Organizing the Global Value Chain*

Pol Antràs
Harvard University

Davin Chor
Singapore Management University

Preliminary and Incomplete

November 18, 2011

Abstract

We develop a property-rights model of the firm in which production entails a continuum of uniquely sequenced production stages. In each stage, a final-good producer contracts with a distinct supplier for the procurement of a customized stage-specific component. Our model yields predictions regarding the optimal allocation of ownership rights along the value chain. We show that whether the incentives to integrate suppliers are relatively higher or lower for downstream stages than for upstream stages depends crucially on the elasticity of demand faced by the final-good producer. Our model allows for various sources of asymmetry across final-good producers and across suppliers within a production line, and we show how it can be taken to the data using international trade statistics. Combining data from the U.S. Census Bureau's Related Party Trade database, U.S. Input-Output Tables, and estimates of import demand elasticities from Broda and Weinstein (2006), we find empirical evidence broadly supportive of our model's predictions. In the process, we propose a new measure of the average position of an industry in the value chain which boils down to a simple function of the square of the Leontief inverse matrix.

*We are grateful to Arnaud Costinot, Don Davis, Ron Findlay, Elhanan Helpman, Kala Krishna, Marc Melitz, Daniel Trefler, Jonathan Vogel, and David Weinstein for their helpful comments, and to Nathan Nunn for generously providing data. Thanks also to audiences at the Econometric Society World Congress (Shanghai), Society for the Advancement of Economic Theory Conference (Singapore), Nanyang Technological University, Singapore Management University, Columbia, Wisconsin, Kiel, Munich, Tübingen and Bonn. Gurmeet Singh Ghumann, Frank Schilbach, and Zhicheng Song provided excellent research assistance. Chor gratefully acknowledges the research funding provided by a Sing Lun Fellowship. All errors are our own.

1 Introduction

Most production processes are sequential in nature. At a broad level, the process of manufacturing cannot commence until the efforts of R&D centers in the development or improvement of goods has proven to be successful, while the sales and distribution of manufactured goods cannot be carried out until their production has taken place. Even within manufacturing processes, there is often a natural sequencing of production stages. First, raw materials are converted into basic components, which are next combined with other components to produce more complicated inputs, which are themselves assembled into final goods. This process very much resembles Henry Ford's original Model T production assembly line, but revolutionary advances in information and communication technology, coupled with a gradual reduction in natural and man-made trade barriers, have allowed the value chain to be 'sliced up' into geographically separated steps.

As emphasized by Grossman and Rossi-Hansberg (2011), Boeing's production of the 787 Dreamliner is an illustrative example of these practices. According to Boeing's website, the 787 "Development Team" encompasses 50 suppliers located in 9 countries (Australia, France, Germany, Italy, Japan, Korea, Sweden, the United Kingdom and the United States).[1] The involvement of foreign suppliers is not anecdotal as it accounts for close to 70 percent of the 787's parts (Newhouse, 2007, p.29). Furthermore, the production process is very much sequential in nature. The manufacturing process only began after Boeing recognized a demand for a more fuel-efficient airliner and after years of design and development of the different variants of the 787. Within the actual manufacturing of the aircraft, the process involves the various suppliers procuring raw materials and components from other producers further upstream and combining them to build and preassemble big pieces of the plane (the wings, different parts of the fuselage, the engine, etc.). The final manufacturing stage consists of the final assembly of the aircraft, which takes place in Boeing's plant in Everett, Washington. After the assembly of the aircraft is completed, it still needs to be subjected to a battery of pre-flight ground tests and to a flight test program.

The implications of sequential production for the workings of open-economy general equilibrium models have been widely explored in the literature. A few papers, most notably Findlay (1978), Dixit and Grossman (1982), Sanyal (1983), Kremer (1993), and Costinot, Vogel and Wang (2011), have emphasized that the pattern of specialization along the value chain has interesting implications for the world income distribution and for how shocks spread across countries. Other papers, including Yi (2003), Harms, Lorz, and Urban (2009), and Baldwin and Venables (2010) have unveiled interesting nonlinearities in the response of trade flows to changes in trade frictions in models in which production entails value being added sequentially along locations around the globe.

The focus of our paper is different. Our aim is to understand how the sequentiality of production shapes the contractual relationships between final-good producers and their various suppliers, and how the allocation of control rights along the value chain can be designed in a way that elicits

[1] The list can be found at http://www.boeing.com/commercial/787family/dev_team.html.

(constrained) optimal effort on the part of suppliers. An obvious premise of our work is that, although absent from most general-equilibrium models, contractual frictions are relevant for the efficiency with which production is carried out, and also for the particular way in which different production processes are organized across borders. We find this to be a natural assumption to make particularly in international trade environments, in which determining which country's laws are applicable to particular contractual disputes is often difficult. The detrimental effects of imperfect contract enforcement on international trade flows are particularly acute in transactions involving intermediate inputs, as those tend to be associated with longer time lags between the time the order is placed (and the contract is signed) and the time the goods or services are delivered (and the contract executed), and they also often entail significant relationship-specific investments and other sources of lock-in on the part of both buyers and suppliers.[2]

The relevance of incomplete contracting for the organization of production also rests on solid empirical underpinnings. A recent literature (see, for instance, Nunn, 2007, and Levchenko, 2007) has convincingly documented that contracting institutions appear to be an important determinant of international specialization, while another branch of the trade literature, to which our paper will contribute, has also shown that the ownership decisions of multinational firms depict various patterns that are consistent with Grossman and Hart's (1986) incomplete-contracting, property-rights theory of firm boundaries (see, among others, Antràs, 2003, Nunn and Trefler, 2008, and Bernard et al., 2010).

Boeing's experience with the 787 Dreamliner again illustrates the significance of these features. The repeated delays experienced during the production process, which is now over three years behind schedule, have been ascribed in part to the fact that multiple suppliers did not stand by their contractual obligations.[3] Boeing responded to these delays by partially reorganizing their sourcing model and bringing some of the problematic upstream production stages within their firm boundaries. For example, during 2008 and 2009, Boeing successively acquired operations from Vought Aircraft Industries, a company that was producing the rear sections of the Dreamliner's fuselage but that had been identified as a problematic supplier. One of those acquisitions entailed forming a 50-50 joint venture with a subsidiary of Italy's Alenia Aeronautica, another key supplier for Boeing, from which it procures the horizontal stabilizer and the center fuselage for the Dreamliner, and with which it has also wrestled in past years. The experience of Boeing illustrates that firms often design the ownership structure along their value chain in a way that attempts to minimize frictions resulting from suppliers not honoring their contractual commitments.

These examples, however, raise the following questions. If increasing its ownership stake in a supplier reduces a final-good producer's (say Boeing's) exposure to production delays caused by that particular supplier, why does this firm not bring the whole production process within their

[2] Suppliers often customize their output to the needs of particular buyers and would find it hard to sell those goods to alternative buyers, should the intended buyer decide not to abide by the terms of the contract. Similarly, buyers often undertake significant investments whose return can be severely diminished by incompatibilities, production line delays or quality debasements associated with suppliers not going through with their contractual obligations.

[3] See http://www.bloomberg.com/apps/news?pid=newsarchive&sid=aF6uWvMb9C08 for a more complete account.

firm boundaries? And if it decides to hold a controlling stake in some suppliers but not in others, what are the key determinants of which suppliers are brought inside the firm and which are not?

In this paper, we develop a property-rights model of firm boundaries that permits an analysis of the optimal allocation of ownership rights in a setting in which production stages are sequential in nature and contracts are incomplete. Our model builds on Acemoglu, Antràs and Helpman (2007). Production of final-goods entails a large number (formally, a continuum) of production stages and each stage is performed by a different producer (or supplier), who needs to undertake a relationship-specific investment in order to produce components that will be compatible with those produced by other suppliers in the value chain. These components are combined according to a constant-elasticity-of substitution (CES) aggregator by a final-good producer that faces an isoelastic demand curve. Contracts between final-good producers and their suppliers are incomplete in the sense that contracts contingent on whether components are compatible or not cannot be enforced by third parties.

The key innovation of our framework relative to the work of Acemoglu, Antràs and Helpman (2007) – and relative to the previous property-rights models of multinational firm boundaries in Antràs (2003, 2005) and Antràs and Helpman (2006, 2008) – is that we introduce a natural (or technological) ordering of production stages, in the sense that production of a stage cannot commence until the inputs or components from upstream stages have been delivered. Absent a binding initial (ex-ante) agreement, the firm and its suppliers are left to sequentially bargain over how the surplus associated with a particular stage is divided between the firm and the particular supplier of that stage. As in Grossman and Hart (1986), in this incomplete-contracting environment, owning a supplier is a source of *power* for the firm because the residual rights of control associated with ownership allow the firm to take actions (or make threats) that enhance their bargaining power vis-à-vis the supplier. Yet the optimal allocation of ownership rights is not always associated with all production stages being integrated, because by reducing the bargaining power of suppliers, integration reduces the incentives of suppliers to invest in the relationship.[4]

We begin in Section 2 by developing a benchmark model of firm behavior that isolates the role of the degree of "donwstreamness" of a supplier in shaping organizational decisions. A key feature of our analysis is that the relationship-specific investments made by suppliers in upstream stages affect the incentives to invest of suppliers in downstream stages. The nature of this dependence is shaped in turn by whether suppliers' investments are *sequential complements* or *sequential substitutes*, in the sense that higher investment levels by prior suppliers increase or decrease the marginal monetary return of a particular supplier. Even though, from a strict technological point of view (i.e., in light of the CES aggregator of inputs), inputs are always complements, suppliers' investments can still prove to be sequential substitutes when the elasticity of demand faced by the final good producer is sufficiently low, because in those cases, marginal revenue falls particularly quickly along the value

[4] Zhang and Zhang (2008, 2011) introduce sequential elements in a standard Grossman and Hart (1986) model but focus on one-supplier environments in which either the firm or the supplier has a first-mover advantage. Other papers that have studied optimal incentive provision in sequential production processes include Winter (2006) and Kim and Shin (2011).

chain. Whether inputs are sequential complements or sequential substitutes turns out to depend *only* on whether the "technological" elasticity of substitution is lower or higher than the price elasticity of demand faced by the final-good producer.

The central result of our model is that the optimal pattern of ownership along the value chain also depends critically on whether production stages are sequential complements or substitutes. When the demand faced by the final-good producer is sufficiently elastic, then there exists a unique cutoff production stage such that all production stages prior to this cutoff stage are outsourced, while all stages (if any) above that threshold are integrated. If demand is sufficiently inelastic, the converse is true: it is optimal to integrate relatively upstream stages and if outsourcing is observed along the value chain, it necessarily involves relatively downstream stages. In Section 3, we show that these results are robust to the introduction of investments by the final-good producer and of productivity or cost heterogeneity across final-good producers and also across suppliers within a production chain. These extensions prove useful in guiding our empirical analysis.

In Sections 4 and 5, we develop an empirical test of the main predictions of our framework. We follow the bulk of the recent empirical literature on multinational firm boundaries in using U.S. Census data on intrafirm trade to measure the relative prevalence of vertical integration in particular products. More specifically, we correlate the share of U.S. intrafirm imports over total U.S. imports in a particular product during the period 2000-10 with the average degree of "downstreamness" of that product, and we study whether this dependence is crucially shaped by the elasticity of demand faced by the average buyer of this product.

We propose two measures of downstreamness, both of which make use of the 2002 U.S. Input-Output Tables. Our first measure is the ratio of the aggregate direct use to the aggregate total use ($DUse\_TUse$) of a particular industry i's goods, where the *direct use* for a pair of industries (i, j) is the value of goods from industry i directly used by firms in industry j to produce goods for final use. The *total use* for (i, j) on the other hand includes the value of goods from industry i used directly or indirectly in producing industry j's output for final use, and thus incorporates the value of goods from industry i that had entered the production chain further upstream via purchases from other industries that also provide inputs to j. A low value of $DUse\_TUse$ for a given industry i thus suggests that most of the contribution of input i to production processes tends to occur at relatively upstream production stages. Conversely, the larger the value of $DUse\_TUse$, the further downstream an industry tends to enter other production processes on average. An advantage of this variable is that it is straightforward to compute given the information provided in Input-Output tables. Nevertheless, it has two main limitations: (i) it does not incorporate the value of industry i's goods that goes to final use (which obviously is correlated with downstreamness); and (ii) it does not necessarily generate a clean ranking for the most upstream inputs, as these may rarely be used as direct inputs into final-good production.

For these reasons, we propose a second measure of downstreamness (*DownIndex*) which is a weighted index of the average position in the value chain at which an industry's output is used (i.e., as final consumption, as direct input to other industries, as direct input to industries serving as direct inputs to other industries, and so on), with the weights being given by the ratio of the use of that industry's output in that position relative to the total output of that industry. Although constructing such a measure would appear to require computing an infinite power series, we show that $DownIndex$ can be succinctly expressed as a simple function of the square of the Leontief inverse matrix.

Our empirical tests also call on us to attempt to distinguish between the cases of sequential complements and substitutes identified in the theory. For that purpose, we use the U.S. import demand elasticities estimated by Broda and Weinstein (2006) and data on U.S. Input-Output Tables to compute a weighted average of the demand elasticity faced by the buyers of goods from each particular industry i.

Our empirical investigation provides support for the theoretical implications of the model. Our empirical specifications and results are discussed in more detail in Sections 4 and 5, but Figure 1 provides a first illustration of the main empirical results of the paper. The two panels in the figure depict a scatterplot of the share of U.S. intrafirm imports in a given sector in 2005 against the ratio of the aggregate direct use to the aggregate total use ($DUse\_TUse$) in that sector. The size of each observation marker is proportional to the value of U.S. imports in that sector. The only difference between the two panels is that the left one only includes sectors with an average buying industry's elasticity of demand above the median industry level for that variable, while the right panel only includes sectors with below-median values. As is clear, the relationship between the intrafirm import share and $DUse\_TUse$ is quite different in the two cases. In the substitutes case, the relationship is negative and statistically significant, regardless of whether we run simple OLS (red dashed line) or weight observations (black solid line) by the value of total import volumes in that sector (a strategy that, as we argue later, helps reduce error-in-variable biases). Conversely, in the complements case, the OLS estimates do not unveil a systematic relationship between downstreamness and the share of intrafirm trade, but when weighting by import volumes, a clear positive and significant relationship emerges between these two variables.

Consistent with the figure, our paper will unveil a robust statistically significant positive relationship between alternative measures of downstreamness and the share of intrafirm trade in a given sector, but such a relationship emerges *only* for high values of the elasticity of demand faced by industries purchasing these sector's goods (i.e., in the complements case). These findings hold even after controlling for other relevant determinants of the intrafirm trade share that have been highlighted by recent empirical work on multinational firm boundaries. We shall also show that our results are robust to specifications that further exploit the cross-country dimension of the data on intrafirm trade, while controlling for unobserved variation in factor costs with country-year fixed effects. The statistical significance of the negative relationship that emerges in the right panel of Figure 1 for low values of the elasticity of demand (the substitutes case) will prove to be less robust to alternative specifications with additional control variables, but as we shall see, the relationship will remain qualitatively and statistically distinct from that in the complements case.

The remainder of this paper is organized as follows. In Section 2, we develop our benchmark

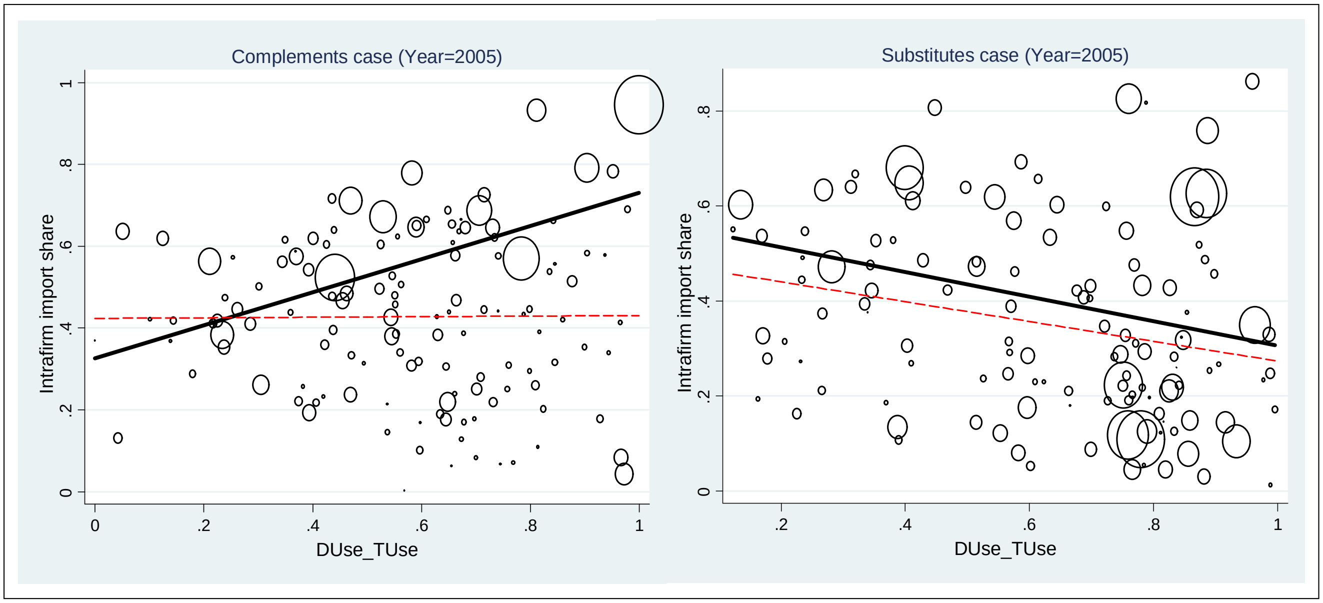

Figure 1: Downstreamness and the Share of Intrafirm Trade

model of sequential production with incomplete contracting and study the optimal ownership structure along the value chain. In Section 3, we develop a few extensions of the model and discuss how we attempt to take it to the data. We describe our data sources and empirical specification in Section 4, and present the results in Section 5. Section 6 offers some concluding remarks.

2 A Model of Sequential Production with Incomplete Contracts

We begin by developing a benchmark model of firm behavior along the lines of Acemoglu et al. (2007), but extended to incorporate a deterministic sequencing of production stages. The model is stylized in order to emphasize the new insights that emerge from considering the sequentiality of production. We will later incorporate more realistic features and will also embed the framework in industry equilibrium to guide the empirical analysis.

2.1 Benchmark Model

Environment We consider the organizational problem of a firm producing a final good. Production requires the completion of a measure one of production stages. We index production stages by $j \in [0, 1]$ with a larger j corresponding to stages further downstream (closer to the final end product) and we let $x(j)$ be the amount of *compatible* intermediate inputs that the supplier of stage j delivers to the firm. The *quality-adjusted* volume of final-good production is then given by:

$$q = \theta \left(\int_0^1 x(j)^\alpha dj \right)^{1/\alpha} \tag{1}$$

where θ is a productivity parameter and the parameter $\alpha \in (0,1)$ captures the (symmetric) degree of substitutability among the different production stages.

There is a large number of profit-maximizing suppliers that can produce the necessary intermediate goods or engage in an alternative activity that delivers them an outside option normalized to 0. We assume that each intermediate input needs to be produced by a different supplier with whom the firm needs to contract. A supplier assigned to the production of an intermediate input needs to undertake a relationship-specific investment in order to produce a *compatible* input. In order to highlight the asymmetries that will arise from the sequencing of production, we assume that production stages are otherwise symmetric: the marginal cost of investment is common for all suppliers and equal to c, and in all stages $j \in [0,1]$, one unit of investment generates one unit of the stage j compatible input when combined with the inputs provided by upstream suppliers. We will relax these symmetry assumptions later in the extensions in Section 3. Incompatible inputs can be produced by all agents (including the firm) at a negligible marginal cost, but they add no value to final-good production.

