

INTERNALIZING GLOBAL VALUE CHAINS: A FIRM-LEVEL ANALYSIS *

Laura Alfaro
Harvard Business School

Pol Antràs
Harvard University

Davin Chor
National University of Singapore

Paola Conconi
Université Libre de Bruxelles (ECARES)

September 2015

Abstract

In recent decades, technological progress in information and communication technology and falling trade barriers have led firms to retain within their boundaries and in their domestic economies only a subset of their production stages. A key decision facing firms worldwide is the extent of control to exert over the different segments of their production processes. Building on Antràs and Chor (2013), we describe a property-rights model of firm boundary choices along the value chain. To assess the evidence, we construct firm-level measures of the upstreamness of integrated and non-integrated inputs by combining information on the production activities of firms operating in more than 100 countries with Input-Output tables. In line with the model's predictions, we find that whether a firm integrates upstream or downstream suppliers depends crucially on the elasticity of demand for its final product. Moreover, a firm's propensity to integrate a given stage of the value chain is shaped by the relative contractibility of the stages located upstream versus downstream from that stage. Our results suggest that contractual frictions play an important role in shaping the integration choices of firms around the world.

JEL classifications: F14, F23, D23, L20.

Keywords: Global value chains, sequential production, incomplete contracts.

*We thank participants at the Global Fragmentation of Production and Trade Policy conference at ECARES, the ERWIT trade conference in Paris, the Barcelona GSE Summer Forum, the Princeton IES Summer Workshop, the NBER ITI meetings, the WEI meetings, the Asia Pacific Trade Seminars, as well as seminar participants at LSE, the Paris Trade Workshop, MIT Sloan, Warwick, University of Ferrara, HKUST, HKU, NUS, and Singapore Management University. We are particularly grateful to Kamran Bilir, Arnaud Costinot, Thibault Fally, Thierry Mayer, Peter Morrow, and David Weinstein for their detailed comments on a previous draft of the paper. Chor thanks colleagues at the Global Production Networks Centre (GPN@NUS) for their engaging discussions. Alfaro: lalfaro@hbs.edu. Antràs: pantras@fas.harvard.edu. Chor: davinchor@nus.edu.sg. Conconi: pconconi@ulb.ac.be.

1 Introduction

Sequential production has been an important feature of modern manufacturing processes at least since Henry Ford introduced his Model T assembly line in 1913. The production of cars, computers, mobile phones and most other manufacturing goods involves a sequencing of stages: raw materials are converted into basic components, which are then combined with other components to produce more complex inputs, before being assembled into final goods.

In recent decades, technological progress in information and communication technology and falling trade barriers have led firms to retain within their boundaries and in their domestic economies only a subset of these production stages. Research and development, design, production of parts, assembly, marketing and branding, previously performed in close proximity, are increasingly fragmented across firms and countries. The semiconductor industry fittingly exemplifies these trends. The first semiconductor chips were manufactured in the United States by vertically integrated firms such as IBM and Texas Instruments. Firms initially kept the design, fabrication, assembly, and testing of integrated circuits within ownership boundaries. The industry has then undergone several reorganization waves in the last fifty years, and many of the production stages are now outsourced to independent contractors in Asia (Brown and Linden, 2005). Another often cited example is the iPhone: while its software and product design are done by Apple, most of its components are produced by independent suppliers around the world (Xing, 2011).

While fragmenting production across firms and countries has become easier, contractual frictions remain a significant obstacle to the globalization of value chains. On top of the inherent difficulties associated with specifying the course of action to be taken by parties in all possible states of nature, international transactions suffer from a disproportionately low level of enforcement of contract clauses and legal remedies (Antràs, 2015). In such an environment, companies are presented with complex organizational decisions. In this paper, we focus on a key decision faced by firms worldwide: the extent of control they choose to exert over the different segments of their production processes.

Although the global fragmentation of production has featured prominently in the trade literature, much less attention has been placed on how the position of a given production stage in the value chain affects firm boundary choices, and firm organizational decisions more broadly. Furthermore, most studies on this topic have been mainly theoretical in nature.¹ To a large extent, this theoretical bias is explained by the challenges one faces when taking models of global value chains to the data. Ideally, researchers would like to access comprehensive datasets that would enable them to track the flow of goods within value chains across borders and organizational forms. Trade statistics are useful in capturing the flows of goods when they cross a particular border, and some countries' customs offices also record whether goods flow in and out of a country within or across firm boundaries. Nevertheless, once a good leaves a country, it is virtually impossible with

¹Recent papers on sequential production include Harms, Lorz, and Urban (2012), Baldwin and Venables (2013), Costinot *et al.* (2013), Antràs and Chor (2013), Kikuchi *et al.* (2014), and Fally and Hillberry (2014). This literature is in turn inspired by earlier contributions in Dixit and Grossman (1982), Sanyal and Jones (1982), Kremer (1993), Yi (2003), and Kohler (2004).

available data sources to trace the subsequent locations (beyond its first immediate destination) in which the good will be combined with other components and services.

A first contribution of this paper is to show how available data on the activities of firms can be combined with information from standard Input-Output tables to study firm boundaries along value chains. A key advantage of approach is that it allows us to study how the integration of stages in a firm’s production process is shaped by the characteristics – in particular, the production line position (or “upstreamness”) – of these different stages. Moreover, the richness of our data allow us to run specifications that exploit variation in organizational features across firms, as well as within firms across their various inputs. Available theoretical frameworks of sequential production are highly stylized and often do not feature asymmetries across production stages other than in their position in the value chain. A second contribution of this paper is to develop a richer framework of firm behavior that can guide our empirical analysis using firm-level data.

On the theoretical side, we extend the property-rights model of the organization of production of Antràs and Chor (2013).² We focus on the problem of a firm controlling the production process of a final-good manufacturing variety, which is associated with a constant price elasticity demand schedule. The production of the final good entails a large number of stages that need to be performed in a predetermined order. The different stage inputs are provided by suppliers, who undertake relationship-specific investments to make their components compatible with those of other suppliers in the value chain. How these supplier investments are transformed into quality-adjusted units of output of the final good is determined by a function that is isomorphic to a constant elasticity of substitution technology, except for the sequential nature of production. The setting is one of incomplete contracting, in the sense that contracts contingent on whether components are compatible or not cannot be enforced by third parties. As a result, the division of surplus between the final-good producer and each supplier is governed by bargaining, after a stage has been completed and the firm has had a chance to inspect the input. The final-good producer must decide which input suppliers (if any) to own along the value chain. As in Grossman and Hart (1986), the integration of suppliers does not change the space of contracts available to the firm and its suppliers, but it affects the relative bargaining power of these agents in their negotiations. A key feature of our model of firm boundaries is that organizational decisions have spillovers along the value chain because relationship-specific investments made by upstream suppliers affect the incentives of suppliers in downstream stages.

Initially, we follow Antràs and Chor (2013) by focusing on a setup in which production stages differ in their position in the value chain, but are symmetric on the technology, contracting and cost sides. This symmetric model delivers the result that a firm’s decision to integrate upstream or downstream suppliers depends crucially on the relative size of the elasticity of demand for its final good and the elasticity of substitution across production stages. When demand is elastic or inputs are not particularly substitutable, inputs are sequential complements, i.e., the marginal

²The property-rights approach builds on the seminal work of Grossman and Hart (1986), and has been fruitfully employed to study the organizational decisions of multinational firms. See Antràs (2015) for a comprehensive overview of this literature.

incentive of a supplier to undertake relationship-specific investments is higher, the larger are the investments by upstream suppliers. In this case, the firm finds it optimal to integrate only the most downstream stages, while contracting at arm's length with upstream suppliers in order to incentivize their investment effort. When instead demand is inelastic or inputs are sufficiently substitutable, inputs are sequential substitutes, i.e., investments by upstream suppliers lower the investment incentives of downstream suppliers. When this is the case, the firm would choose to integrate relatively upstream stages, while engaging in outsourcing to downstream suppliers.

In this paper, we also consider a richer environment, in which we allow for heterogeneity in the marginal productivity of inputs, as well as in the marginal cost of production faced by suppliers at different points in the value chain. Perhaps surprisingly, we show that regardless of how marginal productivity or marginal costs rise or fall along the value chain, the relative propensity to integrate upstream versus downstream suppliers continues to be determined by the same two parameters as described above for our baseline model. The profile of marginal productivities and costs along the value chain does however shape the measure of stages that the firm ends up finding optimal to integrate.

We operationalize this result by mapping differences in marginal input productivity across stages to differences in input contractibility. More specifically, we show that the propensity of a firm to integrate a given stage of its value chain is shaped in subtle ways by the degree of contractibility of upstream and downstream stages. Intuitively, in production processes that feature a high degree of contractibility among upstream relative to downstream inputs, firms need to rely less on the organizational mode to counteract the distortions associated with inefficient investments upstream. As a consequence, high levels of upstream contractibility tend to reduce the set of outsourced stages when inputs are sequential complements, while they tend to reduce the set of integrated stages when inputs are sequential substitutes.

To assess the validity of the model's predictions, we employ the WorldBase dataset of Dun and Bradstreet (D&B), which provides detailed establishment-level information for public and private companies in many countries. For each establishment, the dataset reports a list of up to six production activities. Establishments belonging to the same firm can be linked via information on their global parent using a unique identifier (the DUNS number).³ Our main sample consists of more than 300,000 manufacturing firms in 116 countries.


