

Organizing the Global Value Chain*

Pol Antràs

Harvard University

Davin Chor

Singapore Management University

March 18, 2013

Abstract

We develop a property-rights model of the firm in which production entails a continuum of uniquely sequenced 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 a sharp characterization for the optimal allocation of ownership rights along the value chain. We show that the incentive to integrate suppliers varies systematically with the relative position (upstream versus downstream) at which the supplier enters the production line. Furthermore, the nature of the relationship between integration and “downstreamness” depends crucially on the elasticity of demand faced by the final-good producer. Our model readily accommodates 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 with international trade statistics. Combining data from the U.S. Census Bureau’s Related Party Trade database and estimates of U.S. import demand elasticities from Broda and Weinstein (2006), we find empirical evidence broadly supportive of our key predictions. In the process, we develop two novel measures of the average position of an industry in the value chain, which we construct using U.S. Input-Output Tables.

*We thank the editor and three anonymous referees for their helpful comments and suggestions. We are also grateful to Arnaud Costinot, Don Davis, Ron Findlay, Elhanan Helpman, Kala Krishna, Marc Melitz, Daniel Treffer, Jonathan Vogel, and David Weinstein, as well as audiences at Chicago, Columbia, Harvard, Northeastern, Notre Dame, NYU Stern, Princeton, Stanford, UCLA, UNH, Wisconsin, Kiel, Munich, Tübingen, Bonn, City University of Hong Kong, HKUST, Nanyang Technological University, National University of Singapore, Singapore Management University, the Econometric Society World Congress (Shanghai), the Society for the Advancement of Economic Theory Conference (Singapore), the Asia Pacific Trade Seminars (Honolulu), the Australasian Trade Workshop (UNSW), and the NBER Summer Institute. We thank Nathan Nunn for making available his data. Ruiqing Cao, Mira Frick, Gurmeet Singh Ghumann, Frank Schilbach, and Zhicheng Song provided excellent research assistance. Chor acknowledges the research funding provided by a Sing Lun Fellowship at Singapore Management University. 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 products have 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 stages. First, raw materials are converted into basic components, which are next combined with other components to produce more complicated inputs, before themselves being assembled into final goods. This process very much resembles Henry Ford's original Model T production assembly line, but recent revolutionary advances in information and communication technology, coupled with a gradual reduction in natural and man-made trade barriers, now allow such value chains to be 'sliced up' into geographically separated steps.

The implications of such sequential production for the workings of open-economy general equilibrium models have been widely explored in the literature. Several papers, most notably Findlay (1978), Dixit and Grossman (1982), Sanyal (1983), Kremer (1993), Kohler (2004), and Costinot, Vogel and Wang (2013), have emphasized that the pattern of specialization along the value chain has implications for the world income distribution and for how shocks spread across countries. Others, 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 of production where value is 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 (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 way in which production processes are organized across borders. We find this to be a natural premise 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 these tend to be associated with longer lags between the time an order is placed (and the contract is signed) and the time the goods or services are delivered (and the contract is executed). Such transactions moreover often entail significant relationship-specific investments and other sources of lock-in on the part of both buyers and suppliers.¹ The relevance of contracting frictions for the organization of production also now rests on solid empirical underpinnings.²

¹Suppliers 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.

²A recent literature (see, for instance, Nunn, 2007, and Levchenko, 2007) has convincingly documented that contracting institutions are an important determinant of international specialization. Another branch of the trade

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 where production is 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. Each stage is performed by a different 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. The services of 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 relative to 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, so that production at a stage cannot commence until the inputs or components from all 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 to be divided between the firm and the particular stage supplier. 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. However, the optimal allocation of ownership rights does not always entail all production stages being integrated because by reducing the bargaining power of suppliers, integration reduces the incentives of suppliers to invest in the relationship.³

We begin in Section 2 by developing a benchmark model of firm behavior that isolates the role of the degree of “downstreamness” 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*, according to whether higher investment levels by prior suppliers increase or decrease the value of the marginal product of a particular supplier’s investment. 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 price elasticity of demand faced by the final-good producer is sufficiently low, since in such cases, the value of the marginal product of supplier investments falls particularly quickly along the value chain. Whether inputs are

literature, to which our paper will contribute, has also shown that the ownership decisions of multinational firms exhibit 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 Treffer, 2008, 2012, and Bernard et al., 2010).

³Zhang 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).

sequential complements or sequential substitutes turns out to be determined *only* by whether the elasticity of final-good demand is (respectively) higher or lower than the elasticity of substitution across the services provided by the different suppliers’ investments.

The central result of our model is that the optimal pattern of ownership along the value chain 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 stages prior to this cutoff are outsourced, while all stages (if any) after that threshold are integrated. Intuitively, when inputs are sequential complements, the firm chooses to forgo control rights over upstream suppliers in order to incentivize their investment effort, since this generates positive spillovers on the investment decisions to be made by downstream suppliers. When demand is instead sufficiently inelastic, the converse prediction holds: it is optimal to integrate relatively upstream stages, and if outsourcing is observed along the value chain, it necessarily occurs relatively downstream.

In Section 3, we show that these results are robust to alternative contracting and bargaining assumptions, and stem mainly from the sequential nature of production rather than the sequential nature of bargaining. Furthermore, we show that our framework can easily accommodate a hybrid of sequential and modular production processes (or ‘snakes’ and ‘spiders’ in the terminology of Baldwin and Venables, 2010) as well as several other features which have been built into the recent models of global sourcing cited earlier. These include (headquarter) investments by the final-good producer, productivity heterogeneity across final-good producers, productivity and cost differences across suppliers within a production chain, and partial contractibility. 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. The nature and scope of our test is shaped in significant ways by data availability. Although our model does not explicitly distinguish between domestic and offshore sourcing decisions of firms, data on domestic sourcing decisions are not publicly available. We therefore 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 industries.⁴ More specifically, we correlate the share of U.S. intrafirm imports in total U.S. imports reported during the period 2000-2010 with the average degree of “downstreamness” of that industry, and we study whether this dependence is qualitatively different for the sequential complements versus sequential substitutes cases.

We propose two measures of downstreamness, both of which are constructed from 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, while the *total use* for (i, j) is the value of goods from industry i used either

⁴See, for example, Nunn and Trefler (2008, 2012), Bernard et al. (2010), and Díez (2010). Antràs (2011) contains a comprehensive survey of empirical papers using other datasets, including several firm-level studies, that have similarly used the intrafirm import share to capture the propensity towards integration relative to outsourcing.

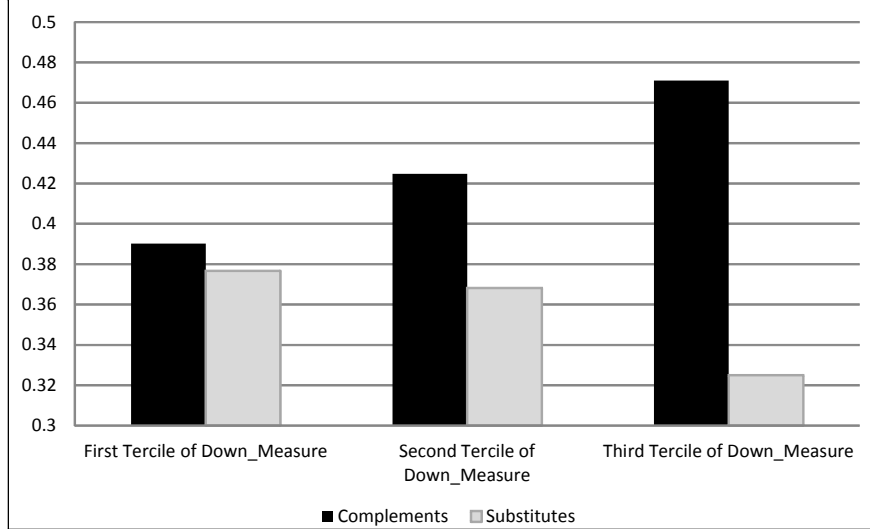


Figure 1: Downstreamness and the Share of Intrafirm Trade

directly or *indirectly* (via purchases from upstream industries) in producing industry j 's output for final use. A high value of $DUse_TUse$ thus suggests that most of the contribution of input i tends to occur at relatively downstream production stages that are close (one stage removed) from final demand. Our second measure of downstreamness ($DownMeasure$) 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 inputs to other industries, as direct inputs 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 $DownMeasure$ can be succinctly expressed as a simple function of the square of the Leontief inverse matrix. As discussed in Antràs et al. (2012), there is a close connection between $DownMeasure$ and the measure of distance to final demand derived independently by Fally (2012).

Our empirical tests also call on us 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 . The idea is that for sufficiently high (respectively, low) values of this average demand elasticity, we can be relatively confident that input substitutability is lower (respectively, higher) than the demand elasticity. Ideally, one would have used direct estimates of cross-input substitutability (and how they compare to demand elasticities) to distinguish between the two theoretical scenarios, but unfortunately these estimates are not readily available in the literature.

Figure 1 provides a preliminary illustration of our key empirical findings, which lend broad support for the theoretical implications of our model. As is apparent from the dark bins, for the subset of industries with above-median average buyer demand elasticities (labeled as "Complements"), the

average U.S. intrafirm import share (for the year 2005) rises as we move from the lowest tercile of *DownMeasure* to the highest. In the light bins, this pattern is exactly reversed when considering those industries facing below-median average buyer demand elasticities (labeled as “Substitutes”), with the intrafirm trade share steadily falling across terciles of *DownMeasure* instead.⁵

Our regression analysis will confirm that the above patterns hold under more formal testing. We uncover a positive and statistically significant relationship between each of the measures of downstreamness and the intrafirm import share in a given sector, with this relationship emerging *only* for high values of the demand elasticity faced by buyer industries (i.e., in the complements case). These findings hold when controlling for other determinants of the intrafirm trade share raised in the literature, and which our theoretical extensions also indicate are important to explicitly consider. They are moreover robust in specifications that further exploit the cross-country dimension of the intrafirm trade data, while controlling for unobserved variation in factor costs with country-year fixed effects. For a wide range of specifications, we will also report a significant *negative* relationship between downstreamness and the intrafirm import share for goods with low average buyer demand elasticities (i.e., in the substitutes case), as predicted by our model.

The remainder of this paper is organized as follows. In Section 2, we develop our benchmark model of sequential production with incomplete contracting and study the optimal ownership structure along the value chain. In Section 3, we develop a few theoretical extensions and discuss how we attempt to take the model 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. All the proofs in the paper are relegated to the Appendix (and an Online Appendix).

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 embed the framework in industry equilibrium to guide the empirical analysis.

2.1 Benchmark Model

Sequential Production. 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 these

⁵It is not hard to find examples of large industries that exhibit similar degrees of downstreamness, but face very different average buyer demand elasticities *and* also very different integration propensities. For instance, Women’s apparel (IO 315230) and Automobiles (IO 336111) are among the ten most downstream manufacturing industries, but buyers tend to be much less price-sensitive in their demand for the former (elasticity=4.90) than for the latter (elasticity=19.02). These two industries are thus classified under the sequential substitutes and complements cases respectively, and consistent with our model, the share of intrafirm trade is low in Women’s apparel (0.108) and very high in Automobiles (0.946). As we shall see later in our econometric analysis, this broad pattern continues to hold when controlling for other industry characteristics that might also affect the propensity toward intrafirm trade.

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 services 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:

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

where θ is a productivity parameter, $\alpha \in (0, 1)$ is a parameter that captures the (symmetric) degree of substitutability among the stage inputs, and $I(j)$ is an indicator function such that:

$$I(j) = \begin{cases} 1 & \text{if input } j \text{ is produced after all inputs } j' < j \text{ have been produced} \\ 0 & \text{otherwise} \end{cases}.$$

We normalize $x(j) = 0$ if an incompatible input is delivered at stage j . Although production requires completion of all stages, note that $\alpha > 0$ ensures that output remains positive even when some stages might be completed with incompatible inputs. In words, although all stages are essential from an engineering point of view, we allow some substitution in how the characteristics of these inputs shape the quality-adjusted volume of final output. For example, producing a car requires four wheels, two headlights, one steering wheel, and so on, but the value of this car for consumers will typically depend on the services obtained from these different components, with a high quality in certain parts partly making up for inferior quality in others.

Our production function in (1) resembles a conventional CES function with a continuum of inputs, but the indicator function $I(j)$ makes the production technology inherently sequential in that downstream stages are useless unless the inputs from upstream stages have been delivered. In fact, the technology in (1) can be expressed in differential form by applying Leibniz's rule as:

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

where $q(m) = \theta \left(\int_0^m x(j)^\alpha I(j) dj \right)^{1/\alpha}$. Thus, the marginal increase in output brought about by the supplier at stage m is given by a simple Cobb-Douglas function of this supplier's (compatible) input production and the quality-adjusted volume of production generated up to that stage (which can be viewed as an intermediate input to the stage- m production process).

Input Production. There is a large number of profit-maximizing suppliers who can either engage in intermediate input production or in an alternative activity that delivers an outside option normalized to 0. We assume that each intermediate input must be produced by a different supplier with whom the firm needs to contract. Each supplier must undertake a relationship-specific investment in order to produce a *compatible* input. For simplicity, we assume that the input is fully customized to the final-good producer, so the value of this input for alternative buyers is equal to 0. To highlight the asymmetries that will arise solely 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 services of the stage j compatible input when combined with the inputs from upstream suppliers. (We will relax these symmetry assumptions later 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 apart from allowing the continuation of the production process.

Preferences. 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 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 accruing from this industry is given by:

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

where Ω denotes the set of varieties. Note that these preferences feature diminishing marginal utility with respect to not only the quantity but also the quality of the goods consumed. As a result, in our previous car example, further quality improvements on a high-end car would not add as much satisfaction to consumers as they would in a low-end car. As is well known, when maximizing (2) 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 a firm that sells variety ω is concave in quality-adjusted output $q(\omega) \equiv \varphi(\omega) \tilde{q}(\omega)$ 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:

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

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 and thus over input services at all stages. 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 input services as stipulated in the contract in exchange for the payment $s(j)$. It is clear that the firm will have an incentive to follow the natural sequencing of production, so that $I(j) = 1$ for all j ,

and the optimal contract simply solves the following maximization program:

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

Solving this problem delivers a common investment level $x = (\rho A^{1-\rho} \theta^\rho / c)^{1/(1-\rho)}$ 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 of zero (i.e., $s = cx$).

Incomplete Contracts. For the above contracts to be enforceable, it is important that a court of law be able to verify the precise value of the input services provided by the suppliers of the different stages. In practice, however, a court of law will generally not be able to verify whether inputs are compatible or not, and whether the services provided by compatible inputs are in accordance with what was stipulated in a written contract. Notice also that the firm might be reluctant to sign binding contracts that are contingent on the quantity of inputs produced but *not* on whether inputs are compatible, because suppliers might 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 remain independent. In Section 3.4, we will briefly discuss the case of partial contractibility, in which some aspects of production (such as the quantity produced) are contractible ex-ante.

Given the lack of a binding contract, a familiar holdup problem emerges. The actual payment to a particular supplier (say for stage m) is negotiated bilaterally only after the stage m input has been produced and the firm has had a chance to inspect it. For the time being, we treat this negotiation independently from the bilateral negotiations that take place at other stages (though we will revisit this assumption in Section 3.1.) Because the intermediate input is assumed compatible only with the firm's output, the supplier's outside option at the bargaining stage is 0. Hence, the quasi-rents over which the firm and the supplier negotiate are given by the incremental contribution to total revenue generated by supplier m at that stage. To compute this incremental contribution, note that the firm has no incentive to approach suppliers in an order different from that dictated by the technological sequencing of production and that it can always unilaterally complete a production stage by producing an incompatible input.⁶ As a result, we have $I(j) = 1$ for all $j < m$, and the

⁶The assumption that the firm is able to complete *any* production stage with incompatible inputs may seem strong but it can be relaxed by considering environments with partial contractibility, as in Grossman and Helpman (2005). For instance, if a fraction of the suppliers' investments is verifiable and contractible, then the firm could use a formal contract to ensure the provision of a minimum amount of compatible input services from the supplier and the production process would never stall.

value of final-good production secured up to stage m is given by:

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

Applying Leibniz's integral rule to this expression, we then have that the incremental contribution of supplier m is given by:

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

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. As in Grossman and Hart (1986), we assume that ownership of suppliers is a source of power, in the sense that the firm is able to extract a higher share of surplus from integrated suppliers than from nonintegrated suppliers. Intuitively, when contracts are incomplete, the fact that an integrating party controls the physical assets used in production will allow that party to dictate a use of these assets that tilts the division of surplus in its favor. To keep our model as tractable as possible, we will not specify in detail the nature of these ex-post negotiations and will simply assume that the firm will obtain a share β_V of the incremental contribution in equation (5) when the supplier is integrated, while only a share $\beta_O < \beta_V$ of that surplus when the supplier is a stand-alone entity.

We now summarize the timeline of the game played by the firm and the continuum of suppliers:⁷

- 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 that stage. After observing the value of this unfinished product (i.e., $r(m)$ in (4)), the supplier chooses an input level, $x(m)$. At the end of the stage, the firm and supplier m bargain over the addition to total revenue that supplier m has contributed at stage m (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 briefly discuss 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 , that these agents are not allowed to exchange lump-sum transfers, and that the terms of exchange are not renegotiated at a

⁷Although we focus throughout on a version of the model with a continuum of production stages, our equilibrium corresponds to the limit $\varepsilon \rightarrow 0$ of a discrete game in which M suppliers each control a measurable range $\varepsilon = 1/M$ of the continuum of intermediate inputs. See Acemoglu et al. (2007) for an analogous derivation.

later stage and do not reflect the outcome of subsequent negotiations between the firm and other suppliers. Although some of these assumptions could be motivated in richer frameworks appealing to the existence of incomplete information and limited commitment frictions (as in Hart and Moore, 1994, or Thomas and Worrall, 1994), these assumptions may admittedly seem special. For these reasons, in Section 3.1, we will explore the robustness of our results to alternative contracting and bargaining assumptions.

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. Denote by $\beta(m)$ the share of the incremental contribution $r'(m)$ that accrues to the firm in its bargaining with supplier m . Our previous discussion implies that:

$$\beta(m) = \begin{cases} \beta_O & \text{if the firm outsources stage } m \\ \beta_V > \beta_O & \text{if the firm integrates stage } m \end{cases}.$$

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:

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

which delivers:

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

The investment made by supplier m is naturally increasing in the demand level, A , the productivity θ of the firm, and the supplier's bargaining share, $1 - \beta(m)$, while it decreases in the investment marginal cost, c . Hence, other things equal, an outsourcing relationship (corresponding to a lower $\beta(m)$) promotes higher investments on the part of supplier m . The effect of the value of production secured 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 prior suppliers, as summarized in $r(m)$, increase the marginal return of supplier m 's own 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 thus refer to $\rho > \alpha$ as the *complements* case and to $\rho < \alpha$ as the *substitutes* case.

