Chapter 4

Contracts and Global Sourcing

In his highly entertaining book *Poorly Made in China*, Paul Midler (2009) describes his experiences as an offshoring consultant in China, where his command of Chinese made him a particularly valuable asset for American companies eager to outsource in that country. The bulk of Midler's book is centered around his work in assisting the chief operating officer of Johnson Carter, a U.S. personal care import company, in his negotiations with King Chemical, a Chinese manufacturer located outside Guangzhou (the names in the book are fictitious but the stories are genuine). Johnson Carter is in the business of supplying large chain stores in the U.S. with house-brand shampoo and soaps.

Midler's book vividly illustrates just how irrelevant formal outsourcing contracts can prove to be in China, reflecting the old Chinese adage that 'signing a contract is just the first step in the real negotiations'. Chinese manufacturers – and King Chemical, in particular – appear to attract clients by offering very low prices in order to secure contracts. After locking in U.S. buyers – in this case, Johnson Carter –, they work to devise numerous ways to increase their profit margins. These maneuvers often involve reductions of material costs, last-minute price increases, or even the use of the client's designs to sell the same products to alternative buyers at higher prices. In one particular instance, and after a contract order had been signed, King Chemical sought to reduce the money they spent on plastic by secretly switching to thinner bottles which were much more likely to collapse when squeezed. In the case of another order, Midler discovered that King Chemical was filling the bottles with significantly less soap than the 850 milliliters they had agreed on. Arbitrary price increases were particularly frustrating to Johnson

Carter, as King Chemical's price hikes were typically poorly justified and were invariably demanded right after the U.S. company had secured a large order from a U.S. retailer, at which point the U.S. company had no time to turn to an alternative Chinese manufacturer. This last-minute price spikes were hardly specific to King Chemical. As Midler writes," 'Price go up!' was the resounding chorus heard across the manufacturing sector" in China.

The goal of this chapter is to initiate the analysis of the implications of weak contract enforcement for the international organization of production. With that goal in mind, I will start by going back to the benchmark two-country model of global sourcing developed in Chapter 2, and I will highlight the contracting assumptions underlying the results of that model. I will then introduce a model of contractual frictions that will shed light on the effect of contracting institutions on the intensive and extensive margins of intermediate input trade. Next, I will develop a series of extensions of the model that will open the door for model-based empirical tests of the effects of contracting considerations on the offshoring decisions of firms. In Chapter 5, I will present suggestive evidence based on U.S. import data of the empirical relevance of the concepts highlighted in this chapter.

A Brief Recap of the Global Sourcing Model

Let us now go back to the model of global sourcing with heterogeneous firms introduced in Chapter 2. To build intuition, I will initially focus on the simple two-stage, two-country offshoring model inspired by the work of Antràs and Helpman (2004). Before transitioning to the empirics, it will prove useful to extend the model to a multi-country environment as outlined towards the end of Chapter 2, but I defer any discussion of multi-country environments until Chapter 5.

I will not review the benchmark global sourcing model in great detail, but it may be useful to remind the reader a few of its key features. We are largely concerned with the behavior of firms in a differentiated good sector in which final goods are produced by combining two stages, headquarter services and manufacturing production, under a Cobb-Douglas technology. Firms differ in their productivity level, which as in Melitz (2003), is only revealed upon paying a fixed cost of entry. There are two countries, the North and the South. The North has comparative advantage in entry and headquarter services, so these stages always occur there. Conversely, the South has comparative advantage in manufacturing production, and absent any fixed costs of offshoring, all firms would want to fragment production and combine Northern headquarter services and Southern manufacturing production. Offshoring is however costly and entails a fixed cost of f_O units of Northern labor, while Northern production of manufacturing goods entails a lower fixed cost equal to $f_D < f_O$ units of Northern labor.

Overall, the total cost of production associated with Domestic sourcing in the North and Offshoring to South are given, respectively, by

$$C_D(q,\varphi) = \left(f_D + \frac{q}{\varphi}\right) w_N,$$

and

$$C_O(q,\varphi) = f_O w_N + \frac{q}{\varphi}(w_N)^{\eta}(\tau w_S)^{1-\eta},$$

where φ is a firm-specific productivity parameter throughout the book, w_i denotes the wage rate in country $i = N, S$, τ reflects variable costs of offshoring (input trade costs, communications costs,...), and η represents the headquarter intensity of production.

Assuming further, and for simplicity, that trade in final goods is free, we then used the demand equation (2.3) to derive the operating profits associated with each of the two sourcing strategies available to firms in the North. In particular, a Northern firm with productivity φ would obtain a profit flow equal to

$$\pi_D(\varphi) = (w_N)^{1-\sigma} B\varphi^{\sigma-1} - f_D w_N \tag{4.1}$$

when sourcing domestically, and a profit flow equal to

$$\pi_O(\varphi) = \left((w_N)^{\eta}(\tau w_S)^{1-\eta}\right)^{1-\sigma} B\varphi^{\sigma-1} - f_O w_N, \tag{4.2}$$

when offshoring the manufacturing stage of production to the South. In these equations, the demand shifter B is given by

$$B = \frac{1}{\sigma}\left(\frac{\sigma}{(\sigma-1)P}\right)^{1-\sigma}\beta(w_N L_N + w_S L_S).$$

As in our discussion of the export profit functions in the Melitz (2003) model, the realization of the entire profit flow (4.2) by the Northern final-good producer rests on strong assumptions, including having information on all parameters of the model, and frictionless contracting between producers (and between producers and workers). Below, I will maintain the assumption that all the parameters of the model are deterministic and common knowledge to all producers, and that labor markets work competitively and efficiently. This will focus our attention on the implications of weak contracting enforcement between Northern headquarters and Southern manufacturing plants. A discussion of contracting issues requires a more detailed discussion of the agents involved in production and the timing of events, a task to which I turn next.

Microeconomic Structure and Contracting

We shall begin by assuming that there are only two agents relevant for contracting in the production of each final-good differentiated variety. On the one hand, there is the final-good producer – agent F – which is in charge of incurring all fixed costs of production (including entry), and who also controls the provision of headquarter services. On the other hand, manufacturing production is controlled by a manager – agent M – in that production facility. Both agents have an outside opportunity that delivers them an income level which for simplicity I normalize to 0.

I next describe the timing of events, which is also illustrated in Figure 4.1. At some initial date t_0, the final-good producer F incurs the fixed cost of entry $w_N f_e$, upon which the productivity level φ is revealed, and F decides whether to have the manufacturing stage of production controlled by a Northern or a Southern manager. At the end of this same period, F approaches a manager M in the chosen location and offers him or her a formal sourcing contract (details below). This initial stage is followed by an investment stage t_1, at which F produces headquarter services and M undertakes manufacturing production. I assume for now that these investments occur simultaneously, but we will contemplate models with sequential production below. Once the investments have been incurred and before the manager M hands over the manufactured goods to the final-good producer F, we shall consider the possibility that the terms in the initial contract are renegotiation and bargained over at stage t_2. Finally, the terms of this renegotiation (or of the initial contract in the absence of renegotiation) are executed at a final stage t_3, when the final good is also produced and sold.

I will first illustrate that a seemingly simple initial contract might suffice for F to be able to attain the 'frictionless' levels of operating profits in (4.1) and (4.2). To fix ideas, let us consider the case in which, at t_0, F has decided to approach a Southern manufacturer (we will later study the loca-

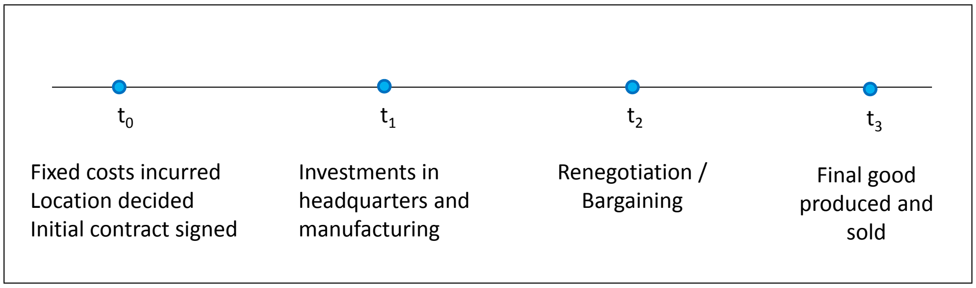

Figure 4.1: Timing of Events

tion choice). Suppose then that in the initial contract, F offers the Southern manager a contract that stipulates a quantity m^c of manufacturing production to be provided by M in exchange for a fee s^c received by M. Assume also that the initial contract includes a clause such that if any party deviates from this initial contract, the other party is entitled to arbitrarily large damages. With that clause, F can be assured that any level of manufacturing services stipulated in the initial contract will be honored by the manager. As a result, F can safely choose $h^c(\varphi)$, $m^c(\varphi)$ and $s^c(\varphi)$ to solve the following problem, where the constraint reflects the participation constraint of the Southern manager:

$$\begin{aligned} \max_{h(\varphi),m(\varphi),s(\varphi)} \quad & p(q(\varphi))\,q(\varphi) - w_N h(\varphi) - w_N f_O - s(\varphi) \\ \text{s.t.} \quad & s(\varphi) - \tau w_S m(\varphi) \geq 0 \end{aligned} \tag{4.3}$$

Naturally, F will set $s^c(\varphi)$ in the initial contract to make the participation constraint of the manager M exactly bind, and as a result, F will choose $h^c(\varphi)$ and $m^c(\varphi)$ to maximize the overall surplus from the relationship, thus resulting in profit-maximizing investments and a profit flow identical to that in equation (4.2) above. Note that the timing of events, investments or payments is quite irrelevant for this result.

Incomplete Contracts and Weak Contract Enforcement

As simple as the above contract may seem, it is hard to imagine that agents will in fact be able to (1) write such a type of contract, and (2) find a court of law to enforce it. More specifically, note that the initial contract needs to specify the level of m, which corresponds to the value or services obtained from the manufacturing stage of production. It does not merely reflect, in particular, the number of physical units of manufacturing goods. This distinction is crucial and carries important consequences. Consider our above anecdotal evidence from Midler's book *Poorly Made in China*. Johnson Carter presumably signed a contract with King Chemical for the provision of a certain number of bottles of soap in exchange for a certain amount of money (or price per bottle). It is hard to imagine, however, that the formal contract specified a variety of characteristics of the product which were clearly relevant for how the combination of Johnson Carter's 'headquarter services' and King Chemical's manufactured goods was to translate into sale revenues in U.S. retail chains. For instance, the contract almost surely did not indicate the plastic content of the bottles or the chemical composition of the soap. Flimsy bottles or abrasive soap would each lead to a low value of m in the model, no matter how large the number of bottles actually manufactured.