In our empirical analysis, we study the determinants of a firm's propensity to integrate upstream versus downstream inputs. To distinguish between integrated and non-integrated inputs, we rely on the methodology of Fan and Lang (2000), combining information on firms' reported activities with Input-Output tables (see also Acemoglu *et al.*, 2009; and Alfaro *et al.*, 2013). To capture the position of different inputs along the value chain, we compute a measure of the upstreamness of each input i in the production of output j using U.S. Input-Output Tables. This extends the measure of

³D&B uses the United States Government Department of Commerce, Office of Management and Budget, Standard Industrial Classification Manual 1987 edition to classify business establishments. The Data Universal Numbering System – the D&B DUNS Number – supports the linking of plants and firms across countries and tracking of plants' histories including name changes.

the upstreamness of an industry with respect to final demand from Fally (2012) and Antràs *et al.* (2012) to the bilateral industry-pair level. To provide a test of the model, we exploit information from WorldBase on the primary activity of each firm, and use estimates of demand elasticities from Broda and Weinstein (2006), as well as measures of contractibility from Nunn (2007).

We first examine how firms’ organizational choices depend on the elasticity of demand for their final good. In line with the first prediction of the model, we find that the higher the elasticity of demand faced by the parent firm, the lower the average upstreamness of its integrated inputs relative to the upstreamness of its non-integrated inputs. This result is illustrated in a simple (unconditional) form in Figure 1, based on different quintiles of the parent firm’s elasticity of demand. As seen in the left panel of the figure, the average upstreamness of integrated inputs is much higher when the parent company belongs to an industry with a low demand elasticity than when it belongs to one associated with a high demand elasticity. Conversely, the right panel shows that the average upstreamness of non-integrated stages is greater the higher the elasticity of demand faced by the parent’s final good.⁴

Figure 1: Average Upstreamness of Production Stages, by Quintile of Parent’s Demand Elasticity



The above pattern is robust in the regression analysis, even when controlling for a comprehensive list of parent firm characteristics (e.g., size, age, employment, sales), using different measures of the demand elasticity, as well as in different samples of firms (e.g., restricting to single-plant firms or multinationals). We also show that our results hold in specifications where the elasticity of demand is replaced by the difference between this same elasticity and a proxy for the degree of input substitutability associated with the firm’s production process. We reach a similar conclusion

⁴To be specific, Figure 1 is plotted using only inputs i that rank within the top 100 manufacturing inputs in terms of total requirements coefficients of the parent’s output industry j . The averages are computed weighting each input by its total requirement coefficient tr_{ij} , while excluding integrated stages belonging to the same industry j as the parent. The figures obtained when considering all manufacturing inputs, when computing unweighted averages, and when considering the output industry j as an input are all qualitatively similar.

when we exploit within-firm variation in integration patterns. In these specifications, we find that a firm’s propensity to integrate is generally larger for downstream inputs, but disproportionately so for firms facing high demand elasticities.

After documenting this fact from a variety of angles, we turn to study how firms’ ownership decisions are shaped by the degree of contractibility of upstream versus downstream inputs. We find that a greater degree of contractibility of upstream inputs increases the likelihood that a firm integrates upstream inputs, when the firm faces a high elasticity of demand (both in absolute terms, as well as relative to our proxy for input substitutability). This result is in line with the second prediction of our theoretical model, according to which greater upstream contractibility reduces a firm’s need to rely on decisions over organizational mode to elicit the right incentives from suppliers positioned at early stages in the value chain. When the elasticity of demand is high, this reduces the propensity to outsource upstream, leading the firm to integrate more upstream inputs.

By conducting our analysis at the firm level, we are able to greatly improve upon the empirical evidence provided in Antràs and Chor (2013), which was based on industry-level data on U.S. intrafirm import shares and lacked direct information on the U.S. entity internalizing these foreign purchases. Moreover, extending the theoretical model to allow for asymmetries across inputs generates a richer set of predictions about firms’ boundary choices that we can bring to the data. Our work is closely related to two contemporaneous papers with similar goals. Del Prete and Rungi (2015) employ a dataset of about 4,000 multinational business groups to explore the correlation between the average “downstreamness” of integrated affiliates (relative to final demand) and that of the parent firm itself (also relative to final demand). They find that this correlation varies depending on the size of the demand elasticity faced by the parent firm, in a manner reminiscent of the predictions in Antràs and Chor (2013). Their work is however silent on the production line position of non-integrated inputs and does not incorporate an industry-pair measure of the upstreamness of affiliates relative to their parents. Luck (2014) reports corroborating evidence based on city-level evidence on the export-import activities of processing firms in China, though his work adopts a value-added notion of production line position (rather than one rooted in actual production staging). As insightful as these contributions are, we view the empirical strategy developed in this paper as a more direct firm-level test of the propositions of the theory. More generally, our paper is related to a recent empirical literature testing various aspects of the property-rights theory of multinational firm boundaries. This includes Yeaple (2006), Nunn and Trefler (2008, 2013), Corcos *et al.* (2013), Defever and Toubal (2013), Díez (2014), and Antràs (2015), among others.⁵

Our analysis suggests that contractual frictions play a crucial role in shaping the integration choices of firms around the world. A recent influential study by Atalay *et al.* (2014) finds little evidence of intrafirm shipments between related plants within the United States. The authors interpret their findings as suggesting that firm boundaries are likely to be shaped by the transfer of intangible inputs, rather than by the transfer of physical goods. Without denying the importance of

⁵Even more broadly, our work is related to the extensive empirical literature on firm boundaries, which is nicely overviewed in Lafontaine and Slade (2007) and Bresnahan and Levin (2012).

intangible inputs in the ownership decisions of firms, we believe that contractual frictions are likely to be more relevant for international production networks than for domestic ones. For instance, the fact that intrafirm imports account for close to 50 percent of U.S. merchandise imports is *prima facie* evidence that U.S. firms choose to internalize a very significant share of their transactions involving physical goods. It is important to stress, however, that intrafirm trade flows (and U.S. intrafirm trade flows in particular) are an imperfect proxy for the extent to which firms react to contractual insecurity by internalizing particular stages of their global value chains. The reason for this is that, at least in our model, firms do not choose to own suppliers to facilitate their vertical trade with them, but rather to better discipline their behavior. In fact, we argue that a simple extension of our model is consistent with an arbitrarily low level of intrafirm trade in the data relative to the overall transaction volume in the value chain. We believe that this helps reconcile our findings with those of Ramondo *et al.* (2015), who find that intrafirm trade between U.S. multinationals and their affiliates abroad is highly concentrated among a small number of large affiliates.

The remainder of the paper is organized as follows. Section 2 presents our model of firm boundaries with sequential production and input asymmetries. Section 3 describes our data. Section 4 outlines our empirical methodology and presents our findings in detail. Section 5 concludes.

2 Theoretical Framework

In this section, we develop our model of sequential production. We begin with a stripped-down version of the model, along the lines of Antràs and Chor (2013), in which all production stages are symmetric except for their position in the value chain. Later, we will introduce technological and contractual asymmetries across stages to enrich the set of empirical predictions emanating from the model.

2.1 Symmetric Model

We focus throughout on the problem of a firm seeking to optimally organize a manufacturing process that culminates in the production of a finished good valued by consumers. The final good is differentiated in the eyes of consumers and belongs to a monopolistically competitive industry with a continuum of active firms, each producing a differentiated variety. Consumer preferences over the industry's varieties feature a constant elasticity of substitution so that the demand faced by the firm in question can be represented by

$$q = Ap^{-1/(1-\rho)}, \tag{1}$$

where $A > 0$ is a term that the firm takes as given, and the parameter $\rho \in (0, 1)$ is positively related to the degree of substitutability across final-good varieties. The parameter A is allowed to vary across firms in the industry – perhaps reflecting differences in quality across firms – while the demand elasticity $1/(1 - \rho)$ is common for all firms in the sector. The latter assumption is

immaterial for our theoretical results, but will be exploited in the empirical implementation of the model, where we rely on sectoral estimates of demand elasticities. Given that we focus on the problem of a representative firm, we abstain from indexing variables by firm or sector to keep the notation tidy.

Obtaining the finished product requires the completion of a unit measure of production stages. These stages are indexed by $i \in [0, 1]$, with a larger i corresponding to stages further downstream and thus closer to the finished product. Denoting by $x(i)$ the value of the services of intermediate inputs that the supplier of stage i delivers to the firm, final-good production is then given by:

$$q = \left(\int_0^1 x(i)^\alpha I(i) di \right)^{1/\alpha}, \quad (2)$$

where $\alpha \in (0, 1)$ is a parameter that captures the (symmetric) degree of substitutability among the stage inputs, and $I(i)$ is an indicator function that takes a value of 1 if input i is produced after all inputs $i' < i$ have been produced, and a value of 0 otherwise. The technology in (2) resembles a conventional symmetric CES production function with a continuum of inputs, but the indicator function $I(i)$ makes the production technology inherently sequential.⁶


Intermediate inputs are produced by a unit measure of suppliers, with the mapping between inputs and suppliers being one-to-one. Inputs are customized to make them compatible with the needs of the firm controlling the finished product. In order to provide a compatible input, each supplier must undertake a relationship-specific investment entailing a marginal cost of c per unit of input services $x(i)$. All agents including the firm are capable of producing *subpar* inputs at a negligible marginal cost, but these inputs add no value to final-good production apart from allowing the continuation of the production process in situations in which a supplier threatens not to deliver his or her input to the firm.