Since $\alpha \in (0, 1)$, it is straightforward to verify that from a purely technological point of view, supplier investments are always (weakly) complementary. More precisely, in light of equation (1), $\partial q/\partial x(m)$ is necessarily nondecreasing in the investment decisions of other suppliers $m' \neq m$. Why

is $x(m)$ then negatively affected by prior investments when $\rho < \alpha$? The reason 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 downstream investments, thus effectively turning supplier investments (in quality-adjusted terms) 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, namely when $\rho < \alpha$.

B. Suppliers' Investments Along the Value Chain

Equation (7) characterizes supplier m 's investment level as a function of $r(m)$, 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:

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

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

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

Equation (9) illustrates how the value of production secured up to stage m depends on all upstream organizational decisions, namely the $\beta(j)$ for $j < m$. Finally, plugging this solution into (7) yields:

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

From this expression, it is clear that the outsourcing of stage m (i.e., choosing $\beta(m) = \beta_O < \beta_V$) enhances investment by that stage's supplier, 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 quasi-rents from suppliers. We next turn to study this optimal organizational structure formally.

C. Optimal Organizational Structure

The firm seeks to maximize the amount of revenue it obtains when the good is sold net of all payments made to suppliers along the value chain. The firm's profits can thus be evaluated as $\pi_F = \int_0^1 \beta(j)r'(j)dj$, which after substituting in the expressions from (8) and (9) is given by:

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

It is in turn easily verified that the payoff $\pi_S(m)$ obtained by suppliers (see equation (6)) is always positive, so their participation constraint can be ignored. The firm's decision problem is then:

$$\begin{aligned} \max_{\{\beta(j)\}_{j \in [0,1]}} \quad & \pi_F \\ \text{s.t.} \quad & \beta(j) \in \{\beta_V, \beta_O\}, \end{aligned}$$

namely to choose the organizational mode for each stage j in order to maximize its profits π_F as given in (11).

In order to determine if integration or outsourcing is optimal at a given stage m , it proves useful to follow the approach in Antràs and Helpman (2004, 2008) and consider first the *relaxed* problem in which the firm could freely choose the function $\beta(m)$ from the whole set of piecewise continuously differentiable real-valued functions rather than from those that only take on values in the set $\{\beta_V, \beta_O\}$. Defining:

$$v(j) \equiv \int_0^j (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk, \quad (12)$$

we can write this problem as that of choosing the real-value function v that maximizes the functional:

$$\pi_F(v) = \kappa \int_0^1 \left(1 - v'(j)^{\frac{1-\alpha}{\alpha}} \right) v'(j) v(j)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj, \quad (13)$$

where $\kappa \equiv A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{\rho}{c} \right)^{\frac{\rho}{1-\rho}}$ is a positive constant.⁸ The profit-maximizing function v must then satisfy the Euler-Lagrange condition, which in light of (13) is given by:

$$v^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (v')^{\frac{1-\alpha}{\alpha}-1} \left[v'' + \frac{\rho-\alpha}{1-\rho} \frac{(v')^2}{v} \right] = 0, \quad (14)$$

provided that v' is at least piecewise differentiable. In the Appendix, we show that the profit-maximizing function v must set the term in the square brackets in (14) to 0. Imposing the initial condition $v(0) = 0$ and the transversality condition $v'(1)^{\frac{1-\alpha}{\alpha}} = \alpha$, and using (12), we can then conclude that the optimal division of surplus at stage m , which we denote by $\beta^*(m)$, is simply given by:

$$\beta^*(m) = 1 - \alpha m^{\frac{\alpha-\rho}{\alpha}}. \quad (15)$$

⁸We thank an anonymous referee for suggesting this approach.

It then follows that:

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

Before describing the implications of this Lemma, it is worth pausing to briefly discuss two technical issues on which we further elaborate in the Online Appendix. First, equation (15) was derived appealing to the Euler-Lagrange condition, which is a necessary condition for optimality. In the Online Appendix, we also solve the Hamilton-Jacobi-Bellman equation associated with our problem and use it to demonstrate that our optimality condition is also *sufficient* for a maximum. Second, we have not constrained the optimal bargaining share $\beta^*(m)$ to be nonnegative, consistent with the notion that the firm might find it optimal to compensate certain suppliers with a payoff that exceeds their marginal contribution. In the Online Appendix, we show how imposing $\beta^*(m) \geq 0$ affects the solution for the optimal bargaining share. Crucially, however, the statement in Lemma 1 remains valid except for the fact that, in this constrained case, $\beta^*(m)$ is only *weakly* increasing in m when $\rho > \alpha$.

The key implication of Lemma 1 is that the relative size of the parameters ρ and α will govern whether the incentive for the firm to retain a larger surplus share increases or decreases along the value chain. Intuitively, when ρ is high relative to α , investments are sequential complements, and integrating early stages of production is particularly costly because this reduces the incentives to invest not only of these early suppliers but also of all suppliers downstream. Furthermore, although integration allows the firm to capture some rents, 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 , investments are sequential substitutes, and outsourcing is particularly costly in upstream stages because high investments early in the value chain lead to reduced incentives to invest for downstream suppliers, whereas the firm would capture a disproportionate amount of surplus by integrating these early stages.

Another way to convey this intuition is by comparing the supplier's investment levels in the complete versus incomplete contracting environments. As we have seen earlier, in the former case, the firm would choose quality-contingent contracts to elicit a common value of input services for all suppliers in the value chain. Instead, with incomplete contracting, if the bargaining weight $\beta(m)$ was common for all stages, investment levels would be increasing along the value chain for $\rho > \alpha$ and decreasing along the value chain for $\rho < \alpha$ (see equation (7)). The optimal choice of $\beta(m)$ in (15) can thus be understood as a second-best instrument that attenuates the distortions arising from incomplete contracting, by *rebalancing* investment levels towards those that would be chosen in the absence of contracting frictions. In the complements case, this involves eliciting more supplier investment in the early stages through outsourcing, and (possibly) integrating the most downstream suppliers to dampen the relative overinvestment in these latter stages; an analogous logic applies in the substitutes case.

Evaluating the function $\beta^*(m)$ in (15) at its extremes, we obtain that $\lim_{m \rightarrow 0} \beta^*(m) = -\infty$ when $\rho > \alpha$ and $\beta^*(0) = 1$ when $\rho < \alpha$, while $\beta^*(1) = 1 - \alpha$ regardless of the relative magnitude of α




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

and ρ . 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 depends solely on the relative size of ρ and α . In the complements case, the firm would 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, the firm necessarily chooses to integrate these initial stages. As for the most downstream stages, 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 stages in the immediate neighborhood of $m = 1$ are integrated or not depends on other parameter restrictions (see the Appendix). 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 organizational mode for stages at both ends of the value chain. In the Appendix, we show that the set of stages under a common organizational form (integration or outsourcing) is necessarily a connected interval in $[0, 1]$, thus implying:

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.

Given that $\beta(m)$ takes on at most two values along the value chain, one can in fact derive a closed-form solution for the cutoff stages, m_C^* and m_S^* , in terms of the parameters β_O , β_V , α and

ρ (see Appendix for details):

$$m_C^* = \min \left\{ \left[1 + \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1 - \frac{\beta_O}{\beta_V}}{1 - \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right]^{-1}, 1 \right\} \quad (16)$$

and

$$m_S^* = \min \left\{ \left[1 + \left(\frac{1 - \beta_V}{1 - \beta_O} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{\left(\frac{1 - \beta_V}{1 - \beta_O} \right)^{-\frac{\alpha}{1-\alpha}} - 1}{\frac{\beta_V}{\beta_O} - 1} \right)^{\frac{\alpha(1-\rho)}{\alpha-\rho}} - 1 \right] \right]^{-1}, 1 \right\}, \quad (17)$$

where remember that $\beta_O < \beta_V$.⁹ With these expressions, 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 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 thus be less concerned with the investment inefficiencies caused by such integration.

3 Extensions and Empirical Implementation

Our Benchmark Model is stylized along several directions and omits many factors that have been shown to be important for the organizational decisions of firms in the global economy. In this section, we develop a few extensions that help us gauge the robustness of our results and also allow us to connect our Benchmark Model to the global sourcing framework in Antràs and Helpman (2004, 2008). These extensions will also serve to justify the regression specification and several control variables that we will adopt later in our empirical analysis. For simplicity, we develop these extensions one at a time, although they could readily be incorporated in a unified framework.

3.1 Alternative Contracting Assumptions

A. Ex-Ante Transfers

We begin by exploring the robustness of our results to alternative contracting assumptions. (To conserve space, we focus on outlining the main results and relegate most mathematical details to

⁹Using (16) and (17), it is straightforward to derive the parameter restrictions that characterize when the cutoff lies strictly in the interior of $(0, 1)$. In the complements case, $m_C^* \in (0, 1)$ if $\beta_V(1 - \beta_V)^{\alpha/(1-\alpha)} > \beta_O(1 - \beta_O)^{\alpha/(1-\alpha)}$, while $m_C^* = 1$ otherwise. In the substitutes case, $m_S^* \in (0, 1)$ if $\beta_V(1 - \beta_V)^{\alpha/(1-\alpha)} < \beta_O(1 - \beta_O)^{\alpha/(1-\alpha)}$, while $m_S^* = 1$ otherwise (see the Appendix).

the Online Appendix.) We first consider the implications of allowing for ex-ante transfers between the firm and suppliers, which naturally affect the *ex-ante* division of surplus between agents. In our Benchmark Model, the optimal choice of ownership structure was partly shaped by the desire of the firm to extract rents from its suppliers. If ex-ante transfers were allowed, the choice of ownership structure would now seek to maximize the *joint surplus* created along the value chain:

$$\pi_F = A^{1-\rho}\theta^\rho \left(\int_0^1 x(j)^\alpha dj \right)^{\rho/\alpha} - \int_0^1 cx(j) dj, \quad (18)$$

rather than just the ex-post surplus obtained by the firm, as in equation (11).

Importantly, the presence of ex-ante transfers has no effect on suppliers' investment levels, which are still given by (10) and thus feature the same distortions as in our Benchmark Model. In particular, supposing that bargaining weights were constant (i.e., $\beta(m) = \beta$), investment levels would continue to increase along the value chain in the complements case ($\rho > \alpha$), while they would continue to decrease along the value chain in the substitutes case ($\rho < \alpha$). As a result, when studying the hypothetical case in which the firm could freely choose $\beta(m)$ from the continuum of values in $[0, 1]$ to maximize (18), we find that the marginal return to raising $\beta(m)$ is once again increasing in m for $\rho > \alpha$ and decreasing in m for $\rho < \alpha$. In words, even in the presence of lump-sum transfers, the central result of our paper remains intact: the incentive to integrate suppliers is highest for downstream suppliers in the complements case, while it is highest for upstream suppliers in the substitutes case.

There is however one key difference that emerges relative to the Benchmark Model. With ex-ante transfers, we find that integration and outsourcing coexist along the value chain only when $\rho < \alpha$, in which case the firm integrates the most upstream stages and outsources the most downstream ones.¹⁰ On the other hand, when $\rho > \alpha$, although the incentive to integrate suppliers is highest for downstream suppliers, the firm nevertheless finds it optimal to outsource all stages of production (including the most downstream ones), regardless of the values of β_O and β_V (see the Online Appendix for details). The intuition is simple: given that the firm can extract surplus from suppliers in a nondistortionary manner via ex-ante transfers, the use of integration for rent extraction purposes is now inefficient. When $\rho < \alpha$, the firm will also use ex-ante transfers to extract surplus from suppliers, but integration of upstream suppliers continues to be attractive because it serves a different role in providing incentives to invest for downstream suppliers, as in our Benchmark Model.

B. Linkages Across Bargaining Rounds

In our Benchmark Model, we have assumed that the firm and the supplier in each stage m bargain only over the marginal addition of that supplier to production value, as captured by $r'(m)$ in (5), independently of the bilateral negotiations that take place at other stages. This seems a sensible

¹⁰The cutoff stage separating the upstream integrated stages from the downstream outsourced stages can in fact be shown to be unique and to lie strictly in the interior of $(0, 1)$. (See the Online Appendix.)

assumption to make in environments in which suppliers might not have precise information over what other suppliers in the value chain do, but formally introducing incomplete information into our model would greatly complicate the analysis. Instead, in this section, we will stick to our assumption that all players have common knowledge of the structure and payoffs of the game, but we will briefly characterize the subgame perfect equilibrium of a more complicated game in which suppliers internalize the effect of their investment levels and their negotiations with the firm on the subsequent negotiations between the firm and downstream suppliers.

In order to do so, it becomes important to specify precisely the implications of an (off-the-equilibrium path) decision by a supplier to refuse to deliver its input to the firm. Remember that we have assumed that the firm has the ability to costlessly produce any type of incompatible input, so such a breach of contract would not drive firm revenues to zero. The key issue is: what is the effect of such a deviation on the productivity of downstream suppliers? In order to consider spillovers from some bargaining stages to others, the simplest case to study is one in which once the production process incorporates an incompatible input (say because a supplier refused to trade with the firm), all downstream inputs are necessarily incompatible as well, and thus their marginal product is zero and firm revenue remains at $r(m)$ if the deviation happened at stage m . (In the Online Appendix, we outline the complications that arise from studying less extreme environments.)

With these assumptions, the supplier at stage m realizes that by not delivering its input, the firm will not only lose an amount of revenue equal to $r'(m)$, but will also lose its share of the value from all subsequent additions of compatible inputs by suppliers positioned downstream of m . This problem clearly takes on a recursive nature, since the negotiations between the firm and supplier at any given stage will be shaped by all negotiations that take place further downstream. To formally characterize the subgame perfect equilibrium of this game, we first develop a discrete-player version of the game in which each of $M > 0$ suppliers controls a measure $1/M$ of the production stages, and then study its behavior in the limit as $M \rightarrow \infty$. In the Online Appendix, we show that the profits obtained by the K -th supplier ($K = 1, \dots, M$) are given by:

$$\pi_S(K) = (1 - \beta(K)) \sum_{i=0}^{M-K} \mu(K, i) (r(K+i) - r(K+i-1)) - \frac{1}{M} cx(K), \quad (19)$$

where:

$$\mu(K, i) = \begin{cases} 1 & \text{if } i = 0 \\ \prod_{l=K+1}^{K+i} \beta(l) & \text{if } i \geq 1 \end{cases}, \quad (20)$$

and the discrete-player analogue of the revenue function is $r(K) = A^{1-\rho} \theta^\rho \left[\sum_{k=1}^K \frac{1}{M} x(k)^\alpha \right]^{\frac{\rho}{\alpha}}$. Note that from a Taylor approximation, we have that the marginal contribution of supplier $K+i$ is given

by:

$$r(K+i) - r(K+i-1) \approx A^{1-\rho} \theta^\rho \frac{\rho}{\alpha} \left[\sum_{k=1}^{K+i-1} \frac{1}{M} x(k)^\alpha \right]^{\frac{\rho-\alpha}{\alpha}} \frac{1}{M} x(K+i)^\alpha \text{ for all } i \geq 0. \quad (21)$$

The key difference relative to our Benchmark Model is that the payoff to a given supplier in equation (19) is now not only a fraction $1 - \beta(K)$ of the supplier's own *direct* contribution to production value, $r(K) - r(K-1)$, but also incorporates a share $\mu(K, i)$ of the direct contribution of each supplier located $i \geq 1$ positions downstream from K , namely $r(K+i) - r(K+i-1)$. Note, however, that the share of supplier $K+i$'s direct contribution captured by K quickly falls in the distance between K and $K+i$ (see equation (20)).

At first glance, it may appear that the introduction of linkages across bargaining stages greatly complicates the choice of investment levels along the value chain. This is for at least two reasons. First, the choice of investment $x(K)$ will now be shaped not just by the marginal return of those investments on supplier K 's own direct contribution, but also by the marginal return to those investments made in subsequent stages. Second, this will in turn lead upstream suppliers to internalize the effect of their investments on the investment decisions of suppliers downstream. These two effects are apparent from inspection of equations (19) and (21).

In the Online Appendix, we show, however, that when considering the limiting case of a continuum of suppliers ($M \rightarrow \infty$), these effects become negligible, and remarkably, the investment choices that maximize $\pi_S(K)$ in equation (19) end up being *identical* to those in the Benchmark Model. This is despite the fact that the *actual* ex-post payoffs obtained by suppliers are distinct and necessarily higher than those in the Benchmark Model. The intuition behind this result is that the effect of a supplier's investment on its own direct contribution is of a different order of magnitude from its effects on other suppliers' direct contributions, as illustrated by equation (21), and only the former remains non-negligible as $M \rightarrow \infty$.¹¹ In sum, investment levels are only relevant insofar as they affect a supplier's own direct contribution, and thus this variant of the model ends up delivering the *exact same* levels of supplier investments as in the Benchmark Model.

Since investment levels are identical to those in the Benchmark Model, the total surplus generated along the value chain will also remain unaltered. Provided that the firm and its suppliers have access to ex-ante transfers in the initial contract, this variant of the model will generate the exact same predictions as our Benchmark Model extended to include lump-sum transfers, as outlined in subsection 3.1.A above. In the absence of such ex-ante transfers, however, the choice of ownership structure becomes significantly more complicated due to the fact that the ex-post rents obtained by the firm in a given stage are now lower than in the Benchmark Model, and more so the more upstream the stage in question. Other things equal, this generates an additional incentive for the firm to integrate relatively upstream suppliers, regardless of the relative size of ρ and α . Unfortunately, an explicit formula for π_S and π_F cannot be obtained in the limiting case $M \rightarrow \infty$, thus

¹¹Notice that as $M \rightarrow \infty$, the direct effect $\frac{1}{M} x(K+i)^\alpha$ goes to zero at the same rate $1/M$ as the cost term in (19). Conversely, the indirect effect goes to zero at rate $1/M^2$.

precluding an analytical characterization of the ownership structure choice along the value chain.

C. On the Sequentiality of Bargaining and Production

We have so far considered environments in which both production and bargaining occur sequentially. How would our results differ if instead the firm were to bargain with all suppliers simultaneously after the entire sequence of production stages has been completed? The answer to this depends naturally on the details of the ex-post bargaining. Consider first the approach in Acemoglu et al. (2007), who use the Shapley value to determine the division of ex-post surplus between the firm and its suppliers, while specifying a symmetric threat point associated with a deviation (e.g., withholding of noncontractible services) on the part *any* given supplier. In such a case, all suppliers would derive the same payoff and the optimal organizational form would be independent of the position of an input in the value chain (see Acemoglu et al., 2007), a prediction that would be starkly at odds with our empirical findings below. We would argue, however, that this symmetric multilateral solution is not a natural one to consider in environments where production is inherently sequential. In particular, this symmetric solution inherently assumes that despite recognizing that their investments have asymmetric effects on profits depending on the stage at which they enter production (i.e., $r'(m)$ varies with m), suppliers are nevertheless able to figure out that, in an equilibrium in which the continuum of all other suppliers cooperates, threatening to withhold her individual noncontractible investments would have the same effect for any supplier. The bargaining solution that would arguably be more natural is one where suppliers perceive their marginal contribution to be given by the increase in production value *at the time their investment takes place*, regardless of whether negotiations occur sequentially or simultaneously at the end of production.¹² Our results, we would then argue, hinge in a more fundamental way on the sequential nature of production than on the sequentiality of bargaining.