The production function in (1) appears mathematically identical to a conventional constant-elasticity-of-substitution production function with a continuum of inputs, but it is important to stress that the production technology is inherently sequential in that downstream stages cannot commence until the inputs or components from upstream stages have been delivered. In fact, the technology in (1) can be expressed in differential form, by noting that the change in output brought about by the supplier at stage m is given by a simple Cobb-Douglas function of this supplier's input production and the quality-adjusted volume of production up to that stage (which can be treated as an input to the stage-m production process). More specifically, the output generated up to stage-m of the production process is: $q(m) = \theta \left(\int_0^m x(j)^\alpha dj \right)^{1/\alpha}$, so that applying Leibniz' rule to (1), we have:

$$q'(m) = \frac{\partial q(m)}{\partial m} = \frac{1}{\alpha} \theta^\alpha x(m)^\alpha q(m)^{(1-\alpha)}.$$

The final good under study is differentiated in the eyes of consumers. The good belongs to an industry in which firms produce a continuum of goods and consumers have preferences that feature a constant elasticity of substitution $1/(1-\rho)$ across these varieties. More specifically, denoting by $\varphi(\omega)$ the quality of a variety and by $\tilde{q}(\omega)$ its consumption in physical units, the sub-utility in the industry is given by

$$U = \left(\int_{\omega \in \Omega} \left(\varphi(\omega) \tilde{q}(\omega) \right)^\rho d\omega \right)^{1/\rho}, \text{ with } \rho \in (0,1). \tag{2}$$

As is well known, when maximizing this function subject to the budget constraint $\int_{\omega \in \Omega} p(\omega) \tilde{q}(\omega) d\omega = E$, where E denotes expenditure, consumer demand for a particular variety features a constant price elasticity equal to $1/(1-\rho)$. Furthermore, the implied revenue function of any firm in the industry is concave in quality-adjusted output $q(\omega) = \varphi(\omega) \tilde{q}(\omega)$ with a constant elasticity ρ. Combining this feature with the production technology in (1), the revenue obtained by the final-good producing firm under study can be written as

$$r = A^{1-\rho}\theta^{\rho}\left(\int_0^1 x(j)^{\alpha}\,dj\right)^{\rho/\alpha}, \tag{3}$$

where $A > 0$ is an industry-wide demand shifter that the firm treats as exogenous.

Complete Contracts Before discussing in detail our contracting assumptions, it is instructive to consider first the case of complete contracts in which the firm has full control over all investments. In such a case, the firm makes a contract offer $[x(j), s(j)]$ for every input $j \in [0,1]$ under which a supplier is obliged to supply $x(j)$ of compatible inputs as stipulated in the contract in exchange for the payments $s(j)$. The optimal contract solves the following maximization problem:

$$\begin{aligned} \max_{\{x(j),s(j)\}_{j\in[0,1]}} \quad & \pi_F = A^{1-\rho}\theta^{\rho}\left(\int_0^1 x(j)^{\alpha}\,dj\right)^{\rho/\alpha} - \int_0^1 s(j)\,x(j)\,di \\ s.t. \quad & s(j) \geq cx(j), \end{aligned}$$

which delivers a common investment level $x = \left(\rho A^{1-\rho}\theta^{\rho}/c\right)^{1/(1-\rho)}$ for all intermediate inputs and associated firm profits equal to $\pi = (1-\rho)\,A(\rho\theta/c)^{\rho/(1-\rho)}$, while leaving suppliers with a net payoff equal to their outside option (i.e., $s = cx$).

Incomplete Contracts Note that for the above contracts to be enforceable, it is important that a court of law be able to verify whether the inputs provided by the suppliers of the different stages are indeed compatible. In practice, a court of law will generally not be able to verify such compatibilities, and the firm will be reluctant to sign binding contracts that are *not* contingent on whether inputs are compatible or not, because suppliers would then have every incentive to produce incompatible inputs at a negligible cost and still demand payment. One could envision that contracts contingent on total revenues could provide investment incentives for suppliers, but in our setting with a continuum of suppliers, these type of contracts have no value as they would elicit zero investment levels. For these reasons, it is natural to study situations in which the terms of exchange between the firm and the suppliers are not disciplined by an ex-ante enforceable contract. In fact, the initial contract is assumed to specify *only* whether suppliers are vertically integrated into the firm or they remain independent (more on this below).

Given the lack of a binding contract, the payment to a particular supplier (say the supplier of stage m) is only negotiated once the stage m intermediate input has been produced and the firm has had a chance to inspect it. For the time being, we treat these bilateral negotiations independently from the bilateral negotiations at other stages, though we will revisit this assumption in section 3.A below. We model the negotiation at each stage m via the generalized Nash bargaining solution, with the share of the ex-post gains from trade going to the firm being a constant $\beta \in [0,1]$. Because the intermediate input produced by the supplier is compatible only with the firm's output, the supplier's outside option at the bargaining stage is 0 regardless of the stage m. Hence, the quasirents over which the firm and the supplier negotiate are given by the incremental contribution to total revenue generated by supplier m at that stage. Up to that stage m, the value of final-good production is given by

$$r(m) = A^{1-\rho}\theta^{\rho}\left[\int_0^m x(j)^{\alpha}dj\right]^{\frac{\rho}{\alpha}}, \tag{4}$$

so applying Leibniz's integral rule to this expression, we have:

$$r'(m) = \frac{\partial r(m)}{\partial m} = \frac{\rho}{\alpha}\left(A^{1-\rho}\theta^{\rho}\right)^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} x(m)^{\alpha}. \tag{5}$$

Following the property-rights theory of firm boundaries, we let the *effective* bargaining power of the firm vis-à-vis a particular supplier depend on whether the firm owns this supplier or not. To see this, note that when input m is outsourced, the supplier is an independent entity and the Nash bargaining solution leaves the firm with a fraction β of the incremental contribution in equation (5). Conversely, when the firm owns the supplier, we assume that in case of a contractual breakdown the firm is able to lay claim and recover a fraction δ of the customized input produced by supplier m, hence obtaining a fraction $\beta + (1-\beta)\,\delta^{\alpha} > \beta$ of the surplus in equation (5).

We now summarize the time line of the game played by the firm and the continuum of suppliers:[5]

- The firm posts contracts for suppliers for each stage $j \in [0,1]$ of the production process. The contract stipulates the organizational form – integration within the boundaries of the firm or arm's-length outsourcing – under which the potential supplier will operate.
- Suppliers apply for each contract, and the firm chooses one supplier for each production stage.
- Production takes place sequentially. At the beginning of each stage m, the supplier is handed the final good completed up to stage m. After observing the value of this unfinished product (i.e., $r(m)$ in (4)), the supplier chooses an input production level $x(m)$. At the end of each stage m, the firm and supplier m bargain over the addition to total revenue that supplier m has contributed (i.e., $r'(m)$ in (5)) and the firm pays the supplier.
- Output of the final good is realized once the final stage is completed. The total revenue, $A^{1-\rho}q^{\rho}$, from the sale of the final good is collected by the firm.

Before describing the equilibrium of this game, it is worth pausing to discuss further our assumptions regarding the sequential nature of contracting and payments. Notice, in particular, that we have assumed that the firm and the supplier bargain only at stage m and that the terms of exchange are not renegotiated at a later stage and do not reflect the outcome of subsequent negotiations. If negotiations with all suppliers occurred after the final stage of production, once all inputs have been produced, then it is straightforward to see that, given our symmetry assumptions, all suppliers would obtain the same payoff and the optimal organizational form would be independent

[5] Although we focus throughout on a version of the model with a continuum of production stages, our equilibrium corresponds to the limit $\varepsilon \to 0$ of a discrete game in which M suppliers each control a measurable range $\varepsilon = 1/M$ of the continuum of intermediate inputs. The details are available upon request. See Acemoglu et al. (2007) for an analogous derivation.

of the position of an input in the value chain (see Acemoglu et al, 2007).[6] Even when bargaining is bilateral and sequential, one could argue that with common knowledge of the technology in (1), suppliers might anticipate that their *ultimate* contribution to sale revenue is different from $r'(m)$ (on account of the effect of their investment on downstream stages), and they might insist that negotiations at stage m be based on a division of their *ultimate* contribution. In section 3.A, we show that our main results are surprisingly robust to specifying alternative bargaining protocols including the possibility of multilateral or forward-looking contracting.[7]

2.2 Equilibrium Firm Behavior

A. Supplier Investment in Stage m

We now characterize the subgame perfect equilibrium of the game described above. We start by solving for the investment level of a particular stage m supplier, taking as given the value of production up to that stage and the chosen organizational mode for that stage m. Denote by $\beta(m)$ the share of incremental revenue $r'(m)$ that accrues to the firm in the bargain. Our previous discussion implies that:

$$\beta(m) = \begin{cases} \beta_O = \beta & \text{if the firm outsources stage } m \\ \beta_V = \beta + (1-\beta)\,\delta^{\alpha} & \text{if the firm integrates stage } m \end{cases} .$$

The stage-m supplier obtains the remaining share $1-\beta(m) \in [0,1]$ of $r'(m)$ and thus chooses an investment level $x(m)$ to solve:

$$\max_{x(m)} \quad \pi_S = (1-\beta(m))\,\frac{\rho}{\alpha}\left(A^{1-\rho}\theta^{\rho}\right)^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} x(m)^{\alpha} - cx(m), \tag{6}$$

which delivers:

$$x(m) = \left[(1-\beta(m))\,\frac{\rho\left(A^{1-\rho}\theta^{\rho}\right)^{\frac{\alpha}{\rho}}}{c}\right]^{\frac{1}{1-\alpha}} r(m)^{\frac{\rho-\alpha}{\rho(1-\alpha)}} . \tag{7}$$

The investment made by supplier m is naturally increasing in the demand level A and productivity θ of the final good, and in the supplier's bargaining share, $1-\beta(m)$, while it decreases in the investment marginal cost c. Hence, other things equal, an outsourcing relationship at stage m promotes higher investments on the part of supplier m. The effect of the value of production up to stage m (and thus of all investment decisions in prior stages, $\{x(j)\}_{j=0}^{m}$) is more subtle. If $\rho > \alpha$, then investment choices are *sequential complements* in the sense that higher investment levels by

[6] When incompatibilities at a given stage m have implications for the productivity of downstream investments (something we have ruled out so far), multilateral negotiations may deliver asymmetric payoffs across suppliers (see section 3.A for details).

[7] One way to justify the sequential nature of contracting that we impose in this section is by appealing to credit constraints stemming from limited commitment frictions (as in Hart and Moore, 1994, or Thomas and Worrall, 1994). Intuitively, suppliers might be reluctant to agree to be paid (long) after they have delivered their inputs if they fear that the firm might at that point renege on any previously agreed transfers.

prior suppliers, as summarized by $r(m)$ increase the marginal return of supplier m's investment. Conversely, if $\rho < \alpha$, investment choices are *sequential substitutes* because high values of upstream investments reduce the marginal return to investing in $x(m)$. Throughout the paper, we shall refer to $\rho > \alpha$ as the *complements* case and to $\rho < \alpha$ as the *substitutes* case.

Because $\alpha \in (0,1)$, it is straightforward to verify that from a purely technological point of view, supplier investments are always complementary. More precisely, in light of equation (1), $\partial q/\partial x(m)$ is necessarily increasing in the investment decisions of other suppliers $m' \neq m$. Why is then $x(m)$ negatively affected by prior investments when $\rho < \alpha$? The key behind this result is that, when $\rho < 1$, the firm faces a downward sloping demand curve for its product and thus prior upstream investments also affect $x(m)$ on account of the induced movements along the demand curve. When ρ is very small, the firm's revenue function is highly concave in quality-adjusted output and thus marginal revenue falls at a relatively fast rate along the value chain. In other words, in industries where firms enjoy significant market power, large upstream investment levels can significantly reduce the value of undertaking donwstream investments, thus effectively turning supplier investments into sequential substitutes. Equation (7) illustrates that this effect will dominate the standard physical output complementarity effect whenever the elasticity of demand faced by the firm is lower than the elasticity of substitution across inputs, or $\rho < \alpha$.

B. Suppliers' Investments Along the Value Chain

Equation (7) characterizes supplier m's investment level as a function of $r(m)$, i.e., the value of production up to stage m. We next solve for $r(m)$ as a function of the primitives of the model and obtain the equilibrium investment levels of all suppliers along the value chain. To achieve this, first plug equation (7) into (5) to obtain:

$$r'(m) = \frac{\rho}{\alpha}\left(\frac{(1-\beta(m))\,\rho\theta}{c}\right)^{\frac{\alpha}{1-\alpha}} A^{\frac{\alpha(1-\rho)}{\rho(1-\alpha)}} r(m)^{\frac{\rho-\alpha}{\rho(1-\alpha)}}. \tag{8}$$

This constitutes a differential equation in $r(m)$, which is easily solved by noting that it is separable in $r(m)$ and $\beta(m)$ and making use of the initial condition $r(0) = 0$:

$$r(m) = A\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}\left(\frac{\rho\theta}{c}\right)^{\frac{\rho}{1-\rho}}\left[\int_0^m (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj\right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}. \tag{9}$$

Equation (9) illustrates how the value of stages up to stage m depends on all upstream organizational decisions (as captured by the $\beta(j)$ for $j < m$). Finally, plugging this solution into (7) produces:

$$x(m) = A\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}}\left(\frac{\rho}{c}\right)^{\frac{1}{1-\rho}}\theta^{\frac{\rho}{1-\rho}}(1-\beta(m))^{\frac{1}{1-\alpha}}\left[\int_0^m (1-\beta(j))^{\frac{\alpha}{1-\alpha}}\, dj\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}. \tag{10}$$

From this expression, it is clear that outsourcing of stage m (i.e., choosing $\beta(m) = \beta_O < \beta_V$)

enhances investment by that stage's suppliers, while the dependence of $x(m)$ on the prior (upstream) organizational choices of the firm crucially depends on whether investment decisions are sequential complements ($\rho > \alpha$) or sequential substitutes ($\rho < \alpha$). In choosing its optimal organizational structure, the firm will weigh these considerations together with the fact that outsourcing of any stage is associated with capturing a lower share of surplus and thus extracting less surplus from suppliers. We next turn to study this optimal organizational structure.

C. Optimal Organizational Structure

The firm seeks to maximize the amount of sale revenue it obtains when the good is sold net of all payments made to suppliers along the value chain, or:

$$\pi_F = A\frac{\rho}{\alpha}\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}}\left(\frac{\rho\theta}{c}\right)^{\frac{\rho}{1-\rho}}\int_0^1 \beta(j)(1-\beta(j))^{\frac{\alpha}{1-\alpha}}\left[\int_0^j (1-\beta(k))^{\frac{\alpha}{1-\alpha}}dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \qquad (11)$$

It is easily verified that the payoff π_S obtained by suppliers (see equation (6)) is always strictly positive, so their participation constraint can be ignored.

The firm's decision problem is then to choose the values of $\beta(j) \in \{\beta_V, \beta_O\}$ for $j \in [0,1]$ that maximize π_F. For a given stage m, in order to determine if integration or outsourcing is optimal, it proves useful to follow the approach in Antràs and Helpman (2004, 2008) and consider the hypothetical case in which the firm could freely choose $\beta(m)$ from the continuum of values in $[0,1]$. After some simplification, the partial derivative of π_F with respect to $\beta(m)$ can be written as

$$\frac{\partial \pi_F}{\partial \beta(m)} = (1-\beta(m))^{\frac{\alpha}{1-\alpha}-1}\,\theta^{\frac{\rho}{1-\rho}}\kappa\Phi(m),$$

where:

$$\Phi(m) = \left(1-\frac{\beta\left(m\right)}{1-\alpha}\right)\left[\int_0^m (1-\beta(k))^{\frac{\alpha}{1-\alpha}}dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} - \frac{\rho-\alpha}{\left(1-\alpha\right)\left(1-\rho\right)}\int_m^1 \beta(i)(1-\beta(i))^{\frac{\alpha}{1-\alpha}}\left[\int_0^i (1-\beta(j))^{\frac{\alpha}{1-\alpha}}dj\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} di. \qquad (12)$$

and κ is a positive constant.[8]

The function $\Phi(m)$ captures how the firm's incentives to raise or lower the bargaining share $\beta\left(m\right)$ depends on the particular production stage $m \in [0,1]$ under consideration. It consists of two terms. The first term captures the balance of the *rent extraction* and *stage-m incentive* effects of a higher bargaining power for the firm (higher $\beta\left(m\right)$). Intuitively, a higher $\beta\left(m\right)$ increases the rents captured by the firm at stage m but also reduces the incentives to invest of the supplier providing stage m. Absent other considerations, the optimal bargaining share would be given by $\beta^*\left(m\right) = 1-\alpha$, and

[8] In particular, $\kappa = A\frac{\rho}{\alpha}\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}}\left(\frac{\rho}{c}\right)^{\frac{\rho}{1-\rho}}$.

would decrease in α due to the fact that the elasticity of a supplier m's investment in response to a reduction in this agent's bargaining share is higher the more substitutable are inputs along the value chain (see equation (10)). As is clear from equation (12), the size of these effects depends on the stage m under consideration because the incremental surplus over which the firm and the supplier negotiate varies along the value chain: as argued above, it increases with m in the complements case and it decreases in m in the substitutes case.

The second term in (12) represents the *downstream incentive* effect of an increase in $\beta(m)$, namely its effects on the incentives to invest of all suppliers that are positioned downstream relative to supplier m. As explained in the discussion of equation (10), the effect is negative in the complements case, but it is positive in the substitutes case because in the latter case, a lower investment at stage m will increase the incentives to invest of downstream suppliers. Importantly, this last effect also interacts with m as the absolute value of the effect will naturally be higher when m is low, i.e., when a given stage precedes a relatively large number of other production stages.

Setting $\Phi(m) = 0$, we then have that the (hypothetical) optimal division of surplus at stage m, which we denote by $\beta^*(m)$, is given by:

$$\beta^*(m) = \max\left\{1 - \alpha - \frac{\rho - \alpha}{1 - \rho} \frac{\int_m^1 \beta(i)(1 - \beta(i))^{\frac{\alpha}{1-\alpha}} \left[\int_0^i (1 - \beta(j))^{\frac{\alpha}{1-\alpha}} dj\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)} - 1} di}{\left[\int_0^m (1 - \beta(k))^{\frac{\alpha}{1-\alpha}} dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}}, 0\right\}. \tag{13}$$

This expression may seem complicated, but straightforward differentiation indicates that (see Appendix for a proof):

Lemma 1 The (unconstrained) optimal bargaining share $\beta^*(m)$ is a weakly increasing function of m in the complements case ($\rho > \alpha$) while it is a decreasing function of m in the substitutes case ($\rho < \alpha$).

In words, Lemma 1 states that whether the incentive for the firm to retain a larger surplus share increases or decreases along the value chain crucially depends on the relative size of the parameters ρ and α, which in turn govern the effects discussed above. Intuitively, although when ρ is high relative to α, integrating early stages of production seems relatively appealing because it leads to relatively small underinvestment problems by these upstream suppliers, this effect is dominated by the fact that when inputs are highly complementary, (i) integration decisions early in the value chain are particularly costly because they reduce the incentives to invest not only of these early suppliers but also of all suppliers downstream, and (ii) the incremental surplus over which the firm and the supplier negotiate is particularly small in these early stages of production. Conversely, when ρ is small relative to α, outsourcing is particularly costly in upstream stages because high investments in the early stages of the value chain lead to reduced incentives to invest for downstream suppliers, while the firm captures a disproportionate amount of surplus by integrating these early stages.

Evaluating the function $\beta^*(m)$ at both extremes of the value chain, we obtain that $\lim_{m \to 0} \beta^*(m) = 0$ when $\rho > \alpha$ and $\lim_{m \to 0} \beta^*(m) = 1$ when $\rho < \alpha$, while $\beta^*(1) = 1 - \alpha$ regardless of the relative

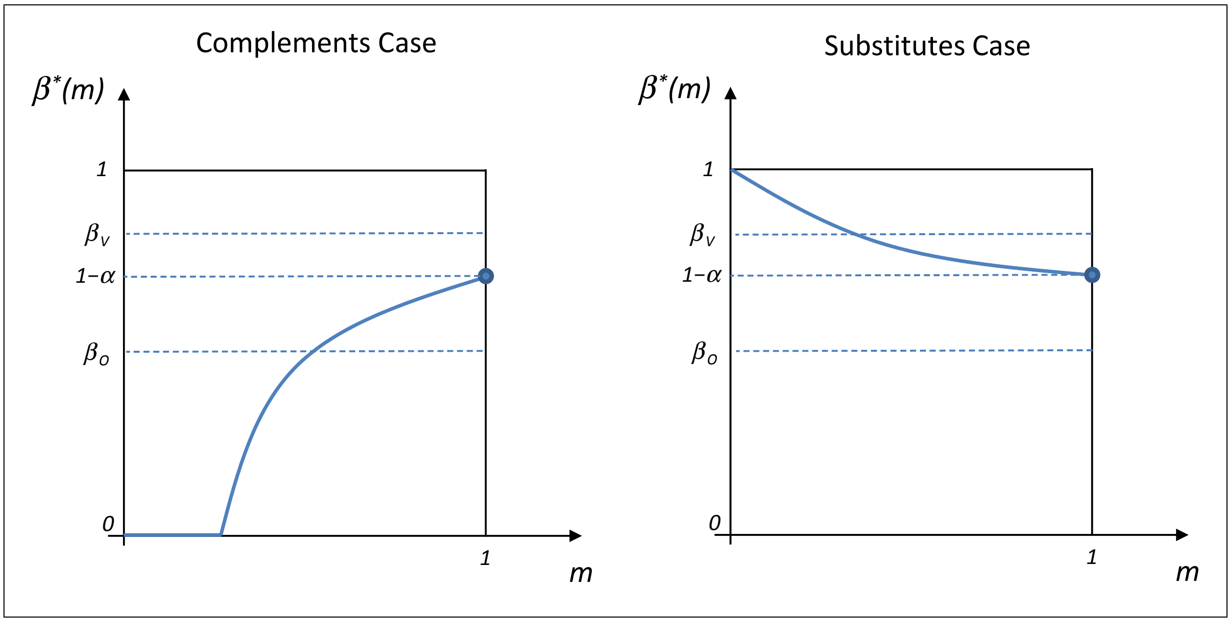

Figure 2: Profit-Maximizing Division of Surplus for Stage m

magnitude of α and ρ (see the proof of Proposition 2 in the Appendix). This implies that when the firm is constrained to choose $\beta(m)$ from the pair of values β_V and β_O, the decision of whether or not to integrate the most upstream stages crucially depends on the relative size of ρ and α. In the complements case ($\rho > \alpha$), the firm would like to select the minimum possible value of $\beta(m)$ at $m = 0$, which corresponds to choosing outsourcing in this initial stage and, by continuity, in a measurable set of the most upstream stages. Conversely, in the substitutes case ($\rho < \alpha$), the firm necessarily chooses to integrate these initial stages. As for the most downstream stages of the value chain, the decision is less clear-cut. In both cases, if $\beta_V < 1 - \alpha = \beta^*(1)$, then it is clear that that last stage will be integrated, while it will necessarily be outsourced if $\beta_O > 1 - \alpha$. When $\beta_V > 1 - \alpha > \beta_O$, whether that stage is integrated or not depends on other parameter values. In sum, when $\rho > \alpha$ the firm will necessarily outsource relatively upstream inputs, while it may (depending on parameter values) find it optimal to integrate the most downstream inputs. Conversely, when $\rho < \alpha$ the firm will necessarily integrate relatively upstream inputs; and may (depending on parameter values) outsource the most downstream inputs. Figure 2 depicts the function $\beta^*(m)$ whenever $\beta_V > 1 - \alpha > \beta_O$, in which case there is the potential for integrated and outsourced stages to coexist along the value chain in both the sequential complements and substitutes cases.