In situations in which the firm could discipline the behavior of suppliers via a comprehensive ex-ante contract, those threats would be irrelevant. For instance, the firm could demand the delivery of a given volume $x(i)$ of input services in exchange for a fee, while including a clause in the contract that would punish the supplier severely when failing to honor such a contractual obligation. 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. For the time being, we will make the stark assumption that none of the aspects of input production can be specified in a binding manner in an initial contract, except for a clause stipulating whether the different suppliers are vertically integrated into the firm or remain independent.

Because the terms of exchange between the firm and the suppliers are not set in stone before production takes place, the actual payment to a particular supplier (say the one controlling stage i) is negotiated bilaterally only after the stage i input has been produced and the firm has had

⁶In fact, Antràs and Chor (2013) show that equation (2) can alternatively be expressed recursively, with value added at each stage i being a Cobb-Douglas function of the volume of production $q(i)$ generated up to that stage and stage- i 's input services $x(i)$.

Figure 2: Timing of Events



a chance to inspect it. At that point, the firm and the supplier negotiate over the division of the incremental contribution to total revenue generated by supplier i . Notice that the lack of an enforceable contract implies that suppliers are free to choose the volume of input services $x(i)$ to maximize their profits conditional on the value of the semi-finished product they are handed by their immediate upstream supplier.

How does integration affect the game played between the firm and the unit measure of suppliers? 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. Under integration, the firm controls the physical assets used in the production of the intermediate input, thus allowing the firm to dictate a use of these assets that tilts the division of surplus in its favor. We capture this central insight of the property-rights theory in a stark manner with the firm obtaining a share β_V of the value of supplier i 's incremental contribution to total revenue when the supplier is integrated, while receiving only a share $\beta_O < \beta_V$ of that surplus when the supplier is a stand-alone entity.

This concludes the description of the assumptions of the model. Figure 2 outlines the timing of events of the game played by the firm and the unit measure of suppliers. Later, we will supplement the analysis by introducing several sources of asymmetry across production stages, most notably incorporating heterogeneity in the degree of contractibility of inputs along the value chain.

Antràs and Chor (2013) describe the subgame equilibrium of the above game in some detail, so we can focus here on outlining its key features. We refer the reader to their paper for detailed derivations, but it should be emphasized that in the next section we will carefully characterize the equilibrium of a richer production environment that encompasses the one considered here.

We begin by noting that if all suppliers provide compatible inputs and the correct technological sequencing of production is followed, equations (1) and (2) imply that the total revenue obtained

by the firm is given by $r(1)$, where the function $r(m)$ is defined by:

$$r(m) = A^{1-\rho} \left[\int_0^m x(i)^\alpha di \right]^{\frac{\rho}{\alpha}}. \quad (3)$$

Because the firm can always unilaterally complete a production stage by producing a subpar input at negligible cost, one can interpret $r(m)$ as the revenue *secured* up to stage m .

Now consider the bargaining between the firm and the supplier at stage m . Because inputs are customized to the needs of the firm, the supplier's outside option at the bargaining stage is 0 and 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.⁷ Applying Leibniz' rule to (3), this is given by

$$r'(m) = \frac{\rho}{\alpha} (A^{1-\rho})^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} x(m)^\alpha. \quad (4)$$

As explained above, in the bargaining, the firm captures a share $\beta(m) \in \{\beta_V, \beta_O\}$ of $r'(m)$, while the supplier obtains the residual share $1 - \beta(m)$. It then follows that the choice of input volume $x(m)$ is characterized by the program

$$x^*(m) = \arg \max_{x(m)} \left\{ (1 - \beta(m)) \frac{\rho}{\alpha} (A^{1-\rho})^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} x(m)^\alpha - cx(m) \right\}. \quad (5)$$

Notice that the marginal return to investing in $x(m)$ is increasing in the demand level A , while it decreases in the marginal cost c . Furthermore, this marginal return is increasing in supplier m 's bargaining share $1 - \beta(m)$, and thus, other things equal, outsourcing provides higher-powered incentives for the supplier to invest. This is a standard feature of property-rights models. The more novel property of program (5) is that a supplier's marginal return to invest at stage m is shaped by all investment decisions in prior stages, i.e., $\{x(i)\}_{i=0}^m$, as captured by the value of production secured up to stage m , i.e., $r(m)$. The nature of such dependence is in turn crucially shaped by the relative size of the demand elasticity parameter ρ and the input substitutability parameter α . When $\rho > \alpha$, investment choices are *sequential complements* in the sense that higher investment levels by upstream suppliers increase the marginal return of supplier m 's own investment. Conversely, when $\rho < \alpha$, investment choices are *sequential substitutes* because high values of upstream investments reduce the marginal return to investing in $x(m)$. As in Antràs and Chor (2013), we shall refer to $\rho > \alpha$ as the *complements* case and to $\rho < \alpha$ as the *substitutes* case.

It is intuitively clear why low values of α will tend to render investments sequential complements. Why might a low value of ρ render investments sequential substitutes? The reason for this is that when ρ is low, the firm's revenue function is highly concave in output and thus marginal revenue falls at a relatively fast rate along the value chain. As a result, the incremental contribution to revenue associated with supplier m – which is what the firm and supplier m bargain over – might

⁷Antràs and Chor (2013) provide an extensive discussion of the robustness of the key results under an alternative bargaining protocol, which allows each supplier m to lay claim over part of the revenues that are realized downstream of m .

be particularly low when upstream suppliers have invested large amounts.

Plugging the first-order condition of program (5) into (4), and solving the resulting separable differential equation, it is possible to express the volume of input $x(m)$ as a function of the whole path of bargaining shares $\{\beta(i)\}_{i \in [0, m]}$ up to stage m (see equation (10) in Antràs and Chor, 2013). With that expression, it is straightforward to see that $x^*(m) > 0$ for all m and thus the firm has every incentive to abide by the proper sequencing of production, hence implying that $I^*(m) = 1$ for all m (consistent with our expressions above).

To complete the description of the equilibrium, we roll back to the initial period prior to any production taking place, in which the firm decides whether the contract associated with a given input m is associated with integration or outsourcing. This amounts to choosing $\{\beta(i)\}_{i \in [0, 1]}$ to maximize $\pi_F = \int_0^1 \beta(i) r'(i) di$, with $r'(m)$ given in equation (4), $x^*(m)$ in equation (5), and $\beta(i) \in \{\beta_V, \beta_O\}$. After several manipulations, the problem of choosing the optimal organizational structure can be reduced to the program:

$$\begin{aligned} \max_{\beta(i)} \quad & \pi_F = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{\rho}{c} \right)^{\frac{\rho}{1-\rho}} \int_0^1 \beta(i) (1-\beta(i))^{\frac{\alpha}{1-\alpha}} \left[\int_0^i (1-\beta(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di \\ \text{s.t.} \quad & \beta(i) \in \{\beta_V, \beta_O\}. \end{aligned} \quad (6)$$

As mentioned above, we will not provide a detailed derivation of the solution of this problem here because in the next section we will work through a more general problem that encompasses the one laid out in (6). Suffice to say, at this point, that Antràs and Chor (2013) derive the following result:

Proposition 1 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.

Figure 3: Firm Boundary Choices along the Value Chain




Figure 3 illustrates the key result of Proposition 1. The optimal pattern of ownership along the value chain depends critically on whether the stage inputs are sequential complements or substitutes. When the demand faced by the final-good producer is sufficiently elastic, then there exists a unique cutoff stage such that all inputs prior to that cutoff are outsourced, and all inputs (if any) downstream of it are integrated. The reason for this is that the firm would like to incentivize upstream suppliers to raise their investment effort, in order to generate positive spillovers on the investment levels of downstream suppliers. The converse prediction holds when demand is sufficiently inelastic (i.e., in the sequential substitutes case): the firm would instead integrate relatively upstream inputs, while outsourcing would take place relatively downstream.

2.2 General Framework with Stage Asymmetries

With the goal of isolating the role of downstreamness in shaping the integration decision, the benchmark model developed above assumed that all inputs enter symmetrically into production and entail a common marginal cost c . Furthermore, the degree of contractibility was also assumed to be symmetric across inputs, simply because it was posited that no aspect of input production could be specified in an enforceable manner in the initial contract.

In this section, we consider an environment in which noncontractible investments related to input production have different effects on output at different stages in the value chain, and in which the marginal cost faced by suppliers may also vary along the value chain. For now, we will take these layers of heterogeneity as exogenously given, but in the next section, we will relate these asymmetries to the ex-ante choices of the firm related to certain contractible aspects of production.

Formally, we consider an environment identical to the one in our benchmark model, except for two features. First, the technology for final-good production is now given by

$$q = \left(\int_0^1 (\psi(i) x(i))^\alpha I(i) di \right)^{1/\alpha}, \quad (7)$$

where $\psi(i)$ captures asymmetries in the marginal product of different inputs' investments. Second, we allow the marginal cost of production of input i , that is $c(i)$, to vary across inputs. For now, we will not place any constraints on the functions $\psi(i)$ and $c(i)$ other than they be non-negative, real-valued and piecewise continuously differentiable.