That said, we have explored several extensions of our framework to allow for richer production structures that feature not just sequential but also modular features. First, we have developed a variant of our model where production resembles a “spider”, following the terminology of Baldwin and Venables (2010). Specifically, the final good combines a continuum of modules or parts, which are put together simultaneously according to a technology featuring a constant elasticity of substitution, $1/(1 - \zeta) > 1$, across the services of the different modules. Prior to final-good assembly, each of these modules is in turn produced by sequentially combining a unit measure of intermediate inputs under the same technology and contracting assumptions as in our Benchmark Model, and where each module producer decides which of its module-specific inputs to integrate. As we demonstrate in the Online Appendix, as long as the negotiations between the final-good producer and the module producers are separate from those between each module producer and its suppliers,

¹²Put differently, remember that the (symmetric) Shapley value of a player is the average of her contributions to all coalitions that consist of players ordered below her in all feasible permutations. When production is simultaneous, one would take an average over *all possible* permutations. But when production is sequential, the order in which suppliers are arranged in any given permutation should matter, with profits accruing only insofar as the coalition players are ordered in the same sequence as production would require.

the incentives to integrate suppliers will be shaped by the same forces as in our Benchmark Model. In particular, the pattern of integration along each module’s chain will now depend crucially on the relative magnitude of the parameter α governing the substitutability of inputs within modules and the parameter ζ which turns out to govern the concavity of each module producer’s revenues with respect to the quality of the module she delivers (as ρ did in our Benchmark Model). If $\zeta > \alpha$, then module inputs will be sequential complements and the propensity to integrate will once again be increasing with the downstreamness of module inputs; the converse statement applies when $\zeta < \alpha$. This modification has some bearing for the empirical interpretation of our results, so we will briefly return to it in Section 5.3.

In the Online Appendix, we also study an alternative hybrid model that resembles more a “snake”: Production is sequential as in our Benchmark Model, but each stage input m is now itself composed of a large number (formally, a unit measure) of components produced *simultaneously*, each by a different supplier, under a symmetric technology featuring a constant elasticity of substitution, $1/(1 - \xi) > 1$, across components. When $\xi \rightarrow 1$, this captures situations in which firms might contract with multiple suppliers to provide essentially the same intermediate input. We model the negotiations between the final-good producer and the set of stage- m suppliers using the Shapley value, as in Acemoglu et al. (2007). As it turns out, in this extension, the incentives to integrate the suppliers of the stage- m components are shaped by downstreamness (i.e., the index m) in qualitatively the same manner as in our Benchmark Model, independently of the value of ξ .

3.2 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 these considerations into our framework. In particular, consider the case in which the production function from (1) is modified to:

$$q = \theta \left(\frac{h}{\eta}\right)^\eta \left(\int_0^1 \left(\frac{x(j)}{1-\eta}\right)^\alpha I(j) dj\right)^{\frac{1-\eta}{\alpha}}, \quad \eta \in (0, 1), \quad (22)$$

where recall that $I(j)$ is an indicator function equal to 1 if and only if input j is produced after all inputs upstream of j have been procured. Suppose also that the provision of headquarter services, h , by the firm is undertaken at marginal cost c_h after suppliers have been hired, but before they have undertaken any stage 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 and we rule out ex-ante transfers to facilitate comparison with our Benchmark Model.

Because the investment in h is sunk by the time suppliers take any actions, the introduction of headquarter services does not alter the above analysis too much. In particular, the value of production generated up to stage m when all inputs are compatible 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 these effects of ρ need to be conditioned on the headquarter intensity of the industry. In particular, we should see a greater propensity towards integrating downstream stages when ρ is high and η is low, with the converse being true when ρ is low and η is high.

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 analogous to those laid out in previous contributions to the property-rights theory. To see 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^* in (16) and (17). 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.¹³

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 .

3.3 Firm Heterogeneity and Prevalence of Integration

Up to now, we have considered the problem of a single firm in isolation. We now show that our model can be readily embedded in an industry equilibrium, in which firms produce a continuum of differentiated final-good varieties that consumers value according to the utility function in (2).

On the technology side, each firm within the industry produces one final-good variety under the same technology and sequencing of production stages in (1). Following Melitz (2003), we let firms differ in their productivity parameter θ . As is commonly done, we assume that θ is drawn

¹³The counterpart of this result is that the unconstrained optimal bargaining share $\beta^*(m)$ spelled out in (15) 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 .

independently for each firm from an underlying Pareto distribution with shape parameter z and minimum threshold $\underline{\theta}$, namely:

$$G(\theta) = 1 - (\underline{\theta}/\theta)^z \quad \text{for } \theta \geq \underline{\theta} > 0, \quad (23)$$

where z is inversely related to the variance of θ and is assumed high enough to ensure a finite variance of the size distribution of firms. We further 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). The values that these fixed costs can take are symmetric for all stages, varying only with the organizational structure chosen by the firm for each given stage. More specifically, and following the arguments in Antràs and Helpman (2004), we assume that:

$$f_V > f_O,$$

reflecting the higher managerial overload associated with running an integrated relationship (f_V) relative to maintaining an arms-length arrangement with an input supplier (f_O).

The introduction of productivity heterogeneity and fixed costs of production enriches the choice 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 ($\rho > \alpha$). As in the Benchmark Model, the incentive for the firm to integrate a given production stage is larger the more downstream the stage, and again there exists a cutoff $m_C \in (0, 1]$ such that all stages before m_C are outsourced and all stages after m_C (if any) are integrated. The presence of fixed costs means however that when $m_C < 1$, this threshold is now implicitly defined by:

$$(m_C)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left[\left(1 - \frac{\beta_O}{\beta_V}\right) - \left(1 - \left(\frac{1-\beta_V}{1-\beta_O}\right)^{\frac{\alpha}{1-\alpha}}\right) \left[1 + \left(\frac{1-\beta_V}{1-\beta_O}\right)^{\frac{\alpha}{1-\alpha}} \left(\frac{1-m_C}{m_C}\right)\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \right] = \frac{f_V - f_O}{\Psi \theta^{\frac{\rho}{1-\rho}}}, \quad (24)$$

where $\Psi \equiv \kappa \frac{\rho(1-\alpha)}{\alpha(1-\rho)} \beta_V (1-\beta_O)^{\frac{\rho}{1-\rho}}$ is a constant, κ being the same constant from equation (13) in the Benchmark Model.

It can be shown that the left-hand side of (24) is increasing in m_C whenever we have an interior solution, and thus the threshold m_C is now a decreasing function of the level of firm productivity θ . Intuitively, relatively more productive firms will find it easier to amortize the extra fixed cost associated with integrating stages, and thus will tend to integrate a larger number of stages. Furthermore, when $\theta \rightarrow \infty$, 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 (16)). Following analogous steps, 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




Figure 3: Firm Heterogeneity and the Integration Decision

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 .¹⁴ As productivity increases, more and more stages become integrated, with these 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.

A key implication of firm heterogeneity is that it 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 $\theta_C(m)$, where $\theta_C(m)$ is the productivity value for which equation (24) holds; the input m will in turn be outsourced by all firms with $\theta < \theta_C(m)$. (Inputs with an index $m < m_C^*$ will not be integrated by any firms.) Appealing to the Pareto distribution in (23), we thus have that the share of firms integrating stage m is given by:

$$\sigma_C(m) = \begin{cases} 0 & \text{if } m \leq m_C^* \\ (\underline{\theta}/\theta_C(m))^z & \text{if } m > m_C^* \end{cases}. \quad (25)$$

From our previous discussion, it is clear that $\theta_C(m)$ is a decreasing function of m , and thus the

¹⁴We assume that f_O is low to ensure that the firms with the lowest productivity level $\underline{\theta}$ will outsource all stages.

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 the dispersion of the productivity distribution, a result that very much resonates with 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 the within-firm variation in the propensity to integrate different stages into predictions regarding the relative prevalence of integration of an input when aggregating over the decisions of all firms within an industry. This is an important step because our empirical application will use industry-level data on intrafirm trade. It is moreover worth stressing that the modeling 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 small. In such a case, our model would 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.

3.4 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, however, different production stages have different effects on output, suppliers differ in their productivity levels, and the widespread process of offshoring also implies that firms undertake different stages of production in various countries where prevailing local factor costs might differ. 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 quality-adjusted final-good production is now given by:

$$q = \theta \left(\int_0^1 (\psi(j) x(j))^\alpha I(j) dj \right)^{1/\alpha}, \quad (26)$$

where $\psi(j)$ captures asymmetries in the marginal product of each input's investments. Furthermore, let the marginal cost of production of input j be given now by $c(j)$, which can vary across inputs due to supplier-specific productivity differences or the heterogeneity in factor costs across the country locations in which inputs are produced.

Following the same steps as in our Benchmark Model, we find that the profits the firm obtains are given more generally by:

$$\pi_F = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (\rho\theta)^{\frac{\rho}{1-\rho}} \int_0^1 \beta(j) \left(\frac{1-\beta(j)}{c(j)/\psi(j)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^j \left(\frac{1-\beta(k)}{c(k)/\psi(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj, \quad (27)$$

This is clearly analogous to equation (11), except for the inclusion of input asymmetries as captured by the term $c(m)/\psi(m)$ for input m . How do these asymmetries affect the firm's choice of ownership structure $\beta(m) \in \{\beta_V, \beta_O\}$ for each stage $m \in [0, 1]$? To build intuition, it is useful once again to consider first the relaxed problem in which the firm could freely choose $\beta(m)$ from the set of piecewise continuously differentiable real-valued functions rather than from $\{\beta_V, \beta_O\}$. After similar derivations to those performed in the Benchmark Model (see the Appendix), we find that the optimal division of surplus must satisfy:

$$\frac{\partial \beta^*(m)}{\partial m} = \frac{(\rho - \alpha)(1 - \alpha)}{(1 - \rho)\alpha} (1 - \beta(m)) \frac{\left(\frac{1-\beta(m)}{c(m)/\psi(m)} \right)^{\frac{\alpha}{1-\alpha}}}{\int_0^m \left(\frac{1-\beta(j)}{c(j)/\psi(j)} \right)^{\frac{\alpha}{1-\alpha}} dj}. \quad (28)$$

It then follows 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 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 remains the same as 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. When embedding the model in the industry equilibrium structure described in Section 3.3, Proposition 5 continues to apply even when firms face heterogenous costs for their inputs. We can thus state:

Proposition 6 Suppose that technology allows for input heterogeneity as in (26), 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.

The result above treats the marginal cost parameters $c(j)$ as exogenous, while in reality they are partly shaped by the endogenous *location* decisions of firms. One might worry that, to the extent that these location decisions are also shaped by downstreamness in a systematic way, the comparative static results regarding the effect of m on the integration decision might become more complex. 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 in a similar manner. Straightforward calculations indicate that although the marginal

incentives for the firm to reduce the marginal cost of a given stage are indeed generally affected by the index of the production stage m , the optimal division of surplus must continue to satisfy the differential equation in (28) and thus the results in Proposition 6 are robust to the endogeneity of cost parameters.

Proposition 6 appears to justify an empirical specification in which the propensity to integrate a particular input is correlated with the degree of downstreamness of that input *regardless* of where that input is produced. From the point of view of U.S. firms importing inputs from abroad, this suggests that one can aggregate or pool observations from several origin countries of these inputs without worrying about variation in country characteristics that might shape the marginal costs faced by producers in those countries. A caveat of this approach, however, is that it ignores the possibility of the existence of heterogeneity in other determinants of integration across locations of input production. In our empirical analysis, we propose two ways to address this caveat. First, we will run specifications that exploit both cross-sectoral and cross-country variation in the prevalence of integration, while introducing country fixed effects to ensure that the effect of downstreamness we identify is not estimated off cross-country variation in unrelated parameters. 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 precisely due to their level of downstreamness. To deal with this concern, we will also experiment with a two-stage Heckman correction specification.

Before we turn to our empirical investigation, we briefly discuss another implication of the result in Proposition 6. In our Benchmark Model, we have assumed that the only enforceable aspect of an ex-ante contract is whether suppliers are vertically integrated or not. Suppose instead that the *quantity* of intermediate inputs produced by a supplier is also contractible and that the services provided by intermediate input m are given by $\psi(m)x(m)$ as in equation (26), with $\psi(m)$ denoting the (contractible) number of units and $x(m)$ denoting the (noncontractible) services per unit (or quality). Our previous discussion of endogenous location decisions suggests that even in situations in which the initial contract specifies heterogeneous quantities $\psi(m)$ for different intermediate input stages (perhaps in an effort to partly correct the inefficient ex-post asymmetries arising from incomplete contracting), our key result in equation (28) characterizing the differential incentives to integrate suppliers along the value chain will continue to hold. Intuitively, as long as quantity and quality are not perfect substitutes in generating value, partial contracting will not be sufficient to restore efficiency, and the second-best role for vertical integration identified in our Benchmark Model will continue to be active.

4 Implementing an Empirical Test

The Benchmark Model that we have developed focuses on firm organizational decisions, and thus firm-level data would appear to be the ideal laboratory for testing it. Nevertheless, firm-level data on integration decisions is not readily available, and while a small number of such datasets have been used to test theories of multinational firm boundaries, these do not provide a sufficiently

rich picture of the heterogeneous sourcing decisions of firms over a large number of inputs.¹⁵ Our approach will instead exploit industry-level variation in the extent to which goods 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 existence of official records of goods crossing borders.

We describe in this section our empirical strategy based on detailed data on U.S. intrafirm imports. Specifically, we will test the prediction in Proposition 6, namely that the relative prevalence of vertical integration of an input when aggregated across the decisions of all final-good producers purchasing that input (see equation (25)), should be a function of the average position of that input’s use in the value chain. Needless to say, implementing such a test requires that we propose appropriate measures for the downstreamness of an input’s use and that we provide a means to distinguish between the sequential complements ($\rho > \alpha$) and substitutes ($\rho < \alpha$) cases. We carefully describe below the construction of these key variables. (Additional details on the industry concordances used and other control variables are documented in the Data Appendix.)

4.1 Intrafirm Import Share

For our 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. We draw this data from the U.S. Census Bureau’s Related Party Trade Database, which reports U.S. trade volumes at the detailed country-industry level, and more importantly, breaks down the value of trade according to whether it was conducted with related versus non-related parties. We focus our analysis on the U.S. import data for manufacturing industries, given the U.S.’ position as a large user of intermediates and consumer of finished goods from the rest of the world. For imports, a related party is defined as a foreign counterpart in which the U.S. importer has at least a 6% equity interest.¹⁶ We work with an extensive amount of data for the years 2000-2010.

For each industry, we use the share of related party imports in total U.S. imports, or (*Related Trade*)/(*Related Trade* + *Non-Related Trade*), to capture the propensity of U.S. firms to integrate foreign suppliers of that particular industrial good. We will refer to this measure as simply the *share* of intrafirm imports, and we can calculate this both at the industry-year and at the exporting country-industry-year levels. The publicly available Census Bureau data is reported at the North American Industry Classification System (NAICS) six-digit level. To facilitate the merging with other industry variables (especially our measures of downstreamness), we converted the related party trade data from NAICS to 2002 Input-Output industry codes (IO2002) using the concordance provided by the Bureau of Economics Analysis (BEA), before calculating the intrafirm import

¹⁵See 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.

¹⁶While this is lower than the conventional 10% cutoff used by the IMF to determine whether a foreign ownership stake qualifies as FDI, extracts from the confidential direct investment dataset collected by the BEA nevertheless suggest that related party trade is generally associated with one of the entities having a controlling stake in the other entity; see Nunn and Treffer (2008).

share. As illustrated by Antràs (2011), there is rich variation in this U.S. intrafirm import share: it varies widely across products and origin countries, and there also exists significant variation across products within exporting countries, as well as across exporting countries within narrowly-defined products. In all, there were 253 IO2002 manufacturing industries for which we had data on intrafirm imports, and that therefore made it into our eventual regression sample.¹⁷

It is further useful to point out that some trade values in the Census Bureau data are recorded under a third category (“unreported”). These are instances where the nature of the transactions – whether it was between related or non-related parties – could not be precisely determined. This constitutes a very small share of total trade flows, and is a reason of concern only for observations with small trade volumes where “unreported” flows might contribute to measurement error in the intrafirm trade shares that we calculate.¹⁸ We will return to this point later below when discussing our empirical findings.

4.2 Downstreamness

Our model emphasizes a novel explanatory variable, namely the relative location of an industry along the value chain. We propose two alternative measures to capture the “downstreamness” of an industry in production processes. As we do not have information on the sequencing of stages for individual technologies, we instead turn to the 2002 Input-Output Tables to obtain *average* measures of the relative position of each industry in U.S. production processes.

To build intuition on these measures, recall the basic input-output identity:

$$Y_i = F_i + Z_i,$$

where Y_i is total output in industry i , F_i is the output of i that goes toward final consumption and investment (“final use”), and Z_i is the use of i ’s output as inputs to other industries (or its “total use” as an input). In a world with N industries, this identity can be expanded as follows:

$$Y_i = F_i + \underbrace{\sum_{j=1}^N d_{ij} F_j}_{\text{direct use of } i \text{ as input}} + \underbrace{\sum_{j=1}^N \sum_{k=1}^N d_{ik} d_{kj} F_j + \sum_{j=1}^N \sum_{k=1}^N \sum_{l=1}^N d_{il} d_{lk} d_{kj} F_j + \dots}_{\text{indirect use of } i \text{ as input}}, \quad (29)$$

where d_{ij} for a pair of industries (i, j) , $1 \leq i, j \leq N$, is the amount of i used as an input in producing one dollar worth of industry j ’s output. Note that the second term on the right-hand side of (29) captures the value of i ’s “direct use” as an input, namely the total value of i purchased by industry j to produce output that immediately goes to final use. The remaining terms that

¹⁷This is out of a maximum possible of 279. Several industries dropped out due to the absence of trade data in the original Census Bureau dataset. A handful of industries were also merged in the process of the mapping from NAICS; see the Data Appendix for more details.

¹⁸The share of U.S. manufacturing imports recorded as “unreported” is below 0.2% in all the years in our sample. Furthermore, using the industry-year observations, we find a raw negative correlation of -0.20 between the log of the share of “unreported” trade and the log of total imports across the IO2002 manufacturing industries.

involve higher-order summations reflect the “indirect use” of i as an input, as these enter further upstream in the value chain, at least two production stages away from final use. The above can be written in compact matrix form by stacking the identity for all industries i :

$$Y = F + DF + D^2F + D^3F + \dots = [I - D]^{-1} F, \quad (30)$$

where Y and F are the $N \times 1$ vectors whose i -th entries are respectively Y_i and F_i , while D is the $N \times N$ direct requirements matrix whose (i, j) -th entry is d_{ij} . Note that $[I - D]^{-1}$ is often called the Leontief inverse matrix.¹⁹

Our first measure of downstreamness, $DUse_TUse$, is the ratio of aggregate direct use to aggregate total use of i as an input.²⁰ Specifically, this is calculated by dividing the i -th element of the column vector DF (i.e., the value of i 's direct use as an input for final-use production, summed over all buyer industries j) by the i -th element of $Y - F$ (which equals the total use value of i as an input, summed over all buyer industries j). The higher is $DUse_TUse$ for a given industry i , the more intensive is its use as a direct input for final-use production, so that the bulk of i 's value enters into production relatively far downstream. Conversely, a low value of $DUse_TUse$ would indicate that most of the contribution of input i to production processes occurs indirectly, namely in more upstream stages.