Our discussion so far has focused on the optimal organization of production for productions stages at both ends of the value chain. In the Appendix, we show that in fact the set of stages under a common organizational form (integration or outsourcing) is necessarily a connected interval in $[0, 1]$, thus implying that:

Proposition 2 In the complements case ($\rho > \alpha$), there exists a unique $m_C^* \in (0,1]$, such that: (i) all production stages $m \in [0, m_C^*)$ are outsourced; and (ii) all stages $m \in [m_C^*, 1]$ are integrated within firm boundaries. In the substitutes case ($\rho < \alpha$), there exists a unique $m_S^* \in (0,1]$, such that: (i) all production stages $m \in [0, m_S^*)$ are integrated within firm boundaries; and (ii) all stages $m \in [m_S^*, 1]$ are outsourced.

Taking into account that $\beta(m)$ only takes at most two values along the value chain, one can easily derive a closed-form solution for the cut-off stages, m_C^* and m_S^*, in terms of the parameters β, δ, α and ρ (see Appendix for details):

$$m_C^* = \left[1 + \left(\frac{1-\beta}{1-\beta_V}\right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1-\frac{\beta}{\beta_V}}{1-\left(\frac{1-\beta}{1-\beta_V}\right)^{-\frac{\alpha}{1-\alpha}}}\right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1\right]\right]^{-1} \tag{14}$$

and

$$m_S^* = \left[1 + \left(\frac{1-\beta_V}{1-\beta}\right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{\left(\frac{1-\beta_V}{1-\beta}\right)^{-\frac{\alpha}{1-\alpha}} - 1}{\frac{\beta_V}{\beta} - 1}\right)^{\frac{\alpha(1-\rho)}{\alpha-\rho}} - 1\right]\right]^{-1}, \tag{15}$$

where remember that $\beta < \beta_V \equiv \beta + (1-\beta)\delta^{\alpha}$. Differentiating these equations, we can then establish:

Proposition 3 Whenever integration and outsourcing coexist along the value chain (i.e., $m_C^* \in (0,1)$ when $\rho > \alpha$ or $m_S^* \in (0,1)$ when $\rho < \alpha$), a decrease in ρ or β will necessarily expand the range of stages that are vertically integrated.

The negative effect of ρ on integration is explained by the fact that when the firm has relatively high market power (low ρ), it will tend to place a relatively high weight on the rent-extraction motive for integration and will be less concerned with the investment inefficiencies caused by such integration. As for the effect of β, note that a higher β gives the firm a stronger bargaining position under both outsourcing and integration, which tends to disincentivize supplier investments across all stages. In order to counteract this effect, it is optimal in both the complements and substitutes cases to increase the range of stages that are outsourced.[9]

3 Extensions and Empirical Implementation

Our benchmark model is special among several directions and omits many factors that have been shown to be important determinants of the organizational decisions of firms in the global economy. In this subsection, we outline four extensions of our benchmark model that help us further connect our framework to the global sourcing framework in Antràs and Helpman (2004, 2008), thus

[9] The comparative statics of the cutoff stages with respect to δ and α are in general ambiguous.

illustrating that the main theoretical predictions continue to hold in richer frameworks. This will also serve as a justification for some of the control variables that we will include in the empirical section. For simplicity, we develop these extensions one at a time, though they could readily be incorporated in a unified framework.

A. Alternative Bargaining Protocols

[TO BE WRITTEN]

B. Headquarter Intensity

We next consider the introduction of investment decisions on the part of the firm. As first discussed by Antràs (2003), to the extent that final-good producers or 'headquarters' undertake significant noncontractible, relationship-specific investments in production, their willingness to give up bargaining power via outsourcing will be tampered by the negative effect of those decisions on the provision of headquarter services. The relative intensity of headquarter services in production thus emerges as a crucial determinant of the integration decision (see also Antràs and Helpman, 2004, 2008). It is straightforward to incorporate headquarter services into our framework. In particular, consider the case in which the production function is modified from (1) to:

$$q = \theta \left(\frac{h}{\eta} \right)^{\eta} \left(\int_0^1 \left(\frac{x(j)}{1-\eta} \right)^{\alpha} dj \right)^{\frac{1-\eta}{\alpha}}, \tag{16}$$

and the provision of headquarter services by the firm is undertaken at marginal cost c_h after suppliers have been hired but before they have undertaken any investments. For instance, one could think of these headquarter services as R&D or managerial inputs that need to be performed before the sourcing of inputs along the supply chain can commence. As in the case of the investments by suppliers, we assume that ex-ante contracts on headquarter services are not enforceable.

Because the investment in h is sunk by the time that the firm sequentially bargains with suppliers, the introduction of headquarter intensity does not alter the above analysis too much. In particular, note that revenue from the final good when only inputs up to stage m are compatible and the rest are incompatible is now given by:

$$r(m) = A^{1-\rho} \theta^{\rho} \left(\frac{h}{\eta} \right)^{\rho\eta} (1-\eta)^{-\tilde{\rho}} \left[\int_0^m x(j)^{\alpha} dj \right]^{\frac{\tilde{\rho}}{\alpha}}$$

where $\tilde{\rho} \equiv (1-\eta)\,\rho$. It is then immediate that one can follow the same steps as in previous sections to conclude that the dependence of the integration decision on the index of a production stage m crucially depends on the relative magnitude of $\tilde{\rho} \equiv (1-\eta)\,\rho$ and α. As before, a high value of ρ relative to α leads to a higher desirability of integrating relatively downstream production stages, while the converse is true when ρ is low relative to α. What this extension illustrates is that the effect of ρ on the desirability of integrating relatively downstream or upstream stages should interact with the headquarter intensity of the industry. In particular, it should be highest for low levels of headquarter intensity.

Beyond this effect, our model also predicts that a higher headquarter intensity (higher η) will also have a positive "level" effect (across all stages) in the integration decision, for reasons very much analogous to those laid out in previous contributions to the property-rights theory. To illustrate this formally, notice that Propositions 2 and 3 will continue to hold with $\tilde{\rho} \equiv (1-\eta)\,\rho$ replacing ρ both in the statements of the Propositions as well as in the formulas for m_C^* and m_S^*. Hence, whenever our model predicts a coexistence of integration and outsourcing along the value chain, an increase in η will necessarily expand the range of stages that are vertically integrated.[10]

We summarize these results as follows (see Appendix for a formal proof):

Proposition 4 In the presence of headquarter services provided by the firm, the results in Propositions 2 and 3 continue to hold except for the fact that: (i) the complements and substitutes cases are now defined by $\tilde{\rho} \equiv (1-\eta)\,\rho > \alpha$ and $\tilde{\rho} \equiv (1-\eta)\,\rho < \alpha$, respectively, and (ii) the range of stages that are vertically integrated is now also (weakly) increasing in η.

C. Firm Heterogeneity and Prevalence of Integration

Up to now, we have considered the problem of a single firm in isolation. We now consider the equilibrium of an industry in which firms produce a continuum of differentiated final good varieties that consumers value according to (2).

On the technology side, each firm produces one final-good variety under the same technology in (1) and the same sequencing of production stages, but we let them differ in their productivity parameter θ, which is drawn from a Pareto distribution with shape parameter z and minimum threshold $\underline{\theta}$, namely:

$$G(\theta) = 1 - (\underline{\theta}/\theta)^z \quad \text{for } \theta \geq \underline{\theta} > 0, \tag{17}$$

where z is inversely related to the variance of $G(\cdot)$ and is assumed high enough to ensure a finite variance of the size distribution of firms. Finally, we introduce a fixed organizational cost $f(j)$ associated with each production stage $j \in [0,1]$. For simplicity, we let the firm pay these fixed costs (or a large enough fraction of them to ensure that no supplier's participation constraint is violated). In the spirit of our symmetry assumptions above, we assume that these fixed costs are symmetric for all stages, but we allow them to vary with the stage's organizational structure. More specifically, and following the arguments in Antràs and Helpman (2004), we assume that:

$$f_V > f_O,$$

reflecting the relatively high managerial overload associated with running an integrated value chain.

The presence of productivity heterogeneity and fixed costs of production enriches the choice

[10] The counterpart of this result is that that the unconstrained optimal bargaining share $\beta^*(m)$ spelled out in (13) is decreasing in $\tilde{\rho}$ in both the complements and substitutes cases, which implies a greater propensity to integrate each stage m the higher is η.

of ownership structure relative to our Benchmark Model. We relegate most mathematical details to the Appendix and focus here on describing the main results. Consider first the complements case in which $\rho > \alpha$. As in the Benchmark Model, the incentive for the firm to integrate a given production stage is larger the more downstream the stage is located, and again there exists a threshold $m_C \in (0, 1]$ such that all stages below m_C are outsourced and all stages above m_C (if any) are integrated. When $m_C < 1$, this threshold is now implicitly defined by:

$$(m_C)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left[\left(1 - \frac{\beta}{\beta_V}\right) - \left(1 - \left(\frac{1-\beta_V}{1-\beta}\right)^{\frac{\alpha}{1-\alpha}}\right) \left[1 + \left(\frac{1-\beta_V}{1-\beta}\right)^{\frac{\alpha}{1-\alpha}} \left(\frac{1-m_C}{m_C}\right)\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \right] = \frac{f_V - f_O}{\Psi \theta^{\frac{\rho}{1-\rho}}}, \tag{18}$$

where Ψ is a constant.[11]

Note that the left-hand-side of (18) is increasing in m_C, and thus the threshold m_C is now a decreasing function of the level of firm productivity θ. Intuitively, relatively more productive firms will be find it relatively easier to amortize the extra fixed cost associated with integrating stages and thus will tend to integrate a larger number of stages. Furthermore, when $\theta \to \infty$, the effect of fixed costs on firm profits becomes negligible and the threshold m_C converges to the one in the Benchmark Model (i.e., m_C^* in equation (14)). Following analogous steps (see the Appendix), it is straightforward to verify that in the substitutes case ($\rho < \alpha$), there exists again a threshold $m_S \in (0, 1]$ such that all stages upstream from m_S are integrated and all stages downstream from m_S (if any) are outsourced. Furthermore, m_S is increasing in firm productivity θ, so again relatively more profitable firms tend to integrate a larger interval of production stages.

Figure 3 illustrates these results. In both panels of the Figure it is assumed that the firms with the lowest values of productivity (in the neighborhood of $\underline{\theta}$) do not find it profitable to integrate any production stage m.[12] As productivity increases, more and more stages become integrated: with those stages being the most downstream ones in the complements case but the most upstream ones in the substitutes case. Furthermore, both panels illustrate that even when productivity becomes arbitrarily large, the firm might want to keep some production stages (the most upstream ones in the complements case and the most downstream ones in the substitutes case) under an outsourcing contract.

The introduction of productivity heterogeneity generates smooth predictions for the prevalence of integration in production stages with different indices m, a feature that will facilitate our transition to the empirical analysis in the next section. More specifically, notice that in the complements case $\rho > \alpha$, we have that input $m > m_C^*$ will be integrated by all firms with productivity higher than the threshold $\theta_C(m)$ such that equation (18) holds, and will be outsourced by all other firms. Inputs with an index $m < m_C^*$ will not be integrated by any firms. Appealing to the Pareto

[11] In particular, $\Psi = \beta_V A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{(1-\beta)\rho\theta}{c}\right)^{\frac{\rho}{1-\rho}}$.

[12] We assume throughout that f_O is low to ensure that the firms with the lowest productivity level $\underline{\theta}$ find it profitable to outsource all stages.

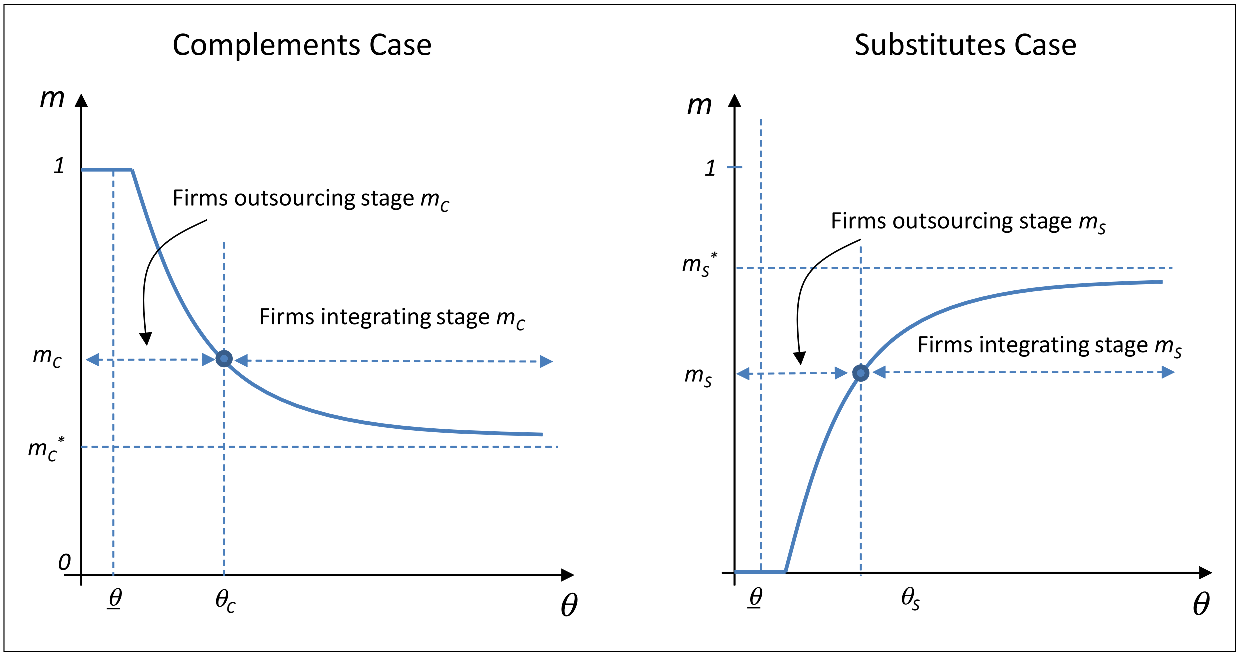

Figure 3: Firm Heterogeneity and the Integration Decision

distribution in (17), we have that the share of firms integrating stage m is given by

$$\sigma_C(m) = \begin{cases} 0 & \text{if } m \leq m_C^* \\ \left(\underline{\theta}/\theta_C(m)\right)^z & \text{if } m > m_C^* \end{cases}. \tag{19}$$

From our previous discussion, $\theta_C(m)$ is a decreasing function of m and thus the share of firms integrating stage m is weakly increasing in the downstreamness of that stage. Notice also that because $\underline{\theta} < \theta_C(m)$, the share of integrating firms is decreasing in z and thus increasing in productivity dispersion, a result that very much resonates those derived by Helpman et al. (2004) and Antràs and Helpman (2004).

Following analogous steps for the substitutes case, we can conclude that:

Proposition 5 The share of firms integrating a particular stage m is weakly increasing in the downstreamness of that stage in the complements case ($\rho > \alpha$), while it is decreasing in the downstreamness of the stage in the substitutes case ($\rho < \alpha$). Furthermore, the share of firms integrating a particular stage m, is weakly increasing in the dispersion of productivity within the industry.

Proposition 5 converts our previous results on within-firm variation on the propensity to integrate different stages into predictions regarding the relative prevalence of integration of an input with a particular level of downstreamness when considering the decisions of all firms within an industry. This is an important step because our empirical application uses product-level data that on intrafirm trade that aggregates various firm level decisions.

It is worth stressing that the modelling of final-good producer heterogeneity highlights that, to the extent that fixed costs of integration are relatively high, the set of stages that will be integrated by final-good producers will be relatively low. In such a case, our model predict that, in the sequential complements case, only a few very downstream stages will be integrated, while in the sequential substitutes case, only a few very upstream stages will be integrated. We will come back to this observation in our empirical section.

D. Input and Supplier Heterogeneity

So far, we have assumed that the only source of asymmetry across production stages is their level of downstreamness. In particular, we have assumed that all inputs enter symmetrically into production and that their production entails a common marginal cost c. In the real world, different production stages have different effects on output, suppliers differ in their productivity levels, and the widespread process of offshoring also implies that firms undertaking different stages of production in different countries will typically face different costs of production, even for a common productivity. For these reasons, it is important to assess the robustness of our results to the existence of asymmetries across suppliers.

To that end, we next consider a situation in which the volume of final-good production is now given by,

$$q = \theta \left(\int_0^1 \left(\psi\left(j\right) x(j) \right)^{\alpha} dj \right)^{1/\alpha}, \tag{20}$$

where $\psi\left(j\right)$ captures asymmetries in the marginal product of different inputs, while the marginal cost of production of input j is given by $c(j)$ and can vary across inputs due to productivity differences or due to the heterogeneity in factor costs across the locations in which inputs are produced.

Following the same steps as in our benchmark model, we find that the profits the firm obtains in such a case are given by

$$\pi_F = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\rho\theta\right)^{\frac{\rho}{1-\rho}} \int_0^1 \beta(j) \left(\frac{1-\beta(j)}{c\left(j\right)/\psi\left(j\right)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^j \left(\frac{1-\beta(k)}{c\left(k\right)/\psi\left(k\right)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj, \tag{21}$$

which is analogous to equation (11) except for the inclusion of input asymmetries as captured by the term $c\left(m\right)/\psi\left(m\right)$ for input m. How do these asymmetries affect the firm's choice of ownership structure $k(m) \in \{V, O\}$ for each stage $m \in [0,1]$? To build intuition it is useful again to treat the choice of $\beta(j)$ as that of a continuous variable in $[0,1]$. After some simplifications analogous to those performed in the benchmark model, setting the partial derivative of π_F with respect to $\beta(m)$

equal to 0 delivers the following expression for the optimal division of surplus $\beta^*(m)$:

$$\beta^*(m) = \max\left\{1-\alpha-\frac{\rho-\alpha}{1-\rho}\frac{\int_m^1 \beta(i)\left(\frac{1-\beta(i)}{c(i)/\psi(i)}\right)^{\frac{\alpha}{1-\alpha}}\left[\int_0^i \left(\frac{1-\beta(j)}{c(j)/\psi(j)}\right)^{\frac{\alpha}{1-\alpha}}dj\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1}di}{\left[\int_0^m \left(\frac{1-\beta(k)}{c(k)/\psi(k)}\right)^{\frac{\alpha}{1-\alpha}}dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}},0\right\}.$$

It can be easily verified that despite the presence of heterogenous marginal products and marginal costs along the value chain, Lemma 1 continues to apply in this richer framework and the sign of sign of the derivative of $\beta^*(m)$ with respect to m is again given by the sign of $\rho-\alpha$. The intuition for how the optimal allocation of bargaining power varies with the stage of production m is also analogous to that in the benchmark model. Furthermore, Proposition 2 continues to apply, though one can no longer solve for the thresholds m_C^* and m_S^* in closed form.