The general lesson here is that while certain aspects of manufacturing, such as the number of units of goods to be produced, the price per unit, or the date of the delivery are relatively easy to incorporate into a formal written contract, many other aspects of production are not. And, crucially, these noncontractible elements of production are often key in shaping the quality of goods or their compatibility with other parts of the production process.

Beyond the obvious incompleteness of real-life commercial contracts, there still remains the issue of contract enforcement. Even if two expert parties were able to design a highly comprehensive contract specifying what each party is supposed to contribute to production and making explicitly what defines high or low quality or compatibility, it is still questionable that a court of law will be able to understand such a contract and enforce it properly. And even when it does, in international transactions there still remains the uncertainty over the cross-border enforcement of damages (see Chapters 1 and 3 for a discussion of this issue). A natural response of the contracting parties is thus to shy away from specifying too harsh penalties for deviations from highly detailed clauses in written contracts. In Midler's soap bottles examples, Johnson Carter could have presumably insisted on the initial contract specifying the plastic content of bottles or that each bottle contained exactly 850 milliliters of soap, but it is somewhat hard to believe that a court of law would be able to verify whether that contract had been honored or not.

In sum, as simple as it might have seemed, the type of contract that we have considered so far is in fact too complicated to realistically discipline the behavior of agents. If such a contract is not feasible or enforceable, then what type of contracts are? One could envision the possibility that the initial contract would at least specify all the characteristics of the contractible aspects of production, while stipulating large penalties for deviations from that contract, thus ruling out the possibility of any ex-post renegotiation. It is not hard to see, however, that these type of contracts will typically deliver unappealing outcomes. For instance, imagine that the initial contract were to stipulate the number of physical units of m to be traded as well as their price in a binding manner, and without allowing M or F to renegotiate those terms upon observing the quality of the manufactured goods. In such a case, M would have every incentive to produce the inputs m in the least-cost possible manner, which would typically result in a low quality level of those goods. Foreseeing this debasements in quality, F would not be willing to offer a particularly large price for those goods in the initial contract, and the overall surplus of such a contractual relationship would end up being quite low and possibly zero, if M were to produce useless goods. In these circumstances, F might be better off by offering a less complete contract at the initial stage, as we will show below.

Another possibility would be for the initial contract to specify a simple sharing rule for the sale revenue obtained by F when selling the final good to consumers. Even if the initial contract did not specify the quality characteristics of the manufactured inputs in a binding manner, it seems natural that the incentives of M to skim on quality will be attenuated when his or her profits are a function of the the willingness of consumers to pay for goods that embody those manufactured inputs. Whenever certain aspects of the investments in h and m remain noncontractible, however, these type of revenue-sharing contracts will not lead to the frictionless levels of investment and profits of the complete-contracting environment, a result reminiscent of Holmstrom's (1982) moral-hazard-in-teams problem. As a result, even when revenue-sharing arrangements are feasible and enforceable, we will see that in some cases it may well be in the interest of the parties to opt out of them in the initial contract. Needless to say, the appeal of these contracts is further diminished whenever they are not perfectly enforceable, perhaps due to manipulation by the agent collecting revenues, who may be tempted to underreport them. Gennaioli (2013) has studied optimal contracting in a model with potentially biased judges, and has shown that the enforcement risk generated by these biases, often leads parties to write simple contracts that are not contingent on revenues.

Borrowing tools from the mechanism design literature, an important body of theoretical work has proposed a variety of ingenious 'mechanisms' to restore efficiency in environments in which the contracting parties (i) have symmetric information, (ii) can commit not to renegotiate an initial contract, and (iii) can resort to a third party (presumably a court of law or arbitrator) to enforce off-the-equilibrium-path penalties (see, for instance, Aghion, Dewatripont and Rey, 1994, or Maskin and Tirole, 1999). This literature is often criticized for suggesting somewhat convoluted contracts that are not observed in the real world. I find that criticism unconvincing: after all, one could have similarly criticized some key contributions to the auction theory literature, and yet they have subsequently had an enormous impact in real-world auctions (see, for instance, Milgrom, 2000). My main reservation with mechanism-design resolutions to incomplete contracting is that they rely on the ability of a third party to enforce contracts, and as argued repeatedly in this book, this is a real sticking point when studying the international organization of production.

'Totally Incomplete' Contracts

Given the discussion above and for pedagogical reasons, I will begin by considering environments in which, when contracts are incomplete, they are so in a rather extreme way. With that in mind, consider the following definition:

Definition *A contract is said to be* 'totally incomplete' *whenever no aspect of the contract is perceived to be enforceable, with the possible exception of a lump-sum transfer exchanged at the time of the agreement.*

For reasons discussed in Chapter 1, it seems natural to assume that certain contracts that are feasible or enforceable in domestic transactions might not be feasible or enforceable in international transactions. To illustrate the implications of this asymmetry I will this consider first the case in which contracting is complete or perfect in the domestic sourcing relationships, while contracting is *totally incomplete* in offshoring relationships. Later in the chapter, I will consider environments with partial contractibility in each of these two sourcing options.

With complete contracting in domestic sourcing, we have demonstrated above that F will be able to design a contract such that the levels headquarter services h and manufacturing production m are set at their joint-profit maximizing level, thus resulting in the frictionless profit flow $\pi_D(\varphi)$ in equation (4.2).

Let us next consider the more interesting implications of incomplete contracting in offshoring relationships. What happens when the initial contract does not stipulate the levels of h or m nor a payment to be paid to the manufacturer M contingent on the volume of m produced or contingent on sale revenues? In that case, the only option left for the parties is to decide on the terms of exchange at t_2 (remember the timing of events in Figure 4.1). The next question is then: how should one model this bargaining/contracting stage? Given that we have assumed above that the final good producer makes a take-it-or-leave it offer to the manager at t_0, it might seem natural to maintain that assumption for stage t_3. This would however ignore what Oliver Williamson famously termed the "fundamental transformation" (see, for instance, Williamson, 1985). This transformation refers to the fact that even though the final good producer might have chosen a particular manager M from a competitive fringe of managers, once the investments h and m have been incurred, a contractual separation is likely to prove costly to both parties. To the extent that parties feel 'locked-in' with each other, the initial competitive environment at t_0 has thus been fundamentally transformed into a bilateral monopoly scenario.

As explained in Chapter 1, in the global sourcing environments that we are considering in this book, there are several natural sources of *lock in* between final-good producers and suppliers. First, manufacturing inputs are often customized to their intended buyers and cannot easily be resold at full price to alternative buyers. Second, certain types of headquarter services are also designed with particular suppliers in mind, and it would prove costly to reuse these services with alternative suppliers. Third, even though we have abstracted from modeling them in the initial period, search frictions are particularly relevant in international environments and they are likely make ex-post separations particularly costly for both final-good producers and suppliers, who would at the very least suffer from delays in obtaining a return on their investments.

In sum, in the presence of lock-in effects, incomplete contracting leads to a situation of bilateral monopoly in which the terms of exchange between F and M will only be determined ex-post (at t_2), after these agents have incurred investments that are by then sunk and have a relatively lower value outside that particular business relationship. The combination of incomplete contracting and lock-in effects, leads to what is often referred to as a *hold-up* *problem*, which in our particular context is two-sided. More specifically, on one end, F will try to push down the price paid for M's manufacturing input, realizing that M might be inclined to accept a reduced price due to the lower value of those inputs for alternative buyers. At the same time, however, M will try to raise the price of m as much as possible, knowing that it might also be in F's best interest to accept a relatively high price if that avoids having to search for a new supplier.

In deciding how assertively to bargain, each party takes into account that a too aggressive offer might lead the other party to refuse to trade, an outcome that is not appealing given the lower value of the sunk investments outside the relationship. As a result, even when bargaining is efficient and trade takes place in equilibrium, the *possibility* of a disagreement and associated failure to trade implies that F and M will tend to have lower incentives to invest in h and m than in the complete contracting case, in which case the initial contract ensures that trade will occur. In more technical terms, with incomplete contracting, the payoff obtained by each party in the ex-post negotiations will put a positive weight on *off-the-equilibrium path* situations in which the return to each party's investments is lower than on the equilibrium path.

In the literature, it is common to characterize the ex-post bargaining at t_2 using the Nash Bargaining solution and assuming symmetric information between F and M with regards to all parameters of the model. In such a case, each party ends up with a payoff equal to the value of their outside option (their payoff under no trade) plus a share of the ex-post gains from trade, which correspond to the difference between the sum of the agent's payoffs under trade and their sum under no trade. For the time being, I will assume that the outside option for each party is equal to 0. In other words, I am assuming that the manufactured input m is fully specialized to F and thus useless to other producers, while headquarters h are also fully tailored to the supplier M and could not be productively combined with inputs provided by other manufacturers. I will also consider the case of *symmetric* Nash bargaining, which implies that F and M share *equally* the ex-post gains, which with zero outside options equal sale revenues. Obviously, these are restrictive assumptions, but I will consider more general environments below.