The inclusion of these sources of heterogeneity has a minor effect on the revenue and marginal contribution functions (3) and (4), which now include extra terms in $\psi(i)^\alpha$. The program related to the optimal investment choice of the supplier at stage m is similarly modified to:

$$x^*(m) = \arg \max_{x(m)} \left\{ (1 - \beta(m)) \frac{\rho}{\alpha} (A^{1-\rho})^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} \psi(m)^\alpha x(m)^\alpha - c(m) x(m) \right\}. \quad (8)$$

The first-order condition associated with (8) relates investment $x(m)$ at stage m with the bargaining share $\beta(m)$, as well as the revenue $r(m)$ secured by suppliers upstream from m . When plugging this expression for $x(m)$ into $r'(m) = \frac{\rho}{\alpha} (A^{1-\rho})^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} \psi(m)^\alpha x(m)^\alpha$, we show in Appendix A-1

that one obtains a separable differential equation for $r(m)$ with solution:

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

With this expression in hand, one can then revert back to the first-order condition of program (8) and solve for the whole path of supplier investments $x(m)$ for all $m \in [0, 1]$ as a function of the path of bargaining shares.

Rolling back to the initial period of the game, we finally tackle the problem of the firm of deciding on the optimal ownership structure – i.e., the path of $\beta(i)$ – along the value chain. As shown in Appendix A-1, this reduces to a simple generalization of program (6):

$$\begin{aligned} \max_{\beta(i)} \quad & \pi_F = \Theta \int_0^1 \beta(i) \left(\frac{(1-\beta(i))\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di, \\ \text{s.t.} \quad & \beta(i) \in \{\beta_V, \beta_O\}. \end{aligned} \quad (10)$$

where $\Theta = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}} > 0$.

Antràs and Chor (2013) derived the same program (10) and showed that the statements in Proposition 1 continue to hold even in the presence of heterogeneity in $\psi(i)$ and $c(i)$ across production stages. Although we will confirm their findings below, our goal in this section is distinct. Instead, we seek to study the particular way in which heterogeneity in the parameters $\psi(i)$ and $c(i)$ affects the optimal ownership structure along the value chain. In particular, we want to examine how the optimal thresholds m_C^* and m_S^* in Proposition 1 are shaped by the paths of input marginal productivity $\psi(i)$ and marginal cost $c(i)$.

With that goal in mind, consider first a relaxed version of program (10) in which rather than constraining $\beta(i)$ to equal β_V or β_O , we allow the firm to freely choose the function $\beta(i)$ from the whole set of piecewise continuously differentiable real-valued functions. Defining

$$v(i) \equiv \int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk, \quad (11)$$

we can then turn this relaxed program (10) into a calculus of variation problem where the firm chooses the real-value function v that maximizes:

$$\pi_F(v) = \Theta \int_0^1 \left(1 - v'(i)^{\frac{1-\alpha}{\alpha}} \frac{c(i)}{\psi(i)} \right) v'(i) v(i)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di. \quad (12)$$

In Appendix A-1, we show that imposing the necessary Euler-Lagrange and transversality conditions, and after a few cumbersome manipulations, the optimal (unrestricted) division of surplus at

stage m can be expressed as:

$$\beta^*(m) = 1 - \alpha \left[\frac{\int_0^m (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk} \right]^{\frac{\alpha-\rho}{\alpha}}. \quad (13)$$

Notice that because the term inside the square brackets is a monotonically increasing function of m , expression (13) confirms the claim in Antràs and Chor (2013) that whether the optimal division of surplus increases or decreases along the value chain is shaped critically by the relative size of the parameters α and ρ .⁸ Hence, as in our stripped-down model in the previous section, the incentive to integrate suppliers will increase as we move downstream in the value chain in the complements case, but will decrease in the substitutes case, regardless of the path of $\psi(k)$ and $c(k)$.

It is worth pausing to explain why this result is not entirely straightforward. Notice that a disproportionately high value of $\psi(m)$ at a given stage m can be interpreted as that stage being relatively important in the production process. Indeed, in a model with complete contracts, the share of m in the total input purchases of the firm would be a monotonically increasing function of $\psi(m)$. According to one of the canonical results of the property-rights literature (cf., Grossman and Hart, 1986, Antràs, 2003), one would then expect the incentive to outsource such a stage to be particularly large. Intuitively, outsourcing provides higher-powered incentives to suppliers, and minimizing underinvestment inefficiencies is particularly beneficial for inputs that are relatively important in production. In terms of the notation of the model, one might have thus expected the optimal division of surplus $\beta^*(m)$ to be decreasing in stage m 's importance $\psi(m)$. For the same reason, and given that input shares are monotonically decreasing in the marginal cost $c(m)$, one might have also expected the share $\beta^*(m)$ to be increasing in $c(m)$. As intuitive as this reasoning might appear, one would then be led to conclude that if the path of $\psi(m)$ were sufficiently increasing in m – or the path of $c(m)$ were sufficiently decreasing in m – then $\beta^*(m)$ would tend to decrease along the value chain, particularly when the difference between ρ and α is small.

Equation (13) demonstrates, however, that this line of reasoning is flawed. No matter by how little ρ and α differ, the slope of $\beta^*(m)$ is uniquely pinned down by the sign of $\rho - \alpha$, regardless of the paths of $\psi(m)$ and $c(m)$. This is not to say, however, that these paths are irrelevant for the incentive to integrate suppliers along the value chain.⁹ Equation (13) illustrates, in particular, that the incentives to integrate a particular input will be notably shaped by the size of the ratio $\psi(k)/c(k)$ for inputs upstream from input m relative to the average size of this ratio along the whole value chain. In particular, in production processes featuring sequential complementarities, the higher is the value of $\psi(k)/c(k)$ for inputs upstream from m relative to its value for inputs downstream from m , the higher will be the incentive of the firm to integrate stage m .

The intuition behind this result is as follows. Remember that when inputs are sequential

⁸Antràs and Chor (2013) failed to derive this explicit formula for $\beta^*(m)$ and simply noted that $\partial\beta^*(m)/\partial m$ inherited the sign of $\rho - \alpha$ (see, in particular, equation (28) in their paper).

⁹This result bears some resemblance to the classic result in consumption theory that the dynamic utility-maximizing level of consumption at any point in time should be a function of the whole path of income (and, in particular, permanent income) rather than just of income in that period.

complements, the marginal incentive of supplier m to invest will be higher, the higher are the levels of investment by suppliers upstream from m . Furthermore, fixing the ownership structure, these upstream investments will also tend to be relatively large whenever stages m' upstream from m are associated with disproportionately large values of $\psi(m')$ or low values of $c(m')$. In those situations, and due to sequential complementarity, the incentives to invest at stage m will also tend to be disproportionately large, and thus the incentive of the firm to integrate stage m will be reduced relative to a situation in which the ratio $\psi(k)/c(k)$ is common for all stages. Conversely, whenever $\rho < \alpha$, investments are sequential substitutes, and thus high upstream investments related to disproportionately high upstream values of $\psi(m')/c(m')$ for $m' < m$ will instead increase the likelihood that stage m is outsourced.

So far, we have focused on a characterization of the optimal bargaining share $\beta^*(m)$, but the above results can easily be turned into statements analogous to those in Proposition 1. In particular, in Appendix A-1 we show that:

Proposition 2 In the presence of input asymmetries in marginal productivity $\psi(m)$ and marginal cost $c(m)$, there continue to exist thresholds $m_C^* \in (0, 1]$ and $m_S^* \in (0, 1]$ such that, in the complements case, all production stages $m \in [0, m_C^*)$ are outsourced and all stages $m \in [m_C^*, 1]$ are integrated, while in the substitutes case, all production stages $m \in [0, m_S^*)$ are integrated, while all stages $m \in [m_S^*, 1]$ are outsourced. Furthermore, both m_C^* and m_S^* are lower, the higher is the ratio $\psi(m)/c(m)$ for upstream inputs relative to downstream inputs.

Although related to our discussion of equation (13), the last statement of Proposition 2 may appear somewhat vague. In order to more formally illustrate it, consider the sequential complements case. As shown in Appendix A-1, integration and outsourcing coexist along the value chain provided that $\beta_V(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} > \beta_O(1 - \beta_O)^{\frac{\alpha}{1-\alpha}}$.¹⁰ In such a case, the threshold m_C^* is given by

$$\frac{\int_0^{m_C^*} (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di} = \left\{ 1 + \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1 - \frac{\beta_O}{\beta_V}}{1 - \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}. \quad (14)$$

Notice then that the larger is the value of $\psi(k)/c(k)$ in upstream production stages (in the numerator of the left-hand-side ratio) relative to downstream production stages, the lower will the value of m_C^* tend to be, and the larger will the set of integrated stages be.