In terms of implementation, we draw on the detailed Use Table issued by the BEA in the 2002 U.S. Input-Output Tables to construct the direct requirements matrix, D . We also constructed the final-use vector, F , by summing over the value of each industry i 's output purchased for consumption and investment by private or government entities (IO2002 codes starting with “F”), but excluding net changes in inventories (F03000), exports (F04000), and imports (F05000). Lastly, the output vector, Y , was obtained by taking the sum of all entries in row i in the Use Table (this being equal to gross output Y_i). We applied an open-economy and inventories adjustment to the entries of D and F , to account for the fact that inter-industry flows across borders (as well as in and out of inventories) are not directly observed.²¹ (For a detailed discussion of this adjustment, see Antràs et al. (2012).)

We supplement our analysis with a second measure of downstreamness, $DownMeasure$, which seeks to make fuller use of the information on indirect input use further upstream. To motivate this, consider the example of IO 331411 (Primary smelting and refining of copper). This is the manufacturing industry with the third lowest value of $DUse_TUse$ (about 0.07), indicating that the

¹⁹This inverse exists so long as $\sum_{i=1}^N d_{ij} < 1$ for all j , a natural assumption given the economic interpretation of the d_{ij} 's as input requirement coefficients.

²⁰See Alfaro and Charlton (2009) and di Giovanni and Levchenko (2010) for measures of production line position that have a similar flavor.

²¹This entailed: (i) multiplying the (i, j) -th entry of D by $Y_j/(Y_i - X_i + M_i - N_i)$, and (ii) multiplying the i -th entry of F by $1/(Y_i - X_i + M_i - N_i)$, where X_i , M_i and N_i denote respectively the value of exports, imports and net changes in inventories reported for industry i in the Use Table. These are the adjustment terms implied by a natural set of proportionality assumptions, namely that the shares of i 's output purchased by other industries j or for final use in domestic transactions are respectively equal to the corresponding shares of i 's various uses both in net exports and net changes in inventories.

vast majority of its use is indirect to final-use production. That said, it is easy to trace production chains of varying lengths which begin with 331411. An example of a short chain with just three stages is: 331411 \rightarrow 336500 (Railroad rolling stock) \rightarrow F02000 (Private fixed investment), while a much longer example with seven stages is: 331411 \rightarrow 331420 (Copper rolling, drawing, extruding and alloying) \rightarrow 332720 (Turned product and screw, nut, and bolt) \rightarrow 33291A (Valve and fittings other than plumbing) \rightarrow 336300 (Motor vehicle parts) \rightarrow 336112 (Automobile) \rightarrow F01000 (Personal consumption).²² To shed light on whether 331411’s use as an input is characterized by short as opposed to long chains, we require a measure that distinguishes the indirect use value according to the number of stages from final-use production at which that input use enters the value chain.

More specifically, referring back to the identity (29), let output for final use (the first term on the right-hand side) be weighted by 1, let the input value used directly in final-use production (the second term on the right-hand side) be weighted by 2, let the third term on the right-hand side be weighted by 3, and so on. In matrix form, this boils down to calculating:

$$F + 2DF + 3D^2F + 4D^3F + \dots = [I - D]^{-2} F. \quad (31)$$

Although evaluating the left-hand side of (31) would appear to require computing an infinite power series, it turns out that this sum is a simple function of the square of the Leontief inverse matrix. For each industry i , we then take the i -th entry of $[I - D]^{-2} F$ and normalize it by Y_i . Since larger weights are applied the further upstream the input enters the production chain, this provides us with a measure of upstreamness, which by construction is greater than or equal to 1. (The value exactly equals 1 if and only if all the output of that industry goes to final use, and it is never used as an input by other industries.) We therefore take the reciprocal to obtain *DownMeasure* for each industry i , where *DownMeasure* now lies in the interval $[0, 1]$.

This second variable has several desirable properties that provide reassurance for its use as a measure of production line position. In Antràs et al. (2012), we established that the upstreamness version of the variable is in fact equivalent to a recursively-defined measure of an industry’s distance to final demand proposed independently by Fally (2012), where Fally’s construction hinges on the idea that industries that purchase a lot of inputs from other upstream industries should themselves be relatively upstream. The upstreamness variable can moreover be interpreted as a measure of cost-push effects or forward linkages – how much the output of all industries in the economy would expand following a one dollar increase in value-added in the industry in question – highlighted in the so-called supply-side branch of the input-output literature (for e.g., see Ghosh (1958), and Miller and Blair (2009)).

We report in Table 1 the ten highest and lowest values of *DUse_TUse* and *DownMeasure* across the IO2002 manufacturing industries. Not surprisingly, the industries that feature low downstreamness values tend to be in the processing of fuel, chemicals or metals, while industries with high values appear to be goods that are near the retail end of the value chain. There is a reassuring

²²In identifying these production chains from the U.S. Input-Output Tables, we selected buying industries at each stage which were among the top ten users by value of the input at that stage.

degree of agreement, with the two measures sharing seven out of the ten bottom industries and six out of the top ten industries. While the correlation between $DUse_TUse$ and $DownMeasure$ is clearly positive (with a Pearson coefficient of 0.60), there are nevertheless useful distinctions between the two measures. For example, Fertilizer (325310) ranks as the 35th most upstream manufacturing industry according to $DUse_TUse$ (with a value of 0.3086), but is actually among the ten most upstream industries according to $DownMeasure$, indicating that a lot of its input value tends to enter early in long production chains. On the other hand, Plastics and rubber industry machinery (333220) is among the ten most upstream industries based on $DUse_TUse$, but only ranks as the 183rd most upstream according to $DownMeasure$ (with a value of 0.6785), consistent with the bulk of its use occurring relatively close to final-good production.²³

4.3 Empirical Specification

We now describe our empirical specifications for uncovering the effect of production line position on the share of intrafirm trade. As a baseline, we work with cross-industry regressions of the form:

$$S_{it} = \beta_1 D_i \times \mathbf{1}(\rho_i < \rho_{med}) + \beta_2 D_i \times \mathbf{1}(\rho_i > \rho_{med}) + \beta_3 \mathbf{1}(\rho_i > \rho_{med}) + \beta_X X_{it} + \alpha_t + \varepsilon_{it}, \quad (32)$$

The dependent variable, S_{it} , is the U.S. intrafirm import share in industry i in a given year t . We seek to explain this as a function of the downstreamness D_i of the industry in question, as captured either by $DUse_TUse$ or $DownMeasure$. Importantly, taking guidance from our model, we seek to distinguish between the effects of downstreamness in the sequential complements and substitutes cases. We do this by interacting D_i with indicator variables, $\mathbf{1}(\rho_i < \rho_{med})$ and $\mathbf{1}(\rho_i > \rho_{med})$, that equal one when the average demand elasticity faced by industries that purchase i as an input is below (respectively above) the cross-industry median value of this variable.

Our theory in fact predicts that the sequential complements and substitutes cases would be delineated by the conditions $\rho > \alpha$ and $\rho < \alpha$ respectively. While it would thus be ideal to empirically capture the degree of technological substitutability across inputs within each industry too, we are unfortunately constrained by the fact that estimates of cross-input substitutability are not readily available in the literature, nor is it clear that these can be obtained from current data sources.²⁴ To make some progress, we therefore 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 an industry's output, so that we can associate the sequential complements case with high values of ρ and the substitutes case with low values of ρ .

We construct this average buyer demand elasticity as follows. We used the U.S. import demand

²³We have further experimented with two other measures of downstreamness: (i) the final plus direct use value divided by total output for that industry; and (ii) the final use value divided by total output for an industry. Our results are reassuringly similar with both of these measures (reported in Online Appendix Tables 3 and 4).

²⁴For example, one might envision estimating α by exploiting time-series variation in the direct requirements coefficients, but comprehensive Input-Output Tables for the U.S. are constructed only every five years. Consequently, it would be challenging to separately identify input substitutability from biased changes in production techniques that might occur over extended periods of time.

elasticities estimated by Broda and Weinstein (2006) from disaggregate ten-digit Harmonized System (HS) product-level trade data. For each IO2002 industry, we then computed a demand elasticity equal to the trade-weighted average elasticity of its constituent HS10 products, using data on U.S. imports as weights. (Details on how this crosswalk between industry codes was implemented are documented in the Data Appendix.) Next, we took a weighted average elasticity across industries that purchase i as an input, with weights proportional to the value of input i used from the 2002 U.S. Input-Output Tables. We included the final-use value of i in this last calculation by assigning it the import demand elasticity of industry i itself. The average buyer demand elasticity that results from these calculations is our empirical proxy for $1/(1 - \rho)$. In our baseline analysis, we split the sample into industries with ρ above the industry median (sequential complements case) and below the median (sequential substitutes case), with our model’s predictions leading us to expect that $\beta_1 < 0$ and $\beta_2 > 0$ in the estimating equation (32). We will later also report estimates using a finer cut of this proxy for ρ by quintiles.

Equation (32) further includes an indicator variable to control for the level effect of the sequential complements case, a control vector of additional industry characteristics, X_{it} (including a constant term), and year fixed effects, α_t . We cluster the standard errors by industry, since the key explanatory variables related to downstreamness and the average buyer demand elasticity vary only at the industry level, and these are being used to explain multiple observations of the intrafirm trade share across years.

The vector X_{it} comprises a set of variables that have been identified previously 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. First, we verify whether measures of headquarter intensity are positively associated with the share of intrafirm trade. This would be consistent with part (ii) of our statement of Proposition 4, but notice further that part (i) of the Proposition also highlights that headquarter intensity can be expected to affect the condition that distinguishes the sequential complements and substitutes case, with ρ being replaced by $\rho(1 - \eta)$. With that in mind, we will also experiment with specifications that include triple interactions between downstreamness, the average buyer demand elasticity proxy, and measures of headquarter intensity, as explained in more detail in the next section. As controls for headquarter intensity, we include industry measures of physical capital per worker (as first suggested by Antràs, 2003) and skill intensity (nonproduction employees over total employment) derived from the NBER-CES Database, as well as a measure of R&D intensity (R&D expenditures divided by sales) computed by Nunn and Trefler (2012) from the Orbis database. In most specifications, we further break down physical capital intensity into equipment capital intensity and plant capital intensity. As pointed out by Nunn and Trefler (2012), capital equipment is much more likely to be relationship-specific than plant structures, and thus we would expect the former to provide a cleaner proxy for headquarter intensity. (We also follow Nunn and Trefler (2012) in including a materials intensity variable, namely materials purchases per worker.) Last but not least, we control for a measure of the within-industry size dispersion from Nunn and Trefler (2008). In light of

Proposition 5, we expect this dispersion variable to have a positive effect on the intrafirm import share. (See the Data Appendix for more details on the construction of these control variables.)

We construct the above factor intensity and dispersion variables in a slightly different way from past papers. The standard practice to date has been to assign to industry i the value of the factor intensity or size dispersion of i itself, namely the industry selling the good in question. A more satisfactory approach that maps more directly into our present model would be to control for the average value over industries that purchase good i . We thus construct “average buyer” industry versions of these variables by taking a weighted average of the characteristic values of industries that purchase good i as an input, using weights derived from the 2002 U.S. Input-Output Tables, in a manner analogous to our construction of the average buyer demand elasticity parameter, ρ . It turns out that using average buyer rather than seller industry variables makes little qualitative difference to our results, but we adopt this approach because it is closer in spirit to the model. (Summary statistics for all the variables can be found in Table 1 in the Online Appendix, while their correlations with our two downstreamness measures are reported there in Table 2.)

As argued in Section 3.4, our model suggests that cross-country variation in the prevalence of integration can be useful for addressing biases that might arise from the endogenous location decisions of firms regarding different stages of production. We will thus also explore specifications that exploit the full country-industry variation in our intrafirm import share data, as follows:

$$S_{ict} = \beta_1 D_i \times \mathbf{1}(\rho_i < \rho_{med}) + \beta_2 D_i \times \mathbf{1}(\rho_i > \rho_{med}) + \beta_3 \mathbf{1}(\rho_i > \rho_{med}) + \beta_X X_{it} + \alpha_{ct} + \varepsilon_{ict}, \quad (33)$$

In words, this seeks to explain the intrafirm import share, S_{ict} , at the exporting country-industry-year level as a function of a similar set of industry variables, while controlling for country-year fixed effects, α_{ct} , and (conservatively) clustering the standard errors by industry. Later on, we will build on equation (33) to discuss tests that make use of this cross-country variation to address concerns related to the bias that might arise as a result of selection into exporting.

Before we turn to our results, it is worth acknowledging and discussing several caveats that apply to our empirical strategy. To understand how the make-or-buy decision over an input is related to that input’s location in a particular industry’s production line, one would ideally like to observe the breakdown of intermediate input imports by the identity of the purchasing industry. Unfortunately, such a level of detail is not available in the U.S. related party trade data. For example, while we observe the share of U.S. intrafirm imports of rubber tires, we do not observe a breakdown vis-à-vis the intrafirm trade shares of rubber tires purchased by automobile versus aircraft makers. We have instead pursued what is arguably the next best possible strategy, which is to correlate the intrafirm trade share of industry i with measures of how far downstream i tends to be used on average in production processes. The lack of detailed information at the level of the purchasing industry also constrains our ability to empirically distinguish between the sequential complements and substitutes cases, so that we have to rely instead on identifying sectors that sell *on average* to industries that feature high versus low demand elasticities. In sum, while we are unable to perform a structural test of our model, we instead work with the data at the level of

aggregation that is available to us to test the implications of our propositions.

The U.S. Census Bureau trade data itself also comes with its limitations, as discussed at length in Antràs (2011). For example, the data do not report which party is owned by whom, namely whether integration is backward or forward, in related party transactions. U.S. intrafirm imports also generally underrepresent the true extent of U.S. multinational firms’ involvement in global sourcing strategies, as we do not get to observe the cross-border shipment of parts and components that takes place before goods are shipped back to the U.S. That said, it is less clear how (if at all) this might systematically bias the empirical results we are about to discuss. On the plus side, it should be emphasized that the U.S. Census Bureau subjects its related party trade data to several quality assurance procedures. The data do offer a complete picture of the sourcing strategies related to transactions that cross U.S. customs, thus making it easier to spot *fundamental* factors that appear to shape the internalization decisions over these transactions at the cross-industry level.

5 Empirical Results

5.1 Core Findings

We begin our empirical analysis in Table 2, by running shorter versions of the industry-year regressions in (32) to replicate and refine some of the key findings from prior studies of the determinants of intrafirm trade (for e.g., Antràs, 2003, Nunn and Treffer, 2008, 2012, and Bernard et al., 2010). Toward this end, columns 1-3 serve to verify whether the characteristics of the “seller” industry systematically explain the propensity to import the good within firm boundaries. In column 1, we indeed find that two commonly-used measures of headquarter intensity – skill intensity ($\log(s/l)$) and physical capital intensity ($\log(k/l)$) – are both positively and significantly correlated with the intrafirm import share. A third such proxy for the importance of headquarter services – R&D intensity ($\log(0.001 + R\&D/Sales)$) – displays a similarly strong positive association when added in column 2.²⁵ Another measure based on factor use – the materials intensity – that one would typically *not* associate with firm headquarters turns out indeed not to have predictive power for the share of intrafirm trade. The last control added in column 2 is the measure of industry size dispersion, which does have a positive and significant effect on the propensity to trade within firm boundaries, as found previously in Nunn and Treffer (2008). Column 3 further highlights the usefulness of distinguishing between equipment capital and plant structures (c.f., Nunn and Treffer 2012). The former is more likely to involve noncontractible, relationship-specific investments by firm headquarters, and thus it is not surprising that the plant capital intensity variable is weakly correlated (actually, with a negative sign) with the share of intrafirm trade.

We repeat these regressions in columns 4-6, but now using the average buyer industry values of the respective industry characteristics, in place of the typically-used “seller” industry values. As we have argued earlier, it would be more consistent with our model of input sourcing to consider

²⁵We add 0.001 to the R&D expenditures over total sales ratio, in order to avoid dropping the industries with zero reported R&D expenditures in the Orbis dataset.

industry characteristics that pertain to the input-purchasing industries. This has some effects on the estimates, although it is reassuring that the role of headquarter intensity still broadly stands. Physical equipment capital and R&D intensity, in particular, continue to have positive and statistically significant effects on the intrafirm trade share. Note however that the effects of skill intensity and the size dispersion are now less significant than in columns 2 and 3.

The novel predictions from our model regarding the role of downstreamness are tested in Table 3 using the $DUse_TUse$ measure of production line position. In column 1, this is introduced directly as an additional explanatory variable in the industry-year regressions.²⁶ When included on its own, the effect of $DUse_TUse$ on the share of intrafirm trade turns out to be statistically insignificant. Following the guidance of our theoretical model, we run our benchmark specification from equation (32) 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.²⁷ The empirical results here are indeed strongly supportive of our model’s central prediction: The effect of downstreamness is positive and significant at the 5% level when the average buyer demand elasticity is above the median for this variable ($\beta_2 > 0$), consistent with a greater propensity toward the integration of input suppliers that enter further downstream in the value chain. Conversely, the negative and significant β_1 coefficient confirms a greater propensity toward integrating upstream production stages in the sequential substitutes case. As indicated in the table, we can comfortably reject the null hypothesis that $\beta_1 = \beta_2 = 0$ (the F-test for joint significance yields a small p-value of 0.0008), while $\beta_2 - \beta_1$ is also significantly different from zero at the 1% level of significance. On a separate note, the effect of the ($Elas > Median$) dummy variable turns out to be negative and statistically significant, which resonates with our comparative static result in Proposition 3 on the effect of ρ .

It is particularly reassuring that these new findings stand even while including the same set of controls for average buyer headquarter intensity and dispersion that were used earlier in Table 2. When we break the physical capital intensity variable down in column 3 into its equipment and plant capital components, we in fact find that equipment intensity and R&D intensity (two natural proxies for headquarter intensity) are positively and significantly correlated with the intrafirm trade share, consistent with part (ii) of Proposition 4.²⁸

The next two columns of Table 3 verify that the effects we have found in the pooled sample are also present when we run our regressions separately on the subsets of industries where the average buyer demand elasticity is below (respectively above) its median value. The effect of $DUse_TUse$ is negative in the substitutes case (column 4) and positive in the complements case (column 5), with both coefficients of interest being significant at the 1% level. In column 6, we return to the

²⁶All our core results reported in Tables 3 and 4 remain very similar if the regressions are instead run year-by-year, namely with the intrafirm trade share of industries in a particular year as the dependent variable, with Huber-White robust standard errors; see Online Appendix Tables 5 and 6.

²⁷Note that there is no need to include the dummy variable for the sequential substitutes case, as the regressions include a constant term.