An implication of this result, is that when studying the global sourcing decisions of a firm, our model continues to predict that the level of downstreamness of an input should be a relevant determinant of the integration decision of that input. Furthermore, when embedding the model in the industry equilibrium structure described in section 3.C., Proposition 5 continues to apply even when firms face heterogenous costs for their inputs and we can thus state (see Appendix):

Proposition 6 Suppose that technology allows for input heterogeneity as in (20) and that marginal costs of production of inputs are also heterogeneous and given by $c(j)$ for $j \in [0,1]$. Then the share of firms integrating a particular stage m is weakly increasing in the downstreamness of that stage in the complements case ($\rho > \alpha$), while it is decreasing in the downstreamness of that stage in the substitutes case ($\rho < \alpha$). Furthermore, the share of firms integrating a particular stage m, is weakly increasing in the dispersion of productivity within the industry.

There is one important caveat to the above result. We have treated the marginal cost parameters, $c(j)$, as exogenous and uncorrelated with m, while in reality they are partly shaped by the endogenous location decisions of firms. To the extent that these location decisions are also shaped by donwstreamness in a systematic way, the above comparative static results regarding the effect of m on the integration decision become more complex because the level of marginal costs might correlate with m. Although this is not the focus of this paper, an analysis of the optimal location of each stage of production and how it varies with the position of that stage in the value chain can be carried out by computing the partial derivative of the profit function in (21) with respect to the marginal cost $c(m)$ of a given stage m. Two results emerge from such an analysis (see the Appendix for details). First, we find that $\partial\pi_F/\partial c(m) < 0$, and thus, other things equal, it is always optimal to choose the minimum cost location for each stage. Second, one can show that the marginal incentives for the firm to reduce the marginal cost of a stage are generally affected by the index of the production stage m, but such dependence is non-monotonic and depends on various aspects of the model in rather subtle ways.

In the empirical section of the paper, we propose two ways to address this caveat. First, we experiment with empirical specifications that exploit both cross-sectoral and cross-country variation in the prevalence of integration, but we introduce country fixed effects to ensure that the effect of downstreamness we identify is not estimated off cross-country variation in production costs. Still, this does not address a potential selection bias related to the fact that certain inputs might not be sourced at all from certain destinations due precisely to their level of downstream. To deal with this concern, we also experiment with a two-stage Heckman correction specification. We will provide more details on our empirical methodology in the next section.

4 Implementing an Empirical Test

The model that we have developed is a model of firm organizational decisions, and thus firm-level data would appear to be the ideal laboratory to use in testing it. Nevertheless, firm-level data on the integration decisions of firms is not readily available, and the datasets that have been employed in testing theories of multinational firm boundaries do not provide a sufficiently rich picture of the heterogeneous sourcing decisions of firms for a large number of inputs.[13] Our approach will instead exploit product-level variation in the extent to which products are transacted across borders within or outside of firm boundaries. Although our framework has implications as well for domestic sourcing decisions, data on international transactions are particularly accessible due to the widespread existence of official records of goods and services crossing borders.

In what follows, we describe our empirical strategy based on detailed data on U.S. intrafirm imports. More precisely, we will use this data to test the prediction highlighted in Proposition 5, namely that the relative prevalence of vertical integration of a particular input when aggregating across the decisions of all final-good producers purchasing that input (see equation (19)), should be a function of the average position of that input's use in the value chain, with the sign of that relationship being crucially determined by the relative size of ρ and α. Needless to say, implementing such a test requires that we propose appropriate proxies for the downstreamness of an input's use, as well as a means to distinguish between the sequential complements ($\rho > \alpha$) and sequential substitutes ($\rho < \alpha$) cases. We carefully describe the construction of these variables below. (Additional details on the industry concordances and other control variables are documented in the Data Appendix.)

A. Intrafirm Import Share

Starting with the key dependent variable, we follow the recent literature in using information on intrafirm trade to capture the propensity to transact a particular input within firm boundaries. While our model applies to closed-economy environments as well, data on the domestic sourcing practices of firms is scant and tends not be representative of the whole manufacturing process.

[13] See Antràs (2011) for a discussion of three firm-level datasets (from Japan, France, and Germany) that have been used to test the property-rights theory of the boundaries of multinational firms.

Instead, we draw our data from the U.S. Census Bureau's Related Party Trade Database.[14] The data report information on U.S. import and export values at the industry- and country-level, and more importantly, breaks down the value of trade according to whether it was conducted with related or non-related parties. A related party is defined as a foreign counterpart in which the U.S. importer has at least a 6 percent equity interest.[15] Some transaction values are recorded under a third category ("unreported"), where the precise nature of the trade flows was not reported or otherwise could not otherwise be determined. But this constitutes a very small share of trade flows and is a reason of concern only for industries with small trade volumes, where there is reason to believe that "unreported" flows might contribute to measurement error in the share of intrafirm trade that we calculate below.[16] We focus the analysis on the U.S. import data, given the position of the U.S. as a large downstream user of intermediates and consumer of finished goods from the rest of the world. We had an extensive amount of data, drawn from the years 2000-2010.

Our measure of the propensity of U.S. producers to integrate foreign suppliers of particular industrial goods is the share of related-party imports in total U.S. imports of that good, or *(Related Trade)/(Related Trade + Non-Related Trade)*. We will often refer to this measure as simply the *share* of intrafirm imports, and we can calculate this both at the industry-year level and at the exporting country-industry-year level. The publicly available U.S. Census data on intrafirm trade is reported at the sixit-digit level North American Industry Classification System (NAICS). In order to facilitate the merging of the data with other industry variables (especially our measures of downstreamness), we convert the related party trade data from NAICS to 2002 BEA Input-Output codes (IO2002) using the concordance provided by the Bureau of Economics Analysis (BEA), before calculating the intrafirm import share. As illustrated by Antràs (2011), there is a large amount of variation in this U.S. intrafirm import share: it varies widely across products and origin countries, and there also exists significant variation in that share across products within exporting countries and across exporting countries within narrowly-defined products.

B. Downstreamness

Our paper emphasizes a novel variable – the relative location of an industry along the value chain – as a determinant of the propensity to source inputs from that industry from affiliated versus unaffiliated parties. We propose two alternative measures that attempt to capture the "downstreamness" of an industry in manufacturing production processes. As we do not have information on the sequencing of stages for specific production technologies, we instead turn to the 2002 Input-Output

[14] This dataset has recently been used by various authors, including Nunn and Trefler (2008a,b), Bernard et al. (2010), and Díez (2010). Antràs (2011) discusses in detail the pros and cons of using this data source in testing theories of firm boundaries. We will overview some of the limitations of our approach at the end of this section.

[15] This number is lower than the conventional 10 percent cutoff used by the IMF to determine whether a foreign ownership stake qualifies as FDI, and is also lower than the cutoff equity stake that would typical convey control. Nevertheless, extracts from the confidential direct investment dataset collected by the BEA suggest that intrafirm trade is generally associated with one of the entities having a controlling stake in the other entity.

[16] The share of U.S. imports with unreported information on whether the parties are related or not is below 3% for all years in our sample. Using the detailed six-digit Harmonized System data from Nunn and Trefler (2008a,b) we find a negative correlation of -0.547 between the log of the share of unreported trade and the log of total imports.

Tables to obtain *average* measures of the relative position of each industry in U.S. manufacturing production processes.

To build intuition on these measures, consider the following identity

$$Y_i = F_i + Z_i,$$

where Y_i is the total output in sector i, F_i is the production of output for "final use" in that sector, and Z_i reflects the use of sector i's output as inputs to other industries (or its "total use' as inputs). Note that in a world with N industries or sectors, this expression can be expanded as follows:

$$Y_i = F_i + \underbrace{\sum_{j=1}^{N} d_{ij} F_j}_{\text{direct use of } i \text{ as input}} + \underbrace{\sum_{j=1}^{N}\sum_{k=1}^{N} d_{ik} d_{kj} F_j + \sum_{j=1}^{N}\sum_{k=1}^{N}\sum_{l=1}^{N} d_{il} d_{lk} d_{kj} F_j + ...}_{\text{indirect use of } i \text{ as input}},$$

where d_{ij} for a pair of industries (i, j) is the amount of sector i's output needed to produce one dollar worth of industry j's output. In the above sum, note that the second term on the right-hand side captures the value of sector i's "direct use" as an input, namely the total value of sector i's products that are used by sector j to produce output that immediately goes to final use. Similarly, the remaining terms that involve higher-order summations reflect the 'indirect' use of sector i as an input, as these enter further upstream in the value chain (at least two production stages away from final use).

This can be summarized in compact matrix form by stacking the above identity for all industries i:

$$Y = F + D\,[I - D]^{-1}\,F, \tag{22}$$

where Y and F are respectively the $N \times 1$ matrices of total output and final-use output, and D is the $N \times N$ matrix of direct requirements (also known as the Direct Requirements Table). The $N \times N$ matrix $[I - D]^{-1}$ is often called the Leontief inverse matrix.

Our first measure of downstreamness, $DUse\_TUse$, is the ratio of aggregate direct use to aggregate total use of industry i's goods in the production of output for final use in all industries.[17] More specifically, this is calculated as the ratio of the i-th element of the column matrix DF (i.e., the value of industry i's direct use as inputs for final-use production, summing over all buying industries j) to the i-th element of the column vector $D\,[I - D]^{-1}\,F$ (which corresponds to the total value of industry i's direct and indirect use as as inputs for final-use production, summing over all buying industries j). As explained in more detail in the Data Appendix, we associate output for 'final use' as all output purchased for consumption and investment purposes (and not for intermediate uses) by either private or government parties.[18] Note that the higher is $DUse\_TUse$ for a given

[17] See Alfaro and Charlton (2009) and di Giovanni and Levchenko (2010) for measures of production line position that have a similar flavor ro $DUse\_TUse$.

[18] Our findings are similar if we use a definition of 'final use' that includes only consumption purchases only, which is not surprisingly since consumption constitutes about 90% of the total value of consumption and investment (available on request). Note that we exclude the value of exports, imports and inventories when computing the final use value

industry i, the more frequently or intensively is i used as a direct input for final-use production, and thus the bulk of i's value enters into production relatively far downstream. Conversely, a low value of $DUse\_TUse$ suggests that most of the contribution of input i to production processes tends to occur in relatively upstream stages. There is in fact one industry IO 331314 (Secondary smelting and alloying of aluminium) that is never used directly as an input to final-use production and has a $DUse\_TUse$ of 0.

As pointed out in the Introduction, the $DUse\_TUse$ measure comes with two key limitations. First, it does not directly factor in the importance of output for final use as a share of total output in an industry, which naturally correlates with downstreamness. Second, it does not necessarily generate a clean ranking among the most upstream inputs, as these may rarely be used as direct inputs into final-use production. To further illustrate this last point, consider the specific case of IO 331411 (Primary smelting and refining of copper). This is the manufacturing industry with the second lowest value of $DUse\_TUse$ of about 0.0421, indicating that the vast majority of its input use is indirect to final-use production. That said, it is easy to trace production chains of different lengths in which IO 331411 is involved in. An example of a short chain with just three stages is: IO 331411 (Primary smelting and refining of copper) → IO 336500 (Railroad rolling stock) → IO F02000 (Private fixed investment), while a much longer example with seven stages is: IO 331411 (Primary smelting and refining of copper) → IO 331420 (Copper rolling, drawing, extruding and alloying) → IO 332720 (Turned product and screw, nut, and bolt) → IO 33291A (Valve and fittings other than plumbing) → IO 336300 (Motor vehicle parts) → IO 336112 (Automobile) → IO F01000 (Personal consumption expenditures).[19] Our first measure, $DUse\_TUse$, unfortunately does not shed light as to whether short or long chains tend to dominate in terms of IO 331411's use of an input (by value).

Our second measure of downstreamness ($DownIndex$) attempts to avoid these limitations by creating a weighted index of the average position in the value chain at which an industry's output is used. More specifically, let output for final use be designated as position 1, inputs directly use in production for final use be in position 2, inputs used directly by industries that are used as direct inputs to other industries be in position 3, and so on. We then calculate the following column matrix

$$Upstream \equiv F + 2DF + 3D^2F + 4D^3F... = [I - D]^{-2} F. \tag{23}$$

Since larger weights are applied the further upstream the input enters the production chain, this provides us with an index of "upstreamness". Although constructing this weighted index would appear to require computing an infinite power series, (23) demonstrates that $DownIndex$ turns out to be a simple function of the square of the Leontief inverse matrix. Our second measure of downstreamness for industry i, $DownIndex$, is thus calculated as the ratio of the i-th element of

of each industry. These categories often contain large negative values that when included in the calculation yield final use share values that lie outside the $[0, 1]$ range and are difficult to interpret.

[19] In identifying these examples of production lines from the Input-Output Tables, we selected buying industries at each stage which were among the top ten users by value of the input in question.

the column matrix Y and the i-th element of the column matrix $Upstream$ in (23).[20]

[TABLE 1 HERE]

Table 1 reports the ten highest and lowest values of both $DUse\_TUse$ and $DownIndex$ across the IO2002 manufacturing industries. By and large, the industries that feature low values of $DUse\_TUse$ tend to be fuel, chemical or metal inputs, while industries with high values of $DUse\_TUse$ and $DownIndex$ appear to be products that are closer to the retail end of the production line. There is some reassuring correspondence between the two measures, with both sharing three out of the ten bottom values and four out of the ten top values reported in Table 1. Nevertheless, although $DUse\_TUse$ and $DownIndex$ are positively correlated, the Pearson coefficient of 0.60 between the two variables indicates that they are far from collinear. To return to the case of IO 331411, the Primary smelting and refining of copper industry turns out to have the lowest $DownIndex$ value (0.1150), indicating that a lot of its input use value tends to occur in long production chains, for which we have given an example earlier. Although it has a similar $DUse_TUse$ value as IO 331314 (Secondary smelting and alloying of aluminium), the latter has a relatively larger $DownIndex$ value of 0.2903 (ranked 50 out of 257 industries), pointing to the fact that the bulk of this industry's input use occurs in short production lines compared to IO 331411.

C. Empirical Specification

Having described the construction of our main variables, we next discuss our empirical specifications. We will begin our analysis by experimenting with purely cross-industry regressions of the form:

$$S_{it} = \beta_1 D_i \times \mathbf{1}(\rho_i < \rho_{med}) + \beta_2 D_i \times \mathbf{1}(\rho_i > \rho_{med}) + \beta_3 \mathbf{1}(\rho_i > \rho_{med}) + \beta_X X_{it} + \alpha_t + \mu_i + \varepsilon_{it}, \quad (24)$$

where the observations are pooled across industries i and years t respectively. The dependent variable, S_{it} is the intrafirm import share in industry i in a given year. The key dependent variables are a proxy for D_i donwstreamness (either $DUse\_TUse$ or $DownIndex$, as described above) interacted with dummy variables capturing the sequential complements and substitutes cases (via an indicator variable of whether the average demand elasticity faced by industries that use i as an input is above or below the median for this variable across industries). We add these two interaction variables to highlight that the model not only predicts a differential impact of D_i on S_{it} depending on the value of ρ_i, but that it actually predicts the sign of the effect to be different depending on the value of ρ_i, and in particular, $\beta_1 > 0$ and $\beta_2 < 0$.

Our theory predicts that the effect of downstreamness on the propensity to integrate input suppliers crucially depends on whether inputs are sequential complements ($\rho > \alpha$) or sequential substitutes ($\rho < \alpha$). Ideally, one would attempt to construct proxies for the elasticity of demand

[20]Since $Y = F + DF + D^2F + D^3F...$, it follows that each element of the column matrix, $F + 2DF + 3D^2F + 4D^3F$ is greater than or equal to the corresponding element of Y, so that $DownIndex$ lies in the interval $[0, 1]$.

$(1/(1-\rho))$ and for the degree of technological substitutability across inputs $(1/(1-\alpha))$ and then create a dummy variable for each of the two cases of interest. Unfortunately, estimates of cross-input substitutability are not available in the literature and it is not clear that they can be obtained with available data sources.[21] To make some progress, we take the agnostic view that any existing cross-sectoral variation in α is largely uncorrelated with the elasticity of demand ρ faced by the average buyer of that industry's output, so that we can associate the sequential complements case with high values of ρ and the sequential substitutes case with low values of ρ. Since ρ is a monotonic function of the elasticity of demand for the final good, we proxy it with a measure of the average demand elasticity faced by industries that use i as an input.

We construct this variable using the following steps. First, we download the U.S. import demand elasticities which Broda and Weinstein (2006) estimated from disaggregate ten-digit Harmonized System (HS) product-level data. For each IO2002 industry, we then compute a demand elasticity that is equal to the trade-weighted average elasticity of each of its constituent HS10 products, using data on total U.S. imports as weights. (Details of how this crosswalk between industry codes was implemented are documented in the Data Appendix.) Next, to obtain an average demand elasticity of industries that purchase i as an input, we take a weighted average elasticity across "using" industries, with weights proportional to the value of input i used as reported in the 2002 Direct Use Table. We include the final-use value of industry i in this last calculation by assigning it the import demand elasticity of industry i itself. The average demand elasticity across "buying" industries that results from this calculations is our empirical proxy for $1/(1-\rho)$. In our baseline analysis, we will split the sample into industries with ρ above the industry median (sequential complements case) and below the median (sequential substitutes case), but we shall also report estimates that consider finer cuts of this variable by quintiles.

In addition, to these interaction terms, equation (24) includes an indicator variable for the sequential complements case, a vector of additional industry characteristics X_{it} (including a constant term), and a vector of year fixed effects, α_t. Since the key explanatory variable related to downstreamness and the elasticity of final-good demand vary at the industry level but not across years, and these are being used to explain multiple observations of the intrafirm trade share at the industry-year level, we cluster the standard errors by industry (as captured by the error term μ_i).

The vector X_{it} comprises a set of variables that the literature has identified as systematic determinants of the propensity to transact within (multinational) firm boundaries, and which our extensions in Section 3 suggest are important to incorporate as additional controls. These variables are described in detail in the Data Appendix, but we briefly overview them here and justify their inclusion in our analysis. First, we verify whether, consistent with the findings of previous papers in the literature, measures of headquarter intensity are positively associated with the share of intrafirm trade. This would be consistent too with part (ii) of our statement of Proposition 4, but notice that this Proposition further highlights that the key interaction terms in our baseline

[21] One could envision estimating α exploiting time-series variation in Input-Output coefficients, but I-O Tables are constructed only every few years and, in any case, it would not be straightforward to separately identify input substitutability from biased changes in production techniques.

specification might interact with our measures of headquarter intensity. With that in mind, we will also experiment with specifications that include triple interactions, as explained in more detail in the next section. Our proxies for headquarter intensity include measures of physical capital intensity (as first suggested by Antràs, 2003) and skill intensity from the NBER-CES Database, and a measure of R&D intensity computed by Nunn and Trefler (2008b) using the Orbis database. Also, in some specifications we use data from the NBER-CES Database to break down physical capital intensity into physical capital equipment intensity and physical capital structures (plant) intensity. As pointed out by Nunn and Trefler (2008b), this is motivated by the fact that the former type of expenditures are much more likely to be relationship-specific than the latter, and thus we would expect the former to be a cleaner proxy for headquarter intensity in the model. We also follow Nunn and Trefler (2008b) in including a measure of materials intensity in the regressions. In addition to these variables, the vector X_{it} includes a measure of the within-industry size dispersion from Nunn and Trefler (2008a), and in light of Proposition 5, we expect this variable to have a positive effect on the share of intrafirm trade.

We construct these variables in the vector X_{it} in a slightly different way from the literature. The common practice to date has been to assign to industry i the value of the capital- or skill-intensity of i itself, namely the industry selling the good in question. A more satisfactory approach that maps more directly into the model we have written is to consider a (weighted) average value of the elements of X_t computed over all industries that purchase inputs from industry i. We thus construct these measures based on "buying" industries using weights derived from the Input-Output Tables, in a manner analogous to our treatment and construction of the proxy for the the demand elasticity parameter ρ. (Summary statistics for all the variables in our analysis can be found in Appendix Table 1, while pairwise correlations between some key variables are reported in Appendix Table 2.)

As argued in sections 3.C. and 3.D., our model suggests that cross-country variation in the prevalence of integration can be fruitfully used to address biases that might arise from the endogenous location decisions of firms regarding different stages of production. With that in mind, we will also explore specifications that exploit the full country-industry dimension of the dependent data:

$$S_{ict} = \beta_1 D_i \times \mathbf{1}(\rho_i < \rho_{med}) + \beta_2 D_i \times \mathbf{1}(\rho_i > \rho_{med}) + \beta_3 \mathbf{1}(\rho_i > \rho_{med}) + \beta_X X_{it} + \alpha_{ct} + \mu_i + \varepsilon_{it}, \quad (25)$$

where the key innovation relative to (24) is the inclusion of a country-year fixed effect α_{ct}. In words, equation (25) seeks to explain the intrafirm trade share at the country-industry level as a function of industry variables, while controlling for country-year fixed effects, and (conservatively) clustering the standard errors by industry. Later on, we will build on (25) to discuss tests that make use of this cross-country variation to address concerns related to selection bias.