In sum, with symmetric Nash bargaining, each party will anticipate obtaining a payoff equal to one-half of sale revenue at t_2, and thus the levels of

$h(\varphi)$ and $m(\varphi)$ will be set at t_1 to solve

$$\max_{h} \quad \frac{1}{2} p(q(\varphi)) q(\varphi) - w_N h \tag{4.4}$$

and

$$\max_{m} \quad \frac{1}{2} p(q(\varphi)) q(\varphi) - \tau w_S m, \tag{4.5}$$

respectively, where remember that

$$q(\varphi) = \varphi \left(\frac{h}{\eta}\right)^{\eta} \left(\frac{m}{1-\eta}\right)^{1-\eta} \tag{4.6}$$

and $p(q(\varphi)) = B^{1/\sigma} \sigma (\sigma - 1)^{(\sigma-1)/\sigma} q(\varphi)^{-1/\sigma}$.[1]

For comparability with the complete-contracting case, I will assume again that at t_0 there is a competitive fringe of potential managers M willing to work for each final-good producer F at a reservation wage equal to 0. Each F then decides the terms of the initial contract and makes a take-it-or-leave it offer to one of those managers. Because the initial contract is allowed to include a lump-sum transfer between parties, F can set the transfer such that the participation constraint of M exactly binds. So, as with complete contracts, F ends up with a profit level equal to

$$\pi_O(\varphi) = p(q(\varphi)) q(\varphi) - w_N h(\varphi) - \tau w_S m(\varphi) - w_N f_O. \tag{4.7}$$

The key difference, however, is that the levels of $h(\varphi)$ and $m(\varphi)$ can no longer be set in the initial contract in an enforceable manner. As a result, $h(\varphi)$ and $m(\varphi)$ are no longer set to maximize $\pi_O(\varphi)$ in (4.7), but instead are chosen simultaneously and non-cooperatively by F and M to solve programs (4.4) and (4.5), respectively.

In analogy to the way we studied contractual frictions in Chapter 3 and to our formulation of the complete-contracting program in (4.3), a compact way to represent the ex-ante problem faced by the final good producer under 'totally incomplete' contracts is:

[1] It should be noted that I am assuming that the agreement in t_2 is always enforced. Enforcement might become an issue if the parties exchange the goods at t_2 and the payment occurs at t_3, but not if the payment occurs at t_2, simultaneously with the exchange of goods.

$$\begin{aligned} \max_{h(\varphi),m(\varphi),s(\varphi)} \quad & \frac{1}{2}p\left(q\left(\varphi\right)\right)q\left(\varphi\right) - w_N h\left(\varphi\right) - w_N f_O - s\left(\varphi\right) \\ \text{s.t.} \quad & s\left(\varphi\right) + \frac{1}{2}p\left(q\left(\varphi\right)\right)q\left(\varphi\right) - \tau w_S m\left(\varphi\right) \geq 0 \\ & h\left(\varphi\right) = \arg\max_h \{\frac{1}{2}p\left(q\left(\varphi\right)\right)q\left(\varphi\right) - w_N h\left(\varphi\right)\} \\ & m\left(\varphi\right) = \arg\max_m \{\frac{1}{2}p\left(q\left(\varphi\right)\right)q\left(\varphi\right) - \tau w_S m\left(\varphi\right)\}. \end{aligned} \tag{4.8}$$

A simple comparison of (4.3) and (4.8) shows that the effect of incomplete contracting is captured by the addition of the last two constraints, which represent incentive compatibility constraints faced by both F and M. Because the bargaining payoffs of each agent put a positive weight (in this case, $1/2$) on an off-the-equilibrium-path zero return to producing in h and m, it is naturally the case that the hold up problem discussed above leads to inefficiently low investment levels at t_1, and consequently, results in depressed overall profits as well.

Solving formally program (4.8) delivers a volume of profits that F anticipates obtaining when choosing foreign sourcing at t_0 equal to

$$\pi_O = \left(\left(w_N\right)^{\eta}\left(\tau w_S\right)^{1-\eta}\right)^{1-\sigma} B\Gamma_O\varphi^{\sigma-1} - w_N f_O, \tag{4.9}$$

where

$$\Gamma_O = \left(\sigma+1\right)\left(\frac{1}{2}\right)^{\sigma} < 1 \quad \text{for } \sigma > 1. \tag{4.10}$$

Note that (4.9) is identical to the complete-contracting expression (4.2) except for the term Γ_O, which is necessarily lower than 1. Hence, this term reflects the loss of efficiency due to incomplete contracting. Furthermore, for the relevant range $\sigma > 1$, Γ_O is a decreasing function of σ, and thus in environments with tougher competition (i.e., lower markups), the profit losses from incomplete-contracting frictions are relatively larger (see the Theoretical Appendix for a formal proof). Intuitively, the last two constraints in (4.8) suggest that the effect of incomplete contracting on the choice of h and m is analogous to a doubling of the marginal cost of each of these stages, and this will tend to reduce profits more, the more price elastic demand is. This marginal cost inflation only becomes irrelevant in the limiting case $\sigma \to 1$, since in that case a firm's market is independent of its cost.

Choice of Location and Prevalence of Offshoring

Having computed the anticipated profits associated with domestic sourcing and offshoring, we can next study the choice of location of the final-good producer in the initial period t_0. Note from equations (4.1) and (4.9) that we can write these profits functions succinctly as

$$\pi_\ell(\varphi) = \psi_\ell B \varphi^{\sigma-1} - w_N f_\ell \quad \text{for } \ell = D, O,$$

with

$$\frac{\psi_D}{\psi_O} = \frac{1}{\Gamma_O}\left(\frac{w_N}{\tau w_S}\right)^{-(1-\eta)(\sigma-1)}. \tag{4.11}$$

As in the benchmark models reviewed in Chapter 2, the profit functions $\pi_D(\varphi)$ and $\pi_O(\varphi)$ are linearly increasing in the transformation of productivity $\varphi^{\sigma-1}$, with the relative slope of the two functions now being governed by the ratio ψ_D/ψ_O. Figure 4.2 depicts these functions for the case in which wage differences across countries are high so that $\psi_D < \psi_O$ (see the line π_O^h), and the case in which wage differences are sufficiently low to imply $\psi_D > \psi_O$ (see the line π_O^l). In the latter case, no firm finds it optimal to offshore in the South, while in the former case, only the most productive firms will. The key difference with the complete-contracting case is that, other things equal, firms will now find it less profitable to offshore to the South due to the term Γ_O in (4.11).

Let us now aggregate the decisions of the various firms in a sector and study the determinants of the relative importance of offshoring in different industries. As in Chapter 2 and with an eye to the empirical applications in Chapter 5, it seems natural to focus on the share of spending on *imported* manufacturing inputs over total manufacturing input purchases in a particular industry as a measure of the prevalence of offshoring. Because contracts are complete in domestic transactions, for all firms sourcing domestically input purchases constitute a share $(1-\eta)(\sigma-1)/\sigma$ of revenue and a multiple $(1-\eta)(\sigma-1)$ of operating profits (defined as revenue minus variable costs)
.

Matters are trickier in the case of foreign intermediate inputs, since one needs to take a stance on how these inputs are priced. One possibility is to assume that the headquarter H's spending on inputs corresponds to the payment obtained by the manufacturing manager M in the ex-post bargaining at t_2, which in the model above is simply $\frac{1}{2}$ of revenue and a share $\sigma/(\sigma+1)$

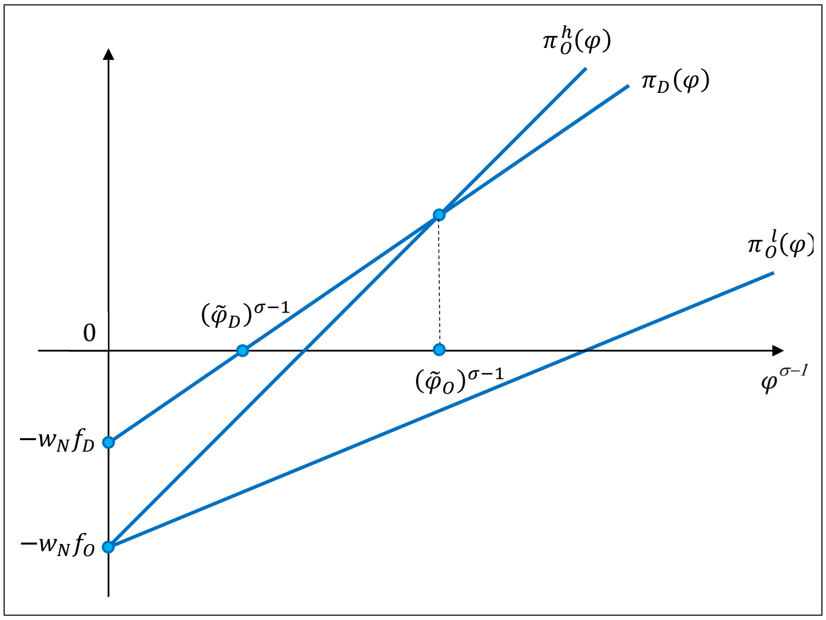

Figure 4.2: Equilibrium Sorting and Contractual Frictions

of operating profits. Alternatively, one could appeal to the existence of the ex-ante lump-sum transfer $s(\varphi)$ at t_0 to argue that H is *effectively* purchasing inputs at marginal cost and thus total spending on foreign inputs is instead given by $\tau w_S m(\varphi)$. From the first-order-condition of the program in the last constraint of program (4.8), this corresponds to a share $\frac{1}{2}(1-\eta)(\sigma-1)/\sigma$ of sale revenue or a share $(1-\eta)(\sigma-1)/(\sigma+1)$ of operating profits.

I do not want to take a strong stance on a particular pricing practice so I will instead take the more agnostic approach of assuming that foreign inputs are priced such that these input expenditures constitute the same multiple $(1-\eta)(\sigma-1)$ of operating profits as in the case of domestic input purchases. I do not believe that this assumption is crucial for the results derived below, but it will significantly simplify the derivations.

Assuming a Pareto distribution of productivity, we can then follow similar steps as in the derivation of equations (2.22) and (2.23) to solve for the share Υ_O of imported manufacturing input purchases in a given industry as

$$\Upsilon_O = \frac{\Gamma_O \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}{\left(\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D}\right)^{\kappa-(\sigma-1)} - 1 + \Gamma_O \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}. \tag{4.12}$$

where

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\Gamma_O \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)} . \tag{4.13}$$

These expressions are analogous to equations (2.22) and (2.23) in Chapter 2 except for the presence of the term $\Gamma_O < 1$.

As is clear from these equations, the prevalence of offshoring is naturally increasing in the term Γ_O, and is thus reduced by the presence of contractual frictions. The remaining comparative statics are analogous to those in the complete-contracting case. The share of imported inputs is increasing in wage differences (w_N/w_S) and productivity dispersion $(1/\kappa)$ and decreasing in (relative) fragmentation barriers $(f_O/f_D, \tau)$ and headquarter intensity (η). Conversely, the overall effect of the elasticity of substitution on the share Υ_O is now ambiguous. As in the complete-contracting case, a higher σ makes offshoring more prevalent by reducing the Pareto shape parameter $\kappa/(\sigma-1)$ for firm sales and by increasing the percentage gain from a reduction in marginal costs of labor from w_N down to τw_S. Nevertheless, as mentioned above in our discussion of equation (4.10), a larger σ now also aggravates the inefficiencies caused by incomplete contracting, hence reducing Υ_O.