²⁸The Online Appendix provides several scatterplots that illustrate the partial correlation between the intrafirm trade share and downstreamness, separately for the substitutes and complements cases.

specification in (32) for the full sample of industries, though we now weight each data point by the value of total imports for that industry-year. This is motivated by our earlier discussion in Section 4.1 on the possible measurement error introduced into the intrafirm trade share by the presence of trade flows whose related party status was not reported to the U.S. Census Bureau, where we also noted that this was likely to pose a greater concern for observations with small trade volumes (see footnote 18 for some evidence). The weighting in column 6 thus attaches more weight to data points that are likely associated with less measurement error. The results clearly reinforce our main findings: The β_2 coefficient increases in magnitude, while remaining highly statistically significant. Based on this point estimate, a one standard deviation increase in $DUse_TUse$ would correspond to an increase in the intrafirm trade share in high ρ industries of $0.482 \times 0.228 = 0.110$, which is over one-half of a standard deviation in the dependent variable, a fairly sizeable effect. While the β_1 coefficient for the substitutes case is now smaller, it remains negative and significant at the 10% level. A similar one standard deviation increase in $DUse_TUse$ in this low ρ case would be associated with a decrease in the intrafirm trade share of about one-fifth of a standard deviation. Note too that the fit of the regression in terms of its R^2 also improves markedly under weighted least squares.²⁹

We exploit the additional variation in the intrafirm import share across source countries in the final two columns of Table 3. Following the specification in (33), column 7 includes country-year fixed effects and clusters the standard errors by industry. Column 8 in turn weights each observation by the total import value for that country-industry-year, for reasons analogous to those discussed above. The unweighted results in column 7 are relatively noisy, while also delivering a low R^2 . The negative effect of downstreamness on integration remains evident in the substitutes case, but that in the complements case is now imprecisely estimated (with a sign opposite to that expected from our model).³⁰ When working with the full country variation though, there is arguably a stronger case for the weighted specification, as there are many more small import flow observations and hence a greater scope for measurement error in the intrafirm trade share to matter. These results in column 8 in fact look very similar to the purely cross-industry findings, with the β_2 coefficient positive and significant at the 1% level, while β_1 remains negative (although not statistically significant). The remaining coefficients are also consistent with the theory, with positive effects for our preferred headquarter intensity measures (equipment capital and R&D intensities), and a negative effect of the dummy variable for the sequential complements case.

In Table 4, we repeat the exercise from Table 3 using our second downstreamness variable, *DownMeasure*, in place of $DUse_TUse$. This reassuringly corroborates our earlier findings on how production line position influences integration outcomes, particularly in the complements case.

²⁹One might be concerned that the largest manufacturing industry by import value, IO 336311 (Automobile), is also the most downstream according to both $DUse_TUse$ and *DownMeasure*, and that this might be driving the significant findings in the weighted regressions. Our results remain qualitatively similar and significant when dropping this industry, in both the cross-industry and the cross-industry, cross-country specifications (available on request).

³⁰Note nevertheless that these point estimates are still consistent with the weaker prediction that downstreamness should have a more negative effect on integration in the substitutes than in the complements case. Formally, we can reject the null hypothesis that the two coefficients are equal ($\beta_1 = \beta_2$) at the 10% level of significance.

We consistently find a positive and significant effect of *DownMeasure* on the propensity to import goods within firm boundaries in the high ρ case, both in the specifications that pool all industries (columns 2, 3, 6) and when we focus on the subset of industries with above-median average buyer demand elasticities (column 5). This effect is especially strong when we weight observations by import volumes: based on column 6, a one standard deviation increase in *DownMeasure* would imply a rise in the intrafirm import share of $0.526 \times 0.222 = 0.117$, which is once again just over half a standard deviation for our dependent variable. We likewise obtain a similar set of results in column 8, which uses the country-industry-year data with weighted least squares. Admittedly, the evidence in Table 4 is more mixed with regard to the role of *DownMeasure* for industries that fall under the substitutes case. The overall results do point toward a negative effect, although the point estimates are much smaller and not statistically distinguishable from zero. Still, the difference between β_1 and β_2 remains statistically significant at conventional levels.

5.2 Alternative Specifications

We turn next to address potential criticisms regarding the use of a median cutoff value for the ρ proxy to distinguish between the sequential substitutes and complements cases. Given that technological substitutability (α in our model) might well vary across sectors, it would be reasonable to expect that the positive effects of downstreamness would be concentrated in the highest ranges of the elasticity demand parameter ρ , while the negative effects of downstreamness might only be evident for particularly low values of ρ . This leads us to consider a more flexible variant of (32) that breaks down our empirical proxy for ρ by quintiles:

$$S_{it} = \sum_{s=1}^5 \beta_s D_i \times \mathbf{1}(\rho_i \in \Omega_{\rho,s}) + \sum_{s=2}^5 \gamma_s \mathbf{1}(\rho_i \in \Omega_{\rho,s}) + \beta_X X_i + \alpha_t + \varepsilon_{it} \quad (34)$$

Here, $\Omega_{\rho,s}$ refers to the subset corresponding to the s -th quintile of ρ (for $s = 1, 2, \dots, 5$), 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.³¹ We report three regressions in Table 5 for each of the downstreamness variables, *DUse_TUse* and *DownMeasure*. For each measure, we estimate (34) both without and with regression weights in the first and second columns respectively; the third column then runs the analogous weighted specification with the country-industry-year data, while controlling for country-year fixed effects. (For all regressions reported in the remainder of this paper, standard errors will be clustered by industry; the full set of average buyer industry variables used in column 3 in Tables 3 and 4 will also be included as controls.)

Table 5 confirms that the effect of downstreamness on the intrafirm import share differs qualitatively depending on the average buyer demand elasticity. The point estimates obtained on the downstreamness variables in the lower (first and second) quintiles of the ρ proxy are negative barring a few exceptions, even being significant at the 10% level in two of the *DUse_TUse* specifications.

³¹We drop the level effect of the first quintile dummy as the regression already includes a constant term.

These coefficients progressively increase as we move from β_1 in the lowest quintile to β_5 in the highest, eventually becoming positive and statistically significant in all columns in the fifth quintile of ρ , precisely for those industries most likely to fall under the sequential complements case. Collectively, we find that the coefficients of downstreamness over these five quintiles are jointly significant, while the difference $\beta_5 - \beta_1$ is also significantly different from zero at the 1% level except in one specification (column 4). These results therefore strengthen our confidence in the empirical relevance of our key theoretical predictions.

We further seek to allay concerns regarding omitted variables bias, specifically those arising from other industry characteristics that could affect the intrafirm trade share which we have not explicitly controlled for so far. Toward this end, we revert to our cross-industry specification in equation (32) and introduce several plausible controls; these results are reported in Table 6 for *DUse_TUse* and in Table 7 for *DownMeasure*.³² We control for the value-added content of each industry – the ratio of value-added to total shipments – in column 1. Column 2 examines whether the importance of a good as an input in production processes might have an effect on the propensity to trade that good within firm boundaries. We capture this through a measure equal to the value of an industry’s total use as an input, divided by the total input purchases made by all its buyer industries, constructed from the 2002 U.S. Input-Output Tables.³³ We incorporate the intermediation variable from Bernard et al. (2010) in column 3, this being a measure of the extent to which wholesalers that serve as intermediaries to transactions are observed to be active in a given industry. Finally, column 4 controls for proxies for industry contractibility, as suggested by Nunn and Treffer (2008). These are based on the underlying share of products from an industry that are transacted on organized exchanges or are reference-priced according to Rauch (1999), and which thus can be regarded as homogeneous and readily contractible goods.³⁴ We control here for both the own contractibility of the good in question as well as an average contractibility taken over industries that purchase the good, in order to distinguish between contracting frictions inherent to the seller and average buyer industries respectively. (Further details on the construction of these additional variables can be found in the Data Appendix.)

Our central finding on the contrasting effects of downstreamness in the substitutes versus complements case turn out to be remarkably robust. This is true when we introduce the aforementioned additional industry variables either individually (columns 1-4) or jointly (column 5), as well as when we run the regressions using total imports as weights (columns 6-10). Throughout Table 6, *DUse_TUse* consistently has a negative correlation with the intrafirm trade share in the low

³²We have also run these same robustness checks on the full country-industry-year data, using the specification in (33). These results are reported in Online Appendix Tables 7 and 8 respectively for *DUse_TUse* and *DownMeasure*, for the weighted least squares regressions. We consistently obtain results akin to our baselines: a positive and significant effect of downstreamness in the complements case, and a negative (albeit insignificant) effect in the substitutes case.

³³The robustness results are very similar if we alternatively divide by the total gross output of buyer industries when constructing this input “importance” variable (available on request).

³⁴We report results using the liberal classification in Rauch (1999). The findings are very similar when using his conservative classification, or when excluding reference-priced products from the definition of what constitutes a contractible good (results available on request).

ρ case, with this coefficient being significant at least at the 10% level except in columns 3 and 8 where we control for the intermediation variable on its own. On the other hand, the positive correlation for the high ρ case is always significant at the 5% level. We also obtain results similar to our baseline regressions when using *DownMeasure* instead in Table 7. Once again, we find that downstreamness is strongly associated with a higher intrafirm trade share in the complements case, although the coefficients for the substitutes case remain for the most part imprecisely estimated.

Of independent interest, the intermediation variable always shows up with a negative and highly significant coefficient when included. This is consistent with Bernard et al. (2010), who interpret the presence of wholesale intermediation as indicative of a reduced need to expand the boundaries of the firm to secure the input in question. Separately, in the columns that control for the role of contracting frictions, we generally find that inputs that are more contractible appear to be transacted more within firm boundaries, while a greater degree of buyer industry contractibility is associated with a reduced propensity toward integration, with these correlations being particularly strong in the weighted specifications.³⁵ In the case of *DownMeasure*, the negative coefficient on downstreamness in the substitutes case even becomes statistically significant in columns 9-10.³⁶

5.3 Extensions

We round off our empirical analysis by testing several auxiliary implications from the extensions that we developed in Sections 3.2-3.4, in order to explore how far the patterns in the data are consistent with these additional predictions of our model of sequential production and sourcing.

Table 8 examines further the role of headquarter intensity. In Section 3.2, we showed that an increase in η not only expands the range of stages that are vertically integrated, 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, recall from Proposition 4 that the complements and substitutes cases are now defined by $\tilde{\rho} \equiv (1 - \eta)\rho > \alpha$ and $\tilde{\rho} < \alpha$ respectively, and thus the larger is η , the less likely it is that downstreamness will have a positive effect on the intrafirm trade share, even for large values of the buyer demand elasticity ρ .

The findings from Table 8 uncover some evidence of this. The key difference relative to the specifications in (32) and (33) is that we now include triple interactions, between $D_i \times \mathbf{1}(\rho_i < \rho_{med})$ and $D_i \times \mathbf{1}(\rho_i > \rho_{med})$ on the one hand, and a set of five dummy variables corresponding to the quintiles of a summary measure of headquarter intensity on the other. The latter is computed as the first principal component of the three main industry measures of headquarter intensity that we have been using, namely the average buyer skill, equipment capital, and R&D intensities. We report three specifications in Table 8 for each downstreamness measure, namely an unweighted cross-industry

³⁵Incidentally, this dovetails with the predictions of the theoretical model in Antràs and Helpman (2008) which introduced a formulation of partial input contractibility; see in particular their Proposition 5.

³⁶In regressions that use the additional source country dimension in the intrafirm import share, Nunn and Trefler (2008) have further interacted the industry contractibility variables with a country index of the rule of law to get a richer proxy for contracting frictions. We have done the same using the most updated version of the rule of law index from Kaufmann, Kraay and Mastruzzi (2010), and verified that this does not affect our main findings at all (see column 5 in Online Appendix Tables 7 and 8).

regression, a weighted cross-industry regression, as well as a weighted cross-country, cross-industry regression. (Throughout, we also control for the main and double interaction effects of the terms in our triple interactions.) As is clear from the table, the positive effect of $D_i \times \mathbf{1}(\rho_i > \rho_{med})$ on the intrafirm trade share is concentrated in the lowest quintile of headquarter intensity, with few systematic patterns apparent for the other quintiles of our composite proxy for η . There is some evidence too that the negative effect of $D_i \times \mathbf{1}(\rho_i < \rho_{med})$ is strongest in a relatively high (fourth) quintile of η ; the coefficients in the fifth quintile are generally negative, though admittedly not significant.³⁷

We turn in Table 9 to an implication of the model that arises from incorporating firm heterogeneity in Section 3.3. As anticipated there, to the extent that the fixed costs of integration are relatively high, only the most downstream of stages would be integrated in the sequential complements case, with the converse prediction (integration of the most upstream stages) applying in the substitutes case. We thus attempt to capture these predictions, by replacing our two main interaction terms in (32) and (33) with quintiles of D_i interacted with each of $\mathbf{1}(\rho_i > \rho_{med})$ and $\mathbf{1}(\rho_i < \rho_{med})$.³⁸ These results are presented in Table 9, where the columns correspond to the same estimation procedures used previously in Table 8 for *DUse_TUse* (columns 1-3) and *DownMeasure* (columns 4-6). We indeed find throughout all six columns that the positive effect of downstreamness on the intrafirm import share is concentrated in the highest quintile of D_i in the high ρ case. Moreover, in the low ρ case, the propensity towards integration is strongest in the lowest quintile of D_i .³⁹ Table 9 thus provides strong evidence consistent with this further prediction from our model.

In Table 10, we seek to correct our cross-country, cross-industry regressions for potential selection bias. As explained in Section 3.4, controlling for country fixed effects may not be sufficient for this purpose when the location of input production is itself affected by the level of downstreamness. For this reason, we carry out a two-stage Heckman selection procedure in Table 10 that seeks to correct for selection into foreign sourcing. In the context of our dataset, observations for which no imports were observed entering the U.S. are necessarily dropped from our regression sample, since the denominator of the intrafirm import share would equal zero in these cases. Any bias that might arise would moreover be more salient in the specifications that incorporate the full cross-country variation, since up to 60% of all potential country-industry-year observations are in fact zero import flows, and hence are dropped. To address this, we adopt as selection variables (to be included only in the first stage) an interaction term between an indicator for whether countries have above sample-median entry costs (taken from the Doing Business dataset) and the selling industry's R&D intensity. This country measure of entry costs is specifically the first principal component of the

³⁷These qualitative patterns should be taken with a pinch of salt. We are not able to consistently reject across all the specifications the hypothesis that the coefficients of $D_i \times \mathbf{1}(\rho_i > \rho_{med})$ in the highest and lowest headquarter intensity quintiles are in fact equal; a similar caveat applies to the coefficients for the substitutes case.

³⁸We drop one of these interactions due to the collinearity with the constant term in the regressions.

³⁹In this low ρ case, the difference between the coefficients of D_i in the highest and lowest quintiles is marginally significant at the 10% level in 3 out of the 6 columns. On the other hand, for the high ρ case, the effect of downstreamness in its highest quintile is always significantly different from that in its lowest quintile (the latter effect having been normalized to zero).

number of procedures, number of days, and the cost (as a percentage of income per capita) incurred to legally start a local business. (Note however that a number of countries have to be dropped due to the lack of entry cost data.) We in turn view the seller industry R&D variable as a proxy for high fixed costs particularly at the entry stage.⁴⁰ Following the logic of Helpman, Melitz and Rubinstein (2008), the level of such country-industry entry costs can be expected to affect the entry decisions of firms, and thus to also be positively correlated with the fixed costs that would be associated with expanding a firm to engage in export activities to the U.S. In contrast, these entry costs are less likely to directly shape the intensive margin of trade.⁴¹

The estimation itself proceeds in two steps. We first run a probit model based on equation (33), in which the dependent variable is a 0-1 variable indicating the presence of positive imports into the U.S. for the country-industry-year observation in question. Our proposed interaction variable indeed turns out to have a negative and significant effect at the 1% level, so that higher entry costs are *ceteris paribus* associated with a lower probability of importing to the U.S. (see column 1 for *DUse_TUse* and column 4 for *DownMeasure*). The second stage involves a weighted least squares specification (using the total import volume as weights) based on (33) once again, in which the inverse Mills ratio calculated from the first stage is included as a further regressor (columns 2 and 5). We once again obtain a positive and significant coefficient on D_i in the sequential complements case, as well as a negative but insignificant point estimate for the effect in the substitutes case. Relative to a weighted least squares specification that excludes the inverse Mills ratio (reported in columns 3 and 6), we find that correcting for selection tends to have relatively little effect on the coefficients of downstreamness in either the complements or the substitutes cases. This provides reassurance that any such selection bias is unlikely to be driving our core results on the relationship between downstreamness and integration decisions in foreign sourcing.

As a final note, we have further considered how our empirical strategy to distinguish the sequential complements and substitutes cases could be modified to accommodate the alternative “spider”-like production setting described in section 3.1.C. As we saw earlier, with symmetric Shapley bargaining across module producers, the sensitivity of each module-producer’s revenues with respect to the services of the module delivered would instead depend on the parameter that governs the cross-module elasticity of substitution (denoted by ζ). Direct estimates of this elasticity of substitution are naturally difficult to obtain, but we have nevertheless experimented with using import demand elasticities that were estimated by Broda and Weinstein (2006) for more aggregate product categories, namely at the SITC Revision 3, three-digit level. As documented in their paper (see in particular their footnote 22), these elasticities were estimated in part off the substitution

⁴⁰In unreported results, we have also experimented with measures of industry scale economies calculated as the employment per establishment, drawn from the U.S. Census Bureau’s County Business Patterns dataset. Our conclusions are largely unaffected with such alternative proxies of industry fixed costs.

⁴¹International tax considerations are often seen as another key factor that influences entry decisions by multinational firms. We have nevertheless verified that our results from specification (33) remain very similar when we restrict our sample to countries whose effective corporate tax rates for U.S. multinationals were within a 5% range of that for these multinationals’ domestic U.S. activities, based on a 2011 PriceWaterhouseCoopers survey canvassing U.S. firms entitled “Global Effective Tax Rates”. We thank Brent Neiman for bringing this survey to our attention.

seen across HS10 product codes that fall under each SITC three-digit heading, and thus would contain information on the degree of substitution across production modules under the further assumption that the constituent HS10 products in each SITC three-digit category are typically used together as module inputs in final-good production. We do not present these results in detail in our main paper given the rough nature of this proxy for ζ , but the patterns we find are qualitatively similar to our baseline findings, albeit with slightly lower levels of significance.⁴²

6 Conclusion

In this paper, we have developed a model of the organizational decisions of firms in which production entails a continuum of sequential stages. We have shown that, for each stage, the firm’s make-or-buy decision depends on that stage’s position in the value chain, and that dependence is crucially shaped by the relative magnitude of the average buyer demand elasticity faced by the firm and the degree of complementarity between inputs in production. When the average buyer demand elasticity is high relative to input substitutability, stage inputs are sequential complements and the firm finds it optimal to outsource relatively upstream stages and vertically integrate relatively downstream stages. In the converse case of a low demand elasticity relative to input substitutability, stage inputs are sequential substitutes and the firm instead finds it optimal to integrate relatively upstream stages and outsource relatively downstream stages. We have shown that our framework can be readily embedded into existing theoretical frameworks of global sourcing, which motivates our use of international trade data to test the model. Using data on U.S. related-party trade shares, we have shown that the evidence is broadly consistent with the model’s main predictions, as well as with several auxiliary predictions of our framework.

