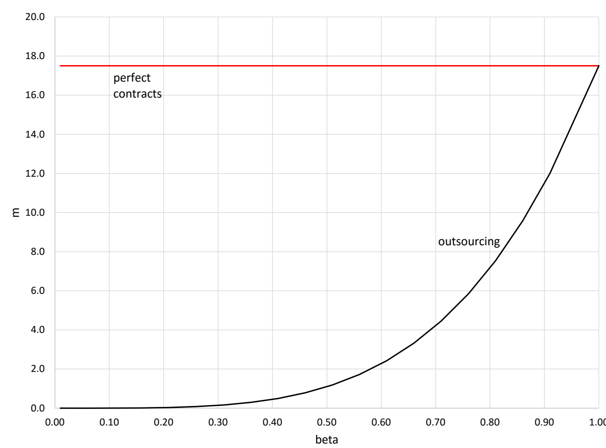

Figure 2: Joint profit.

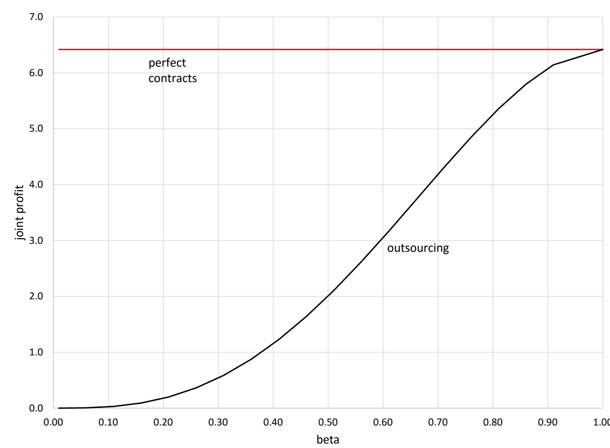

In Figure 1, we plot, as we change β, the production of components by the supplier firm in the outsourcing scenario. As β increases, the supplier firm's incentive problem diminishes. When $\beta = 1$ the supplier completely internalizes the marginal return from producing a unit of components, so the supplier produces m^*. In Figure 2, we plot the joint profit of the two firms for different values of β. As β increases, m^B gets closer to m^*, and the joint profit of the two firms gets closer to the joint profit in the best-case scenario.

Figure 3: Supplier firm profit.

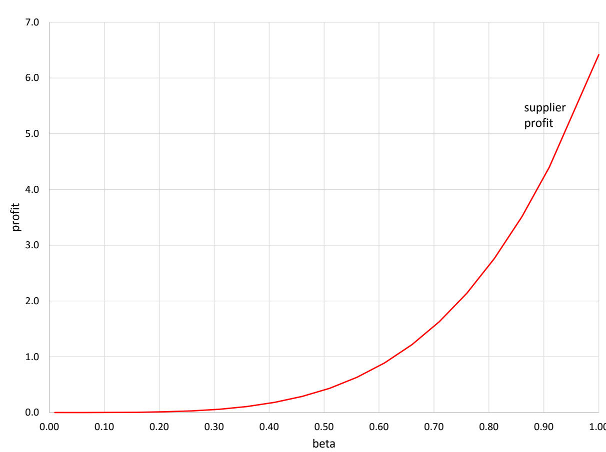

Figure 4: Final-good firm profit.

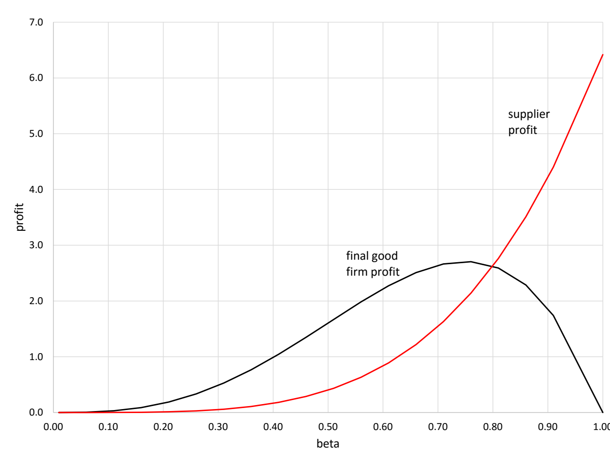

Notice that the joint profit function in Figure 2 is relatively flat for low values of β, steep for intermediate values of β, and flat again as β nears one. In Figure 3, we plot the profits for the supplier. The supplier's profit always increases as we increase β: As β increases, both the joint profit increases (because the supplier produces more m) and the revenue share of the supplier increases.

The final-good firm's profits, however, are not always increasing in β. For very low values of β, the final-good firm receives almost all of the revenues, but the revenues are very small — the low β severely distorts the supplier's decision and very little m is produced. For very large β, the supplier's decision is not very distorted and revenues are close to their ideal level — but, since β is close to one, the final-good firm receives almost none of the revenues. For intermediate values of β, the supplier's incentive problem is reduced and the final-good firm still receives a reasonable share of it. In our example, the final-good firm does best when $\beta \approx 0.75$.

When β is "too high" or "too low," the final-good firm's profits in the outsourcing scenario will be low, and the firm will be more likely to integrate.

Final-good firm productivity

If the final-good firm chooses to integrate, it must pay a fixed cost f^I to do so. When final-good firms differ in their productivities, this fixed cost induces selection among the firms. The idea here is the same as in the horizontal FDI models we studied earlier in the semester. Firms with better productivity (higher A_i) will earn more profits. The larger profits earned by more productive final-good firms make them able to pay for the fixed cost of integrating.[3] We plot the final-good firm profits for outsourcing and integrating for various A_i in Figure 5.

The prediction is that — all else equal — more productive firms are more likely to integrate production.

Figure 5: Profits and final-good firm productivity.

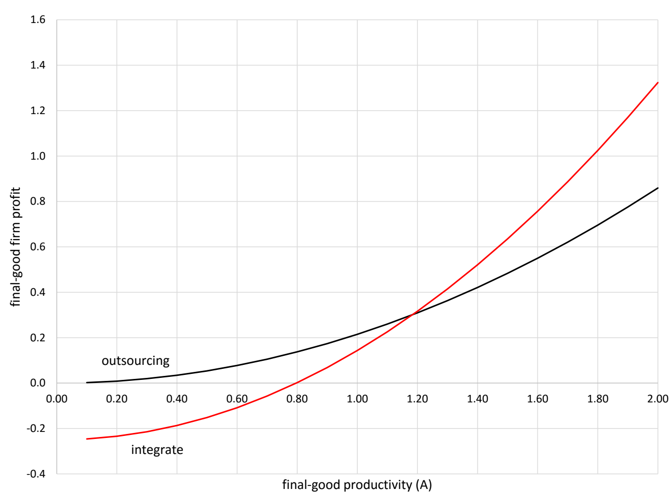

[3] We have simplified our model by assuming that there is no fixed cost to pay when the final-good firm outsources. Everything works the same in the model if we introduce a fixed cost for outsourcing, as long as it is smaller than the fixed cost of producing components in the final-good firm.