Extensions of the Basic Model

So far, our analysis has illustrated that weak contract enforcement across countries will tend to reduce the profitability associated with firms engaging in global sourcing strategies and will thus lead to a larger reliance on domestic intermediate inputs. We have illustrated this insight with a highly stylized framework with many simplifying assumptions, and as a result the model has not delivered any particularly valuable empirical predictions other than the intuitive negative effect of weak contracting on offshoring. I will next turn to studying more general environments that relax some of the strong assumptions above. This will serve to verify the robustness of the key comparative statics emphasized so far, and also to develop a richer set of comparative statics that are better suited to guide empirical work on the contractual determinants of the global sourcing decisions of firms.

I will consider six basic extensions of the model: (i) a generalization of the bargaining process at t_2, (ii) the possibility of restrictions on ex-ante transfers at t_0, (iii) environments with partial contractibility at t_0, (iv) investments at

t_1 that are only partially relationship-specific, (v) global sourcing decisions with multiple suppliers, and (vi) sequential production in which the investments of different suppliers occurs at different points in time. In each of these cases, we will confirm that the basic insights obtained so far continue to apply to those more general environments, but we will obtain sharper predictions about the differential effect of contractual institutions on firm profitability across firms, sectors and countries.

To be as succinct as possible, I will focus on outlining how each of these extensions affects and enriches the determination of the term Γ_O in (4.9) capturing the inefficiencies associated with incomplete contracting in offshore relationships. I will also develop variants of the model in which contractual frictions reduce the profitability of domestic sourcing via an analogous term Γ_D which is also shaped by the parameters of the model in ways to be discussed below. In the next chapter, I will return to the aggregation of firms' decisions to illustrate how Γ_O and Γ_D jointly affect the share of imported manufacturing input purchases in a given industry as in equation (4.12). This will serve to motivate my discussion of empirical evidence based on U.S. import data.

Generalized Nash Bargaining and Revenue-Sharing Contracts

In our basic model above, we have assumed that F and M share the ex-post gains from trade equally at t_2. In some circumstances, it may make sense to assume that the primitive bargaining power of final-good producers might be higher (or perhaps lower) than that of supplying managers. The large literature on non-cooperative models of bargaining emanating from the seminal work of Rubinstein (1982), has uncovered several potential determinants of primitive bargaining power. It is well-known, for instance, that relatively impatient or risk averse agents will tend to have relatively low bargaining power, and the same will be true about agents for which a bargaining delay might be particularly costly for reasons other than impatience, such as credit constraints (see, for instance, Rubinstein, 1982, or Roth, 1985).

Rather than developing any of these microfoundations in great detail, let me just assume that the final-good producer F obtains a share β of the ex-post gains from trade, with the manager M obtaining the remaining share $1-\beta$. This is often referred to as the *generalized* Nash bargaining solution. With the maintained assumption that the investments incurred to produce h and m are fully relationship-specific, β and $1-\beta$ will also correspond to the shares of revenue obtained by F and M, respectively, in their ex-post negotiations at t_2. Below, we will develop variants of the model in which the division of revenue will be shaped by factors other than primitive bargaining power.

It should be clear by now that with complete contracting, the particular bargaining solution adopted to characterize the t_2 stage is irrelevant, and thus the profits associated with domestic sourcing are still given by (4.1). Conversely, when solving for the equilibrium associated with offshoring, one needs to replace the term 1/2 in (4.4) and the second constraint in (4.8) with β, and the term 1/2 in (4.5) and the third constraint in (4.8) with $1-\beta$. Naturally, the larger is β, the lower will tend to be the underinvestment in the provision of headquarter services, but the larger will be the underinvestment in the provision of manufacturing services.

Solving for the equilibrium profits obtained by F under generalized Nash bargaining, delivers a profit flow identical to that in equation (4.9), but with

$$\Gamma_O = \Gamma_\beta \equiv \left(\sigma - (\sigma-1)\left(\beta\eta + (1-\beta)(1-\eta)\right)\right)\left(\beta^\eta (1-\beta)^{1-\eta}\right)^{\sigma-1}. \quad (4.14)$$

It is straightforward (though somewhat tedious) to verify that, regardless of the value of the primitive bargaining power β, incomplete contracting still necessarily reduces the profitability of offshoring, i.e., $\Gamma_\beta < 1$ for $\sigma > 1$.[2] Furthermore, in the Technical Appendix I show that the size of these contractual distortions continues to be increasing in the elasticity of demand (or $\partial\Gamma_\beta/\partial\sigma < 0$). The main novelty in (4.14) is that the level of contractual frictions is no longer only shaped by the elasticity of substitution σ, but now also depends on headquarter intensity η and on the bargaining power parameter β. The effects of these parameters on the level of Γ_β are non-monotonic and interact closely with each other. More specifically, it can be shown that Γ_β is decreasing in η when $\beta < 1/2$, while it is increasing in η when $\beta > 1/2$.[3] Notice also that when $\eta \to 0$, $\Gamma_\beta \to (\sigma - (\sigma-1)(1-\beta))\left((1-\beta)\right)^{\sigma-1}$ and thus Γ_β is a decreasing function of β, while when $\eta \to 1$, $\Gamma_\beta \to (\sigma - (\sigma-1)\beta)\beta^{\sigma-1}$, and Γ_β is instead an increasing function of β. In other words, whether increases or decreases in the bargaining power of the final good producer increase or

[2] This is a special case of the proof of Proposition 1 in Antràs and Helpman (2008). The result can also be proven more directly by noting that Γ_β was derived from an optimization problem that is more constrained than the one that delivered the profit flow in (4.2), and the latter profit flow can be obtained from (4.9) by setting $\Gamma_\beta = 1$.

[3] The proof is straightforward but cumbersome, so I relegate it to the Theoretical Appendix.

decrease profits also crucially depends on the level of headquarter intensity. Intuitively, and as argued above, when β increases, the underinvestment in headquarter services is alleviated, while the underinvestment in manufacturing production is aggravated. Whether the net effect is positive or negative naturally depends on the intensity with which these two stages are combined in production. In fact, straightforward calculations show that there exists a unique value β^* that maximizes Γ_β, and it satisfies

$$\frac{\beta^*}{1-\beta^*} = \sqrt{\frac{\eta}{1-\eta}\frac{\sigma-(\sigma-1)(1-\eta)}{\sigma-(\sigma-1)\eta}}. \tag{4.15}$$

In line with our intuition above, this profit-maximizing level of β^* is an increasing function of headquarter intensity η.

This result also shed lights on the implications of the model when allowing for revenue-sharing contracts to be signed at t_0. In particular, imagine a situation in which the ex-ante contract was not 'totally incomplete' but rather was allowed to include a division rule contingent on the volume of revenue generated at t_3. Denoting revenue by $R = p(q)q$ and the sharing rule by $\beta(R)$, the optimal initial contract would now solve (I am omitting the argument φ in all functions for simplicity):

$$\begin{aligned}
\max_{h,m,s,\beta(R)} \quad & \beta(R)R - w_N h - w_N f_O - s \\
\text{s.t.} \quad & s + (1-\beta(R))R - \tau w_S m \geq 0 \\
& h = \arg\max_h\{\beta(R)R - w_N h\} \\
& m = \arg\max_m\{(1-\beta(R))R - \tau w_S m\}.
\end{aligned}$$

If one restricts attention to linear sharing rules where $\beta(R)$ is independent of R, then our above discussion indicates that the optimal contract will set $\beta(R) = \beta^*$, where β^* is given in (4.15). Even with this more complete initial contract, the frictionless profit flow in (4.2) cannot possibly be attained because remember that Γ_β in (4.14) is less than 1 for any $\beta \in (0,1)$. Importantly, this conclusion is not specific to the case of linear sharing rules. As shown by Holmstrom (1982), for general sharing rules $\beta(R)$ satisfying budget-balance, the resulting investment levels h and m will continue to differ from the efficient ones, and the equilibrium profits associated with offshoring will necessarily fall short of those under complete contracts.

Limitations on Ex-Ante Transfers: Financial Constraints

So far, I have assumed that F and M are allowed to freely exchange lump-sum transfers when signing the initial contract at t_0. This transfer can be inferred from the participation constraint of the manager M, which implies

$$s(\varphi) = \tau w_S m(\varphi) - \frac{1}{2} p(q(\varphi)) q(\varphi).$$

Plugging the equilibrium values of $m(\varphi)$ and $q(\varphi)$ it is straightforward to show that $s(\varphi) \leq 0$ and thus the optimal contract calls for the manager M to post a bond in order to be able to transact with the final good producer F. In practice, it is not obvious that supplying firms will be willing or able to make that initial transfer. This is due to at least two reasons. First, managers might worry that the final good producer will disappear after period t_0 without incurring the fixed cost of offshoring or investing in headquarter services.[4] Second, depending on the financial environment in the manager's country, it may be hard for M to raise from financiers the full amount of cash $s(\varphi)$ stipulated in the contract we have considered so far. We next explore the implications of the existence of constraints on these ex-ante transfers.

To fix ideas, consider the case in which M can pledge to external financiers in his domestic economy at most a share ϕ of the net income it receives from transacting with F. I will not specify the source of these financial frictions, but they could stem from a limited commitment friction on the part of M along the lines of the models we have explored in Chapter 3. The equilibrium under financial constraints can again be reduced to a program analogous to that in (4.8), but with the additional constraint

$$-s(\varphi) \leq \phi \left[\frac{1}{2} p(q(\varphi)) q(\varphi) - \tau w_S m(\varphi) \right].$$

This constraint is tighter than the original participation constraint in (4.8) and will bind in equilibrium. Solving the program, we then find that F again obtains a payoff equal to that in equation (4.9), but now with

$$\Gamma_O = \Gamma_\phi \equiv (\sigma + \phi - (\sigma - 1)(1 - \phi)\eta) \left(\frac{1}{2} \right)^{\sigma}. \tag{4.16}$$

[4] It can be shown, however, that if the manager insists that the fixed cost of offshoring be incurred prior to the payment $s(\varphi)$, the final-good producer would no longer have an incentive to abscond with the transfer.