Although our empirical results are suggestive of the empirical relevance of our theory, we acknowledge the existence of a tension between our “firm-level” theoretical model and our “industry-level” empirical analysis. We are limited, however, by the data that is currently available to researchers in our field. It is our hope that in the near future, new firm-level datasets featuring detailed information on the sourcing decisions of firms for different inputs will become available.

⁴²The interested reader is directed to Online Appendix Tables 9 and 10, where these results are reported using *DUse_TUse* and *DownMeasure* respectively. To construct this proxy for ζ , we associated each SITC three-digit elasticity to each of its constituent HS10 product codes, before following analogous steps as with our earlier proxy for ρ to construct a weighted-average buyer elasticity. A median cutoff was once again used to delineate the complements case from the substitutes case.

References

- Acemoglu, Daron, Pol Antràs, and Elhanan Helpman, (2007), “Contracts and Technology Adoption,” *American Economic Review* 97(3): 916-943.
- Alfaro, Laura, and Andrew Charlton, (2009), “Intra-Industry Foreign Direct Investment,” *American Economic Review* 99(5): 2096-2119.
- Antràs, Pol, (2003), “Firms, Contracts, and Trade Structure,” *Quarterly Journal of Economics* 118(4): 1375-1418.
- Antràs, Pol, (2005), “Incomplete Contracts and the Product Cycle,” *American Economic Review* 95(4): 1054-1073.
- Antràs, Pol, (2011), “Grossman-Hart (1986) Goes Global: Incomplete Contracts, Property Rights, and the International Organization of Production,” *NBER Working Paper* 17470. *Journal of Law, Economics and Organization*, forthcoming.
- Antràs, Pol, and Elhanan Helpman, (2004), “Global Sourcing,” *Journal of Political Economy* 112(3): 552-580.
- Antràs, Pol, and Elhanan Helpman, (2008), “Contractual Frictions and Global Sourcing,” in E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*, Harvard University Press.
- Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry, (2012), “Measuring the Upstreamness of Production and Trade Flows,” *American Economic Review Papers & Proceedings* 102(3): 412-416.
- Baldwin, Richard, and Anthony Venables, (2010), “Spiders and Snakes: Offshoring and Agglomeration in the Global Economy,” *NBER Working Paper* 16611.
- Becker, Randy A., and Wayne B. Gray, (2009), “NBER-CES Manufacturing Industry Database (1958-2005)”.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, (2010), “Intra-Firm Trade and Product Contractibility,” *American Economic Review Papers & Proceedings* 100(2): 444-448.
- Broda, Christian, and David Weinstein, (2006), “Globalization and the Gains from Variety,” *Quarterly Journal of Economics* 121(2): 541-585.
- Costinot, Arnaud, Jonathan Vogel, and Su Wang, (2013), “An Elementary Theory of Global Supply Chains,” *Review of Economic Studies* 80(1): 109-144.
- Díez, Federico, (2010), “The Asymmetric Effects of Tariffs on Offshoring Industries: How North/South Tariffs Affect Intra-Firm Trade,” 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(4): 583-594.
- Fally, Thibault, (2012), “On the Fragmentation of Production in the U.S.,” mimeo.

- Feenstra, Robert C., John Romalis, Peter K. Schott, (2002), "U.S. Imports, Exports and Tariff Data, 1989-2001," *NBER Working Paper* 9387.
- Findlay, Ronald, (1978), "An 'Austrian' Model of International Trade and Interest Rate Equalization," *Journal of Political Economy* 86(6): 989-1007.
- Ghosh, Ambica, (1958), "Input-Output Approach in an Allocation System," *Economica* 25(97): 58-64.
- Grossman, Gene M., and Elhanan Helpman, (2005), "Outsourcing in a Global Economy," *Review of Economic Studies*, 72(1): 135-159.
- 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(4): 691-719.
- Harms, Philipp, Oliver Lorz, and Dieter Urban, (2012), "Offshoring along the Production Chain," *Canadian Journal of Economics* 45(1): 93-106.
- Hart, Oliver, and John Moore, (1994), "A Theory of Debt Based on the Inalienability of Human Capital," *Quarterly Journal of Economics* 109(4): 841-879.
- Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, (2004), "Exports versus FDI with Heterogeneous Firms," *American Economic Review* 94(1): 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.
- Kaufmann, Daniel, Aart Kraay, and Massimo Mastruzzi, (2010), "The Worldwide Governance Indicators: Methodology and Analytical Issues," *World Bank Policy Research Working Paper* 5430.
- Kim, Se-Jik and Hyun Song Shin, (2012), "Sustaining Production Chains Through Financial Linkages," *American Economic Review Papers & Proceedings* 102(3): 402-406.
- Kohler, Wilhelm, (2004), "International Outsourcing and Factor Prices with Multistage Production," *Economic Journal* 114(494): C166-C185.
- Kremer, Michael, (1993), "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics* 108(3): 551-575.
- Levchenko, Andrei, (2007), "Institutional Quality and International Trade," *Review of Economic Studies* 74(3): 791-819.
- Melitz, Marc J., (2003), "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71(6): 1695-1725.
- Miller, Ronald E., and Peter D. Blair, (2009), *Input-Output Analysis: Foundations and Extensions*, second edition, Cambridge University Press.
- Nunn, Nathan, (2007), "Relationship-Specificity, Incomplete Contracts and the Pattern of Trade," *Quarterly Journal of Economics* 122(2): 569-600.
- Nunn, Nathan, and Daniel Treffer, (2008), "The Boundaries of the Multinational Firm: An Empirical Analysis," in E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*, Harvard University Press.
- Nunn, Nathan, and Daniel Treffer, (2012), "Incomplete Contracts and the Boundaries of the Multinational Firm," *Journal of Economic Behavior and Organization*, forthcoming.

- 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.
- PriceWaterhouseCoopers LLP, (2011), "Global Effective Tax Rates."
- Rauch, James E., (1999), "Networks versus Markets in International Trade," *Journal of International Economics* 48(1): 7-35.
- Sanyal, Kalyan K., (1983), "Vertical Specialization in a Ricardian Model with a Continuum of Stages of Production," *Economica* 50(197): 71-78.
- Sanyal, Kalyan K., and Ronald W. Jones, (1982), "The Theory of Trade in Middle Products," *American Economic Review* 72(1): 16-31.
- Thomas, Jonathan, and Tim Worrall, (1994), "Foreign Direct Investment and the Risk of Expropriation," *Review of Economic Studies* 61(1): 81-108.
- Winter, Eyal (2006), "Optimal Incentives for Sequential Production Processes," *RAND Journal of Economics* 37(2): 376-390.
- Yi, Kei-Mu, (2003), "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111(1): 52-102.
- Zhang, Juyan, and Yi Zhang, (2008), "Sequential Hold-up, and Ownership Structure," mimeo.
- Zhang, Juyan, and Yi Zhang, (2011), "Sequential Hold-up, and Strategic Delay," mimeo.

A Mathematical Appendix

Proof of Lemma 1: The result is immediate given the optimal bargaining share $\beta^*(m)$ in (15), so we focus here on providing more details on the derivation of this equation. Let $L \equiv \left(1 - (v')^{\frac{1-\alpha}{\alpha}}\right) v' v^{\frac{\rho-\alpha}{\alpha(1-\rho)}}$ and note that:

$$\begin{aligned}\frac{\partial L}{\partial v} &= \frac{\rho - \alpha}{\alpha(1 - \rho)} \left(1 - (v')^{\frac{1-\alpha}{\alpha}}\right) v' v^{\frac{\rho-\alpha}{\alpha(1-\rho)} - 1}; \text{ and} \\ \frac{\partial L}{\partial v'} &= \left(1 - \frac{1}{\alpha}(v')^{\frac{1-\alpha}{\alpha}}\right) v^{\frac{\rho-\alpha}{\alpha(1-\rho)}}.\end{aligned}$$

The Euler-Lagrange condition associated with profit maximization is then:

$$\frac{\rho - \alpha}{\alpha(1 - \rho)} \left(1 - (v')^{\frac{1-\alpha}{\alpha}}\right) v' v^{\frac{\rho-\alpha}{\alpha(1-\rho)} - 1} = \left(1 - \frac{1}{\alpha}(v')^{\frac{1-\alpha}{\alpha}}\right) \frac{\rho - \alpha}{\alpha(1 - \rho)} v^{\frac{\rho-\alpha}{\alpha(1-\rho)} - 1} v' - \frac{1}{\alpha} \frac{1 - \alpha}{\alpha} (v')^{\frac{1-\alpha}{\alpha} - 1} v'' v^{\frac{\rho-\alpha}{\alpha(1-\rho)}}$$

where we have assumed (for the time being) that v' is at least piecewise differentiable. Simplifying the above yields:

$$v^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (v')^{\frac{1-\alpha}{\alpha} - 1} \left[v'' + \frac{\rho - \alpha}{1 - \rho} \frac{(v')^2}{v} \right] = 0.$$

To rule out discontinuous jumps in v' , we appeal to the Weierstrass-Erdmann condition. Suppose that v' has a discontinuous jump (and hence v has a corner) at some $\tilde{j} \in (0, 1)$. In other words, $\lim_{j \rightarrow \tilde{j}^-} v(j) = \lim_{j \rightarrow \tilde{j}^+} v(j)$, but $\lim_{j \rightarrow \tilde{j}^-} v'(j) \neq \lim_{j \rightarrow \tilde{j}^+} v'(j)$. By the Weierstrass-Erdmann condition, we must have, however, that $\lim_{j \rightarrow \tilde{j}^-} L_{v'}(j) = \lim_{j \rightarrow \tilde{j}^+} L_{v'}(j)$. Using the above expression for $L_{v'}$ and the fact that $\lim_{j \rightarrow \tilde{j}^-} v(j) = \lim_{j \rightarrow \tilde{j}^+} v(j)$, this immediately implies that $\lim_{j \rightarrow \tilde{j}^-} v'(j) = \lim_{j \rightarrow \tilde{j}^+} v'(j)$, which contradicts our assumption of a point of discontinuity for v' at \tilde{j} . In sum, v' is continuous.

There are three types of solutions to the Euler-Lagrange equation above:

$$\begin{aligned}(v(j))^{\frac{\rho-\alpha}{\alpha(1-\rho)}} &= 0 \text{ for all } j; \\ (v'(j))^{\frac{1-\alpha}{\alpha} - 1} &= 0 \text{ for all } j; \\ v'' + \frac{\rho - \alpha}{1 - \rho} \frac{(v')^2}{v} &= 0 \text{ for all } j.\end{aligned}\tag{35}$$

The first solution would generate a value of the problem equal to 0, while the second one implies $v'(j) = 0$ ($\beta(j) = 1$) for $\alpha < 1/2$ and $v'(j) \rightarrow +\infty$ ($\beta(j) \rightarrow -\infty$) for $\alpha > 1/2$. In the former case, the functional attains a value of 0, and in the latter case it goes to $-\infty$, so neither of these can constitute a maximum for the problem at hand. We will thus focus hereafter on the third case, which generates a strictly positive profit for the firm.

The second-order differential equation implied by the Euler-Lagrange necessary equation is straightforward to solve. In particular, define $z \equiv v'$ and note that we can write (35) as:

$$\frac{\rho - \alpha}{1 - \rho} \frac{(z)^2}{v} = -\frac{dz}{dj} = -\frac{dz}{dv} \frac{dv}{dj} = -\frac{dz}{dv} z,$$

so that:

$$\frac{\alpha - \rho}{1 - \rho} \frac{z}{v} = \frac{dz}{dv}.$$

It is simple to see that the solution to this first-order differential equation satisfies:

$$z = v' = C_1 v^{\frac{\alpha-\rho}{1-\rho}},$$

where C_1 is a positive constant. This leaves us with a new first-order differential equation, which again is straightforward to solve and yields:

$$v(j) = \left(\frac{(1-\alpha)C_1}{1-\rho} (j - C_2) \right)^{\frac{1-\rho}{1-\alpha}},$$

where C_2 is a second constant of integration. Now, imposing the initial condition $v(0) = 0$ and the transversality condition $v'(1)^{\frac{1-\alpha}{\alpha}} = \alpha$ at the right-boundary of the unit interval, we finally obtain:

$$v(j) = \left(\frac{1-\alpha}{1-\rho} \right) \alpha^{\frac{\alpha}{1-\alpha}} j^{\frac{1-\rho}{1-\alpha}},$$

from which $\beta^*(m)$ can be derived by recalling that $\beta(j) = 1 - v'(j)^{\frac{1-\alpha}{\alpha}}$.

In the Online Appendix we show that this solution also satisfies a sufficient condition for a maximum, and we also characterize the solution when $\beta^*(m)$ is constrained to take nonnegative values.

Proof of Proposition 2: As discussed in the main text, 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, 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 the 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 that up to a positive multiplicative constant:

$$\begin{aligned} \Pi_1 - \Pi_2 &\propto \int_{\tilde{m}-\varepsilon}^{\tilde{m}} \beta_V (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} [B + (j - \tilde{m} + \varepsilon)(1 - \beta_V)^{\frac{\alpha}{1-\alpha}}]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\ &\quad + \int_{\tilde{m}}^{\tilde{m}+\varepsilon} \beta_O (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} [B + \varepsilon(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} + (j - \tilde{m})(1 - \beta_O)^{\frac{\alpha}{1-\alpha}}]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\ &\quad - \int_{\tilde{m}-\varepsilon}^{\tilde{m}} \beta_O (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} [B + (j - \tilde{m} + \varepsilon)(1 - \beta_O)^{\frac{\alpha}{1-\alpha}}]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \\ &\quad - \int_{\tilde{m}}^{\tilde{m}+\varepsilon} \beta_V (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} [B + \varepsilon(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} + (j - \tilde{m})(1 - \beta_V)^{\frac{\alpha}{1-\alpha}}]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj. \end{aligned}$$

where we define $B \equiv \int_0^{\tilde{m}-\varepsilon} (1 - \beta(k))^{\frac{\alpha}{1-\alpha}} dk$. That the difference in profits depends only on profits in the interval $(\tilde{m} - \varepsilon, \tilde{m} + \varepsilon)$ and is not affected by decisions downstream follows from the fact that we have chosen the width ε to be common for both sub-intervals. Evaluating the integrals above with respect to j and simplifying, we obtain after some tedious algebra:

$$\begin{aligned} \Pi_1 - \Pi_2 \propto (\beta_V - \beta_O) & \left[(B + \varepsilon(1 - \beta_V))^{\frac{\alpha}{1-\alpha}} \frac{\rho(1-\alpha)}{\alpha(1-\rho)} + (B + \varepsilon(1 - \beta_O))^{\frac{\alpha}{1-\alpha}} \frac{\rho(1-\alpha)}{\alpha(1-\rho)} \right. \\ & \left. - (B + \varepsilon(1 - \beta_V))^{\frac{\alpha}{1-\alpha}} + \varepsilon(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \frac{\rho(1-\alpha)}{\alpha(1-\rho)} - B \frac{\rho(1-\alpha)}{\alpha(1-\rho)} \right]. \end{aligned}$$

Since $\beta_V - \beta_O > 0$, it suffices to show that the expression in square parentheses is negative. To see this, consider the function $f(y) = y^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}$. Simple differentiation will show that for $y, a > 0$ and $b \geq 0$, $f(y + a + b) - f(y + b)$ is an increasing function in b when $\rho > \alpha$. Hence, $(y + a + b)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - (y + b)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} > (y + a)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - (y)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}$. Setting $y = B$, $a = \varepsilon(1 - \beta_O)^{\frac{\alpha}{1-\alpha}}$ 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 can be established using an analogous proof by contradiction. The limit values in this case imply 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.

Proof of Proposition 3: We begin by deriving equations (16) and (17). For each case, this is achieved by first plugging the optimal values of $\beta(m) \in \{\beta_V, \beta_O\}$ for all $m \in [0, 1]$ implied by Proposition 2 into the firm's maximand in (11), and then solving:

$$\begin{aligned} m_C^* &= \arg \max_m \left\{ \beta_O(1 - \beta_O)^{\frac{\rho}{1-\rho}} \int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj + \beta_V(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_m^1 [(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} m + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} (j - m)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \right\}; \\ m_S^* &= \arg \max_m \left\{ \beta_V(1 - \beta_V)^{\frac{\rho}{1-\rho}} \int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj + \beta_O(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_m^1 [(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} m + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} (j - m)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \right\}. \end{aligned}$$

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

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

which after a few simplifications can be written as:

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

from which the formula for m_C^* in (16) is obtained. (One can moreover verify that the second-order condition when evaluated at m_C^* is indeed negative.) The formula for m_S^* in (17) can be derived in an analogous manner.

Using the expression in (16), one can show that $m_C^* < 1$ when $\beta_V(1 - \beta_V)^{\alpha/(1-\alpha)} > \beta_O(1 - \beta_O)^{\alpha/(1-\alpha)}$. When this inequality is satisfied, one can readily check that $1 - \frac{\beta_O}{\beta_V} > 1 - \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\alpha/(1-\alpha)}$, which from (16)

implies that $m_C^* > 0$, so that the cutoff stage lies strictly in the interior of $(0, 1)$. Conversely, using (17) for the substitutes case, we have that $m_S^* < 1$ whenever $\beta_V(1 - \beta_V)^{\alpha/(1-\alpha)} < \beta_O(1 - \beta_O)^{\alpha/(1-\alpha)}$, and that this condition is also sufficient to ensure that $m_S^* > 0$, so that $m_S^* \in (0, 1)$.

Next, consider the effect of ρ on these thresholds. For m_C^* , the exponent $\frac{\alpha(1-\rho)}{\rho-\alpha}$ is positive and decreasing in ρ . Moreover, when the parameter restrictions for $m_C^* \in (0, 1)$ apply, the fraction that is the base of this exponent is strictly greater than 1. It thus follows that $dm_C^*/d\rho > 0$, as claimed in the Proposition, namely that when ρ falls, m_C^* also falls, expanding the range of stages downstream of m_C^* that are integrated. An analogous set of arguments can be applied to show that $dm_S^*/d\rho < 0$.