Before we start discussing our results, it is worth acknowledging and discussing several caveats that apply to our empirical strategy. First, the approach is by its nature limited in regard to how refined a test of our theory it provides. To understand how an input's location in a particular industry's production line affects the make-or-buy decision related to the procurement of that input, one would ideally want to observe the breakdown of intermediate input imports by the identity of the purchasing industry. Unfortunately, U.S. intrafirm trade data are reported based only on the sector or industry category of the good being transacted. For example, while we observe the share of U.S. intrafirm imports of rubber tires, we do not observe a more disaggregate breakdown vis-à-vis the share of intrafirm imports of rubber tires purchased by automobile versus aircraft makers. In the absence of more detailed firm- or transactions-level data, our empirical strategy pursues what is arguably the next best possible strategy, namely to correlate the intrafirm trade share of industry i with average measure of how far downstream input i tends to be used in all production processes that use this input. The lack of detailed data on the "buying" industry also constrains our ability to empirically distinguish between the sequential complements and sequential substitutes cases, as we can only hope to identify sectors that sell *on average* to industries that feature high or low demand elasticities.

The intrafirm trade data also does have its limitations, as argued by Antràs (2011). The data do not typically report which firm is owned by whom, i.e., whether integration is backward or forward, and also do not provide information on the extent of control (or ownership share) of the parent company, though as argued above, in most cases the integrating party is in the U.S. and the seller is majority owned (see Nunn and Trefler , 2008, for more on this). Also, using only data on U.S. imports may be problematic as it allows us to capture only those sourcing decisions that entail goods being shipped back to the U.S., while in practice some large firms have production networks in which parts and components are shipped across foreign locations and then only shipped back to the home country after being assembled abroad. For this reason, U.S. intrafirm imports generally underrepresent the involvement of U.S. multinational firms in global sourcing strategies, though it is not obvious how this phenomenon biases the empirical results we are about to discuss. On the plus side, it should be emphasized that the U.S. Census data on intrafirm trade data are subjected to several quality assurance procedures and the data offer a complete picture of the sourcing strategies of firms in the U.S., thus making it easier to spot *fundamental* factors that appear to shape whether international transactions are internalized or not independently of the sector one studies.

5 Empirical Results

We begin our empirical analysis in Table 2 by running versions of (24) to replicate and refine some of the key findings obtained in prior studies of the determinants of intrafirm trade (see, for instance, Antràs, 2003, Nunn and Trefler, 2008a, and Bernard et al., 2010). Using the industry-year observations of the dependent variable, columns 1 and 2 provides baselines in which we examine if the characteristics of the "selling" industry systematically explain the propensity to import the product within firm boundaries. In column 1, we indeed find that the various alternative measures of headquarter intensity (skill intensity, physical capital intensity, and R&D intensity) are all positively and statistically significantly correlated with the share of intrafirm trade, as is a measure of productivity dispersion within an industry. Another measure of input intensity, materials intensity, that one would typically *not* associate with headquarters instead has indeed no predictive power for the share of intrafirm trade. Column 2 highlights the importance of distinguishing between physical capital equipment and physical capital structures (c.f., Nunn and Trefler, 2008a); the latter is less likely to involve noncontractible, relationship-specific investments by headquarters, and thus it is not surprising that this variable is only weakly correlated (actually, with a negative sign) with the share of intrafirm trade. In columns 3 and 4, we re-run the same specification as in columns 1 and 2 but now using the average "buying" industry values of the industry characteristics (which we have argued are more consistent with our model of input sourcing), in place of the typically used "selling" industry values. This has some effects on the estimates, although the importance of headquarter intensity still broadly stands. Physical capital equipment and R&D intensity continue to have a positive and statistically significant effect on the share of intrafirm trade, although skill intensity no longer does. The effect of productivity dispersion is also less statistically significant than in columns 1 and 2.

[TABLE 2 HERE]

The novel predictions from our model, as outlined in the estimating equation (24), are tested in Table 3. In column 1, we begin by introducing our first proxy for downstreamness, the ratio of aggregate direct use to aggregate total use ($DUse\_TUse$) to the industry-year regressions.[22] When included on its own, the effect of $DUse\_TUse$ on the share of intrafirm trade is statistically insignificant. Following the guidance of our theoretical model, we run our benchmark specification in equation (24) in column 2, which includes the interactions of $DUse\_TUse$ with our proxies for the sequential substitutes ($Elas < Median$) and complements ($Elas > Median$) cases, as well as a dummy variable for the sequential complements case.[23] Our empirical results are broadly supportive of our model's predictions. The effect of downstreamness is positive and highly statistically significant when the elasticity of demand faced by final-good producers purchasing that sector's inputs is above the median for this variable ($\beta_2 > 0$), while the effect is of the opposite sign (and also significant at the 10% level) when the elasticity is below its median value ($\beta_1 < 0$). The results thus confirm that the effect of downstreamness is qualitatively different depending on the elasticity of final-good demand, which is the key insight of our model. Furthermore, the effect of the "high elasticity" dummy variable is negative and statistically significant, which also resonates with our comparative static result in Proposition 3. It is particularly reassuring that these new findings are robust even while including the same set of controls for "buying" industry headquarter intensity and productivity dispersion that we used in our Table 2 baseline. Note too that R&D intensity and physical equipment intensity (two natural proxies for headquarter intensity) remain important in shaping the integration decisions of firms.[24]

[22] The results are qualitatively similar if the regressions are run year-by year, namely with the intrafirm trade share of each industry in a particular year as the dependent variable, with Huber-White robust standard errors.

[23] There is no need to include a the dummy variable for the sequential substitutes case, as the regressions include a constant term.

[24] Similarly, adding a measure of the ratio of value-added to total shipments averaged across buying industries does not alter our results.

[TABLE 3 HERE]

The next two columns of Table 3 split the regression sample according to whether the total imports corresponding to each observation (the denominator of the intrafirm share) fall below or above the median (columns 3 and 4, respectively). This is motivated by the existence of measurement error in the reporting of whether trade flows involve related or unrelated parties, and by the fact that this is more likely to be a concern for industries with small trade volumes (see footnote 16 for some evidence). Consistent with this conjecture, the predictions of our theory appear to be more salient in the larger trade volumes, with the coefficient β_2 being positive and significant in column 4 but not in column 3, as well as the coefficient β_1 being more negative in column 4 than in column 3. Motivated by these findings, in column 5 we report results of running the specification in (24) for the full sample of industries while weighting observations by the total imports for that industry-year, thereby attaching more weight (and less measurement error) to data points that are associated with lower reporting problems. This clearly reinforces our main finding that integration tends to rise with downstreamness precisely in the sequential complements case, with the effect for the substitutes case being negative though not statistically significant. The magnitude of our key β_2 coefficient moreover increases, while the fit of the regression improves markedly in terms of its R^2. Based on the point estimate in column 5, a one standard deviation increase in $DUse\_TUse$ would correspond to an increase in the share of intrafirm trade in high ρ industries of $0.481 \times 0.235 = 0.113$, which is over one-half of a standard deviation in the intrafirm trade share, a fairly sizeable effect.[25]

The final two columns of Table 3 presents results from exploiting the variation across import source countries available from the intrafirm trade data. Following the specification in (25), column 6 includes country-year fixed effects and clusters the standard errors by industry. Furthermore, in column 7 we present results that weight the observations by total import volumes, for reasons analogous to those discussed above.[26] The results in column 7 look very similar to the purely cross-industry ones, with the coefficient β_2 being positive and significant at the 1% level, while β_1 being negative (though not statistically significant). The remaining coefficients are also consistent with the theory, with positive effects of our preferred headquarter intensity measures (physical capital equipment and R&D intensities), and a negative effect of the dummy variable for the sequential complements case. The unweighted results in column 6 are instead quite noisy (the R-square is much lower), and the positive effect of downstreamness on integration in the complements case disappears. It should be noted however that the results continue to be consistent with the model's *weaker* prediction that downstreamness should have a more negative effect on integration in the substitutes case than in the complements case.

---

[25] One might be concerned that the industry with the largest total import value, IO 336311 (Automobile), is also the most downstream manufacturing industry according to both $DUse_TUse$ and $DownIndex$, and that this might in turn be driving our significant findings in the sequential complements case. Reassuringly, our results are robust to dropping this industry for the sample (available on request).

[26] In fact, when working with country-industry-year observations, there are many more small import flow observations that may be particularly prone to measurement errors in terms of the level of relatedness of the parties, and thus weighting by import volumes might be even more justified in this case.

In Table 4, we replicate the results in Table 3, but using our second measure $DownIndex$ as an empirical proxy for downstreamness. Reassuringly, we obtain similar conclusions with this alternative measure of downstreamness. As can be seen from the positive and significant coefficient estimates of $DownIndex_i \times \mathbf{1}(\rho_i > \rho_{med})$, downstreamness tends to raise the intrafirm trade share for those products whose buyer industries face high elasticities of demand. Once again, this effect is especially strong when weighting observations by import volumes (Columns 3-4). The magnitude of this effect is moreover consistent with the previous evidence based on $DUse\_TUse$; based on Column 5, for the industries that fall under the sequential complements case, a one standard deviation increase in $DownIndex$ would imply an increase in the intrafirm trade share of $0.490 \times 0.238 = 0.117$, which is again just over one-half of a one standard deviation in the share of intrafirm imports. Admittedly, the results are a bit more mixed when it comes to the effects of $DownIndex$ in industries selling to sectors facing low elasticities of demand, with some point estimates being positive, although statistically indistinguishable from zero.

[**TABLE 4 HERE**]

In Table 5, we address potential criticisms regarding the use of a median cutoff value for the ρ proxy to differentiate the sequential complements and sequential substitutes case. Given that technological substitutability across sectors (α in the model) might very well vary across sectors, it seems plausible that the positive effect of downstreamness will be concentrated in the highest values of the elasticity of demand parameter ρ_i, while the negative effect of downstreamness might only appear for particularly low values of ρ_i. With that in mind, in Table 5 we estimate the following more flexible variant of (24) that breaks the data into quintiles of ρ_i:

$$S_{it} = \sum_{s=1}^{5} \beta_{1s} DUse\_TUse_i \times \mathbf{1}(\rho_i \in \Omega_{\rho,s}) + \sum_{s=2}^{5} \beta_{2s} \mathbf{1}(\rho_i \in \Omega_{\rho,s}) + \beta_X X_i + \alpha_t + \mu_i + \varepsilon_{it} \qquad (26)$$

Here, $\Omega_{\rho,s}$ refers to the interval of values corresponding to the s-th quintile of ρ ($s = 1, 2, \ldots, 5$), with $\mathbf{1}(\rho_i \in \Omega_{\rho,s})$ being a dummy variable equal to 1 if industry i falls within this s-th quintile. Note that we include the main effects of the quintile dummies, although we drop the first quintile dummy due to the collinearity with the full set of year fixed effects, α_t. We report three regressions in Table 5 for each of the downstreamness measures, $DUse\_TUse$ and $DownIndex$. For each measure, we estimate (26) both without and with regression weights in the first and second columns respectively; the third column runs the analogous specification with the country-industry-year observations, while controlling for country-year fixed effects. As before, all standard errors are clustered by industry.) All columns in Table 5 include the key control variables in Tables 3 and 4.

[**TABLE 5 HERE**]

The findings in Table 5 reveal that downstreamness appears to matter for the intrafirm trade share only in the highest quintiles of the proxy for ρ. This gives strong confirmation that the effect of downstreamness on the propensity to integrate is positive exactly in the range of values of ρ for which the industry is most likely to fall within the sequential complements case. The point estimates obtained for the lower quintiles tend in particular to be of a negative sign, even being statistically significant at the 10% level in one case. Nevertheless, what is striking is that the point estimates of the effect of downstreamness appear to progressive increase as we move from the lowest quintile of ρ to the highest. Overall, the results in the Table 5 strengthen our confidence in the empirical relevance of our theoretical results.

In Table 6, we test an auxiliary prediction of the model that arises when introducing headquarter intensity. In section 3.B., we have shown that an increase in η not only expands the range of stages that are vertically integrated (an effect analogous to that unveiled by previous property-rights models and consistent with the positive effect of our proxies of headquarter intensity on the share of intrafirm trade), but also affects the range of parameter values for which downstreamness is predicted to have a positive effect on the share of intrafirm trade. In particular, remember that the complements and substitutes cases are now defined by $\tilde{\rho} \equiv (1-\eta)\,\rho > \alpha$ and $\tilde{\rho} \equiv (1-\eta)\,\rho < \alpha$, and thus the larger is η, the less likely it is that downstreamness will have a positive effect on the share of intrafirm trade, even for large values of the elasticity of demand ρ. Table 6 illustrates that this is precisely what one observes in the data. The key innovation relative to specifications (24) and (25) is that we now include a full set of triple interactions of each of our key interaction variables – $D_i \times \mathbf{1}(\rho_i < \rho_{med})$ and $D_i \times \mathbf{1}(\rho_i > \rho_{med})$ – with each of five dummy variables corresponding to the quintiles of an aggregate measure of headquarter intensity. The latter is computed as the first principal component of the three main measures of headquarter intensity in our specifications, namely skill, physical capital and R&D intensity. As is clear from the table, the positive effect of $D_i \times \mathbf{1}(\rho_i > \rho_{med})$ on the integration decision is largely concentrated in the lowest quintile of headquarter intensity, with few results being significant at other quintiles of the distribution of headquarter intensity.

[TABLE 6 HERE]

A similar auxiliary prediction of the model is tested in Table 7. The motivation for this test stems from our modeling of final-good producer heterogeneity in section 3.C. As already anticipated in that section, to the extent that fixed costs of integration are relatively high, in the sequential complements case, only a few very downstream stages will be integrated, while in the sequential substitutes case, only a few very upstream stages will be integrated. This implies that when assessing the effects of downstreamness on the share of intrafirm trade, it is important to recognize that the effect of this variable might be concentrated in the highest quintiles of this variable in the sequential complements case, and in the lowest quintiles of this variable in the sequential substitutes case. The specification we employ in Table 7 attempts to capture these predictions by breaking up our two main interaction variables in specifications (24) and (25) into various interaction terms depending on the quintile of D_i. The specification is analogous to that in (26), in which we broke these benchmark interactions for different quintile values of ρ.[27] The fact that in all six columns the positive effect of downstreamness is concentrated in the highest quintiles of D_i whenever ρ is above its median value is again consistent with the model. Furthermore, we find evidence (statistically significantly too for our preferred measure of downstreamness $DUse\_TUse$) that the share of intrafirm trade is highest for the least downstream stages whenever the value of ρ is below its median, again confirming this auxiliary prediction from our model.

[TABLE 7 HERE]

The remaining two Tables 8 and 9 present further robustness tests. In Table 8, we correct our industry-country regressions for potential selection biases. As explained in section 3.D., controlling for country fixed effects may not be sufficient when the location of input production is itself affected by the level of downstreameness (as studied in our Mathematical Appendix). For this reason, in Table 8, we carry out a two-stage Heckman selection procedure in which the excluded variable in the second stage is an interaction of a dummy variable for countries with above sample-median entry costs (from the World Bank's Doing Business database) and the selling industry's R&D intensity. Following the logic of Helpman, Melitz and Rubinstein (2008), the level of entry costs is likely to affect the location decision of firms, but it is less clear how it will shape the intensive margin of trade. In order to enhance the fit of the first stage, we interact this variable with R&D intensity, but do so using the selling industry classification, in order to minimize endogeneity concerns. As shown in columns 1 and 4 of Table 8, this interaction variable has a negative effect on the probability of observing imports from a particular country-industry pair. The inverse Mills ratio appears to be statistically significant in the second stage, but when comparing columns 2 and 3, and also 5 and 6, it is clear that the correction for selection does not have a quantitatively important effect on our results. For computational reasons, our results in Table 8 focus on a single year in our sample, and we choose 2005 since it falls right in the middle of our sample, but the results are very similar for other years.

[TABLE 8 HERE]

Finally, in Table 9 we experiment with specifications that include as controls two measures of contractibility. In our robustness tests, we incorporate the two measures of contractibility suggested by Nunn and Trefler (2008a) and Bernard et al. (2010), which have been shown to correlate significantly with the share of intrafirm trade. This is motivated by the theoretical results in Antràs and Helpman (2008) and by previous empirical findings in Nunn and Trefler (2008a,b) and Bernard et al. (2010), from which we borrow the proxies for contractibility and intermediation respectively (see the Data Appendix for details). Consistently with those previous papers, we find that the share of intrafirm trade appears to be lower in more contractible industries. The inclusion of these variables affects the magnitude of our key estimates of interest and in particular the size of the coefficient β_2 is markedly lower, though it maintains its statistically significance at the 1% level in

[27] As in the that case, here we also need to drop one of these interactions due to the collinearity.

most specifications. The results for β_2 appear to be more supportive for the theory when using the proxy $DUse\_TUse$ for downstreamness, though it is interesting to note that when using the alternative proxy $DownIndex$, we find evidence of a statistically significant (at the 10% level) negative effect of downstreamness on integration in the sequential substitutes case.

[TABLE 9 HERE]

6 Conclusion

[TO BE ADDED]

References

Acemoglu, Daron, Pol Antràs, and Elhanan Helpman, (2007), "Contracts and Technology Adoption," *American Economic Review* 97: 916-943.

Alfaro, Laura, and Andrew Charlton, (2009), "Intra-Industry Foreign Direct Investment," *American Economic Review* 99: 2096-2119.

Antràs, Pol, (2003), "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics* 118: 1375-1418.

Antràs, Pol (2005), "Incomplete Contracts and the Product Cycle," *American Economic Review*, 95:4, pp. 1054-1073.

Antràs, Pol, (2011) "Grossman-Hart (1986) Goes Global: Incomplete Contracts, Property Rights, and the International Organization of Production," mimeo Harvard University.

Antràs, Pol, and Elhanan Helpman, (2004), "Global Sourcing," *Journal of Political Economy* 112: 552-580.

Antràs, Pol, and Elhanan Helpman, (2008), "Contractual Frictions and Global Sourcing," forthcoming in E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*, Harvard University Press.

Baldwin, Richard, and Anthony Venables, (2010), "Relocating the Value Chain: Offshoring and Agglomeration in the Global Economy," *NBER Working Paper* No. 16611.

Bartelsman, Eric J., Randy A. Becker, Wayne B. Gray, (2000), NBER-CES Manufacturing Industry Database.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, (2010), "Intra-Firm Trade and Product Contractibility (Long Version)," *NBER Working Paper* 15881.

Broda, Christian, and David Weinstein, (2006), "Globalization and the Gains from Variety," *Quarterly Journal of Economics* 121(2): 541-585.

Corcos, Gregory, Delphine M. Irac, Giordano Mion, and Thierry Verdier, (2008), "The Determinants of Intra-Firm Trade," mimeo.

Costinot, Arnaud, Jonathan Vogel, and Su Wang, (2011), "An Elementary Theory of Global Supply Chains," MIT mimeo.

Debeare, Peter, Hongshik Lee, and Joon Hyung Lee, (2009), "Explaining International Intra-firm Trade: A Firm-Level Perspective," mimeo.

Defever, Fabrice, and Farid Toubal, (2007), "Productivity and the Sourcing Modes of International Firms: Evidence from French Firm-Level Data," *CEP Discussion Paper* 842.

Díez, Federico, (2010), "The Asymmetric Effects of Tariffs on Offshoring Industries: How North/South Tariffs Affect Intra-Firm Trade," University of Wisconsin-Madison mimeo.

di Giovanni, Julian, and Andrei Levchenko, (2010), "Putting the Parts Together: Trade, Vertical Linkages, and Business Cycle Comovement," *American Economic Journal: Macroeconomics* 2(2): 95-124.

Dixit, Avinash, and Gene Grossman, (1982), "Trade and Protection with Multistage Production," *Review of Economic Studies* 49: 583-594.

Feenstra, Robert C., John Romalis, Peter K. Schott, (2002), “U.S. Imports, Exports and Tariff Data, 1989-2001,” *NBER Working Paper* 9387.

Fernandes, Ana, and Heiwai Tang, (2011), “Determinants of Vertical Integration in Export Processing: Theory and Evidence from China,” Tufts University mimeo.

Findlay, Ronald (1978), “An ‘Austrian’ Model of International Trade and Interest Rate Equalization,” *Journal of Political Economy,* Vol. 86, No. 6, pp. 989-1007.

Grossman, Sanford J., and Hart, Oliver D., (1986), “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy* 94: 691-719.

Grossman, Gene M., and Esteban Rossi-Hansberg, (2008), “Trading Tasks: A Simple Theory of Offshoring,” *American Economic Review* 98(5): 1978–1997.

Hanson, Gordon, Raymond Mataloni, and Matthew Slaughter, (2005), “Vertical Production Networks in Multinational Firms,” *Review of Economics and Statistics* 87: 664-678.

Harms, Philipp, Oliver Lorz, and Dieter Urban, (2009), “Offshoring along the Production Chain,” *CESifo Working Paper* 2564.

Helpman, Elhanan, Marc J. Melitz, and Stephen R.Yeaple (2004), “Exports versus FDI with Heterogeneous Firms,” *American Economic Review*, 94:1, pp. 300-316.