It is straightforward to see that Γ_ϕ is increasing in ϕ and attains the symmetric Nash bargaining value of $(\sigma + 1)(1/2)^\sigma$ when financial constraints disappear, i.e., $\phi = 1$. A first implication is thus that the profitability of offshoring will be increasing in the quality of financial contracting in the manager's foreign country, as summarized by the parameter ϕ. Intuitively, offshoring is now not only associated with distorted investments, but it also entails a loss of rents for the final good producer F. It is interesting to note that Γ_ϕ is also decreasing in headquarter intensity. This is due to the fact that headquarter services are complementary to manufacturing production, and the higher is η, the larger are the rents that M obtains in the ex-post bargaining at t_2 relative to the costs of production her or she incurs at t_1. Consequently, the larger is η, the larger is the loss of rents for F associated with a low ϕ. For the same reason, we observe in (4.16) that an increase in improvement in the quality of financial contracting will have a differentially large positive effect on the profitability of offshoring in production processes with high headquarter intensity (i.e., $\partial (\partial \Gamma_\phi / \partial \phi) / \partial \eta > 0$).[5]

Partial Contractibility

It is obviously unrealistic to assume, as we have done so far, that contracts in international transactions are 'totally incomplete'. It seems natural that some aspects of production can be specified in a contract in a manner that contracting parties feel confident that those aspects of the contract will be enforced. Moreover, it is also unrealistic to assume, as we have done so far, that contracts in domestic transactions are complete. Surely some aspects of production are nonverifiable or certain contracts are perceived to be hard to enforce in domestic transactions. I next incorporate partial contractibility into our global sourcing model following the approach in Antràs and Helpman (2008).

It will prove useful to assume that the production of headquarter services and manufacturing inputs now entails a continuum of processes or activities all of them carried out at t_1. A fraction of these processes is assumed to be ex-ante contractible, in the sense that contracts specifying how those processes should be carried out can be designed in a way that a court of law can verify their fulfillment and penalize any deviation from what was stipulated in

[5] It is also worth pointing out that, under plausible parametric restrictions, in this extension of the model too, higher demand elasticities are associated with larger incomplete-contracting distortions (see the Technical Appendix for details).

the contract. Conversely, the complementary fraction of processes is non-contractible and contracts specifying these activities would fail to discipline their provision at t_1.

There are two main determinants of the degree to which the overall production process is contractible or not. First, a low fraction of contractible activities could reflect technological factors that make it particularly hard to write down enforceable contracts disciplining the behavior of the agents engaged in production. For instance, the production of new and high-tech goods is more contractually demanding than that of more traditional and standardized goods. Second, even focusing on the same production process, it seems reasonable to assume that the fraction of contractible activities will vary across countries reflecting international variation in the quality of contracting institutions. In other words, certain types of contracts are perceived to be enforceable in some environments but perhaps not in others.

To further illustrate these two determinants, it may be useful to consider the following analogy to refereeing in (European) football. There are some rules in football that are almost trivial to enforce, such as ensuring that each team has no more than eleven players on the pitch or preventing a team from performing more than three substitutions. On the other hand, other rules are much trickier to enforce, such as calling a close offside infraction or deciding whether the entire ball crossed the goal line in a ghost-goal situation. It is thus not surprising that players anticipate the former rules to be properly enforced with a much higher likelihood than the latter ones. At the same time, the quality of the referees is also obviously critical in predicting whether the rules will be correctly applied or not. Skilled referees (or linesmen) make fewer mistakes in enforcing these rules than incompetent ones.

In order to formally introduce partial contractibility into the framework, we now let headquarter services h and manufacturing production m be a Cobb-Douglas aggregate of the services of a continuum of measure one of activities, so

$$h = \exp\left[\int_0^1 \log h\left(i\right) di\right] \tag{4.17}$$

and

$$m = \exp\left[\int_0^1 \log m\left(i\right) di\right]. \tag{4.18}$$

Our key new assumption is that activities related to input $k = h, m$ in the range $\left[0, \mu_{kj}\right]$ (with $0 \leq \mu_{kj} \leq 1$) are contractible in country $j = N, S$, in the sense that the characteristics of these activities can be fully specified in advance in an enforceable ex-ante contract involving a manufacturer M from country j. Hence, the initial contract is no longer 'totally incomplete' because in addition to a lump-sum transfer between F and M, it also specifies the level of contractible activities to be carried out at t_1. The remaining activities in the range $\left(\mu_{kj}, 1\right]$ continue to be noncontractible as in our benchmark model and F and M decide on the terms of exchange for those activities only after they have been produced. Because the initial contract does not compel any of the two parties to provide a positive amount of these noncontractible tasks, the threat point for each party in the negotiations at t_2 is to withhold the services from those activities, which in light of the Cobb-Douglas production technologies (4.6), (4.17), and (4.18), and our maintained assumption that all investments are fully relationship-specific, would lead to a zero payoff for both parties.[6] Thus each agent ends up capturing a constant share of sale revenues, and for simplicity, I will now revert back to the assumption of symmetric Nash bargaining, so that the two parties end up sharing evenly total sale revenues.

The symmetry assumptions on technology built into (4.17) and (4.18) allows us to simplify the problem of the firm conditional on having selected a location $j = N, S$, to the choice of an ex-ante transfer s, a common value h_c for all contractible headquarter activities, a common value h_n for all non-contractible headquarter services, and analogous values m_c and m_n for contractible and noncontractible manufacturing tasks, respectively. Formally, we can now write the problem (ignoring fixed costs) as

$$\begin{aligned} \max_{h_c,h_n,m_c,m_n,s} \quad & \frac{1}{2}R - w_N\left(\mu_{hj}h_c + \left(1-\mu_{hj}\right)h_n\right) - s \\ \text{s.t.} \quad & s + \frac{1}{2}R - c_j\left(\mu_{mj}m_c + \left(1-\mu_{mj}\right)m_n\right) \geq 0 \\ & h_n = \arg\max_h \{\frac{1}{2}R - w_N\left(1-\mu_{hj}\right)h_n\} \\ & m_n = \arg\max_m \{\frac{1}{2}R - c_j\left(1-\mu_{mj}\right)m_n\}. \end{aligned} \tag{4.19}$$

[6] An implicit assumption in the analysis is that these noncontractible tasks are not yet fully embodied in into the manufactured inputs at the time of bargaining.

where revenue is given by

$$\begin{aligned} R &= B^{1/\sigma}\sigma\left(\sigma-1\right)^{-(\sigma-1)/\sigma}\varphi^{(\sigma-1)/\sigma} \\ &\quad\times\left(\frac{(h_c)^{\mu_{hj}}\left(h_n\right)^{1-\mu_{hj}}}{\eta}\right)^{(\sigma-1)\eta/\sigma}\left(\frac{(m_c)^{\mu_{mj}}\left(m_n\right)^{1-\mu_{mj}}}{1-\eta}\right)^{(\sigma-1)(1-\eta)/\sigma} \end{aligned} \quad (4.20)$$

and where $c_j = w_N$ when $j = N$ and $c_j = \tau w_S$ when $j = S$.

This problem is somewhat tedious to solve, so I will not go over the derivations here. The interested reader can find the details in the Theoretical Appendix, where I reproduce the derivations in Antràs and Helpman (2008), where we solved the same problem for a general division of revenue (β_h, β_m), rather than $(1/2, 1/2)$. Using these results, we find that, in the case of domestic sourcing, the profits obtained by F are now given by

$$\pi_D\left(\varphi\right) = \left(c_N\right)^{1-\sigma} B\Gamma_D\left(\mu_N\right)\varphi^{\sigma-1} - f_D w_N,$$

where

$$\Gamma_D\left(\mu_N\right) = \left(\frac{\sigma}{\sigma-\left(\sigma-1\right)\left(1-\mu_N\right)}+1\right)^{\sigma-(\sigma-1)(1-\mu_N)}\left(\frac{1}{2}\right)^{\sigma} \quad (4.21)$$

and

$$\mu_N \equiv \eta\mu_{hN} + \left(1-\eta\right)\mu_{mN}.$$

The derived parameter μ_N measures the average contractibility associated with domestic sourcing and is a weighted sum of the contractibility of headquarter services and manufacturing.

F's profits under foreign sourcing can be similarly computed, resulting in

$$\pi_O = \left(\left(w_N\right)^{\eta}\left(\tau w_S\right)^{1-\eta}\right)^{1-\sigma} B\Gamma_O\left(\mu_S\right)\varphi^{\sigma-1} - w_N f_O,$$

with

$$\Gamma_O\left(\mu_S\right) = \left(\frac{\sigma}{\sigma-\left(\sigma-1\right)\left(1-\mu_S\right)}+1\right)^{\sigma-(\sigma-1)(1-\mu_S)}\left(\frac{1}{2}\right)^{\sigma} \quad (4.22)$$

and

$$\mu_S \equiv \eta\mu_{hS} + \left(1-\eta\right)\mu_{mS}.$$

As in the simpler models developed above, the term $\Gamma_\ell\left(\mu_j\right)$ captures the contractual frictions associated with the sourcing options $\ell = D$ and

$\ell = O$, which entail manufacturing in country $j = N$ and country $j = S$, respectively. Differentiation of (4.21) and (4.22) demonstrates that each of these terms is increasing in their associated index of contractibility (see the Theoretical Appendix). Hence, as in our simpler model above, contract incompleteness reduces the profitability of production but the effect is now smoothly shaped by the partial contractibility parameters μ_{hj} and μ_{mj} for $j = N, S$. In fact, our initial model is a special case of the current one, with complete contracting in domestic sourcing (so $\mu_{hN} = \mu_{mN} = 1$), and 'totally incomplete contracts' in foreign sourcing (or $\mu_{hS} = \mu_{mS} = 0$).[7]

Our notation associates the relevant degree of contractibility in foreign sourcing with the quality of contractual institutions in the South. In particular, agents engaged in this type of sourcing strategy perceive that the quality of Southern institutions will be the key one determining the extent to which contracts specifying certain aspects of production, including headquarter service provision, will be enforced. This is a strong assumption to make. One would expect that the contractual insecurity of offshoring relationships would be a function of both the Northern and Southern institutions and perhaps their legal similarity, as argued in Chapter 3 when studying the exporting decision. Nevertheless, by using this notation, I seek to stress the notion that the quality of contracting institutions in the country where manufacturing takes place will be an important determinant of the profitability of offshore transactions. I will later appeal to this result when discussing the empirical evidence in Chapter 5. Of course, as I discussed in Chapter 3, parties can seek to insulate a given transaction from weak contract enforcement in 'the South' by including choice-of-law and forum-of-law clauses (see Chapter 1). Still, Southern institutions will likely remain crucial in determining the degree to which damages set by international courts of law or arbitrators are enforced.