Proof of Proposition 4: From the discussion in Section 3.2 and equation (10), we have that investments by suppliers in this extension with headquarter intensity will be given by:

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

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

It follows from equation (11) that the final good producer will capture revenues given by:

$$R_F = \left(A^{1-\rho} \theta^\rho \left(\frac{\tilde{\rho}}{c} \right)^{\tilde{\rho}} \left(\frac{h}{\eta} \right)^{\rho\eta} (1-\eta)^{-\tilde{\rho}} \right)^{\frac{1}{1-\tilde{\rho}}} \frac{\tilde{\rho}}{\alpha} \left(\frac{1-\tilde{\rho}}{1-\alpha} \right)^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}} \int_0^1 \beta(j)(1-\beta(j))^{\frac{\alpha}{1-\alpha}} \left[\int_0^j (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}} dj \quad (36)$$

Before suppliers make any investments, the firm will choose h to maximize $R_F - c_h h$. From equation (36), it is clear that this 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 $\pi_F = \left(\frac{1-\rho}{1-\tilde{\rho}} \right) R_F$. One can then substitute the optimal value of h from the above first-order condition back into the expression for R_F in (36), to solve for π_F as a function of the model parameters and the organizational decisions (the $\beta(j)$'s) only. From this, it will be straightforward to see that the sequence of organizational forms that maximizes profits will be that which maximizes:

$$\int_0^1 \beta(j)(1-\beta(j))^{\frac{\alpha}{1-\alpha}} \left[\int_0^j (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\tilde{\rho}-\alpha}{\alpha(1-\tilde{\rho})}} dj,$$

but this is precisely analogous to the objective function in the Benchmark Model, except with ρ replaced by $\tilde{\rho}$. 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 (16) and (17) respectively, with ρ replaced by $\tilde{\rho}$. Differentiating (16) and (17) 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:

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

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 choice of a (hypothetical) unconstrained optimal division of surplus $\beta^*(m)$ for stage m is not affected by the fixed costs terms, 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:

$$m_C^* = \arg \max_m \left\{ \begin{array}{l} \kappa \theta^{\frac{\rho}{1-\rho}} \left\{ \beta_O (1 - \beta_O)^{\frac{\rho}{1-\rho}} \int_0^m j^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj + \right. \\ \left. \beta_V (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_m^1 [(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} m + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} (j - m)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj \right\} - m f_O - (1 - m) f_V \end{array} \right\}$$

where recall that $\kappa = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{\rho}{c} \right)^{\frac{\rho}{1-\rho}}$. Solving this in a manner analogous to the proof of Proposition 3 delivers equation (24). The rest of the proof follows from the discussion in Section 3.3.

Proof of Proposition 6: We focus here on deriving equation (28), with the rest of the proof following from the same steps as the case with homogeneity in input costs and productivity. Notice first that we can write equation (27) as:

$$\pi_F(v) = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (\rho \theta)^{\frac{\rho}{1-\rho}} \int_0^1 \left[1 - v'(j)^{\frac{1-\alpha}{\alpha}} \gamma(j) \right] v'(j) [v(j)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} dj,$$

where:

$$v(j) \equiv \int_0^j \left(\frac{1 - \beta(k)}{\gamma(k)} \right)^{\frac{1-\alpha}{\alpha}} dk \quad (37)$$

and $\gamma(j) \equiv c(j)/\psi(j)$. The Euler-Lagrange equation associated with maximizing $\pi_F(v)$ is given by:

$$\frac{\rho - \alpha}{\alpha(1 - \rho)} \left[1 - v'(j)^{\frac{1-\alpha}{\alpha}} \gamma(j) \right] v'(j) [v(j)]^{\frac{\rho-\alpha}{\alpha(1-\rho)} - 1} = \frac{d}{dj} \left[[v(j)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(1 - \frac{1}{\alpha} v'(j)^{\frac{1-\alpha}{\alpha}} \gamma(j) \right) \right],$$

which after a couple of manipulations can be reduced to:

$$\frac{\rho - \alpha}{1 - \rho} \frac{[v'(j)]^2}{v(j)} + v'(j) \frac{\alpha}{1 - \alpha} \frac{\gamma'(j)}{\gamma(j)} = -v''(j). \quad (38)$$

Plugging (37) and its two first derivatives into (38) and simplifying produces (28).

B Data Appendix

Intrafirm trade: From the U.S. Census Bureau’s Related Party Trade Database, for the years 2000-2010. The data in NAICS industry codes were mapped to six-digit IO2002 industries using the correspondence provided by the Bureau Economic Analysis (BEA) as a supplement to the 2002 U.S. Input-Output (I-O) Tables. This is a straightforward many-to-one mapping for the manufacturing industries (NAICS first digit = 3). Two industries required a separate treatment as the Census Bureau data was at a coarser level of aggregation than could be mapped into six-digit IO2002 codes. A synthetic code 31131X was created to merge IO 311313 (Beet sugar manufacturing) and 31131A (Sugar cane mills and refining), while a separate code 33641X merged IO 336311, 336412, 336413, 336414, 33641A (all related to the manufacture of aircraft and related components). All other industry variables described below were also constructed for these two synthetic IO2002 codes. After converting the related party and non-related party import data to the IO2002 codes, the share of intrafirm imports was calculated for each industry-year or country-industry-year as: $(Related\ Trade)/(Related\ Trade + Non-Related\ Trade)$.

DUse_TUse: Calculated from the 2002 U.S. I-O Tables, as described in Section 4.2, using the detailed Supplementary Use Table after redefinitions issued by the BEA. For the synthetic codes 33131X and 33641X, we took a weighted average of the $DUse_TUse$ values of the component IO2002 industries, using the output of these component industries as weights.

DownMeasure: Calculated from the 2002 U.S. I-O Tables, as described in Section 4.2. In particular, *DownMeasure* is the reciprocal of the upstreamness measure discussed in detail in Antràs et al. (2012). A treatment analogous to that described above for $DUse_TUse$ was used to obtain *DownMeasure* for the synthetic codes 33131X and 33641X.

Import demand elasticities: U.S. import demand elasticities for HS10 products were from Broda and Weinstein (2006). This was merged with a comprehensive list of HS10 codes from Pierce and Schott (2009). For each HS10 code missing an elasticity value, we assigned a value equal to the trade-weighted average elasticity of the available HS10 codes with which it shared the same first nine digits. This was done successively up to codes that shared the same first two digits, to fill in as many HS10 elasticities as possible. Using the IO-HS concordance provided by the BEA with the 2002 U.S. I-O Tables, we then took the trade-weighted average of the HS10 elasticities within each IO2002 category. At each stage, the weights used were the total value of U.S. imports by HS10 code from 1989-2006, calculated from Feenstra et al. (2002). There remained 13 IO2002 industries without elasticity values after the above procedure. For these, we assigned a value equal to the weighted average elasticity of the IO2002 codes with which the industry shared the same first four digits, or (if the value was still missing) the same first three digits, using industry output values as weights. This yielded import elasticities for the industry that sells the input in question. For the average buyer elasticity, we took a weighted average of the elasticities of industries that purchase the input in question, with weights equal to these input purchase values as reported in the 2002 U.S. I-O Tables.

Factor intensities: From the NBER-CES Manufacturing Industry Database (Becker and Gray, 2009). Skill intensity is the log of the number of non-production workers divided by total employment. Physical capital intensity is the log of the real capital stock per worker. Equipment capital intensity and plant capital intensity are respectively the log of the equipment and plant capital stock per worker. Materials intensity is the log of materials purchases per worker. The NBER-CES data for NAICS industries were mapped to IO2002 codes using the procedure described above for the related party trade data. For each factor intensity variable, a simple average of the annual values from 2000-2005 was taken to obtain the seller industry

measures. The factor intensities for the average buyer were then calculated using the same procedure as described for the average buyer import demand elasticity.

R&D intensity: From Nunn and Treffer (2012), who calculated R&D expenditures to total sales on an annual basis for IO1997 industries using the U.S. firms in the Orbis dataset. We constructed a crosswalk from IO1997 to IO2002 through the NAICS industry codes. The R&D intensity for each IO2002 industry was then calculated as the weighted average value of $\log(0.001 + R\&D/Sales)$ over that of its constituent IO1997 industries over the years 2000-2005, using the industry output values in the 1997 U.S. I-O Tables as weights. A similar procedure to that described above for the import demand elasticity was used to obtain the R&D intensity for the remaining 13 IO2002 codes. The R&D intensity for the average buyer was then calculated using the same procedure as described for the average buyer import demand elasticity.

Dispersion: From Nunn and Treffer (2008), who constructed dispersion for each HS6 code as the standard deviation of log exports for its HS10 sub-codes across U.S. port locations and destination countries in the year 2000, from U.S. Department of Commerce data. We associated the dispersion value of each HS6 code to each of its HS10 sub-codes. These were mapped into IO2002 industries using the IO-HS concordance, taking a trade-weighted average of the dispersion value over HS10 constituent codes; the weights used were the total value of U.S. imports for each HS10 code from 1989-2006, from Feenstra et al. (2002). A similar procedure to that described above for the import demand elasticity was used to obtain the dispersion measure for the remaining 13 IO2002 codes. The dispersion for the average buyer was then calculated using the same procedure described for the average buyer import demand elasticity.

Other industry controls: Value-added over the value of shipments was calculated from the NBER-CES Manufacturing Industry Database, as described for the factor intensity variables; an average over 2000-2005 was used. Input “importance” was computed from the 2002 U.S. I-O Tables, as the industry’s total use value as an input divided by the total input purchases made by all of its buyer industries.

Contractibility was computed from the 2002 U.S. I-O Tables, following the methodology of Nunn (2007). For each IO2002 industry, we first calculated the fraction of HS10 constituent codes classified by Rauch (1999) as neither reference-priced nor traded on an organized exchange, under Rauch’s “liberal” classification. (The original Rauch classification was for SITC Rev. 2 products; these were associated with HS10 codes using a mapping derived from U.S. imports in Feenstra et al. (2002).) We took one minus this value as a measure of the own contractibility of each IO2002 industry. The average buyer contractibility was then calculated using the same procedure described for computing the average buyer import demand elasticity.

Intermediation was from Bernard et al. (2010), who calculated this from U.S. establishment-level data as the weighted average of the wholesale employment share of firms in 1997, using the import share of each firm as weights. This variable was reported at the HS2 level in the NBER long version of their paper. We associated the intermediation value for each HS2 code to each of its HS10 sub-codes. These were mapped into IO2002 industries using the IO-HS concordance, taking a trade-weighted average of the intermediation value over HS2 constituent codes; the weights used were the total value of U.S. imports for each HS10 code from 1989-2006, from Feenstra et al. (2002). A similar procedure to that described above for the import demand elasticity was used to obtain the intermediation measure for the remaining 13 IO2002 codes.

Country variables: Country entry costs were taken from the Doing Business dataset. Data on the number of procedures, number of days, and cost (as a percentage of income per capita) required to start a business were used. These were averaged over 2003-2005 for each variable. Country rule of law was from the Worldwide Governance Indicators (Kaufmann et al., 2010). The annual index was linearly rescaled from its original range of -2.5 to 2.5 , to lie between 0 and 1.

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

IO2002	Industry	DUse_TUse	IO2002	Industry	DownMeasure
<i>Lowest 10 values</i>			<i>Lowest 10 values</i>		
331314	Secondary smelting and alloying of aluminum	0.0000	325110	Petrochemical	0.2150
325110	Petrochemical	0.0599	331411	Primary smelting and refining of copper	0.2296
331411	Primary smelting and refining of copper	0.0741	331314	Secondary smelting and alloying of aluminum	0.2461
325211	Plastics material and resin	0.1205	325190	Other basic organic chemical	0.2595
325910	Printing ink	0.1325	33131A	Alumina refining and primary aluminum	0.2622
311119	Other animal food	0.1385	325310	Fertilizer	0.2658
333220	Plastics and rubber industry machinery	0.1420	335991	Carbon and graphite product	0.2668
33131A	Alumina refining and primary aluminum	0.1447	325181	Alkalies and chlorine	0.2769
335991	Carbon and graphite product	0.1615	331420	Copper rolling, drawing, extruding, and alloying	0.2769
331420	Copper rolling, drawing, extruding, and alloying	0.1804	325211	Plastics material and resin	0.2800
<i>Highest 10 values</i>			<i>Highest 10 values</i>		
334517	Irradiation apparatus	0.9669	339930	Doll, Toy, and Game	0.9705
339930	Doll, Toy, and Game	0.9686	311111	Dog and cat food	0.9717
337910	Mattress	0.9779	337910	Mattress	0.9720
322291	Sanitary paper product	0.9790	315230	Women's and girl's cut and sew apparel	0.9762
337121	Upholstered household furniture	0.9864	321991	Manufactured home (mobile home)	0.9810
337212	Office furniture and custom woodwork & millwork	0.9868	336212	Truck trailer	0.9837
336213	Motor home	0.9879	336213	Motor home	0.9879
33299A	Ammunition	0.9956	316200	Footwear	0.9927
316200	Footwear	0.9967	337121	Upholstered household furniture	0.9928
336111	Automobile	0.9997	336111	Automobile	0.9997

Notes: Tabulated based on the set of 253 IO2002 manufacturing industries for which data on intrafirm import shares was available.

Table 2
Baseline Determinants of the Intrafirm Import Share

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Log (s/l)	0.200*** [0.031]	0.109*** [0.038]	0.117*** [0.037]	0.143*** [0.039]	0.004 [0.043]	0.019 [0.043]
Log (k/l)	0.068*** [0.014]	0.037* [0.021]		0.096*** [0.018]	0.047 [0.028]	
Log (equipment k / l)			0.070*** [0.024]			0.083** [0.033]
Log (plant k / l)			-0.051 [0.032]			-0.063 [0.045]
Log (materials/l)		0.018 [0.026]	0.023 [0.026]		0.056 [0.035]	0.063* [0.035]
Log (0.001+ R&D/Sales)		0.029*** [0.007]	0.030*** [0.006]		0.055*** [0.009]	0.054*** [0.009]
Dispersion		0.130** [0.063]	0.150** [0.061]		0.083 [0.070]	0.120 [0.075]
Industry controls for: Year fixed effects?	Seller Yes	Seller Yes	Seller Yes	Buyer Yes	Buyer Yes	Buyer Yes
Observations	2783	2783	2783	2783	2783	2783
R-squared	0.23	0.30	0.31	0.17	0.27	0.28

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. All columns use industry-year observations controlling for year fixed effects. Estimation is by OLS. Industry factor intensity and dispersion variables in Columns 1-3 are that of the seller industry (namely, the industry that sells the input in question), while in Columns 4-6, 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)	(8)
				Elas < Median	Elas >= Median	Weighted		Weighted
Log (s/l)	0.005 [0.044]	0.039 [0.043]	0.056 [0.042]	0.112* [0.064]	0.038 [0.055]	-0.098 [0.079]	0.005 [0.020]	-0.068 [0.076]
Log (k/l)	0.044 [0.029]	0.034 [0.027]						
Log (equipment k / l)			0.085** [0.034]	0.022 [0.047]	0.153*** [0.043]	0.188*** [0.061]	0.026 [0.016]	0.134*** [0.051]
Log (plant k / l)			-0.077* [0.045]	-0.011 [0.057]	-0.159** [0.064]	-0.151** [0.070]	-0.056*** [0.019]	-0.142*** [0.050]
Log (materials/l)	0.058* [0.035]	0.060* [0.034]	0.065* [0.033]	0.049 [0.049]	0.072 [0.047]	0.060 [0.058]	0.025* [0.014]	0.080 [0.049]
Log (0.001+ R&D/Sales)	0.055*** [0.009]	0.054*** [0.009]	0.053*** [0.009]	0.050*** [0.013]	0.054*** [0.013]	0.090*** [0.018]	0.031*** [0.004]	0.073*** [0.016]
Dispersion	0.081 [0.070]	0.061 [0.070]	0.103 [0.075]	0.034 [0.108]	0.188* [0.100]	0.160 [0.124]	0.108*** [0.038]	0.083 [0.108]
DUse_TUse	-0.018 [0.054]			-0.216*** [0.075]	0.225*** [0.069]			
DUse_TUse X 1(Elas < Median), β_1		-0.196*** [0.071]	-0.174** [0.072]			-0.166* [0.089]	-0.115*** [0.033]	-0.075 [0.073]
DUse_TUse X 1(Elas > Median), β_2		0.171** [0.067]	0.198*** [0.068]			0.482*** [0.123]	-0.035 [0.030]	0.352*** [0.118]
1(Elas > Median)		-0.191*** [0.062]	-0.191*** [0.061]			-0.410*** [0.085]	-0.049* [0.029]	-0.291*** [0.075]
p-value: Joint significance of β_1 and β_2		[0.0008]	[0.0005]			[0.0000]	[0.0021]	[0.0013]
p-value: Test of $\beta_2 - \beta_1 = 0$		[0.0002]	[0.0001]			[0.0000]	[0.0595]	[0.0003]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects?	No	No	No	No	No	No	Yes	Yes
Observations	2783	2783	2783	1375	1408	2783	207991	207991
R-squared	0.27	0.32	0.33	0.37	0.28	0.61	0.18	0.59

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-6 use industry-year observations controlling for year fixed effects, while Columns 7-8 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 are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 4 and 5 restrict the sample to observations where the buyer industry elasticity is smaller (respectively larger) than the industry 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: DownMeasure

	Dependent variable: Intrafirm Import Share							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
				Elas < Median	Elas >= Median	Weighted		Weighted
Log (s/l)	-0.011 [0.045]	0.019 [0.042]	0.037 [0.041]	0.088 [0.067]	0.032 [0.051]	-0.143** [0.058]	0.000 [0.020]	-0.097* [0.054]
Log (k/l)	0.062** [0.027]	0.058** [0.026]						
Log (equipment k / l)			0.125*** [0.036]	0.061 [0.047]	0.192*** [0.048]	0.157** [0.066]	0.039** [0.017]	0.139*** [0.047]
Log (plant k / l)			-0.100** [0.049]	-0.027 [0.061]	-0.192*** [0.073]	-0.091 [0.080]	-0.062*** [0.020]	-0.117** [0.052]
Log (materials/l)	0.050 [0.033]	0.043 [0.033]	0.049 [0.032]	0.022 [0.050]	0.064 [0.043]	0.032 [0.055]	0.017 [0.014]	0.046 [0.043]
Log (0.001+ R&D/Sales)	0.058*** [0.010]	0.054*** [0.009]	0.054*** [0.009]	0.053*** [0.014]	0.050*** [0.014]	0.089*** [0.017]	0.032*** [0.004]	0.072*** [0.014]
Dispersion	0.087 [0.072]	0.092 [0.076]	0.150* [0.079]	0.082 [0.114]	0.258** [0.105]	0.249* [0.147]	0.116*** [0.043]	0.163 [0.114]
DownMeasure	0.101* [0.055]			-0.036 [0.069]	0.342*** [0.081]			
DownMeasure X 1(Elas < Median), β_1		-0.024 [0.064]	0.025 [0.065]			-0.119 [0.107]	-0.002 [0.035]	-0.002 [0.091]
DownMeasure X 1(Elas > Median), β_2		0.249*** [0.085]	0.298*** [0.081]			0.526*** [0.100]	-0.039 [0.033]	0.440*** [0.089]
1(Elas > Median)		-0.115* [0.064]	-0.110* [0.062]			-0.386*** [0.081]	0.022 [0.030]	-0.279*** [0.072]
p-value: Joint significance of β_1 and β_2		[0.0138]	[0.0013]			[0.0000]	[0.4995]	[0.0000]
p-value: Test of $\beta_2 - \beta_1 = 0$		[0.0103]	[0.0067]			[0.0000]	[0.4180]	[0.0001]
Industry controls for:	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer	Buyer
Year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Country-Year fixed effects?	No	No	No	No	No	No	Yes	Yes
Observations	2783	2783	2783	1375	1408	2783	207991	207991
R-squared	0.28	0.31	0.33	0.33	0.33	0.64	0.18	0.61