2.3 Contractibility and Integration Decisions

In the previous section, we have taken the sources of input heterogeneity as exogenously given. In order to develop empirical tests of Proposition 2 – and especially its last statement – it is important to map variation in the ratio $\psi(m)/c(m)$ along the value chain to certain observables. As mentioned in the last section, in the absence of contractual frictions, $\psi(m)/c(m)$ would be

¹⁰When instead $\beta_V(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} < \beta_O(1 - \beta_O)^{\frac{\alpha}{1-\alpha}}$, $m_C^* = 1$ and the firm finds it optimal to outsource *all* stages.

positively related to the relative use of input m in the production of the firm's good and one could presumably use information from Input-Output tables to construct empirical proxies for this ratio. Unfortunately, such a mapping between $\psi(m)/c(m)$ and input m 's share in the total input purchases of firms is blurred by incomplete contracting and sequential production. More specifically, program (8) indicates that leaving aside variation in $\beta(m)$ and $\psi(m)/c(m)$, the volume of input $x(m)$ will tend to increase along the value chain due to the positive complementarity effect of upstream investments. The endogenous choice of $\beta(m)$ as well as the path of $\psi(m)/c(m)$ may of course affect the monotonic path of $x(m)$, but it is clear that nothing guarantees that $c(m)x(m)$ will positively co-move with $\psi(m)/c(m)$ along the value chain.

With that in mind, in this section we explore the link between $\psi(m)/c(m)$ and other characteristics of production. Our emphasis will be on mapping variation in $\psi(m)$ to differences in the degree of contractibility of different stage inputs. In the Appendix, we shall also briefly relate marginal cost $c(m)$ variation along the value chain to the sourcing location decisions of the firm.

Consider first a variant of our model in which the marginal productivity parameter $\psi(m)$ at any stage m is not exogenously given, but instead is related to the services of certain contractible investments at that stage. We continue to assume that $x(m)$ captures the services related to the noncontractible aspects of input production. The volume $x(m)$ cannot be disciplined via an initial contract and is chosen unilaterally by suppliers. Conversely, the features of production encapsulated by $\psi(m)$ can be specified in the initial contract in a way that precludes any deviation from that agreed level. In light of equation (7), our assumptions imply that input production is a symmetric Cobb-Douglas function of contractible and non-contractible aspects of production. To capture differential contractibility along the value chain, we let stages differ in the (legal) costs associated with specifying these contractible aspects of production. More specifically, we denote these contracting costs by $(\psi(m))^\phi/\mu(m)$ per unit of $\psi(m)$. We shall refer to $\mu(m)$ as the level of *contractibility* of stage m .¹¹ The parameter $\phi > 1$ captures the intuitive notion that it becomes increasingly costly to render additional aspects of production contractible. We shall assume that the firm bears the full cost of these contractible investments (perhaps by compensating suppliers for them upfront), but our results would not be affected if the firm bore only a fraction of these costs. To simplify matters, we let the marginal cost $c(m)$ of noncontractible investments be constant along the value chain, i.e., $c(m) = c$ for all m .

In terms of the timing of events summarized in Figure 2, notice that nothing has changed except for the fact that the initial contract also specifies the profit-maximizing choice of $\psi(m)$ along the value chain. Furthermore, once the levels of $\psi(m)$ have been set at stage t_0 , the subgame perfect equilibrium is identical to that in our previous model in which $\psi(m)$ was assumed exogenous. This implies that the firm's optimal ownership structure along the value chain will seek to maximize the program in (10), and the solution of this problem will be characterized by Proposition 2.

¹¹Acemoglu, Antràs and Helpman (2007) also model input production as involving a Cobb-Douglas function of contractible and noncontractible inputs, but they capture the degree of contractibility by the elasticity of input production to the contractible components of production. In our setup with sequential production, however, such an approach precludes an analytical solution of the differential equations characterizing the equilibrium.

As shown in Appendix A-1, after solving for the optimal choice of $\beta(m) \in \{\beta_V, \beta_O\}$, one can express firm profits net of contracting costs as

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} c^{\frac{-\rho}{1-\rho}} \Gamma(\beta_O, \beta_V) \left[\int_0^1 \psi(i)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - \int_0^1 \frac{(\psi(i))^\phi}{\mu(i)} di, \quad (15)$$

where remember that $\Theta = A \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}} > 0$ and where $\Gamma(\beta_O, \beta_V) > 0$ is a function of β_O and β_V , as well as of α and ρ (see Appendix A-1). The choice of the profit-maximizing path of $\psi(m)$ will thus seek to maximize π_F in (15).

A notable feature of equation (15) is that, leaving aside variation in the costs of contracting $\mu(i)$, the marginal incentive to invest in the contractible components of input production is independent of the position of the input in the value chain. This result is not entirely intuitive because, relative to a complete contracting benchmark, the degree of underinvestment in noncontractible inputs varies along the value chain and the endogenous (but coarse) choice of ownership structure does not fully correct these distortions. One might have then imagined (as Antràs and Chor, 2013, incorrectly hypothesized) that the choice of $\psi(i)$ would have partly sought to remedy these remaining inefficiencies. Instead, variation in the firm's choice of contractible investments $\psi(i)$ is solely shaped by variation in contractibility $\mu(i)$. More precisely, the first-order conditions associated with problem (15) imply that for any two inputs at stages m and m' , we have that

$$\frac{\psi(m)}{\psi(m')} = \left(\frac{\mu(m)}{\mu(m')} \right)^{\phi - \frac{\alpha}{1-\alpha}}. \quad (16)$$

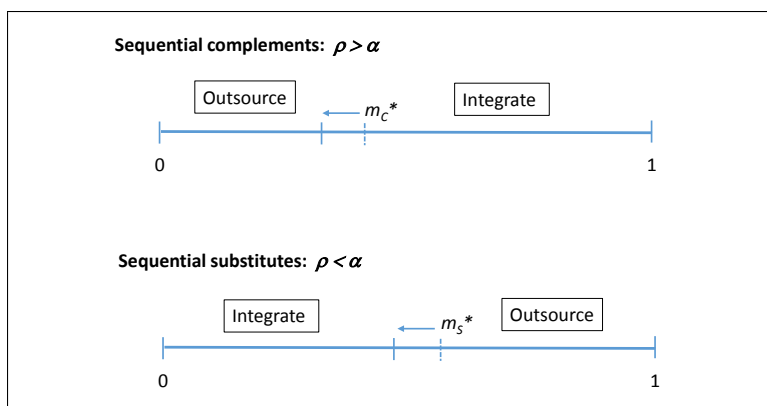
For the second-order conditions of problem (15) to be satisfied, we need to assume that $\phi > \alpha/(1-\alpha)$, and thus the path of $\psi(m)$ along the value chain is inversely related to the path of the exogenous contracting costs $1/\mu(m)$. In light of our discussion in the last section, this implies:

Proposition 3 There exist thresholds $m_C^* \in (0, 1]$ and $m_S^* \in (0, 1]$ such that, in the complements case, all production stages $m \in [0, m_C^*)$ are outsourced and all stages $m \in [m_C^*, 1]$ are integrated, while in the substitutes case, all production stages $m \in [0, m_S^*)$ are integrated, while all stages $m \in [m_S^*, 1]$ are outsourced. Furthermore, both m_C^* and m_S^* are lower, the higher is contractibility $\mu(m)$ for upstream inputs relative to downstream inputs.

The key result of Proposition 3 is illustrated in Figure 4. The intuition for this result is that, the higher the contractibility of the inputs, the less firms need to rely on upstream organizational decisions as a way to counteract the distortions associated with inefficient investments by upstream suppliers. As a consequence, high levels of upstream contractibility tend to reduce the set of outsourced stages whenever final-good demand is elastic or inputs are not too substitutable, while they tend to reduce the set of integrated stages whenever final-good demand is inelastic or inputs are highly substitutable.

By mapping variation in $\psi(m)$ to the degree of input contractibility, Proposition 3 helps oper-

Figure 4: The Effect of an Increase in Upstream Contractibility on Integration Decisions



ationalize our previous, more abstract Proposition 2. More specifically, in our empirical analysis, we will employ empirical proxies for input contractibility to develop a sector-level measure of the extent to which non-contractibilities feature disproportionately in upstream versus downstream stages in the production of that sector’s output. We will then study how firm-level ownership decisions are shaped by this relative importance of upstream versus downstream contractibilities in both the complements and substitutes cases. Before diving into our empirical analysis, however, a brief theoretical digression is necessary.

2.4 Sparse Integration and Intrafirm Trade

Our framework has the strong implication that the sets of integrated and outsourced stages are both connected and jointly constitute a partition of $[0, 1]$. Obviously, this strong prediction of the model is not borne out in the data. In fact, integrated stages are very sparse in our dataset and the overwhelming majority of them *border* with outsourced stages immediately upstream and downstream from them.¹² This sparsity of integration might be due to technological or regulatory factors. Regardless of its source, the goal of this section is to briefly outline an extension that can accommodate such sparsity while maintaining the key predictions of the model that we will take to the data.