Helpman, Elhanan, Marc J. Melitz, and Yona Rubinstein (2008), “Estimating Trade Flows: Trading Partners and Trading Volumes,” *Quarterly Journal of Economics*, 123(2), 441-487.

Johnson, Robert C., and Guillermo Noguera, (2010), “Accounting for Intermediates: Production Sharing and Trade in Value Added,” Dartmouth mimeo.

Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi, (2009), “Governance Matters VIII: Aggregate and Individual Governance Indicators 1996-2008,” *World Bank Policy Research Working Paper* 4978.

Kim, Se-Jik and Hyun Song Shin (2011), “Sustaining Production Chains Through Financial Linkages,” mimeo Princeton University.

Kohler, Wilhelm, (2004), “International Outsourcing and Factor Prices with Multistage Production,” *Economic Journal* 114: C166-C185.

Levine, David K., (2010), “Production Chains,” *NBER Working Paper* No. 16571.

Nunn, Nathan, and Daniel Trefler, (2008), “The Boundaries of the Multinational Firm: An Empirical Analysis," forthcoming in E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*, Harvard University Press.

Pierce, Justin R., and Peter K. Schott, (2009), “A Concordance between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries,” *NBER Working Paper* 15548.

Rauch, James E., (1999), “Networks versus Markets in International Trade,” *Journal of International Economics* 48: 7-35.

Sanyal, Kalyan K., (1983), “Vertical Specialization in a Ricardian Model with a Continuum of Stages of Production,” *Economica* 50: 71-78.

Sanyal, Kalyan K., and Ronald W. Jones, (1982), “The Theory of Trade in Middle Products,” *American Economic Review* 72: 16-31.

Williamson, Oliver E., (1975), *Markets, Hierarchies: Analysis, Antitrust Implications*, Free Press: New York.

Williamson, Oliver E., (1985), *The Economic Institutions of Capitalism*, Free Press: New York.

Winter, Eyal (2006), "Optimal Incentives for Sequential Production Processes, RAND Journal of Economics, Vol. 37, No. 2, pp. 376-390.

Yeaple, Stephen R., (2006), "Offshoring, Foreign Direct Investment, and the Structure of US Trade," *Journal of the European Economic Association* 4(2-3): 602-611.

Yi, Kei-Mu, (2003), "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111: 52-102.

Zhang, Juyan, and Yi Zhang, (2008), "Sequential Hold-up, and Ownership Structure," Singapore Management University mimeo.

Zhang, Juyan, and Yi Zhang, (2011), "Sequential Hold-up, and Strategic Delay," Singapore Management University mimeo.

A Mathematical Appendix (Incomplete)

Proof of Lemma 1: The result follows from studying how the ratio in the expression for $\beta^*(m)$ in (13) depends on m. When $\rho > \alpha$, it is clear that the numerator is decreasing in m, while the denominator is increasing in m. Due to the negative sign in front of the ratio, it follows that $\beta^*(m)$ is increasing in m whenever $\beta^*(m) > 0$. The case $\rho < \alpha$ is more cumbersome to study. It proves useful to re-write $\beta^*(m)$ as:

$$\frac{\beta^*(m) - (1-\alpha)}{\frac{\alpha-\rho}{(1-\rho)}} = \left(\int_{i=m}^{1} \beta(i)(1-\beta(i))^{\frac{\alpha}{1-\alpha}} \left[\int_0^i (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} di \right) \left[\int_0^m (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\alpha-\rho}{\alpha(1-\rho)}} . \tag{27}$$

Straightforward differentiation shows that the sign of the derivative of the right-hand side of this expression with respect to m is determined by that of:

$$-\beta(m) + \frac{\alpha-\rho}{\alpha(1-\rho)} \left(\int_m^1 \beta(i)(1-\beta(i))^{\frac{\alpha}{1-\alpha}} \left[\int_0^i (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} di \right) \left[\int_0^m (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj \right]^{\frac{\alpha-\rho}{\alpha(1-\rho)}}$$

which using the above expression for $\beta^*(m)$ simplifies to:

$$-\frac{1-\alpha}{\alpha}(1-\beta^*(m)) \leq 0,$$

and thus $\beta^*(m)$ is decreasing in m when $\rho < \alpha$. Note that we do not allow the firm to set $\beta^*(m) > 1$, but in the proof of Proposition 2 below we will verify that the firm never has an incentive to do so.

Proof of Proposition 2: We begin by examining the limiting values of $\beta^*(m)$. First take the case $\rho > \alpha$. When $m \to 0$, the numerator of the ratio in the expression for $\beta^*(m)$ in (13) converges to a positive, finite value, while the denominator goes to 0. Hence, $\lim_{m\to 0}\beta^*(m) = 0$ (remember that $\beta^*(m) \geq 0$). When $m \to 1$, the numerator goes to 0 while the denominator converges to a positive, finite value, implying $\lim_{m\to 1}\beta^*(1) = 1-\alpha$. For the case $\rho < \alpha$, we again appeal to alternative formula for $\beta^*(m)$ in (27). Note that

$$\lim_{m\to 0} \left\{ \frac{\frac{\beta^*(m)}{(1-\alpha)} - 1}{\frac{\alpha-\rho}{(1-\alpha)(1-\rho)}} \right\} = \lim_{m\to 0} \left(\int_{i=m}^{1} \beta(i)(1-\beta(i))^{\frac{\alpha}{1-\alpha}} \left[\int_{j=0}^{i} (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} di \right)$$
$$\div \lim_{m\to 0} \left[\int_{k=0}^{m} (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}$$

Both the limits in the expression on the right-hand side tend to $+\infty$, so we need to use l'Hospital rule:

$$\lim_{m\to 0} \left\{ \frac{\beta^*(m) - (1-\alpha)}{\frac{\alpha-\rho}{(1-\rho)}} \right\} = \lim_{m\to 0} \left\{ \frac{-\beta(m)(1-\beta(m))^{\frac{\alpha}{1-\alpha}} \left[\int_0^m (1-\beta(j))^{\frac{\alpha}{1-\alpha}} dj \right]^{-\frac{\alpha-\rho}{\alpha(1-\rho)}-1}}{\left(-\frac{\alpha-\rho}{\alpha(1-\rho)} \right) \left[\int_0^m (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{-\frac{\alpha-\rho}{\alpha(1-\rho)}-1} (1-\beta(m))^{\frac{\alpha}{1-\alpha}}} \right\},$$

which simplifies to

$$\frac{\beta^*(0) - (1-\alpha)}{\frac{\alpha-\rho}{(1-\rho)}} = \frac{\beta^*(0)}{\frac{\alpha-\rho}{\alpha(1-\rho)}},$$

from which we obtain $\beta^*(0) = 1$.

Together with Lemma 1, the above limit values imply that when $\rho > \alpha$, it is optimal for the firm to choose β_O (namely outsourcing) for stages with a small index in the neighborhood of $m = 0$, since $0 < \beta_O < \beta_V$. Conversely, in the $\rho < \alpha$ case, $\beta^*(m)$ it is optimal for the firm to choose β_V (namely integration) for stages in a neighborhood of $m = 0$.

To fully establish Proposition 2 for case $\rho > \alpha$, we proceed to show that we cannot have a positive measure of integrated stages located upstream relative to a positive measure of outsourced stages in the optimal organizational structure. Since the limit values above indicate that stage 0 will be outsourced, it follows that if any stages are to be integrated, they have to be downstream relative to all outsourced stages. In other words, there exists an optimal cutoff $m_C^* \in (0, 1]$ such that all stages in $[0, m_C^*)$ are outsourced and stages in $[m_C^*, 1]$ are integrated. (If $m_C^* = 1$, then all stages along the production line are outsourced.)

We establish the above claim by contradiction. Suppose that there exists a stage $\tilde{m} \in (0, 1)$ and a positive constant $\varepsilon > 0$ such that stages in $(\tilde{m} - \varepsilon, \tilde{m})$ are integrated, while stages in $(\tilde{m}, \tilde{m} + \varepsilon)$ are outsourced. The width of both of these sub-intervals, ε, can clearly be chosen to be equal without loss of generality. Let profits from this mode of organization be Π_1. On the other hand, consider an alternative organizational mode which instead outsources the stages in $(\tilde{m} - \varepsilon, \tilde{m})$ and integrates the stages in $(\tilde{m}, \tilde{m} + \varepsilon)$, while retaining the same organizational decision for all other stages. Let profits from this alternative be Π_2. Using the expression for the firm's profits from (11), one can show after some tedious algebra that up to a positive multiplicative constant:

$$
\begin{aligned}
\Pi_1 - \Pi_2 \propto & \int_{\tilde{m}-\varepsilon}^{\tilde{m}} \beta_V (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + (j - \tilde{m} + \varepsilon)(1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\
& + \int_{\tilde{m}}^{\tilde{m}+\varepsilon} \beta (1-\beta)^{\frac{\alpha}{1-\alpha}} \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + \varepsilon(1-\beta_V)^{\frac{\alpha}{1-\alpha}} + (j - \tilde{m})(1-\beta)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\
& - \int_{\tilde{m}-\varepsilon}^{\tilde{m}} \beta (1-\beta)^{\frac{\alpha}{1-\alpha}} \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + (j - \tilde{m} + \varepsilon)(1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\
& - \int_{\tilde{m}}^{\tilde{m}+\varepsilon} \beta_V (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + \varepsilon(1-\beta)^{\frac{\alpha}{1-\alpha}} + (j - \tilde{m})(1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj
\end{aligned}
$$

and thus

$$
\begin{aligned}
\Pi_1 - \Pi_2 \propto (\beta_V - \beta) \Bigg[& \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + \varepsilon(1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} + \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + \varepsilon(1-\beta)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} \\
& - \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di + \varepsilon(1-\beta_V)^{\frac{\alpha}{1-\alpha}} + \varepsilon(1-\beta)^{\frac{\alpha}{1-\alpha}} \right)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} - \left(\int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di \right)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} \Bigg]
\end{aligned}
$$

where the second step follows from evaluating the integrals with respect to j.

Since $\beta_V - \beta > 0$, it suffices to show that the expression in square parentheses is negative. To see this, consider the function $f(y) = y^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}}$. Simple differentiation will show that for $y, a > 0$ and $b \geq 0$, that $f(y + a + b) - f(y + b)$ is an increasing function in b when $\rho > \alpha$ and hence $\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho} > 1$. Hence, $(y+a+b)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} - (y+b)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} > (y+a)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}} - (y)^{\frac{\rho}{\alpha}\frac{1-\alpha}{1-\rho}}$. Setting $y = \int_0^{\tilde{m}-\varepsilon} (1-\beta(i))^{\frac{\alpha}{1-\alpha}} di$, $a = \varepsilon(1-\beta)^{\frac{\alpha}{1-\alpha}}$ and $b = \varepsilon(1-\beta_V)^{\frac{\alpha}{1-\alpha}}$, it follows that the last term in square brackets is negative and that $\Pi_1 - \Pi_2 < 0$. This yields the desired contradiction as profits can be strictly increased by switching to the organizational mode that yields profits Π_2.

The full proof for the $\rho < \alpha$ case is analogous. The limit values show that it is optimal to integrate stage 0. One can then show that if any stages are to be outsourced, they occur downstream to all the integrated stages, so that there is a unique cutoff $m_S^* \in (0,1)$ with all stages prior to m_S^* being integrated and all stages after m_S^* being outsourced. This can be established using a similar proof by contradiction.

Proof of Proposition 3: We begin by deriving equations (14) and (15). For each case, this is achieved by first plugging the optimal values of $\beta(m) \in \{\beta_V, \beta_O\}$ for all $m \in [0,1]$ in the firm's maximand in (11) and then solving

$$m_C^* = \arg\max_m \left\{ \beta(1-\beta)^{\frac{\rho}{1-\rho}} \int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj + \beta_V (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_m^1 \left[(1-\beta)^{\frac{\alpha}{1-\alpha}} m + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} (j-m) \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \right\};$$

$$m_S^* = \arg\max_m \left\{ \beta_V(1-\beta_V)^{\frac{\rho}{1-\rho}} \int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj + \beta (1-\beta)^{\frac{\alpha}{1-\alpha}} \int_m^1 \left[(1-\beta_V)^{\frac{\alpha}{1-\alpha}} m + (1-\beta)^{\frac{\alpha}{1-\alpha}} (j-m) \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \right\}.$$

Let us illustrate it for the case $\rho > \alpha$. The first-order condition associated with the optimal choice of m is given by:

$$\beta(1-\beta)^{\frac{\rho}{1-\rho}} m^{\frac{\rho-\alpha}{\alpha(1-\rho)}} - \beta_V(1-\beta_V)^{\frac{\alpha}{1-\alpha}} (1-\beta)^{\frac{\alpha}{1-\alpha}\frac{\rho-\alpha}{\alpha(1-\rho)}} m^{\frac{\rho-\alpha}{\alpha(1-\rho)}}$$
$$+ \beta_V(1-\beta_V)^{\frac{\alpha}{1-\alpha}} \frac{\rho-\alpha}{\alpha(1-\rho)} \left((1-\beta)^{\frac{\alpha}{1-\alpha}} - (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right) \int_m^1 \left[(1-\beta)^{\frac{\alpha}{1-\alpha}} m + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} (j-m) \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} dj = 0$$

which after a few simplifications can be written as

$$\beta_V - \beta = \beta_V \left(1 - \left(\frac{1-\beta_V}{1-\beta} \right)^{\frac{\alpha}{1-\alpha}} \right) \left(\left[1 + \left(\frac{1-\beta_V}{1-\beta} \right)^{\frac{\alpha}{1-\alpha}} \left(\frac{1-m}{m} \right) \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \right),$$

from which the formula for m_C^* in (14) is obtained. The formula for m_S^* in (15) can be derived in an analogous manner.

Next, the effect of ρ on the thresholds follows immediate from noting that $m_C^* < 1$ only if $\left(\frac{1-\beta}{1-\beta_V}\right)^{-\frac{\alpha}{1-\alpha}} > \frac{\beta}{\beta_V}$, while $m_S^* < 1$ only if $\left(\frac{1-\beta_V}{1-\beta}\right)^{-\frac{\alpha}{1-\alpha}} > \frac{\beta_V}{\beta}$. Hence, m_C^* increases in ρ and m_S^* decreases in ρ whenever the thresholds are in $(0,1)$. As for the effect of β, using $\beta_V = \beta + (1-\beta)\delta^\alpha$ we can rewrite the thresholds as

$$m_C^* = \left[1 + (1-\delta^\alpha)^{\frac{-\alpha}{1-\alpha}} \left[\left(\frac{1 - \frac{1}{1+\left(\frac{1-\beta}{\beta}\right)\delta^\alpha}}{(1-\delta^\alpha)^{\frac{\alpha}{1-\alpha}}} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right]^{-1}$$

and

$$m_S^* = \left[1 + (1-\delta^\alpha)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{(1-\delta^\alpha)^{\frac{-\alpha}{1-\alpha}} - 1}{\frac{(1-\beta)\delta^\alpha}{\beta}} \right)^{\frac{\alpha(1-\rho)}{\alpha-\rho}} - 1 \right] \right]^{-1},$$

from which it is clear that m_C^* increases in β and m_S^* decreases in β whenever the thresholds are in $(0,1)$.

Proof of Proposition 4: From the discussion in the main text, and equation (10), we have that investments by suppliers will be given by:

$$x(m)=\left(\frac{1-\tilde{\rho}}{1-\alpha}\right)^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}}\left(\tilde{\rho}\frac{A^{1-\rho}\theta^{\rho}\left(\frac{h}{\eta}\right)^{\rho\eta}(1-\eta)^{-\tilde{\rho}}}{c}\right)^{\frac{1}{1-\tilde{\rho}}}(1-\beta(m))^{\frac{1}{1-\alpha}}\left[\int_{j=0}^{m}(1-\beta(j))^{\frac{\alpha}{1-\alpha}}dj\right]^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}},$$

and the final good producer will capture revenues given by

$$R_F=\left(A^{1-\rho}\theta^{\rho}\left(\frac{h}{\eta}\right)^{\rho\eta}(1-\eta)^{-\tilde{\rho}}\right)^{\frac{1}{1-\tilde{\rho}}}\frac{\tilde{\rho}}{\alpha}\left(\frac{1-\tilde{\rho}}{1-\alpha}\right)^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}}\left(\frac{\tilde{\rho}}{c}\right)^{\frac{\tilde{\rho}}{1-\tilde{\rho}}}\int_{j=0}^{1}\beta(j)(1-\beta(j))^{\frac{\alpha}{1-\alpha}}\left[\int_{k=0}^{j}(1-\beta(k))^{\frac{\alpha}{1-\alpha}}dk\right]^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}}dj \tag{28}$$

where remember that $\tilde{\rho}\equiv(1-\eta)\rho$.

Before suppliers makes any investments, the firm will choose h to maximize $R_F-c_h h$. From equation (28) and the definition of $\tilde{A}$, it is clear that the optimal choice of h will satisfy

$$\frac{\rho\eta}{1-\tilde{\rho}}R_F=c_h h,$$

and thus the firm obtains profits equal to

$$\pi_F=\left(\frac{1-\rho}{1-\tilde{\rho}}\right)R_F.$$

Hence, the profit-maximizing sequence of organizational forms will also maximize the expression in (28), which is analogous to the objective function in the benchmark model, but with ρ and A being replaced by $\tilde{\rho}$ and $\tilde{A}$. This establishes part (i) of Proposition 4.

For part (ii) of the proposition, the cutoff expressions for the two cases, m_C^* and m_S^*, are now given by (14) and (15) respectively, with ρ replaced by $\tilde{\rho}$. Differentiating (14) and (15) with respect to ρ (as in Proposition 3) and bearing in mind that $\tilde{\rho}$ is decreasing in η, yields the desired comparative static results.

Proof of Proposition 5: A firm with productivity parameter θ now chooses its organizational structure along the value chain to maximize

$$\pi_F=A\frac{\rho}{\alpha}\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}}\left(\frac{\rho\theta}{c}\right)^{\frac{\rho}{1-\rho}}\int_0^1\beta(j)(1-\beta(j))^{\frac{\alpha}{1-\alpha}}\left[\int_0^j(1-\beta(k))^{\frac{\alpha}{1-\alpha}}dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}dj-\int_0^1 f(j),$$

where $(\beta(m),f(m))=(\beta_V,f_V)$ when stage m is integrated and $(\beta(m),f(m))=(\beta_O,f_O)$ when it is outsourced. It should be clear that the study of the choice of a (hypothetical) unconstrained optimal division of surplus $\beta^*(m)$ for stage m is not affected by the presence of fixed costs since these do not impact the derivative and the fixed costs are independent of the stage of production being considered.

For the complements case, we thus have

$$m_C^*=\arg\max_m\left\{\begin{array}{c}\Omega\beta(1-\beta)^{\frac{\rho}{1-\rho}}\int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}}dj+\beta_V(1-\beta_V)^{\frac{\alpha}{1-\alpha}}\int_m^1\left[(1-\beta)^{\frac{\alpha}{1-\alpha}}m+(1-\beta_V)^{\frac{\alpha}{1-\alpha}}(j-m)\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}dj\\-mf_O-(1-m)f_V\end{array}\right\}$$

where $\Omega=A\frac{\rho}{\alpha}\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}}\left(\frac{\rho\theta}{c}\right)^{\frac{\rho}{1-\rho}}$. This delivers equation (18).

B Data Appendix

Intrafirm trade: From the US Census Bureau's Related Party Trade Database, for the years 2000-2010. The US Census data on related party and non-related party trade volumes are reported for NAICS industries. We map these into six-digit IO2002 industries using the correspondence provided by the Bureau of Economic Analysis (BEA) as a supplement to the 2002 Input-Output Tables. This is a straightforward many-to-one mapping for the manufacturing industries (NAICS first digit = 3), except for one exception. For NAICS code 31131, the US Census data is at a coarser level of aggregation than can be mapped into six-digit IO2002 industries: NAICS 31131 maps into both IO2002 311313 and 31131A. We therefore create a synthetic category called 31131X that merges IO2002 311313 and 31131A; all other industry variables below are also constructed for this synthetic IO2002 category. For NAICS code 33641X, the US Census does not report any import trade value for this industry, so we drop it altogether. After mapping to IO2002 industries, the share of intrafirm imports is then calculated for each industry-year or exporting country-industry-year as: *(Related Trade)/(Related Trade + Non-Related Trade)*.

DUse_TUse: From the 2002 US Input-Output Tables, we obtain the Direct Use Table (D) which is the square matrix whose (i, j)-th entry reports the value of inputs from industry i used directly in producing each dollar of industry j output. The Total Use Table (T) is calculated as: $T = D[I - D]^{-1}$, where I is the identity matrix; the (i, j)-th entry of this table reports the value of industry i input used in total (both directly and indirectly) in the production of each dollar of industry j output. From the Input-Output Tables, we also obtain the vector F of final-use values for each IO2002 industry, where final use is defined as all output from the industry purchased for consumption and investment purposes by either private or government parties. This is the sum over all final use categories in the IO Tables beginning with "F", except the change in private inventories (F03000), exports of goods and services (F04000), and imports of goods and services (F05000). The i-th entry of the column vector DF thus reports the value of industry i input used directly in the economy, after aggregating over all buying industries j; the column vector TF gives us the analogous aggregate values for the total use of industry i inputs summed over all buyer industries j.