Notice that equations (4.21) and (4.22) not only illustrate the positive effect of better contract enforcement on profitability, but they also shed light on the differential effect of such an improvement on institutions depending on other features of the environment. For instance, tedious differentiation of these expressions delivers the intuitive result that an increase in the reduced-form aggregate contractibility μ_j for $j = N, S$, will have a disproportionately larger effect on profitability whenever σ is high, that is when

[7] This is easily verifed by plugging these values of μ_{hj} and μ_{mj} for $j = N, S$ into the profit functions above and comparing these with (4.1) and (4.9).

the final-good producer faces a particularly competitive environment.[8] Intuitively, the higher the price elasticity faced by the final-good producer, the costlier will the investment inefficiencies associated with weak contracting prove to be. Apart from this interaction effect, in the Theoretical Appendix, we also show that the elasticity of demand continues to have an unambiguous negative effect on $\Gamma_\ell\left(\mu_j\right)$ in this more general framework.

When inspecting how the terms μ_N and μ_S are shaped by the contractibility of the different process of production, it is also evident that improvements in contractibility will interact with the headquarter intensity of production depending on the source of these changes in contractibility. For instance, if improvements in Southern institutions affect disproportionately the contractibility of manufacturing, then this version of the model predicts that these improvements will have a disproportionate effect on profitability in sectors with low headquarter intensity. Conversely, if Southern institutions disproportionately affect the extent to which F will capture the full marginal return from his or her investments in headquarter services, then the model predicts a larger impact of improved contracting on headquarter intensive sectors.

Partial Relationship Specificity

Although relationship-specific investments are pervasive in economic transactions, the assumption of *full* relationship-specificity in our basic model is extreme. Even when particular transactions end up not occurring, suppliers can generally recoup part of the cost of their investment, perhaps by reselling their goods to alternative buyers. Similarly, contractual breaches by suppliers may reduce the overall profitability of headquarter services, but will generally not render them useless. A proper modeling of partial-relationship-specificity would require the introduction of a secondary market for inputs as well as of the negotiations between final-good producers and suppliers in that market, which in turn might depend on the outside options of agents in a tertiary market, and so on (see, for instance, Grossman and Helpman, 2002). The main idea would then be that the lower is the degree of specificity, the larger is the value of inputs in the secondary markets and thus the lower should be the incentive of agents to underinvest. I will next consider a reduced-form version of such a model.

[8] More precisely, in the Theoretical Appendix we establish that $\partial\left(\partial \ln \Gamma_\ell / \partial \mu_j\right) / \partial \sigma > 0$.

In particular, assume that there indeed exists a second market for inputs in which the manager M can obtain a price $p_m^s(\varphi)$ for each unit of m, while the final-good producer F anticipates obtaining a monetary return $p_h^s(\varphi)$ per unit of headquarter services. These constitute the outside options for each party at the bargaining stage t_2 we have studied above. I will assume that the agents perceive these secondary market transaction prices $p_h^s(\varphi)$ and $p_m^s(\varphi)$ as unaffected by their actions and in particular their investment levels (more on this below).

Assuming again symmetric Nash bargaining, the payoff obtained by final-good producer will now be given $p_h^s(\varphi)\,h+\frac{1}{2}\left(R-p_h^s(\varphi)\,h-p_m^s(\varphi)\,m\right)$, while the supplier will obtain $p_m^s(\varphi)\,m+\frac{1}{2}\left(R(\varphi)-p_h^s(\varphi)\,h-p_m^s(\varphi)\,m\right)$. As a result, the levels of investments $h(\varphi)$ and $m(\varphi)$ at t_1 will satisfy the following first-order conditions:

$$\begin{aligned}\frac{1}{2}\left(\frac{\partial R(\varphi)}{\partial h}+p_h^s(\varphi)\right) &= w_N\\ \frac{1}{2}\left(\frac{\partial R(\varphi)}{\partial m}+p_m^s(\varphi)\right) &= \tau w_S.\end{aligned} \tag{4.23}$$

Consider next the determination of the prices $p_h^s(\varphi)$ and $p_m^s(\varphi)$. In a frictionless environment *without* any relationship-specificity, one would expect that this secondary market would provide a thick market for each input and that, in equilibrium, the price commanded by these inputs would correspond to the monetary value of their marginal product. In that case, we would have $p_k^s(\varphi)=\partial R(\varphi)/\partial k$ for $k=h,m$, and the corresponding investments in (4.23) would coincide with the efficient ones under complete contracts. In other words, in the absence of relationship-specificity of investments, weak contract enforcement is irrelevant as the hold-up problem disappears. Conversely, in the other extreme case with full relationship specificity, we instead have $p_h^s(\varphi)=p_m^s(\varphi)=0$, and the model collapses back to our basic model.[9]

In order to consider environments with partial relationship specificity assume then that the secondary market price commanded by each input is a share $1-\epsilon$ of the actual value of the marginal product of this input, so that

[9]The setup I am developing is admittedly special in that I am allowing the value of the marginal product of the manufacturer's investment m in the secondary market to be a function of the productivity level φ of the final-good producer with whom it initially contracted. This might reflect the fact that the secondary market is thick for any level of φ or perhaps that the supplier is able to assimilate F's technology while producing m.

larger values of ϵ are associated with larger degrees of customization or relationship specificity. As is clear from equation (4.23), this then corresponds to the case in which F and M choose investments h and m while anticipating obtaining a share $\beta_h = \beta_m = 1 - \epsilon/2$ of the actual value of the marginal return to these investments.

The rest of the equilibrium of this variant of the model is as in our basic model. Notice in particular that the parties will still find it efficient to reach an agreement at t_2 and thus the secondary market is never used in equilibrium.[10] Overall, the equilibrium can be solved in a manner analogous to the program in (4.19), but with $1 - \epsilon/2$ replacing $1/2$ in the second and third constraints.[11]

As mentioned before, Antràs and Helpman (2008) solved this program for a general division of revenue (β_h, β_m), so we can just plug $\beta_h = \beta_m = 1 - \epsilon/2$ into their equilibrium equations (see the Theoretical Appendix). This yields a level of profits associated with domestic sourcing ($\ell = D$ and $j = N$) and foreign sourcing ($\ell = O$ and $j = S$) equal to

$$\pi_\ell = (c_j)^{1-\sigma} B\Gamma_\ell\left(\mu_j, \epsilon\right)\varphi^{\sigma-1} - w_N f_\ell,$$

with

$$\Gamma_\ell\left(\mu_j, \epsilon\right) = \left(1 + \frac{\frac{\epsilon}{2}}{1-\frac{\epsilon}{2}}\frac{\sigma}{\sigma - (\sigma-1)\left(1-\mu_j\right)}\right)^{\sigma-(\sigma-1)\left(1-\mu_j\right)}\left(1-\frac{\epsilon}{2}\right)^{\sigma}, \tag{4.24}$$

and where remember that $\mu_j = \eta\mu_{hj} + (1-\eta)\,\mu_{mj}$, $c_j = w_N$ when $j = N$ and $c_j = \tau w_S$ when $j = S$. As in the model with full relationship specificity, it continues to be the case that improvements in contractibility are associated with larger values of $\Gamma_\ell\left(\mu_j, \epsilon\right)$. Similarly, the negative effect of σ on $\Gamma_\ell\left(\mu_j, \epsilon\right)$ and the positive "interaction" effect $\partial\left(\partial \ln \Gamma_\ell / \partial\mu_j\right)/\partial\sigma > 0$ continue to apply in this more general environment (see the Theoretical Appendix).

The main new feature of expression (4.24) is that the inefficiencies derived from incomplete contracting are now increasing in the degree of specificity ϵ in the sense that $\Gamma_\ell\left(\mu_j, \epsilon\right)$ decreases in ϵ. This intuitive result is not immediate from inspection of equation (4.24), but it can be verified by analyzing the

[10] More specifically, given the concavity of the revenue function, we necessarily have that $R(\varphi) > (1-\epsilon/2)\left(\frac{\partial R(\varphi)}{\partial h}h + \frac{\partial R(\varphi)}{\partial m}m\right)$.

[11] Despite the fact that F and M do not each receive a share $1 - \epsilon/2$ of revenue, their investments are determined *as if they did*.

partial derivative $\partial \ln \Gamma_\ell \left(\mu_j, \epsilon\right) / \partial\epsilon$. The reader is referred to the Theoretical Appendix for the mathematical derivations, which also demonstrate that the cross-partial derivative $\partial \left(\ln \partial\Gamma_\ell \left(\mu_j, \epsilon\right) / \partial\epsilon\right) / \partial\mu_j$ is positive. In words, the positive effect of higher quality of contracting institutions on firm profitability is predicted to be disproportionately higher in production processes with high degrees of specificity ϵ. Or, put differently, the model seems to be consistent with the fact that countries with weak contracting environments appear to export manufactured goods featuring relatively low levels of specificity, as empirically shown by Nunn (2007). I will further illustrate this result in Chapter 5, when I develop a multi-country version of the model.