Consider then the case in which the firm cannot possibly integrate certain segments of the value chain. Although the analysis below could be generalized (see the Online Appendix), we shall focus on a particular example in which all stages in the sets $O_1 = [0, \omega]$, $O_2 = [\frac{1}{2} - \omega, \frac{1}{2} + \omega]$, and $O_3 = [1 - \omega, 1]$ are necessarily outsourced, with $\omega \in [0, \frac{1}{4}]$. The firm is however still free to choose the optimal ownership structure in the sets $I_1 = (\omega, \frac{1}{2} - \omega)$ and $I_2 = (\frac{1}{2} + \omega, 1 - \omega)$. When $\omega = 0$, we revert to our benchmark model without technological or regulatory constraints,

¹²For our full sample of firms, the median number of integrated stages is 2, while the median number of non-integrated stages – i.e., all inputs with positive total requirements coefficients – is 906. Furthermore, even when restricting the sample to the top 100 manufacturing inputs ranked by the total requirements coefficients of the associated output industry, a mere 0.11 percent of all integrated stages are immediately preceded or succeeded by another integrated stage. In the next section, we will discuss in detail how we identify integrated and non-integrated stages, and their position in the value chain.

but as ω increases, the set of potentially integrated stages becomes increasingly sparse, and when $\omega = \frac{1}{4}$, no stage can possibly be integrated. As we demonstrate in the Online Appendix, in this variant of the model it continues to be the case that the average upstreamness of integrated stages is higher than that of outsourced stages whenever $\rho < \alpha$, with the converse being true for $\rho > \alpha$.¹³ Intuitively, despite the presence of exogenously outsourced stages, it is still the case that the firm cannot possibly find it optimal, when $\rho > \alpha$, to integrate a positive measure of stages located upstream from a positive measure of outsourced stages belonging to $I_1 \cup I_2$. This provides a natural generalization of our Propositions 1 and 2 to situations in which the set of integrated stages can be made arbitrarily sparse. In the limit in which $\omega \rightarrow 1/4$, any positive measure of integrated stages is arbitrarily small and is necessarily surrounded by outsourced stages in the value chain.

An interesting implication of the sparsity of integrated stages in the value chain is that, as ω increases in the example above, the volume of intrafirm trade in the value chain becomes smaller and smaller. Intuitively, in such a case, each interval of integrated stages becomes increasingly isolated, and necessarily trades at arm’s-length with their immediate neighbors in the value chain. This confirms our claim in the Introduction that in sequential production processes in which physical goods flow through both integrated and non-integrated plants, and in which the former are largely outnumbered by the latter, the volume of intrafirm trade flows may be a poor proxy of the extent to which firms’ integration decisions are shaped by contractual incompleteness.

3 Dataset and Key Variables

We turn now to our empirical analysis. We aim to measure the relative propensity of firms to integrate or outsource inputs at different positions in the value chain. For that purpose, we need firm-specific information on input integration and outsourcing, as well as a measure of the “upstreamness” of these various inputs. To assess the validity of our model, we also require proxies for whether a final-good industry falls into the complements or substitutes case, and a measure of input contractibility. In this section, we discuss the dataset that we employ to identify integrated inputs, together with the construction of several key variables.

3.1 The WorldBase Dataset

Our core firm-level dataset is Dun & Bradstreet’s (D&B) WorldBase, which provides comprehensive coverage of public and private companies across more than 100 countries and territories. WorldBase has been used extensively in the literature, in particular to explore research questions related to the organizational practices of firms around the world.¹⁴ Cross-country studies at the firm level

¹³This strong result also relies on the fact that the average upstreamness of the *forced* outsourced stages m in $O_1 \cup O_2 \cup O_3$ is exactly equal to $1/2$. When relaxing this assumption, the following weaker result still applies (see the Online Appendix): the ratio of the average upstreamness of integrated stages relative to the average upstreamness of outsourced stages is lower when $\rho > \alpha$ than when $\rho < \alpha$. This is precisely the type of implication from our model that we take to the data.

¹⁴An early example was Caves’ (1975) analysis of size and diversification patterns between Canadian and U.S. plants. More recent uses include Harrison *et al.* (2004), Acemoglu *et al.* (2009), Alfaro and Charlton (2009), Alfaro

are challenging, as there are few high-quality datasets that are comparable across countries; when such data are available, these tend to be limited to advanced countries. One of the advantages of WorldBase is thus the inclusion of a wide set of countries at different levels of development.¹⁵

Another key advantage is that the unit of observation is the establishment, namely a single physical location where industrial operations or services are performed, or business is conducted. Each establishment in WorldBase is assigned a unique identifier, called a DUNS number. Where applicable, the DUNS number of the global ultimate owner is also reported, which allows us to keep track of ownership linkages within the dataset. In addition, we will make use of the following information from WorldBase, namely: (i) the location (address) of each establishment; (ii) the 4-digit SIC code (1987 version) of its primary industry, and the SIC codes of up to five secondary industries; (iii) the year it was started or in which current ownership took control; and (iv) basic information on employment and sales.

We use the 2004/2005 WorldBase vintage. For our analysis, we focus on firms in the manufacturing sector – i.e., whose primary SIC code lies between 2000 and 3999 – with a minimum total employment of 20. Note that each firm in our sample is either: (i) a single-establishment firm; or (ii) is identified in WorldBase as a “global ultimate”. The former refers to a business entity whose entire activity is in one location, and which does not report ownership links with other firms in WorldBase. The set of integrated SIC activities for a single-establishment plant is thus simply the list of up to six SIC codes associated with it. On the other hand, a global ultimate is defined as a business entity with legal and financial responsibility over another establishment. We link each global ultimate to its identified majority-owned subsidiaries, both in manufacturing and in non-manufacturing, by using the DUNS number of the global ultimate that is reported for establishments. The set of integrated SIC codes for a global ultimate is then the complete list of SIC activities that is performed either in its headquarters or by one of its subsidiaries.

In all, our sample contains 320,254 firms from 116 countries; 259,312 of these are single-establishment firms, while 60,942 are global ultimates. Among the global ultimates, 6,370 observations have subsidiaries in more than one country, and can thus be labeled as multinational firms. Panel A of Table A-1 provides some descriptive statistics for our full sample, as well as for the subset of multinationals. Not surprisingly, multinationals are on average larger in terms of employment, sales and number of integrated SIC codes, as compared to the typical firm in our data. That said, we will show that our core findings concerning the relationship between “upstreamness” and integration patterns are stable when we look separately at the subsamples of single-establishment firms, global ultimates, or even multinationals. Moving forward, we will refer for simplicity to each observation in our sample as a “parent” firm, indexed by p .

et al. (2013), Alfaro and Chen (2014), and Fajgelbaum *et al.* (2015).

¹⁵The data in WorldBase is compiled from a large number of underlying sources, including partner firms, business registers, telephone directory records, company websites, and even self-registration. See Alfaro and Charlton (2009) for a more detailed discussion, and comparisons with other data sources such as the Bureau of Economic Analysis (BEA) data on U.S. multinational activity.

3.2 Key Variables

Integrated and Outsourced Inputs

For each parent, WorldBase provides us with information on the inputs that are integrated within the firm’s ownership boundaries. In order to further identify which inputs are outsourced, we combine the above with information from U.S. Input-Output (I-O) Tables, following the methodology of Fan and Lang (2000). (See also Acemoglu *et al.*, 2009 and Alfaro *et al.*, 2013).

To fix ideas, consider an economy with $N > 1$ industries. In what follows, we refer to *output* industries by j and *input* industries by i . For each industry pair, $1 \leq i, j \leq N$, the I-O Tables report the dollar value of i used directly as an input in the production of \$1 of j , also known as the direct requirements coefficient, dr_{ij} . Denote with D the corresponding square matrix that has dr_{ij} as its (i, j) -th entry. In practice, each input i can be used not just directly, but could also enter further upstream, i.e., more than one stage prior to the actual production of j . The total dollar value of i used either directly or indirectly to produce \$1 of j is called the total requirements coefficient, tr_{ij} , and this reflects the overall importance of the input for the production of j . As is well known, tr_{ij} is given by the (i, j) -th entry of $[I - D]^{-1}D$, where I is the identity matrix and $[I - D]^{-1}$ is the Leontief inverse matrix.

In our baseline analysis, we designate the primary SIC code reported in WorldBase for each parent p as its output industry j . We first use the I-O Tables to deduce the set of 4-digit SIC inputs $S(j)$ that are used (directly or indirectly) in the production of j , namely: $S(j) = \{i : tr_{ij} > 0\}$. We then identify which inputs are integrated and which are outsourced as follows. Define $I(p) \subseteq S(j)$ to be the set of integrated inputs of parent p . From WorldBase, this consists of the primary and secondary SIC codes of p and all its subsidiaries (if any), these being inputs that the parent can in principle obtain within its ownership boundaries. We then define the complement set, $NI(p) = S(j) \setminus I(p)$, to be the set of non-integrated SICs for parent p , these being the inputs required in the production of j that have not been identified as integrated in $I(p)$. Note that with this construction, the primary SIC activity j of the parent is automatically classified as an element of $I(p)$, so we will later explore the robustness of our results to dropping this “self-SIC” code. (We will also consider several alternative treatments of what constitutes the output industry j for those parent firms that feature multiple manufacturing SIC codes.)

To implement the above, we turn to the 1992 U.S. Benchmark I-O Tables from the Bureau of Economic Analysis (BEA). The U.S. Tables are one of the few publicly-available I-O accounts that provide a level of industry detail close to the 4-digit SIC codes used in WorldBase, while the 1992 vintage is the most recent year for which the BEA provides a concordance from its 6-digit I-O industry classification to the 1987 SIC system.¹⁶ Readers familiar with these tables will be aware that this concordance is not a one-to-one key. This is not a major problem given that our focus is on parents whose primary output j is in manufacturing, as the key assigns a unique I-O industry to each 4-digit SIC code between 2000 and 3999. The complication arises only with those inputs i

¹⁶This concordance is available from: <http://www.bea.gov/industry/exe/ndn0017.exe>.

whose 6-digit I-O industry code maps to multiple 4-digit SIC codes; when this occurs, we split the total requirements value tr_{ij} equally across the multiple SIC codes that the input i maps to.