The direct use to total use ratio ($DUse\_TUse$) for industry i is then calculated as the i-th entry in DF divided by the i-the entry in TF.

DownIndex: From the 2002 US Input-Output Tables. We first calculate a weighted sum equal to the value of industry i inputs use multiplied by the number of stages from final use that the input is used at. Specifically, the final use value is multiplied by 1, the input use value that then directly goes to final use is multiplied by 2, and so on. Mathematically, this is given by the i-th entry of $[I - D]^{-2}$. Dividing this by the total value of industry i output yields a normalized index measure that is increasing in the upstreamness of an input's use. The $DownIndex$ measure is the reciprocal of this measure of upstreamness, and by construction lies between 0 and 1.

Final Use Share: From the 2002 Input-Output Tables. Calculated as: *(Final Use Value)/(Final Use Value + Intermediate Use Value)*, for each IO2002 industry. The definition of what constitutes final use is as above; intermediate use is the sum by value of that which is purchased for use as inputs by all industries (including the industry itself).

Import demand elasticities: US import demand elasticities for HS10 products are from Broda and Weinstein (2006). Compared against the comprehensive list of HS10 codes in Pierce and Schott (2009), there are a significant number of codes for which the import demand elasticity is not available. For each missing HS10 elasticity, we assign an elasticity value equal to the weighted average of the available elasticities for all HS10 codes with which it shares the same first nine digits. For HS10 elasticities that are still missing, we then assign a value equal to the weighted average of the available elasticities for all HS10 codes with which it shares the same first eight digits. We do this successively up to codes that share the same first two digits, namely at the HS2 level. We map these elasticities to six-digit IO2002 categories using the IO-HS concordance provided by the BEA with the 2002 Input-Output Tables. When multiple HS10 codes map into the same IO2002 industry, we take a weighted average of the HS10 elasticities. At each stage, the weights used in the averaging are proportional to the total US imports by HS10 code from 1989-2006, calculated from Feenstra et al. (2002).

The above procedure yields import demand elasticities for 265 of the 278 IO2002 manufacturing industry codes (first digit = 3). For the remaining 13 industries, we assign a value equal to the weighted average elasticity of IO2002 codes with which it shares the same first 4 digits. If the elasticity is still missing, we successively assign it to be equal to the weighted average elasticity of IO2002 codes with which it shares the same first 3 digits. The weights used here are proportional to the total recorded US output by industry in the 2002 Input-Output Tables.

"Buying" industry variables: For several of the industry variables, most notably the import demand elasticity, we require a measure for each industry i of the average elasticity of "buying" industries, namely industries j that purchase i as an input. We calculate this as a weighted average of the elasticities over all buying industries j, where the weights are proportional to the value of input i used by each industry j, taken from the 2002 IO Direct Use Table. We include the final use value of industry i in this calculation, by assigning it the import demand elasticity of industry i itself. For other industry variables, such as the factor intensities and productivity dispersion, we calculate the variable for "buying" industries in an analogous way.

Factor intensities: From the NBER-CES Manufacturing Industry Database. Skill intensity and physical capital intensity are respectively the log of non-production employment per worker and the log of real capital stock per worker. Equipment capital intensity and plant capital intensity are respectively the log of equipment capital and plant capital per worker. Materials intensity is equal to the log of materials purchases per worker. The latest version of the NBER-CES dataset reports this information for NAICS industries from 1958-2005. We map NAICS industries into IO2002 industries using the same procedure described above for the intrafirm trade flows. For each factor intensity measure, we then take a simple average of the annual values across 2000-2005 to obtain a measure for the input "selling" industry. The factor intensities for "buying industries" are calculated as described above.

R&D intensity: From Nunn and Trefler (2008b), who construct this as the value of R&D expenditures to total sales in an industry, as reported in the Orbis firm-level dataset. We use the measure constructed from US firms in Orbis on an annual basis for IO1997 industries, and transform it to $\log(0.001 + R\&D/Sales)$, so that zero values do not drop out from the dataset. We average this over the years 2000-2005 to obtain an R&D intensity measure for each IO1997 industry. These are mapped to IO2002 industries by constructing a crosswalk through NAICS industries, where the R&D intensity for a given IO2002 industry is equal to weighted average intensity over its constituent IO1997 industries, using the total input use value in the 1997 IO Input-Output Tables as weights. This yields R&D intensity values for 265 of the 278 IO2002 manufacturing industries. A similar procedure to that described above for the import demand elasticity is then used to obtain the dispersion value for the remaining 13 IO2002 codes.

Dispersion: From Nunn and Trefler (2008a), who construct this for each HS6 category as the standard deviation of log exports reported by the US Department of Commerce for each of the constituent HS10 products from each US location to each destination country in the year 2005. We associate the dispersion measure for each HS6 industry to each of its constituent HS10 products. This is then mapped into IO2002 industries using the IO-HS concordance, taking a trade-weighted average of the dispersion value over HS10 constituent codes where necessary; the weights used are the total value of US imports for each HS10 code from 1989-2006, from Feenstra et al. (2002). This yields dispersion values for 265 of the 278 IO2002 manufacturing industries. A similar procedure to that described above for the import demand elasticity is then used to obtain the dispersion value for the remaining 13 IO2002 codes. The dispersion measure for "buying industries" is then calculated as described above.

Contractibility: From the 2002 Input-Output Tables, computed following the methodology of Nunn (2007). We first calculate the fraction of HS10 constituent codes classified by Rauch (1999) as homogeneous goods using Rauch's "liberal" classification. The original Rauch classification is for SITC Rev. 2 products. These are associated with HS10 products by using the mapping in Feenstra et al. (2002), and then finally converted to IO2002 format using the IO-HS concordance. We take one minus this calculated value, to obtain a measure of the own contractibility of each industry. The upstream contractibility faced by each IO2002 industry is the average contractibility of the inputs used, calculated with the input use values from the Direct Use Table as weights. This is thus one minus the zrs^{1l} measure in Nunn (2007), where only homogeneous goods are treated as non-relationship specific. (The results are not sensitive to whether reference-priced goods are treated as relationship-specific, nor to the use of the alternative "conservative" classification in Rauch (1999).)

Following Nunn and Trefler (2008), we also construct a contractibility measure that takes into account differences in country enforcement. For this, we interact the upstream contractibility measure with a rule of law index from Kaufmann et al. (2011). This rule of law index is first linearly rescaled from its original range of -2.5 to 2.5, to lie between 0 and 1. No rule of law data is available for 2001, so we drop this year from the analysis.

Intermediation: From Bernard et al. (2010), who calculate this from micro data as the weighted average of the retail and wholesale employment share at the firm level, using the import share of each firm as weights. This variable is reported at the HS2 level by Bernard et al. We associate the intermediation measure at the HS2 level to each of the HS10 sub-codes. Using the IO-HS concordance, we then compute intermediation for each IO2002 industry as the weighted average over its constituent HS10 products; the weights used are as described above for the factor intensity variables. This yields intermediation values for 265 of the 278 IO2002 manufacturing industries. A similar procedure to that described above for the import demand elasticity is then used to obtain the intermediation value for the remaining 13 IO2002 codes.

Table 1
Tail Values of Measures of Production Line Position (Downstreamness)

IO2002	Industry	DUse_TUse
Lowest 10 values		
331314	Secondary smelting and alloying of aluminum	0.0000
331411	Primary smelting and refining of copper	0.0421
325110	Petrochemical	0.0510
335991	Carbon and graphite product	0.1008
325910	Printing ink	0.1226
33131A	Alumina refining and primary aluminum	0.1249
325211	Plastics material and resin	0.1347
311119	Other animal food	0.1383
333220	Plastics and rubber industry machinery	0.1442
327992	Ground or treated mineral and earth	0.1632
Highest 10 values		
336211	Motor vehicle body	0.9655
336411	Aircraft	0.9669
316200	Footwear	0.9728
337910	Mattress	0.9777
322291	Sanitary paper product	0.9784
337121	Upholstered household furniture	0.9861
337212	Office furniture and custom architectural woodwork & millwork	0.9882
336213	Motor home	0.9891
33299A	Ammunition	0.9961
336111	Automobile	0.9995

IO2002	Industry	DownIndex
Lowest 10 values		
331411	Primary smelting and refining of copper	0.1150
33721A	Wood television, radio and sewing machine cabinet	0.1243
334411	Electron tube	0.1321
313240	Knit fabric mills	0.1436
316100	Leather and hide tanning and finishing	0.1634
331419	Primary smelting and refining of nonferrous metal	0.1687
335991	Carbon and graphite product	0.1697
325110	Petrochemical	0.1699
313220	Narrow fabric mills and schiffli machine embroidery	0.1946
313310	Textile and fabric finishing mills	0.1948
Highest 10 values		
335222	Household refrigerator and home freezer	0.9483
337920	Blind and shade	0.9560
337910	Mattress	0.9703
311111	Dog and cat food	0.9722
321991	Manufactured home (mobile home)	0.9787
336612	Boat building	0.9794
336212	Truck trailer	0.9816
336213	Motor home	0.9903
337121	Upholstered household furniture	0.9922
336111	Automobile	0.9993

Notes: Tabulated based on the 257 IO2002 manufacturing industries for which data on intrafirm import shares is available.

Table 2
Baseline Determinants of the Intrafirm Import Share

	Dependent variable:	**Intrafirm Import Share**		
	(1)	(2)	(3)	(4)
Industry controls for:	Seller	Seller	Buyer	Buyer
Log (s/l)	0.096**	0.107***	-0.011	0.013
	[0.038]	[0.037]	[0.042]	[0.042]
Log (k/l)	0.035*		0.049*	
	[0.021]		[0.028]	
Log (equipment k / l)		0.075***		0.091***
		[0.024]		[0.031]
Log (plant k / l)		-0.059*		-0.074*
		[0.032]		[0.045]
Log (materials/l)	0.025	0.031	0.053	0.062*
	[0.027]	[0.026]	[0.034]	[0.034]
Log (0.001+ R&D/Sales)	0.032***	0.033***	0.056***	0.055***
	[0.007]	[0.006]	[0.009]	[0.009]
Dispersion	0.109*	0.130**	0.090	0.131*
	[0.060]	[0.058]	[0.070]	[0.074]
Year fixed effects	Yes	Yes	Yes	Yes
Observations	2823	2823	2823	2823
R-squared	0.28	0.29	0.26	0.27

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. All columns use industry-year observatons controlling for year fixed effects. Estimation is by OLS. The industry factor intensity and dispersion variables in Columns 1-2 are that of the seller industry (namely, the industry that sells the input in question), while in Columns 3-4, these variables are a weighted average of the characteristics of buyer industries (the industries that buy the input in question).

Table 3
Downstreamness and the Intrafirm Import Share: DUse_TUse

	Dependent variable: Intrafirm Import Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Imp < Median	Imp >= Median	Weighted		Weighted
Log (s/l)	0.012	0.045	0.010	0.051	-0.073	0.002	-0.050
	[0.042]	[0.042]	[0.054]	[0.054]	[0.070]	[0.019]	[0.067]
Log (equipment k / l)	0.095***	0.099***	0.088**	0.158***	0.193***	0.029*	0.155***
	[0.035]	[0.034]	[0.042]	[0.047]	[0.062]	[0.016]	[0.052]
Log (plant k / l)	-0.077*	-0.092**	-0.064	-0.140**	-0.141**	-0.061***	-0.139***
	[0.046]	[0.046]	[0.061]	[0.054]	[0.068]	[0.019]	[0.046]
Log (materials/l)	0.061*	0.058*	0.023	0.058	0.036	0.024*	0.044
	[0.034]	[0.034]	[0.041]	[0.046]	[0.057]	[0.015]	[0.047]
Log (0.001+ R&D/Sales)	0.055***	0.053***	0.033***	0.058***	0.079***	0.031***	0.063***
	[0.009]	[0.009]	[0.012]	[0.013]	[0.017]	[0.004]	[0.014]
Dispersion	0.133*	0.126	0.020	0.239*	0.225	0.107***	0.185
	[0.075]	[0.077]	[0.092]	[0.127]	[0.147]	[0.040]	[0.125]
DUse_TUse	0.014						
	[0.055]						
DUse_TUse X 1(Elas < Median)		-0.128*	-0.111	-0.128	-0.141	-0.090***	-0.036
		[0.070]	[0.080]	[0.092]	[0.100]	[0.033]	[0.086]
DUse_TUse X 1(Elas > Median)		0.175***	0.069	0.298***	0.481***	-0.049	0.362***
		[0.067]	[0.074]	[0.094]	[0.121]	[0.030]	[0.112]
1(Elas > Median)		-0.143**	-0.072	-0.231***	-0.372***	-0.019	-0.257***
		[0.059]	[0.069]	[0.080]	[0.083]	[0.028]	[0.072]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects	No	No	No	No	No	Yes	Yes
Observations	2823	2823	1411	1412	2823	209592	209592
R-squared	0.28	0.31	0.14	0.46	0.60	0.18	0.59

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-5 use industry-year observations controlling for year fixed effects, while Columns 6-7 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. In all columns, the industry factor intensity and dispersion variables used are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 3 and 4 restrict the sample to observations where total imports (the denominator of the intrafirm import share) are smaller (respectively larger) than the sample median value. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 4
Downstreamness and the Intrafirm Import Share: DownIndex

	Dependent variable: Intrafirm Import Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Imp < Median	Imp >= Median	Weighted		Weighted
Log (s/l)	-0.003	0.022	-0.006	0.017	-0.141**	-0.001	-0.095**
	[0.042]	[0.041]	[0.052]	[0.053]	[0.055]	[0.020]	[0.046]
Log (equipment k / l)	0.126***	0.136***	0.121***	0.156***	0.172***	0.043***	0.157***
	[0.033]	[0.034]	[0.043]	[0.041]	[0.056]	[0.015]	[0.040]
Log (plant k / l)	-0.099**	-0.114**	-0.087	-0.118**	-0.098	-0.067***	-0.120***
	[0.046]	[0.048]	[0.063]	[0.056]	[0.077]	[0.020]	[0.045]
Log (materials/l)	0.057*	0.046	0.017	0.035	0.009	0.017	0.018
	[0.032]	[0.032]	[0.041]	[0.042]	[0.050]	[0.014]	[0.038]
Log (0.001+ R&D/Sales)	0.059***	0.057***	0.040***	0.055***	0.086***	0.033***	0.068***
	[0.010]	[0.010]	[0.012]	[0.014]	[0.017]	[0.004]	[0.013]
Dispersion	0.157**	0.166**	0.034	0.337**	0.309*	0.113***	0.257**
	[0.075]	[0.078]	[0.091]	[0.145]	[0.160]	[0.041]	[0.125]
DownIndex	0.144***						
	[0.051]						
DownIndex X 1(Elas < Median)		0.059	0.071	-0.018	-0.055	0.012	0.042
		[0.055]	[0.064]	[0.077]	[0.091]	[0.029]	[0.080]
DownIndex X 1(Elas > Median)		0.257***	0.140	0.356***	0.490***	-0.054	0.414***
		[0.078]	[0.086]	[0.103]	[0.089]	[0.033]	[0.077]
1(Elas > Median)		-0.060	0.005	-0.173**	-0.283***	0.042	-0.209***
		[0.054]	[0.061]	[0.075]	[0.074]	[0.027]	[0.066]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects	No	No	No	No	No	Yes	Yes
Observations	2823	2823	1411	1412	2823	209592	209592
R-squared	0.30	0.32	0.14	0.47	0.63	0.18	0.61

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-5 use industry-year observations controlling for year fixed effects, while Columns 6-7 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. In all columns, the industry factor intensity and dispersion variables used are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 3 and 4 restrict the sample to observations where total imports (the denominator of the intrafirm import share) are smaller (respectively larger) than the sample median value. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 5
Effect of Downstreamness: By Import Elasticity Quintiles

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse Weighted	DUse_TUse Weighted	DownIndex	DownIndex Weighted	DownIndex Weighted
Log (s/l)	0.031	-0.054	-0.042	0.013	-0.141***	-0.109***
	[0.042]	[0.072]	[0.067]	[0.041]	[0.054]	[0.040]
Log (equipment k / l)	0.096***	0.171***	0.132**	0.129***	0.128**	0.111**
	[0.035]	[0.063]	[0.054]	[0.035]	[0.056]	[0.043]
Log (plant k / l)	-0.079*	-0.125**	-0.124***	-0.103**	-0.067	-0.089**
	[0.045]	[0.058]	[0.040]	[0.048]	[0.059]	[0.036]
Log (materials/l)	0.063*	0.046	0.058	0.055	0.030	0.045
	[0.034]	[0.054]	[0.046]	[0.033]	[0.047]	[0.040]
Log (0.001+ R&D/Sales)	0.056***	0.077***	0.062***	0.058***	0.086***	0.070***
	[0.009]	[0.016]	[0.014]	[0.010]	[0.017]	[0.013]
Dispersion	0.134*	0.185	0.150	0.175**	0.268*	0.225*
	[0.075]	[0.151]	[0.128]	[0.077]	[0.152]	[0.117]
Downstream X 1(Elas Quintile 1)	-0.180*	-0.305**	-0.155	0.035	-0.283	-0.099
	[0.093]	[0.140]	[0.100]	[0.115]	[0.173]	[0.122]
Downstream X 1(Elas Quintile 2)	-0.118	-0.048	0.027	0.039	-0.001	0.079
	[0.110]	[0.143]	[0.128]	[0.073]	[0.118]	[0.101]
Downstream X 1(Elas Quintile 3)	-0.067	0.030	0.066	0.032	0.097	0.171
	[0.119]	[0.158]	[0.149]	[0.114]	[0.164]	[0.164]
Downstream X 1(Elas Quintile 4)	0.226**	0.300**	0.252***	0.342***	0.189	0.211*
	[0.102]	[0.152]	[0.094]	[0.119]	[0.186]	[0.120]
Downstream X 1(Elas Quintile 5)	0.173*	0.693***	0.546***	0.267**	0.654***	0.553***
	[0.100]	[0.190]	[0.190]	[0.125]	[0.101]	[0.083]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Elas Quintile dummies?	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	No	Yes	Yes	No
Country-Year fixed effects	No	No	Yes	No	No	Yes
Observations	2823	2823	209592	2823	2823	209592
R-squared	0.32	0.63	0.60	0.33	0.66	0.62

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-2 and 4-5 use industry-year observations controlling for year fixed effects, while Columns 3 and 6 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. The downstreamness measure used is indicated under each column number. In all columns, the industry factor intensity and dispersion variables are a weighted average of the characteristics of buyer industries (namely, the industries that buy the input in question). The main effects of the import elasticity quintile dummies are controlled for, but not reported. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 6
Extension: Effect of Headquarter Intensity