Multiple Inputs and Multilateral Contracting

So far, I have focused on situations in which F is concerned only with the provision of one input. In modern manufacturing processes final-good producers instead combine intermediate inputs provided by various suppliers. I will next return to the version of the global sourcing model introduced in Chapter 2, in which the manufacturing stage of production entails the procurement of a continuum of measure one of inputs indexed by v, all produced simultaneously at t_1. Assuming that the services from these stages are imperfectly substitutable with each other with a constant and symmetric elasticity of substitution equal to $\sigma_\rho \equiv 1/\left(1-\rho\right)$, we can now write the production function as

$$q\left(\varphi\right) = \varphi \left(\frac{h}{\eta}\right)^{\eta} \left(\frac{\left[\int_0^1 m\left(v\right)^{\rho} dv\right]^{1/\rho}}{1-\eta}\right)^{1-\eta}. \tag{4.25}$$

Note that if one interprets $q\left(\varphi\right)$ as the *quality-adjusted* volume of output, this formulation is perfectly consistent with the notion that, from an engineering point of view, all stages might be essential. For example, producing a car requires four wheels, two headlights, one steering wheel, and so on, but the value of this car for consumers will typically depend on the services obtained from these different components, with a high quality in certain parts potentially making up for inferior quality in others.

I will assume that the continuum of inputs are not only symmetric in technology but are also produced with the same marginal cost in a given location. Manufacturing also continues to entail fixed costs that depend on the location of this activity but I assume that these fixed costs are independent of the number of inputs produced in a location. For this reason, it is natural to focus on symmetric equilibria in which all manufacturing inputs are produced in the same location. In Chapter 5 we will consider a more interesting and realistic environment in which firms source inputs from various locations.

Headquarter service provision continues to be controlled by the final good producer, agent F. To obtain the various intermediate inputs, F now needs to contract with a continuum of managers $M(v)$ each controlling one input. If all the aspects associated with the production of the different inputs could be specified in an enforceable manner in an initial contract, then it is straightforward to show that the resulting profit functions for the final-good producer associated with domestic sourcing and offshoring would be exactly identical to those of the single manufacturing input model. These profits flows are given by equations (4.1) and (4.2). Note that, given our symmetry assumptions and complete contracting, these profit flows are independent of the value of the input substitution parameter ρ. As we will next demonstrate, this parameter will play a much more relevant role in the presence of contractual frictions.

Consider now the case of partial contractibility introduced above, in which some of the characteristics of production are contractible, while others are not. Specifically, headquarter services h and each manufacturing input $m(v)$ are a Cobb-Douglas aggregate of the services of a continuum of measure one of activities, as in equations (4.17) and (4.18), and only a share μ_{hj} and μ_{mj} of those activities are contractible when manufacturing takes place in country $j = N, S$. Note that, for simplicity, the share μ_{mj} is common for all inputs v. The terms of exchange related to the noncontractible activities are only decided at t_2, after they have been performed but not yet embodied into production. The threat point for each party in the negotiations at t_2 is to withhold the services from those activities.

The key novel feature of this richer environment is that the ex-post negotiations at t_2 are now multilateral, rather than bilateral. How should one model these ex-post bargaining? One possible way would be to apply Nash bargaining to our multilateral setup with each agent obtaining their outside option plus a share of the difference between joint surplus under cooperation and the sum of outside options (see, for instance, Osborne and Rubinstein, 1990, p. 23). With zero outside options, and hence full relationship specificity, this would amount to all agents obtaining a constant share of revenues. This would lead, however, to a situation analogous to a moral hazard in teams problem (see Holmstrom, 1982) with an arbitrarily large number of agents. In such a case, the agents would have no incentive to invest in noncontractible tasks and revenue would be zero. In sum, a minimal amount of contractual frictions would be sufficient to drive production efficiency to zero.

This extreme result is in part due to our Cobb-Douglas assumptions in (4.17) and (4.18) but it also reflects the limitations of the Nash bargaining solution in multi-agent environments. In particular, this solution does not allow for situations of *partial cooperation* in which even if one supplier rejects an agreement, the other agents are still allowed to cooperate with each other and obtain some surplus. For this reason, in multilateral bargaining setups it is customary to adopt the Shapley value as the solution concept characterizing the equilibrium of these negotiations. In a bargaining game with a finite number of players, each player's Shapley value is the average of her contributions to all coalitions that consist of players ordered below her in all feasible permutations.[12]

A complication arises from the fact that, in our environment, we have a continuum of agents bargaining over surplus. Acemoglu, Antràs and Helpman (2007) resolve this issue by considering a discrete-player version of the game and computing the *asymptotic* Shapley value of Aumann and Shapley (1974). I will next develop an alternative, heuristic derivation of this Shapley value.

First note that agent F is an *essential* player in the bargaining game and thus a supplier $M(v)$'s marginal contribution is equal to zero when being added to a coalition that does not include the firm. When that coalition does include the firm and a measure n of suppliers, the marginal contribution of supplier v is equal to $\Delta R(v,n) = \partial R(\varphi,n)/\partial n$, where $R(\varphi,n) = p(q(\varphi,n))\,q(\varphi,n)$ and $q(\varphi,n)$ is as in (4.25) but with the integral running up to n rather than 1. Using Leibniz' rule and invoking symmetry,

[12] More formally, in a game with M players, let $g = \{g(1),...,g(M)\}$ be a permutation of $1,2,...,M$, and let $z_g^j = \{j' \mid g(j) > g(j')\}$ be the set of players ordered below j in the permutation g. Denoting by G the set of feasible permutations, and by $v : G \to \mathbb{R}$ the value (or surplus generated) of the coalition consisting of any subset of the M players, the Shapley value of player j is then

$$s_j = \frac{1}{(M+1)!}\sum_{g\in G}\left[v\left(z_g^j \cup j\right) - v\left(z_g^j\right)\right].$$

this marginal contribution can be succinctly written as

$$\Delta R(v,n) = \frac{(\sigma-1)(1-\eta)}{\sigma\rho} R(\varphi) \left(\frac{m_n(v)}{m_n(-v)}\right)^{\rho} n^{\frac{(\sigma-1)}{\sigma\rho}-1},$$

where $m(-v)$ represents the (symmetric) investments of all suppliers other than v. Note that in deriving this expression we also have imposed, without loss of generality, a symmetric choice for contractible manufacturing tasks, or $m_c(v) = m_c$ for all v.

The Shapley value of $M(v)$ is the average of her marginal contributions to coalitions that consist of players ordered below her in all feasible orderings. A supplier that has a measure n of players ordered below her has a marginal contribution of $\Delta R(v,n)$ if the firm is ordered below her – which occurs with probability n –, and 0 otherwise. Averaging over all possible orderings of the players and using the above formula for $\Delta R(v,n)$ we obtain the following payoff for supplier $M(v)$:

$$P_m(v) = \int_0^1 n\Delta R(v',n)\,dv' = \frac{(\sigma-1)(1-\eta)}{(\sigma-1)(1-\eta)+\sigma\rho} R(\varphi) \left(\frac{m_n(v)}{m_n(-v)}\right)^{\rho}. \tag{4.26}$$

A number of features of (4.26) are worth noting. First, in equilibrium, all suppliers invest equally in all the noncontractible activities, and thus each receives a share $(\sigma-1)(1-\eta)/(\sigma-1+\sigma\rho)$ of revenue, leaving F with the residual share $\sigma\rho/((\sigma-1)(1-\eta)+\sigma\rho)$. The bargaining power of the firm is thus naturally increasing in the substitutability of inputs as governed by ρ, since the suppliers' bargaining threats are less effective in that case. Second, and although in equilibrium suppliers end up with an equal share of sale revenue, equation (4.26) indicates that suppliers perceive their noncontractible investments to have a non-negligible (i.e., measurable) effect on their payoffs, and thus the moral-hazard-in-teams, zero-investment result mentioned above does not apply here. Third, the degree of substitutability ρ crucially impacts the marginal return to suppliers' investments by shaping the degree to which increases in the investments of a given supplier affect output. Intuitively, when inputs are highly complementary (low ρ), the marginal return to increasing the production of one input v while holding the rest fixed is particularly low.

Having solved for the division of surplus at t_2, the rest of the equilibrium is as in previous models. In particular, the program is analogous to that in (4.19), but with $\beta_h = \sigma\rho/((\sigma-1)(1-\eta)+\sigma\rho)$ replacing 1/2 in

F's choice of h and with each supplier $M(v)$ choosing $m_n(v)$ to maximize $P_m(v) - c_j(1-\mu_{mj})m_n(v)$. In equilibrium, the latter choice is isomorphic to that of a single supplier choosing a *common* m_n for all v to maximize $\beta_m R(\varphi) - c_j(1-\mu_{mj})m_n$ with $\beta_m = \rho\sigma/((\sigma-1)(1-\eta)+\sigma\rho)$.[13] Thus we can apply the general formula in Antràs and Helpman (2008) (see the Theoretical Appendix) to express the level of contractual frictions $\Gamma_\ell(\mu_j,\rho)$ associated with manufacturing taking place in country j as

$$\Gamma_\ell(\mu_j,\rho) = \left(1+\frac{1}{\rho}\frac{(\sigma-1)(1-\eta)}{\sigma-(\sigma-1)(1-\mu_j)}\right)^{\sigma-(\sigma-1)(1-\mu_j)}\left(\frac{\rho\sigma}{\rho\sigma+(\sigma-1)(1-\eta)}\right)^{\sigma}, \tag{4.27}$$

with again $\mu_j = \eta\mu_{hj} + (1-\eta)\mu_{mj}$ (where $j = N$ when $\ell = D$ and $j = S$ when $\ell = O$).

The generality of the results in Antràs and Helpman (2008) allows us to conclude, without having to differentiate this expression, that $\Gamma_\ell(\mu_j,\rho)$ is again increasing in the degree of contractibility μ_j and decreasing in the elasticity of demand σ.[14] In addition, the positive effect of contract enforcement affects high-substitutability sectors disproportionately, or $\partial(\partial\ln\Gamma_\ell/\partial\mu_j)/\partial\sigma > 0$.

The main novelty of equation (4.27) is that the degree of input substitutability is now a key determinant of the extent to which contractual frictions depress the profitability of production. Straightforward differentiation demonstrates that $\Gamma_\ell(\mu_j,\rho)$ is increasing in ρ, and thus contractual frictions are lower, the more substitutable inputs are. Intuitively, investments tend to be less distorted in that case because a higher level of ρ (i) provides more ex-post surplus to F thus enhancing the investments in headquarter services by F, and (ii) increases the sensitivity of suppliers' ex-post payoffs to their own investments. Naturally, a high ρ also reduces the share of ex-post surplus accruing to suppliers, but given the functional forms in the model, this is a dominated effect.