As mentioned before and reported in Panel A of Table A-1, firms tend to integrate very few of the inputs necessary to produce their final good. The median number of integrated stages is 2, compared to a median number of non-integrated stages equal to 906. The corresponding 90th, 95th and 99th percentiles of the number of integrated stages are 3, 4, and 6. As discussed below, however, integrated inputs tend to be “bunched” together along the value chain, consistent with our model.

Upstreamness

We make further use of the information on production linkages contained in I-O Tables, to obtain a measure of the “upstreamness” of an input i in the production of output j . To capture this, we build on the methodology in Fally (2012) and Antràs *et al.* (2012), and define the following:

$$upst_{ij} = \frac{dr_{ij} + 2 \sum_{k=1}^N dr_{ik} dr_{kj} + 3 \sum_{k=1}^N \sum_{l=1}^N dr_{ik} dr_{kl} dr_{lj} + \dots}{dr_{ij} + \sum_{k=1}^N dr_{ik} dr_{kj} + \sum_{k=1}^N \sum_{l=1}^N dr_{ik} dr_{kl} dr_{lj} + \dots}. \quad (17)$$

Recall that dr_{ij} is the value of i that enters directly in (i.e., exactly one stage prior to) the production of j , that $\sum_{k=1}^N dr_{ik} dr_{kj}$ is the value of i that enters two stages prior to production of j , and so on and so forth. The denominator in (17) is therefore equal to tr_{ij} , written as an infinite sum over the value of i 's use that enters exactly n stages removed from the production of j (where $n = 1, 2, \dots, \infty$). The numerator is similarly an infinite sum, but there each input use term is multiplied by an integer equal to the number of stages upstream at which the input value enters the production process. Looking then at (17), $upst_{ij}$ is a weighted average of the number of stages it takes for i to enter in j 's production, where the weights correspond to the share of tr_{ij} that enters at that corresponding upstream stage. In particular, a larger $upst_{ij}$ means that a greater share of the total input use value of i is accrued further upstream in the production process for j . We thus refer to $upst_{ij}$ simply as the “upstreamness” of i in the production of j .

Note that $upst_{ij} \geq 1$ by construction, with equality if and only if $tr_{ij} = dr_{ij}$, namely when the entirety of the input use of i goes directly into the production of j via one stage. With some matrix algebra, one can see that the numerator of (17) is equal to the (i, j) -th entry of $[I - D]^{-2}D$. Together with the formula for tr_{ij} noted earlier (i.e., the (i, j) -th entry of $[I - D]^{-1}D$), one can then calculate $upst_{ij}$ when provided with the direct requirements matrix, D .

Two additional remarks are in order. First, we should stress the distinction between $upst_{ij}$ and the upstreamness measure put forward previously in Fally (2012) and Antràs *et al.* (2012). The measure in this earlier work served to capture the average production line position of each industry i with respect to final demand (i.e., consumption and investment), whereas our current $upst_{ij}$ instead reflects the position of input i with respect to output industry j . This is therefore a measure of production staging specific to each input-output industry pair, which we can directly map to the firm-level observations in our dataset to assess the validity of the model's predictions.

Second, $upst_{ij}$ also has the interpretation of an “average propagation length”, a concept introduced in Dietzenbacher *et al.* (2005) to capture the average number of stages taken by a shock in i to spread to industry j . Dietzenbacher *et al.* (2005) in fact show that this average propagation length has the nice property that it is invariant to whether one adopts a forward or backward linkage perspective when computing the average number of stages between a pair of industries.

We use the 1992 U.S. I-O Tables – specifically, the direct requirements matrix derived from it – to calculate $upst_{ij}$.¹⁷ We first obtain $upst_{ij}$ for each 6-digit I-O industry pair, before mapping these to 4-digit SIC codes. As mentioned earlier, each 4-digit manufacturing SIC code is associated with a single 6-digit I-O code; this means that we can uniquely assign an $upst_{ij}$ value to SIC code pairs where both the input i and output j are in manufacturing. The mapping concerns apply only when we have a non-manufacturing input i which is associated with multiple 6-digit I-O codes. We adopt a variety of approaches in such cases, by taking either: (i) the simple mean of $upst_{ij}$ over constituent I-O codes of the SIC input industry; (ii) the median value; (iii) a random pick; or (iv) the tr_{ij} -weighted average value. Reassuringly, the pairwise correlation of the upstreamness measures obtained under these different treatments is very high (> 0.98), and our regression results will not depend on which specific approach we adopt, so we will focus on the version that uses a simple mean as our baseline. To be clear, what this yields is a measure of the average number of production stages based on the I-O classification system that are traversed between a given pair of SIC industries.


Panel B of Table A-1 presents some basic information on the total requirements and upstreamness variables after the mapping to SIC codes. We can further report that 98.0% of the (i, j) pairs in our WorldBase sample, namely inputs i that are integrated by a parent firm with output industry j , are relevant for production in the sense that $tr_{ij} > 0$.¹⁸ Figure 5 provides an illustration of the variation contained in the $upst_{ij}$ measure, even when focusing on one particular input industry, in this case Tires and Inner Tubes (SIC 3011). Notice that $upst_{ij}$ is indeed smaller for j sectors such as Mobile Homes (2451), Lawn and Garden Equipment (3524), Industrial Trucks and Tractors (3537), Motorcycles, Bicycles, and Parts (3751), and Transportation Equipment, n.e.c. (3799), this being industries that use tires almost exclusively as a direct input. For comparison and to illustrate the difference, Figure 5 also depicts the upstreamness of Tires with respect to final demand (from Antràs *et al.* 2012); this is the horizontal line with value 2.0954.

As mentioned before, firms tend to integrate few inputs. Our upstreamness measure allows us to study the extent to which — though sparse — integrated inputs tend to be “bunched” together. To do so, we focus on firms that report at least two secondary SIC codes (on top of their primary output industry code j) and identify all the relevant manufacturing inputs i (all those for which

¹⁷We apply an open-economy and net-inventories correction to the direct requirements matrix D , before calculating tr_{ij} and $upst_{ij}$. This involves a simple adjustment to each dr_{ij} to take into account input flows across borders, as well as into and out of inventories, on the assumption that these flows occur in proportion to what is observed in domestic input-output transactions; see Antràs *et al.* (2012) for details.

¹⁸85.6% of these pairs actually exceed the median positive tr_{ij} value of 0.000163 (where this median is taken over all input-output pairs in which j is in manufacturing). We obtain similarly high relevance rates when restricting the count to manufacturing inputs only, or if we drop the self-SIC of the parent firm (i.e., pairs where $i = j$).

Figure 5: Upstreamness of Tires (SIC 3011) in the Production of all Other Manufacturing Industries



$tr_{ij} > 0$). Using (17), we then measure the position of these inputs along the value chain, ranking them into quintiles of upstreamness. Table A-2 reports the probability that any two randomly drawn integrated manufacturing input SICs of a given firm belong to any two quintiles of $upst_{ij}$, where j is the SIC output industry of the firm and the quintiles are taken over all SIC manufacturing inputs i . A first noteworthy feature of Table A-2 is that firms appear more likely to integrate inputs in the first quintile of upstreamness than in the other quintiles. Leaving aside this first quintile, we observe however that the probability that a firm integrates an input is significantly higher when it already owns an input in the same quintile, and furthermore these probabilities fall monotonically as one moves to more distant quintiles. These patterns are very much suggestive of the existence of “bunching” along the value chain in the integration decisions of firms.

Ratio-Upstreamness

To test whether the variation *across* parent firms in integration decisions is consistent with our theory, we first explore specifications with a dependent variable that summarizes the extent to which a firm’s integrated inputs tend to be more upstream compared to its non-integrated inputs. For this purpose, we construct the following R_{jp} measure for each parent:

$$R_{jp} = \frac{\sum_{i \in I(p)} \theta_{ijp}^I upst_{ij}}{\sum_{i \in NI(p)} \theta_{ijp}^{NI} upst_{ij}}, \quad (18)$$

where $\theta_{ijp}^I = tr_{ij} / \sum_{i \in I(p)} tr_{ij}$ and $\theta_{ijp}^{NI} = tr_{ij} / \sum_{i \in NI(p)} tr_{ij}$. This takes the ratio of a weighted-average upstreamness of p 's integrated inputs relative to that of its non-integrated inputs; the weights here are proportional to the total requirements coefficients to capture the relative importance of each input in the production of j . Thus, by design, R_{jp} increases the greater is the propensity of p to integrate relatively upstream inputs, while outsourcing its more downstream inputs. For convenience, we will refer to R_{jp} simply as the “ratio-upstreamness” of parent p . We will later consider several variants of R_{jp} to assess the robustness of our results under different constructions. These include: (i) restricting $S(j)$ to the set of “ever-integrated” inputs, namely inputs i for which we actually observe at least one parent in industry j that integrates i within firm boundaries; (ii) restricting $S(j)$ to the set of manufacturing inputs; and (iii) excluding the self-SIC from $S(j)$.

Panel C of Table A-1 presents summary statistics for the different “ratio-upstreamness” measures. Note that R_{jp} tends to take on values smaller than one for the constructions that include the self-SIC of the parent in the set $I(p)$. This is because the upstreamness of j 's use of itself as an input ($upst_{jj}$) tends to be relatively small in value, and this acts to lower the numerator of R_{jp} . When we drop the self-SIC, this results in a R_{jp} measure with a median value closer to 1. The pairwise correlation between the different versions of R_{jp} is high (typically > 0.8), except when the self-SIC is excluded, in which case the correlation with the baseline measure drops to about 0.15.