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse	DUse_TUse	DownIndex	DownIndex	DownIndex
Buyer industry hq intensity:	First Principal Component Log (s/l), Log (equipment k/l), and Log (0.001+R&D/Sales)					
		Weighted	Weighted		Weighted	Weighted
Log (s/l)	0.191***	0.146*	0.163**	0.211***	0.174**	0.169**
	[0.069]	[0.088]	[0.079]	[0.069]	[0.079]	[0.069]
Log (equipment k / l)	0.152***	0.172**	0.155***	0.194***	0.202***	0.187***
	[0.036]	[0.080]	[0.058]	[0.035]	[0.072]	[0.050]
Log (plant k / l)	-0.108**	-0.101	-0.127**	-0.150***	-0.106	-0.141***
	[0.046]	[0.089]	[0.059]	[0.046]	[0.082]	[0.053]
Log (materials/l)	0.046	0.035	0.036	0.045	0.032	0.035
	[0.031]	[0.047]	[0.039]	[0.029]	[0.041]	[0.034]
Log (0.001+ R&D/Sales)	0.084***	0.110***	0.088***	0.092***	0.135***	0.106***
	[0.018]	[0.026]	[0.022]	[0.018]	[0.025]	[0.020]
Dispersion	0.090	0.140	0.093	0.114	0.111	0.080
	[0.073]	[0.091]	[0.084]	[0.076]	[0.102]	[0.088]
Downstream X 1 (Elas < Med) X (HQ Quintile 1)	0.101	0.072	0.071	0.251**	0.253**	0.208**
	[0.120]	[0.133]	[0.121]	[0.109]	[0.124]	[0.104]
Downstream X 1 (Elas < Med) X (HQ Quintile 2)	-0.236**	-0.482**	-0.308*	-0.080	-0.163	-0.031
	[0.107]	[0.196]	[0.169]	[0.079]	[0.177]	[0.141]
Downstream X 1 (Elas < Med) X (HQ Quintile 3)	-0.044	-0.169	-0.162	0.129	-0.017	-0.045
	[0.177]	[0.148]	[0.134]	[0.133]	[0.152]	[0.139]
Downstream X 1 (Elas < Med) X (HQ Quintile 4)	-0.394***	-0.224	-0.049	-0.077	-0.130	0.014
	[0.110]	[0.185]	[0.156]	[0.117]	[0.183]	[0.158]
Downstream X 1 (Elas < Med) X (HQ Quintile 5)	0.207	-0.180	-0.053	0.130	-0.303**	-0.142
	[0.173]	[0.205]	[0.183]	[0.129]	[0.124]	[0.139]
Downstream X 1 (Elas > Med) X (HQ Quintile 1)	0.172	0.831**	0.807***	0.369**	0.570***	0.539***
	[0.230]	[0.325]	[0.279]	[0.151]	[0.083]	[0.064]
Downstream X 1 (Elas > Med) X (HQ Quintile 2)	-0.033	-0.088	-0.188	0.002	-0.055	-0.056
	[0.141]	[0.157]	[0.152]	[0.110]	[0.205]	[0.173]
Downstream X 1 (Elas > Med) X (HQ Quintile 3)	0.201*	0.046	-0.046	0.351***	0.295	0.100
	[0.119]	[0.214]	[0.146]	[0.123]	[0.238]	[0.181]
Downstream X 1 (Elas > Med) X (HQ Quintile 4)	0.170	0.714*	0.389	0.142	0.234	-0.036
	[0.152]	[0.368]	[0.305]	[0.185]	[0.372]	[0.295]
Downstream X 1 (Elas > Med) X (HQ Quintile 5)	0.099	0.174	0.106	0.115	0.157	0.152
	[0.103]	[0.150]	[0.105]	[0.127]	[0.172]	[0.100]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Main and double interaction effects?	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	No	Yes	Yes	No
Country-Year fixed effects	No	No	Yes	No	No	Yes
Observations	2823	2823	209592	2823	2823	209592
R-squared	0.40	0.68	0.63	0.40	0.71	0.64

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-2 and 4-5 use industry-year observations controlling for year fixed effects, while Columns 3 and 6 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. The downstreamness measure used is indicated under each column number. The hq intensity measure is the first principal component of the buyer industry Log (s/l), Log (equipment k/l) and Log (R&D/Sales). In all columns, the industry factor intensity and dispersion variables are a weighted average of the characteristics of buyer industries (namely, the industries that buy the input in question). The main effects and double interaction effects of the import elasticity median and hq intensity quintile dummies are controlled for, but not reported. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 7
Extension: Implications of Firm Heterogeneity

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse Weighted	DUse_TUse Weighted	DownIndex	DownIndex Weighted	DownIndex Weighted
Log (s/l)	0.049	-0.075	-0.049	0.028	-0.086*	-0.060
	[0.040]	[0.057]	[0.051]	[0.042]	[0.050]	[0.043]
Log (equipment k / l)	0.106***	0.163**	0.127**	0.114***	0.115*	0.111***
	[0.034]	[0.068]	[0.055]	[0.035]	[0.061]	[0.042]
Log (plant k / l)	-0.105**	-0.159**	-0.153***	-0.100**	-0.075	-0.107**
	[0.046]	[0.075]	[0.054]	[0.048]	[0.079]	[0.054]
Log (materials/l)	0.060*	0.064	0.067	0.056*	0.049	0.048
	[0.034]	[0.053]	[0.045]	[0.033]	[0.047]	[0.038]
Log (R&D/Sales)	0.053***	0.075***	0.058***	0.058***	0.089***	0.076***
	[0.009]	[0.016]	[0.013]	[0.009]	[0.015]	[0.012]
Dispersion	0.129*	0.292*	0.240*	0.140*	0.237	0.192*
	[0.073]	[0.156]	[0.127]	[0.075]	[0.148]	[0.112]
(Downstream Quin 1) X 1(Elas < Median)	0.068	0.205***	0.130***	-0.029	0.102	0.082
	[0.045]	[0.060]	[0.048]	[0.048]	[0.067]	[0.059]
(Downstream Quin 2) X 1(Elas < Median)	0.011	0.137**	0.110**	0.042	0.086	0.044
	[0.043]	[0.059]	[0.051]	[0.049]	[0.078]	[0.061]
(Downstream Quin 3) X 1(Elas < Median)	-0.006	0.081	0.033	-0.044	0.063	0.049
	[0.043]	[0.059]	[0.047]	[0.052]	[0.066]	[0.051]
(Downstream Quin 4) X 1(Elas < Median)	-0.065	0.026	0.012	0.011	0.022	0.022
	[0.045]	[0.086]	[0.067]	[0.045]	[0.071]	[0.057]
(Downstream Quin 5) X 1(Elas < Median)	0.012	0.113*	0.114**	0.005	0.003	0.054
	[0.042]	[0.063]	[0.052]	[0.045]	[0.068]	[0.051]
(Downstream Quin 2) X 1(Elas > Median)	0.013	0.053	0.025	-0.005	-0.005	-0.038
	[0.036]	[0.062]	[0.051]	[0.041]	[0.057]	[0.046]
(Downstream Quin 3) X 1(Elas > Median)	0.032	0.123*	0.047	-0.009	-0.058	-0.022
	[0.045]	[0.074]	[0.053]	[0.052]	[0.083]	[0.074]
(Downstream Quin 4) X 1(Elas > Median)	0.057	0.129*	0.071	0.040	0.047	0.007
	[0.050]	[0.067]	[0.060]	[0.045]	[0.064]	[0.047]
(Downstream Quin 5) X 1(Elas > Median)	0.150***	0.337***	0.248***	0.170***	0.349***	0.289***
	[0.048]	[0.074]	[0.068]	[0.060]	[0.083]	[0.068]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	No	Yes	Yes	No
Country-Year fixed effects	No	No	Yes	No	No	Yes
Observations	2823	2823	209592	2823	2823	209592
R-squared	0.33	0.62	0.60	0.33	0.65	0.62

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-2 and 4-5 use industry-year observations controlling for year fixed effects, while Columns 3 and 6 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. The downstreamness measure used is indicated under each column number. In all columns, the industry factor intensity and dispersion variables are a weighted average of the characteristics of buyer industries (namely, the industries that buy the input in question). "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 8
Heckman selection (Year = 2005)

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse	DUse_TUse	DownIndex	DownIndex	DownIndex
	Probit	Weighted	Weighted	Probit	Weighted	Weighted
Log (s/l)	0.307*	-0.054	-0.023	0.245	-0.096*	-0.073
	[0.170]	[0.085]	[0.074]	[0.171]	[0.050]	[0.048]
Log (equipment k / l)	-0.010	0.161***	0.164***	-0.048	0.175***	0.169***
	[0.190]	[0.056]	[0.058]	[0.164]	[0.044]	[0.045]
Log (plant k / l)	-0.169	-0.127***	-0.155***	-0.117	-0.123***	-0.140***
	[0.226]	[0.047]	[0.048]	[0.210]	[0.046]	[0.046]
Log (materials/l)	-0.049	0.060	0.044	-0.051	0.029	0.016
	[0.135]	[0.050]	[0.049]	[0.133]	[0.040]	[0.040]
Buyer Log (0.001+ R&D/Sales)	0.022	0.051***	0.058***	0.036	0.057***	0.064***
	[0.065]	[0.014]	[0.015]	[0.061]	[0.013]	[0.013]
Dispersion	-0.307	0.250*	0.225	-0.406	0.342**	0.306**
	[0.391]	[0.142]	[0.143]	[0.394]	[0.137]	[0.142]
Downstream X 1(Elas < Median)	0.090	-0.054	-0.031	0.495*	-0.026	0.040
	[0.301]	[0.081]	[0.085]	[0.299]	[0.087]	[0.086]
Downstream X 1(Elas > Median)	0.109	0.358***	0.371***	-0.493**	0.513***	0.442***
	[0.252]	[0.116]	[0.118]	[0.232]	[0.085]	[0.075]
1(Imp Elas > Median)	0.114	-0.259***	-0.237***	0.635***	-0.294***	-0.206***
	[0.228]	[0.074]	[0.072]	[0.206]	[0.078]	[0.068]
Entry Costs X Seller R&D int.	-0.015*			-0.015*		
	[0.008]			[0.008]		
Seller Log (0.001+ R&D/Sales)	0.017			0.012		
	[0.043]			[0.038]		
Inverse Mills Ratio		-0.919**			-0.773***	
		[0.442]			[0.276]	
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Country fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	44461	18088	18088	44461	18088	18088
R-squared / Psuedo R-squared	0.42	0.59	0.59	0.42	0.62	0.62

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1 and 4 report the first-stage probits; our exclusion restriction variable is the seller industry's R&D intensity and its interaction with a dummy variable for countries with above sample-median entry costs. Columns 2 and 5 report the second stage of the Heckman seletion, where observations have been weighted by the value of total imports for the country-industry in question. Columns 3 and 6 report the weighted OLS regression excluding the inverse Mills ratio term, to provide a comparison. The downstreamness measure used is indicated under each column number. In all columns, all other industry factor intensity and dispersion variables are a weighted average of the characteristics of buyer industries (namely, the industries that buy the input in question).

Table 9
Robustness: Contractibility and Intermediation

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse Weighted	DUse_TUse Weighted	DownIndex	DownIndex Weighted	DownIndex Weighted
Log (s/l)	0.016	-0.061	-0.039	0.000	-0.081*	-0.060
	[0.041]	[0.055]	[0.058]	[0.040]	[0.042]	[0.044]
Log (equipment k / l)	0.116***	0.146***	0.147***	0.133***	0.143***	0.150***
	[0.034]	[0.050]	[0.044]	[0.035]	[0.050]	[0.041]
Log (plant k / l)	-0.056	-0.008	-0.057	-0.068	0.002	-0.061
	[0.045]	[0.053]	[0.039]	[0.048]	[0.060]	[0.041]
Log (materials/l)	0.023	-0.033	-0.020	0.017	-0.046	-0.028
	[0.030]	[0.040]	[0.036]	[0.030]	[0.039]	[0.033]
Log (0.001+ R&D/Sales)	0.037***	0.051***	0.046***	0.039***	0.053***	0.049***
	[0.009]	[0.013]	[0.013]	[0.010]	[0.012]	[0.011]
Dispersion	-0.020	0.049	0.072	0.010	0.093	0.120
	[0.077]	[0.127]	[0.115]	[0.079]	[0.138]	[0.118]
Downstream X 1(Elas < Median)	-0.104	-0.204**	-0.113	0.037	-0.116	-0.026
	[0.071]	[0.089]	[0.080]	[0.057]	[0.087]	[0.085]
Downstream X 1(Elas > Median)	0.148**	0.128	0.116*	0.194***	0.196**	0.229***
	[0.062]	[0.082]	[0.069]	[0.070]	[0.088]	[0.076]
1(Imp Elas > Median)	-0.112**	-0.198***	-0.150**	-0.041	-0.164**	-0.146**
	[0.055]	[0.072]	[0.065]	[0.049]	[0.068]	[0.063]
Upstream Contractability	-0.293***	-0.531***	-0.677***	-0.240***	-0.463***	-0.601**
	[0.081]	[0.090]	[0.220]	[0.081]	[0.107]	[0.247]
Upstream Contractability X Rule of Law			0.312			0.367
			[0.348]			[0.346]
Intermediation	-0.431***	-0.532***	-0.311**	-0.453***	-0.533***	-0.321**
	[0.098]	[0.145]	[0.124]	[0.100]	[0.152]	[0.124]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	No	Yes	Yes	No
Country-Year fixed effects	No	No	Yes	No	No	Yes
Observations	2823	2823	184819	2823	2823	184819
R-squared	0.39	0.72	0.62	0.39	0.72	0.63

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-2 and 4-5 use industry-year observations controlling for year fixed effects, while Columns 3 and 6 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. The downstreamness measure used is indicated under each column number. In all columns, the industry factor intensity and dispersion variables are a weighted average of the characteristics of buyer industries (namely, the industries that buy the input in question). "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights. Observations for 2001 are dropped due to the absence of rule of law data for that year.

Appendix Table 1
Summary Statistics: By IO2002 Manufacturing Categories (257 Industries)

Variable	10th	25th	Median	75th	90th	Mean	Std. Dev.
Share of Intrafirm trade (year=2000)	0.105	0.205	0.365	0.535	0.659	0.382	0.210
Share of Intrafirm trade (year=2005)	0.126	0.221	0.384	0.554	0.650	0.390	0.203
Share of Intrafirm trade (year=2010)	0.131	0.231	0.403	0.559	0.663	0.401	0.209
Of Seller Industries:							
DUse_TUse	0.239	0.426	0.645	0.788	0.887	0.606	0.235
DownIndex	0.242	0.323	0.460	0.709	0.898	0.519	0.238
Skill Intensity (Log(s/l))	-1.723	-1.541	-1.292	-0.984	-0.744	-1.267	0.388
Physical Capital Intensity (Log(k/l))	3.875	4.246	4.747	5.263	6.091	4.836	0.819
Log(equipment k / l)	3.271	3.785	4.311	4.836	5.664	4.365	0.898
Log(plant k / l)	2.930	3.283	3.671	4.206	4.855	3.802	0.754
Materials intensity (Log(materials/l))	4.054	4.311	4.734	5.258	5.711	4.839	0.716
R&D intensity (Log(0.001+R&D/Sales))	-6.908	-6.908	-6.150	-4.300	-2.912	-5.430	1.755
Dispersion	1.636	1.745	1.847	1.990	2.186	1.888	0.232
Of Buyer Industries:							
Import elasticity (ρ)	3.154	4.938	7.758	10.532	18.465	10.252	10.926
Skill Intensity (Log(s/l))	-1.708	-1.485	-1.293	-1.020	-0.732	-1.250	0.356
Physical Capital Intensity (Log(k/l))	3.999	4.393	4.751	5.131	5.574	4.769	0.626
Log(equipment k / l)	3.411	3.863	4.322	4.682	5.142	4.281	0.699
Log(plant k / l)	3.054	3.372	3.682	4.049	4.566	3.755	0.572
Materials intensity (Log(materials/l))	4.212	4.546	4.868	5.222	5.643	4.904	0.575
R&D intensity (Log(0.001+R&D/Sales))	-6.904	-6.639	-5.600	-4.543	-3.328	-5.396	1.353
Dispersion	1.708	1.787	1.905	2.011	2.120	1.910	0.181

Notes: For details on the construction of the data variables, please see the Data Appendix.

Appendix Table 2
Industry Variables: Correlation Coefficients

	Wtd. DUse	Buyer Industry Import Elasticity (ρ)	Buyer Industry Log(s/l)	Buyer Industry Log(k/l)	Buyer Industry Log(equip k/l)	Buyer Industry Log(plant k/l)	Buyer Industry Log(mat/l)	Buyer Industry R&D Intensity	Buyer Industry Dispersion
DUse_TUse	0.603***	0.054	-0.046	-0.250***	-0.286***	-0.158**	-0.090	-0.150**	-0.143**
Wtd. DUse		0.116*	0.047	-0.290***	-0.331***	-0.186***	-0.172***	-0.117*	-0.148**
Buyer Industry Import Elasticity (ρ)			-0.161***	0.017	-0.032	0.078	0.227***	-0.085	0.182***
Buyer Industry Log(s/l)				0.159**	0.101	0.268***	0.050	0.610***	0.325***
Buyer Industry Log(k/l)					0.980***	0.944***	0.755***	0.111*	0.393***
Buyer Industry Log(equip k/l)						0.864***	0.725***	0.081	0.321***
Buyer Industry Log(plant k/l)							0.746***	0.167***	0.471***
Buyer Industry Log(mat/l)								0.065	0.375***
Buyer Industry R&D Intensity									0.316***

Notes: ***, ** and * indicate significance at the 1%, 5% and 10% levels respectively.

Appendix Table 3
Downstreamness and the Intrafirm Import Share: FDUse

	Dependent variable: Intrafirm Import Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Imp < Median	Imp >= Median	Weighted		Weighted
Log (s/l)	0.013	0.040	0.001	0.050	-0.133	-0.005	-0.097
	[0.041]	[0.041]	[0.052]	[0.055]	[0.089]	[0.019]	[0.077]
Log (equipment k / l)	0.097***	0.105***	0.099**	0.129***	0.142**	0.036**	0.106**
	[0.034]	[0.034]	[0.045]	[0.043]	[0.055]	[0.016]	[0.045]
Log (plant k / l)	-0.079*	-0.084*	-0.066	-0.101*	-0.089	-0.064***	-0.093**
	[0.046]	[0.046]	[0.064]	[0.052]	[0.063]	[0.020]	[0.043]
Log (materials/l)	0.061*	0.050	0.015	0.056	0.050	0.019	0.065
	[0.034]	[0.034]	[0.041]	[0.048]	[0.058]	[0.014]	[0.047]
Log (0.001+ R&D/Sales)	0.054***	0.052***	0.036***	0.051***	0.078***	0.033***	0.062***
	[0.009]	[0.009]	[0.012]	[0.013]	[0.018]	[0.004]	[0.015]
Dispersion	0.135*	0.132*	0.023	0.247*	0.227	0.115***	0.174
	[0.075]	[0.076]	[0.093]	[0.127]	[0.141]	[0.041]	[0.120]
FDUse	0.015						
	[0.036]						
FDUse X 1(Elas < Median)		-0.044	-0.021	-0.124	-0.145	-0.019	-0.073
		[0.034]	[0.035]	[0.087]	[0.089]	[0.021]	[0.078]
FDUse X 1(Elas > Median)		0.165***	0.071	0.223**	0.342***	-0.054**	0.225**
		[0.059]	[0.068]	[0.089]	[0.105]	[0.027]	[0.093]
1(Elas > Median)		-0.120**	-0.030	-0.250***	-0.376***	0.033	-0.255***
		[0.055]	[0.062]	[0.095]	[0.095]	[0.029]	[0.079]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects	No	No	No	No	No	Yes	Yes
Observations	2823	2823	1411	1412	2823	209592	209592
R-squared	0.28	0.31	0.13	0.45	0.56	0.18	0.57

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-5 use industry-year observations controlling for year fixed effects, while Columns 6-7 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. In all columns, the industry factor intensity and dispersion variables used are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 3 and 4 restrict the sample to observations where total imports (the denominator of the intrafirm import share) are smaller (respectively larger) than the sample median value. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Appendix Table 4
Downstreamness and the Intrafirm Import Share: FUse

	Dependent variable: Intrafirm Import Share						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
			Imp < Median	Imp >= Median	Weighted		Weighted
Log (s/l)	-0.005	0.017	-0.008	0.013	-0.185***	-0.006	-0.126**
	[0.042]	[0.041]	[0.052]	[0.054]	[0.069]	[0.020]	[0.055]
Log (equipment k / l)	0.127***	0.131***	0.129***	0.129***	0.131**	0.052***	0.123***
	[0.031]	[0.032]	[0.041]	[0.040]	[0.057]	[0.014]	[0.039]
Log (plant k / l)	-0.102**	-0.109**	-0.096	-0.101*	-0.073	-0.076***	-0.096**
	[0.046]	[0.047]	[0.061]	[0.055]	[0.077]	[0.020]	[0.044]
Log (materials/l)	0.059*	0.048	0.018	0.049	0.031	0.016	0.036
	[0.032]	[0.033]	[0.042]	[0.045]	[0.054]	[0.013]	[0.040]
Log (0.001+ R&D/Sales)	0.058***	0.056***	0.040***	0.055***	0.088***	0.034***	0.068***
	[0.010]	[0.009]	[0.012]	[0.014]	[0.019]	[0.004]	[0.014]
Dispersion	0.160**	0.162**	0.043	0.308**	0.295*	0.117***	0.238*
	[0.077]	[0.078]	[0.091]	[0.145]	[0.159]	[0.043]	[0.125]
FUse	0.097***						
	[0.032]						
FUse X 1(Elas < Median)		0.059	0.075*	-0.015	-0.058	0.033	0.011
		[0.039]	[0.045]	[0.055]	[0.067]	[0.020]	[0.058]
FUse X 1(Elas > Median)		0.144***	0.094*	0.173***	0.277***	-0.023	0.241***
		[0.045]	[0.056]	[0.060]	[0.070]	[0.019]	[0.056]
1(Elas > Median)		0.009	0.033	-0.061	-0.144***	0.031*	-0.117**
		[0.033]	[0.039]	[0.047]	[0.052]	[0.017]	[0.046]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects	No	No	No	No	No	Yes	Yes
Observations	2823	2823	1411	1412	2823	209592	209592
R-squared	0.30	0.31	0.15	0.44	0.60	0.18	0.60

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-5 use industry-year observations controlling for year fixed effects, while Columns 6-7 use country-industry-year observations controlling for country-year fixed effects. Estimation is by OLS. In all columns, the industry factor intensity and dispersion variables used are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 3 and 4 restrict the sample to observations where total imports (the denominator of the intrafirm import share) are smaller (respectively larger) than the sample median value. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.