Differentiation of (4.27) also demonstrates (see the Theoretical Appendix) that $\partial(\partial\ln\Gamma_\ell(\mu_j,\rho)/\partial\mu_j)/\partial\rho < 0$ and thus the effect of an improvement

[13] This follows from noting that β_m must be such that $\rho P_m(v) = \beta_m\frac{(\sigma-1)(1-\eta)}{\sigma}R(\varphi)$ whenever $m_n(v) = m_n$ for all v.

[14] This latter comparative static result would appear to be complicated by the fact that the bargaining weights β_h and β_m are now endogenous and a function of σ. But since $\Gamma_\ell(\mu_j,\rho)$ in (8.9) is increasing in β_h and β_m, and each of these two shares is decreasing in σ, this does not affect the sign of the derivative $\partial\ln\Gamma_\ell/\partial\sigma$.

in contractual institutions has a differentially larger effect in sectors featuring higher input complementarities. The model thus suggests that, other things equal, foreign sourcing to countries with particularly weak contract enforcement should be more prevalent in sectors with higher substitutability between inputs. We will further formalize and test this result in Chapter 5.

Sequential Production

The variant of the model with multiple suppliers that I have developed above assumes that all stages of production are performed simultaneously. In real-life manufacturing processes, there is often a natural sequencing of stages. First, raw materials are converted into basic components, which are next combined with other components to produce more complicated inputs, before themselves being assembled into final goods. Antràs and Chor (2013) develop a sequential production variant of the model with a continuum of inputs we have just studied. The key new feature of their analysis is that the relationship-specific investments made by suppliers in upstream stages can affect the incentives of suppliers involved in downstream stages thereby generating investment inefficiencies that vary systematically along the value chain.

The model developed by Antràs and Chor (2013) turns out to be very tractable but some of the details of the analysis are somewhat intricate, so I refer the reader to the paper and its Supplemental Appendix for many details. Antràs and Chor (2013) assume a production technology analogous to (4.25) but with $v \in [0, 1]$ indexing the position of an input in the value chain, with a larger v corresponding to stages further downstream (closer to the final end product). Although they develop extensions with headquarter services and partial contractibility, I will focus below on their benchmark model in which $\eta = 0$ and in which all investments are noncontractible.

The final-good producer F plays two roles in the model. On the one hand, it is in charge of assembling the measure one of sequentially produced inputs into a final good valued by consumers. Second, it sequentially negotiates with suppliers once their stage input has been produced and the firm has had a chance to inspect it. It is simplest to consider the case in which this negotiation at stage v is treated independently from the bilateral negotiations that take place at other stages (see Antràs and Chor, 2013, for alternative formulations). Because each intermediate input v is assumed compatible only with the firm's output, the supplier's outside option at the bargaining stage is 0. Hence, the quasi-rents over which the firm and the supplier negotiate are given by the incremental contribution to total revenue generated by supplier v at that stage. In light of (4.20) and (4.25), this incremental contribution is given by

$$\Delta R(v) = \frac{(\sigma-1)}{\sigma\rho} B^{1/\sigma}\sigma(\sigma-1)^{(\sigma-1)/\sigma}\varphi^{(\sigma-1)/\sigma} \times \left(\int_0^v m(u)^{(\sigma_\rho-1)/\sigma_\rho} du\right)^{(1-\sigma_\rho/\sigma)/(\sigma_\rho-1)} m(v)^\rho, \quad (4.28)$$

where remember that $\sigma_\rho = 1/(1-\rho)$. Assume that the share of these quasi-rents accruing to F are given by $\beta(v)$. Below, I will allow this share to be affected by the location of manufacturing production.

Notice that if $\sigma > \sigma_\rho$, then the investment choices of suppliers are *sequential complements* in the sense that higher investment levels by prior suppliers increase the marginal return of supplier v's own investment $m(v)$. Conversely, if $\sigma < \sigma_\rho$, investment choices are *sequential substitutes* because high values of upstream investments reduce the marginal return to investing in $m(v)$. Because the supplier at position v chooses $m(v)$ to maximize $(1-\beta(v))\Delta R(v) - c_j m(v)$, equation (4.28) illustrates the trickle down effect that upstream investment inefficiencies can have on downstream stages.

Exploiting the recursive structure of the model, Antràs and Chor (2013) show that if agent F is able to use ex-ante transfers to extract all surplus from suppliers, then the overall profits obtained by a final-good producer with productivity φ when all inputs are produced under a marginal cost equal to j are given by

$$\pi_j = (c_j)^{1-\sigma} B\Gamma_j\left(\{\beta(v)\}_{v=0}^1\right)\varphi^{\sigma-1}$$

where

$$\Gamma_\ell\left(\{\beta(v)\}_{v=0}^1\right) = \frac{(\sigma-1)}{(\sigma_\rho-1)}\left(\frac{\sigma_\rho}{\sigma}\right)^{\frac{\sigma-\sigma_\rho}{\sigma_\rho-1}} \int_0^1 \{\left(\frac{\sigma_\rho}{1-\beta(v)} - (\sigma_\rho-1)\right) \times (1-\beta(v))^{\sigma_\rho}\left[\int_0^v (1-\beta(u))^{\sigma_\rho-1} du\right]^{\frac{\sigma-\sigma_\rho}{\sigma_\rho-1}}\} dv. \quad (4.29)$$

In the case of a symmetric bargaining power at all stages, so $\beta(v) = \beta$ for all v, equation (4.29) reduces to

$$\Gamma_\ell\left(\{\beta(v)\}_{v=0}^1\right) = \left(\frac{\sigma_\rho}{\sigma}\right)^{\frac{\sigma-\sigma_\rho}{\sigma_\rho-1}}\left(\frac{\sigma_\rho}{1-\beta} - (\sigma_\rho-1)\right)(1-\beta)^\sigma .$$

This expression in turns collapses to the single-supplier index of contractual frictions in equation (4.14) when $\sigma = \sigma_\rho$ (and $\eta = 0$, of course). This is intuitive since in that knife-edge case, the payoff to a supplier is independent of other suppliers' investments, and the trickle-down effects mentioned above become irrelevant.

More interesting implications from the modeling of sequential production can be obtained when allowing the bargaining share $\beta(v)$ to vary along the value chain and across manufacturing locations. To build intuition, it is instructive to consider first the case in which the (infinite-dimensional) vector of $\beta(v)$'s is chosen to maximize F's profits, and thus $\Gamma_j\left(\{\beta(v)\}_{v=0}^{1}\right)$. Antràs and Chor (2013) show that this seemingly complicated problem can be reduced to a standard calculus of variation problem which delivers the surprisingly simple Euler-Lagrange condition

$$\frac{\partial\beta^*(v)}{\partial v} = \frac{1-\sigma_\rho/\sigma}{\sigma_\rho - 1} v^{\frac{\sigma_\rho(\sigma-1)}{(\sigma_\rho-1)\sigma}}.$$

The key implication of this expression is that the relative size of the input and final-good elasticities of substitution σ_ρ and σ governs whether the incentive for F to retain a larger surplus share increases or decreases along the value chain. Intuitively, when σ is high relative to σ_ρ, investments are sequential complements, and high upstream values of $\beta(v)$ are particularly costly since they reduces the incentives to invest not only of these early suppliers but also of all suppliers downstream. Conversely, when σ is small relative to σ_ρ, investments are sequential substitutes, and low values of $\beta(v)$ in upstream stages are now relatively detrimental, since they reduce the incentives to invest for downstream suppliers, who are already underinvesting to begin with.[15]

This result has interesting implications for the choice between domestic and foreign sourcing whenever these sourcing strategies are associated with

[15] Imposing two boundary conditions on the Euler-Lagrange equation, the optimal stage-v bargaining share can in fact be solved in closed form, and is given by $\beta^*(v) = 1 - v^{\frac{\sigma_\rho/\sigma-1}{\sigma_\rho-1}}$. In the sequential complements case ($\sigma > \sigma_\rho$) this implies $\beta^*(v) < 0$ for all v, so F has an incentive to allocate to suppliers more than their entire incremental contribution. This extreme result does not apply when F cannot extract all surplus from suppliers via ex-ante lump-sum transfer or when the model includes headquarter service provision (see Antràs and Chor, 2013, for details). Importantly, in those cases, it continues to be the case that the sign of $\partial\beta^*(v)\,\partial v$ is determined by the relative size of σ and σ_ρ, and $\partial\beta^*(v)\,\partial v > 0$ whenever $\sigma > \sigma_\rho$.

different levels of contract enforcement or with different bargaining shares for F in its negotiations with suppliers. To see this, consider first the case in which contracting in domestic Northern transactions is complete, while foreign sourcing is associated with totally incomplete contracting. Our results above then suggest that, in the sequential complements case ($\sigma > \sigma_\rho$), foreign sourcing is particularly unappealing in upstream stages. Thus, if domestic and foreign sourcing coexist along the value chain, then only relatively downstream inputs will be offshored.[16] Conversely, in the sequential substitutes case, ($\sigma < \sigma_\rho$) one would expect relatively upstream stages to be offshored. In sum, the model predicts that the 'upstreamness' of an input should be a relevant determinant of the extent to which it is procured from foreign suppliers, with the sign of that dependence being crucially shaped by the relative size of σ and σ_ρ.

Note, however, that very different results might arise if domestic and foreign sourcing do not differ significantly in their contractibility, but are associated with F obtaining a higher share of surplus under domestic sourcing than under offshoring, i.e., $\beta_D(v) > \beta_O(v)$, as suggested for instance by Antràs and Helpman (2008). In such a case, offshoring would be relatively more appealing in upstream stages in the sequential complements case, and relatively more appealing in downstream stages in the sequential substitutes case. We will explore the empirical relevance of these different scenarios in the next chapter.

Summary and Next Steps

This chapter has explored the determinants of the global sourcing decisions of firms in the presence of incomplete contracting frictions in vertical relationships. The different variants of our global sourcing model have delivered a rich set of comparative statics and have also provided tools for testing these predictions with data on intermediate input trade. In the next chapter, I will test the empirical success of the model with detailed data on U.S. imports by product and source country. In the process, and given the cross-country dimension of the data, it will prove necessary to develop a multi-country version of the model.

---

[16] This result can be proved in a manner analogous to Proposition 2 in Antràs and Chor (2013).