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. Columns 1-6 use industry-year observations controlling for year fixed effects, while Columns 7-8 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 are a weighted average of the characteristics of buyer industries (the industries that buy the input in question). Columns 4 and 5 restrict the sample to observations where the buyer industry elasticity is smaller (respectively larger) than the industry 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

Downstreamness:	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
	DUse_TUse	DUse_TUse Weighted	DUse_TUse Weighted	DownMeasure	DownMeasure Weighted	DownMeasure Weighted
Downstream X 1(Elas Quintile 1), β_1	-0.165* [0.093]	-0.255* [0.149]	-0.138 [0.100]	0.049 [0.119]	-0.283 [0.202]	-0.089 [0.142]
Downstream X 1(Elas Quintile 2), β_2	-0.173 [0.108]	-0.100 [0.136]	-0.037 [0.115]	-0.042 [0.099]	-0.154 [0.139]	-0.040 [0.122]
Downstream X 1(Elas Quintile 3), β_3	-0.145 [0.130]	0.007 [0.166]	0.051 [0.153]	0.019 [0.124]	0.114 [0.166]	0.224 [0.172]
Downstream X 1(Elas Quintile 4), β_4	0.215** [0.100]	0.156 [0.143]	0.066 [0.119]	0.332*** [0.126]	0.066 [0.174]	0.073 [0.120]
Downstream X 1(Elas Quintile 5), β_5	0.198* [0.103]	0.785*** [0.194]	0.637*** [0.199]	0.312** [0.130]	0.736*** [0.110]	0.621*** [0.096]
p-value: Joint significance of $\beta_1, \beta_2, \dots, \beta_5$	[0.0104]	[0.0010]	[0.0317]	[0.0338]	[0.0000]	[0.0000]
p-value: Test of $\beta_5 - \beta_1 = 0$	[0.0092]	[0.0000]	[0.0006]	[0.1344]	[0.0000]	[0.0001]
	Additional buyer industry controls included: Main effects of Elasticity Quintile dummies, Log (s/l), Log (equipment k / l), Log (plant k / l), Log (materials/l), Log (0.001+ R&D/Sales), Dispersion					
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	2783	2783	207991	2783	2783	207991
R-squared	0.34	0.64	0.61	0.34	0.67	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. Columns 1-3 use DUse_TUse and Columns 4-6 use DownMeasure as the downstreamness variable respectively. All columns include additional control variables whose coefficients are not reported, namely: (i) the main effects of the buyer industry elasticity quintile dummies, and (ii) buyer industry factor intensity and dispersion variables. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 6
Robustness Checks: DUse_TUse

	Dependent variable: Intrafirm Import Share									
	(1)	(2)	(3)	(4)	(5)	(6) Weighted	(7) Weighted	(8) Weighted	(9) Weighted	(10) Weighted
DUse_TUse X 1(Elas < Median), β_1	-0.187** [0.074]	-0.177** [0.072]	-0.103 [0.071]	-0.180** [0.077]	-0.130* [0.079]	-0.171* [0.093]	-0.188** [0.093]	-0.119 [0.080]	-0.199** [0.091]	-0.185** [0.084]
DUse_TUse X 1(Elas < Median), β_2	0.157** [0.077]	0.198*** [0.068]	0.236*** [0.065]	0.159** [0.068]	0.158** [0.074]	0.472*** [0.144]	0.504*** [0.111]	0.451*** [0.125]	0.391*** [0.103]	0.374*** [0.090]
Value-added / Value shipments	0.187 [0.135]				0.216* [0.125]	0.054 [0.275]				0.107 [0.167]
Input "Importance"		-1.687 [1.231]			-1.453 [1.302]		-3.484*** [1.021]			-3.975*** [0.802]
Intermediation			-0.464*** [0.106]		-0.413*** [0.102]			-0.673*** [0.168]		-0.654*** [0.149]
Own Contractibility				0.019 [0.048]	0.024 [0.047]				0.184** [0.076]	0.198*** [0.075]
Buyer contractibility				-0.199*** [0.067]	-0.169** [0.065]				-0.499*** [0.108]	-0.508*** [0.109]
p-value: Joint significance of β_1 and β_2	[0.0022]	[0.0004]	[0.0003]	[0.0012]	[0.0074]	[0.0001]	[0.0000]	[0.0002]	[0.0000]	[0.0000]
p-value: Test of $\beta_2 - \beta_1 = 0$	[0.0005]	[0.0001]	[0.0002]	[0.0003]	[0.0019]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
	Additional buyer industry controls included: 1(Elas > Median), Log (s/l), Log (equipment k / l), Log (plant k / l), Log (materials/l), Log (0.001+ R&D/Sales), Dispersion									
Year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2783	2783	2783	2783	2783	2783	2783	2783	2783	2783
R-squared	0.33	0.33	0.38	0.37	0.41	0.61	0.63	0.65	0.66	0.73

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. All columns use industry-year observations controlling for year fixed effects. Estimation is by OLS. The value-added / value shipments, intermediation, input "importance", and own contractibility variables are characteristics of the seller industry (namely, the industry that sells the input in question), while the buyer contractibility variable is a weighted average of the contractibility of buyer industries (the industries that buy the input in question). All columns include additional control variables whose coefficients are not reported, namely: (i) the level effect of the buyer industry elasticity dummy, and (ii) buyer industry factor intensity and dispersion variables. "Weighted" columns use the value of total imports for the industry-year as regression weights.

Table 7
Robustness Checks: DownMeasure

	Dependent variable: Intrafirm Import Share									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
						Weighted	Weighted	Weighted	Weighted	Weighted
DownMeasure X 1(Elas < Median), β_1	0.014 [0.067]	0.021 [0.065]	0.083 [0.062]	0.007 [0.067]	0.046 [0.065]	-0.137 [0.110]	-0.152 [0.110]	-0.040 [0.101]	-0.213** [0.103]	-0.216** [0.099]
DownMeasure X 1(Elas < Median), β_2	0.277*** [0.084]	0.294*** [0.081]	0.310*** [0.076]	0.284*** [0.075]	0.270*** [0.075]	0.515*** [0.098]	0.505*** [0.098]	0.500*** [0.097]	0.442*** [0.084]	0.371*** [0.066]
Value-added / Value shipments	0.168 [0.130]				0.235* [0.120]	0.331 [0.208]				-0.599*** [0.150]
Input "Importance"		-0.954 [1.351]			-0.866 [1.403]		-1.565* [0.806]			-2.805*** [0.633]
Intermediation			-0.488*** [0.104]		-0.443*** [0.103]			-0.652*** [0.161]		-0.599*** [0.150]
Own Contractibility				0.062 [0.046]	0.059 [0.044]				0.204*** [0.077]	0.211*** [0.064]
Buyer contractibility				-0.238*** [0.065]	-0.197*** [0.063]				-0.509*** [0.110]	-0.514*** [0.102]
p-value: Joint significance of β_1 and β_2	[0.0050]	[0.0015]	[0.0002]	[0.0009]	[0.0018]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
p-value: Test of $\beta_2 - \beta_1 = 0$	[0.0102]	[0.0066]	[0.0170]	[0.0039]	[0.0169]	[0.0000]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
	Additional buyer industry controls included: 1(Elas > Median), Log (s/l), Log (equipment k / l), Log (plant k / l), Log (materials/l), Log (0.001+ R&D/Sales), Dispersion									
Year fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2783	2783	2783	2783	2783	2783	2783	2783	2783	2783
R-squared	0.33	0.39	0.33	0.37	0.42	0.65	0.68	0.65	0.69	0.75

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Standard errors are clustered by industry. All columns use industry-year observations controlling for year fixed effects. Estimation is by OLS. The value-added / value shipments, intermediation, input "importance", and own contractibility variables are characteristics of the seller industry (namely, the industry that sells the input in question), while the buyer contractibility variable is a weighted average of the contractibility of buyer industries (the industries that buy the input in question). All columns include additional control variables whose coefficients are not reported, namely: (i) the level effect of the buyer industry elasticity dummy, and (ii) buyer industry factor intensity and dispersion variables. "Weighted" columns use the value of total imports for the industry-year as regression weights.

Table 8
Extension: Effect of Headquarter Intensity

Downstreamness:	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
	DUse_TUse	DUse_TUse Weighted	DUse_TUse Weighted	DownMeasure	DownMeasure Weighted	DownMeasure Weighted
	HQ intensity: First Principal Component of Buyer Industry Log (s/l), Log (equipment k/l), and Log (0.001+R&D/Sales)					
Downstream X 1 (Elas < Med) X (HQ Quin 1), β_{11}	0.049 [0.129]	0.009 [0.139]	-0.016 [0.118]	0.351*** [0.128]	0.301** [0.134]	0.290** [0.129]
Downstream X 1 (Elas < Med) X (HQ Quin 2), β_{12}	-0.268** [0.118]	-0.477*** [0.162]	-0.331** [0.148]	-0.161 [0.103]	-0.312* [0.186]	-0.169 [0.150]
Downstream X 1 (Elas < Med) X (HQ Quin 3), β_{13}	-0.175 [0.148]	-0.162 [0.138]	-0.158 [0.124]	0.072 [0.131]	-0.020 [0.160]	-0.058 [0.145]
Downstream X 1 (Elas < Med) X (HQ Quin 4), β_{14}	-0.377*** [0.134]	-0.658*** [0.108]	-0.455*** [0.069]	-0.121 [0.128]	-0.394** [0.172]	-0.280** [0.125]
Downstream X 1 (Elas < Med) X (HQ Quin 5), β_{15}	0.177 [0.171]	-0.046 [0.192]	0.024 [0.172]	0.229 [0.157]	-0.181 [0.160]	-0.050 [0.158]
Downstream X 1 (Elas > Med) X (HQ Quin 1), β_{21}	0.196 [0.248]	0.805** [0.325]	0.798*** [0.278]	0.465*** [0.163]	0.624*** [0.113]	0.593*** [0.087]
Downstream X 1 (Elas > Med) X (HQ Quin 2), β_{22}	-0.037 [0.146]	-0.099 [0.181]	-0.180 [0.175]	0.032 [0.130]	-0.001 [0.238]	-0.061 [0.189]
Downstream X 1 (Elas > Med) X (HQ Quin 3), β_{23}	0.216* [0.119]	0.172 [0.207]	0.027 [0.145]	0.358*** [0.135]	0.357 [0.257]	0.099 [0.199]
Downstream X 1 (Elas > Med) X (HQ Quin 4), β_{24}	0.130 [0.161]	0.819** [0.390]	0.427 [0.341]	0.280* [0.159]	0.519 [0.427]	0.079 [0.352]
Downstream X 1 (Elas > Med) X (HQ Quin 5), β_{25}	0.183** [0.082]	0.271* [0.144]	0.162 [0.114]	0.179 [0.117]	0.274 [0.197]	0.225* [0.117]
p-value: Joint significance of $\beta_{11}, \dots, \beta_{15}, \beta_{21}, \dots, \beta_{25}$	[0.0065]	[0.0000]	[0.0000]	[0.0004]	[0.0000]	[0.0000]
p-value: Test of $\beta_{11} - \beta_{15} = 0$	[0.5271]	[0.8062]	[0.8438]	[0.5397]	[0.0182]	[0.0833]
p-value: Test of $\beta_{21} - \beta_{25} = 0$	[0.9599]	[0.1790]	[0.0477]	[0.1699]	[0.1921]	[0.0262]
	Additional buyer industry controls included: Main and double interaction effects, Log (s/l), Log (equipment k / l), Log (plant k / l), Log (materials/l), Log (0.001+ R&D/Sales), Dispersion					
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	2783	2783	207991	2783	2783	207991
R-squared	0.40	0.69	0.63	0.40	0.71	0.65

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. Columns 1-3 use DUse_TUse and Columns 4-6 use DownMeasure as the downstreamness variable respectively. The hq intensity measure is the first principal component of the buyer industry's Log (s/l), Log (equipment k/l), and Log (0.001 + R&D/Sales). All columns include additional control variables whose coefficients are not reported, namely: (i) the main and double interaction effects of the buyer industry elasticity dummies and hq intensity quintile dummies, and (ii) buyer industry factor intensity and dispersion variables. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 9
Extension: Implications of Firm Heterogeneity

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
Downstreamness:	DUse_TUse	DUse_TUse	DUse_TUse	DownMeasure	DownMeasure	DownMeasure
		Weighted	Weighted		Weighted	Weighted
(Downstream Quin 1) X 1(Elas < Median), β_{11}	0.101** [0.045]	0.210*** [0.058]	0.145*** [0.045]	0.060 [0.051]	0.184*** [0.067]	0.145** [0.059]
(Downstream Quin 2) X 1(Elas < Median), β_{12}	0.041 [0.043]	0.131** [0.062]	0.096* [0.054]	-0.005 [0.045]	0.062 [0.069]	0.019 [0.055]
(Downstream Quin 3) X 1(Elas < Median), β_{13}	0.008 [0.046]	0.063 [0.057]	0.010 [0.046]	-0.016 [0.049]	0.089 [0.076]	0.049 [0.060]
(Downstream Quin 4) X 1(Elas < Median), β_{14}	-0.047 [0.043]	0.092 [0.065]	0.047 [0.051]	0.024 [0.045]	0.064 [0.075]	0.032 [0.058]
(Downstream Quin 5) X 1(Elas < Median), β_{15}	0.017 [0.042]	0.084 [0.063]	0.082 [0.053]	0.023 [0.042]	0.039 [0.081]	0.058 [0.064]
(Downstream Quin 2) X 1(Elas > Median), β_{22}	0.021 [0.036]	0.031 [0.067]	0.016 [0.056]	0.013 [0.039]	-0.054 [0.069]	-0.068 [0.052]
(Downstream Quin 3) X 1(Elas > Median), β_{23}	0.086** [0.044]	0.150** [0.065]	0.072 [0.052]	0.035 [0.047]	0.076 [0.069]	0.018 [0.066]
(Downstream Quin 4) X 1(Elas > Median), β_{24}	0.030 [0.048]	0.086 [0.056]	0.017 [0.050]	0.073* [0.043]	0.056 [0.072]	0.013 [0.060]
(Downstream Quin 5) X 1(Elas > Median), β_{25}	0.181*** [0.048]	0.352*** [0.067]	0.257*** [0.062]	0.171*** [0.059]	0.380*** [0.085]	0.300*** [0.074]
p-value: Joint significance of $\beta_{11}, \dots, \beta_{15}, \beta_{21}, \dots, \beta_{25}$	[0.0002]	[0.0000]	[0.0000]	[0.0607]	[0.0000]	[0.0000]
p-value: Test of $\beta_{11} - \beta_{15} = 0$	[0.0644]	[0.0530]	[0.1986]	[0.4766]	[0.0278]	[0.1006]
p-value: Test of $\beta_{25} = 0$	[0.0002]	[0.0000]	[0.0000]	[0.0043]	[0.0000]	[0.0001]
	Additional buyer industry controls included: Main effects of Downstreamness Quintiles and Elasticity dummies, Log (s/l), Log (equipment k / l), Log (plant k / l), Log (materials/l), Log (0.001+ R&D/Sales), Dispersion					
Downstream Quin and Elasticity 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	2783	2783	207991	2783	2783	207991
R-squared	0.35	0.63	0.61	0.33	0.67	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. Columns 1-3 use DUse_TUse and Columns 4-6 use DownMeasure as the downstreamness variable respectively. All columns include additional control variables whose coefficients are not reported, namely: (i) the main effects of the downstreamness quintile dummies and the buyer industry elasticity dummies, and (ii) buyer industry factor intensity and dispersion variables. "Weighted" columns use the value of total imports for the industry-year or country-industry-year respectively as regression weights.

Table 10
Selection into Intrafirm Trade

	Dependent variable: Intrafirm Import Share					
	(1)	(2)	(3)	(4)	(5)	(6)
	DUse_TUse Probit	DUse_TUse Weighted	DUse_TUse Weighted	DownMeasure Probit	DownMeasure Weighted	DownMeasure Weighted
Downstreamness:						
Log (s/l)	0.424** [0.166]	-0.079 [0.065]	-0.072 [0.059]	0.343** [0.167]	-0.107** [0.048]	-0.104** [0.050]
Log (equipment k / l)	-0.084 [0.203]	0.131*** [0.040]	0.130*** [0.041]	-0.056 [0.176]	0.125*** [0.040]	0.124*** [0.041]
Log (plant k / l)	-0.026 [0.234]	-0.139*** [0.046]	-0.140*** [0.046]	-0.030 [0.217]	-0.113** [0.049]	-0.113** [0.048]
Log (materials/l)	-0.083 [0.136]	0.063 [0.043]	0.061 [0.042]	-0.098 [0.133]	0.042 [0.042]	0.040 [0.041]
Buyer Log (0.001+ R&D/Sales)	-0.015 [0.064]	0.080*** [0.014]	0.080*** [0.014]	0.001 [0.061]	0.080*** [0.014]	0.080*** [0.014]
Dispersion	-0.234 [0.405]	0.117 [0.104]	0.113 [0.104]	-0.284 [0.392]	0.178 [0.115]	0.175 [0.117]
Downstream X 1(Elas < Median), β_1	0.047 [0.310]	-0.073 [0.074]	-0.072 [0.074]	0.634** [0.317]	-0.020 [0.097]	-0.011 [0.088]
Downstream X 1(Elas > Median), β_2	0.274 [0.252]	0.310*** [0.106]	0.314*** [0.102]	-0.291 [0.246]	0.353*** [0.078]	0.347*** [0.084]
1(Elas > Median)	0.003 [0.231]	-0.268*** [0.075]	-0.268*** [0.075]	0.661*** [0.233]	-0.244*** [0.076]	-0.234*** [0.073]
Seller Log (0.001+ R&D/Sales) X Country Entry Costs				-0.028*** [0.007]		
Seller Log (0.001+ R&D/Sales)				0.026 [0.045]		
Inverse Mills Ratio		-0.120 [0.250]			-0.082 [0.197]	
p-value: Joint significance of β_1 and β_2		[0.0036]	[0.0024]		[0.0000]	[0.0001]
p-value: Test of $\beta_2 - \beta_1 = 0$		[0.0009]	[0.0006]		[0.0007]	[0.0008]
Industry controls for: Country-year fixed effects?	Buyer Yes	Buyer Yes	Buyer Yes	Buyer Yes	Buyer Yes	Buyer Yes
Sample:	Total imports >=0	Total imports > 0	Total imports > 0	Total imports >=0	Total imports > 0	Total imports > 0
Observations	462990	186029	186029	462990	186029	186029
R-squared	-	0.63	0.63	-	0.64	0.64

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. The exclusion restriction variable is the seller industry's R&D intensity, Log (0.001+ R&D/Sales), and its interaction with a dummy variable for countries with above sample-median entry costs. The latter is constructed from the first principal component of the number of procedures, number of days, and monetary cost of starting a business, from the Doing Business dataset. Columns 2 and 5 report the second stage, which is run using weighted least squares; the observation weights are the value of total imports for the country-industry-year in question. Columns 3 and 6 report the second-stage regression excluding the inverse Mills ratio, to provide a comparison. Columns 1-3 use DUse_TUse and Columns 4-6 use DownMeasure as the downstreamness variable respectively. All columns also control for the buyer industry factor intensity and dispersion variables, as well as country-year fixed effects.