Our first set of regression specifications will use “ratio-upstreamness” as the dependent variable, and thus seek to exploit the variation *across firms* in this measure. Our theory has predictions at the input level as well, so we will also present evidence based on variation *within firms* in integration decisions across inputs. For this second set of specifications, we will adopt as the dependent variable a 0-1 indicator for whether the input in question is integrated within the parent's ownership structure, i.e., whether $i \in I(p)$.

Our dataset does not allow us to directly observe whether plants that are related in an ownership sense actually contribute inputs and components to a common production process. It is important to stress that any potential misclassification of integrated versus non-integrated inputs (in the sets $I(p)$ and $NI(p)$) would give rise to measurement error in the dependent variable in our regressions. To the extent that this is classical measurement error, it would make our coefficient estimates less precise, making it harder to find empirical support for the model's predictions.

Demand Elasticity

As highlighted in our theory, the incentives to integrate upstream or downstream suppliers are crucially affected by whether the elasticity of demand faced by the firm (ρ_j) is higher or lower than the elasticity of technological substitution across its inputs (α_j). For practical reasons, we focus on variation in the former in most of the regressions, since detailed estimates of the demand elasticity are available from standard sources. More specifically, to capture ρ_j , we use the U.S. import demand elasticities from Broda and Weinstein (2006). The original estimates are for detailed HS10 products, and we aggregate these to the SIC industry level using U.S. import trade values as weights

(see Appendix A-2 for further details). Although this measure of ρ_j is an average of product-level rather than firm-level elasticities, the HS10 codes are highly disaggregated and should thus provide a reasonable approximation to the firm-level elasticity in the model. We will also pursue several refinements of ρ_j , by using only elasticities for those HS10 codes deemed as consumption and capital goods in the United Nations' Classification by Broad Economic Categories (BEC). (The omitted category is goods classified as intermediates.) As the model arguably applies more to final-goods industries, a demand elasticity constructed based on such products should in principle yield a cleaner proxy for ρ_j . Note that when refining the construction in this manner, about half of the 459 SIC manufacturing industries are dropped, namely those industries composed entirely of intermediate goods.

The UN BEC classification also provides a basis for constructing a proxy for α_j . From the model, α_j is closely related to the elasticity of demand for each intermediate input by firms in industry j . We therefore begin by computing the average demand elasticity for each 4-digit SIC code using now only those HS10 elasticities that correspond to products classified as *intermediates*, in an analogous fashion to the construction of the ρ_j refinements above. We proceed to construct our proxy for α_j by taking a weighted average of the intermediate-good demand elasticity across inputs i used in j 's production, with weights proportional to the total requirements coefficients, tr_{ij} . We then use the value of $\rho_j - \alpha_j$ for the industry in question to distinguish between the complements and substitutes cases. Nevertheless, because we find our estimates of ρ_j to be more reliable than those of α_j , in our benchmark regressions we will associate the sequential complements case with high values of ρ_j and the substitutes case with low values of ρ_j . This approach is valid insofar as the demand elasticity and input substitutability parameters are relatively uncorrelated across industries.¹⁹

Input Contractibility

The model further predicts that patterns of integration will depend on the extent to which contractible inputs tend to be “front-loaded” or located in the early stages of the production process. We therefore construct an “upstream contractibility” variable, $UpstCont_j$, to reflect the tendency for high-contractibility inputs to enter the production of output j at relatively upstream stages, for use in the cross-firm regressions.

We start by following Nunn (2007) in constructing a measure of input contractibility for each SIC industry. The basis for this measure is the Rauch (1999) classification of products into whether they are: (i) homogeneous; (ii) reference-priced; or (iii) differentiated in nature. The “contract-intensity” of an industry is then the share of the constituent HS product codes in the composition of the industry’s input use that is classified as differentiated (i.e., neither homogeneous nor reference-priced), on the premise that it is inherently more difficult to specify and enforce the terms of contractual agreements for such products. As our interest is in the converse concept of contractibil-

¹⁹Indeed, the pairwise correlation between the constructed proxy for α_j and the measures of ρ_j (both the baseline measure and its refinements) is low, ranging between -0.026 and 0.083 .

ity, we use instead one minus the Nunn measure of contract-intensity.²⁰ Denote this metric of input contractibility for industry i by $cont_i$. Then, for each output industry j , we calculate $UpstCont_j$ as a weighted covariance between the upstreamness and the contractibility of its manufacturing inputs, namely:

$$UpstCont_j = \sum_{i \in S^m(j)} \theta_{ij}^m (upst_{ij} - \overline{upst}_{ij}) (cont_i - \overline{cont}_i), \quad (19)$$

where $S^m(j)$ is the set of all manufacturing inputs used in the production of j (i.e., with $tr_{ij} > 0$). The weights are given explicitly by: $\theta_{ij}^m = tr_{ij} / \sum_{k \in S^m(j)} tr_{kj}$, while $\overline{upst}_{ij} = \sum_{i \in S^m(j)} \theta_{ij}^m upst_{ij}$ and $\overline{cont}_i = \sum_{i \in S^m(j)} \theta_{ij}^m cont_i$ are total requirements weighted averages of the upstreamness and contractibility variables respectively. Therefore, if high-contractibility inputs tend to be located at earlier production stages, this will lead to a larger (more positive) covariance and hence a higher value of $UpstCont_j$; we will refer to such an industry as exhibiting a higher degree of “upstream contractibility”.²¹

In the within-firm regressions, we can perform a more detailed test of the role of contractibility in explaining the propensity to integrate particular inputs. Motivated by the theory, we construct the variable $ContUpT oi_{ij}$, which is an input-output industry pair-specific measure of the “contractibility up to i in the production of j ”. This is computed as:

$$ContUpT oi_{ij} = \frac{\sum_{k \in S_i^m(j)} tr_{kj} cont_k}{\sum_{k \in S^m(j)} tr_{kj} cont_k}, \quad (20)$$

where the relevant set of inputs, $S_i^m(j)$, for the sum in the numerator is those manufacturing inputs that are located upstream of and including i itself in the production of j , i.e., $S_i^m(j) = \{k \in S^m(j) : upst_{kj} \geq upst_{ij}\}$. The denominator thus sums up the product of the total requirements coefficients and contractibility values across all manufacturing inputs used in the production of j , with the numerator being the partial sum excluding all inputs downstream of i . The construction of (20) is intended to mimic the $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ terms, which appear in equation (13) of the extension of the model with stage asymmetries. There, it was shown that “contractibility up to i ” plays a central role in the expression for the optimal β^* , and hence the propensity towards integration of each input in the production of output j .

²⁰In Nunn’s (2007) notation, the measure of input contractibility that we use is equal to $1 - z^{rs1}$. The results reported are based on the “conservative” Rauch (1999) classification, but are robust when using the alternative “liberal” classification instead.

²¹We have also experimented with alternative measures of $UpstCont_j$, by taking a ratio of the tr_{ij} -weighted upstreamness of inputs classified as being of high contractibility relative to those classified as low contractibility, in a manner analogous to the construction of the ratio-upstreamness measure in (18). To distinguish high- versus low-contractibility inputs, we have adopted either the first tercile, median, or second tercile values of $cont_i$ across the 459 SIC manufacturing industries as cutoffs. The results with these alternative versions of $UpstCont_j$ all continue to lend strong support to the model (results available on request).

4 Empirical Methodology and Results

We translate the propositions from the theoretical model into a series of empirical predictions that we then take to the data. According to Proposition 1, integration patterns along the value chain should vary systematically for industries that fall under the sequential complements versus substitutes cases. As mentioned, we pursue two approaches to distinguish between these cases in the data. The baseline approach focuses on variation in ρ_j , although we will also report specifications that use the difference between ρ_j and a proxy for α_j . However, the limitations inherent in how we construct ρ_j and α_j mean that we cannot use them to precisely delineate where the cutoff between the complements and substitutes cases lies. Consequently, what we test in our regressions is a milder version of Proposition 1, that examines whether the propensity to integrate upstream stages falls as ρ_j (respectively, $\rho_j - \alpha_j$) increases, and we move towards the complements case. We thus formulate the first cross-firm prediction of our model as follows:

P.1 (Cross): A firm’s propensity to integrate upstream versus downstream inputs should fall with ρ_j (alternatively, $\rho_j - \alpha_j$), where j is the final-good industry of the firm.

Our data also affords us the opportunity to explore integration decisions made across different inputs at the firm level, through specifications in which the unit of observation is a parent firm by input SIC pair. In this within-firm setting, we can restate this first prediction as:

P.1 (Within): The upstreamness of an input should have a lower effect (i.e., less positive, or more negative) on the propensity of a firm to integrate that input, the larger is ρ_j (alternatively, $\rho_j - \alpha_j$), where j is the final-good industry of the firm.

The extension of the model developed in Section 2.3 provides us with further predictions that emerge from considering heterogeneity in the contractibility of inputs. In particular, Proposition 3 suggests that the relative propensity to integrate upstream (as opposed to downstream) inputs depends on the extent to which contractible inputs tend to be located in the early stages of production. Moreover, the effect of “upstream contractibility” varies subtly across the sequential complements and substitutes cases. The second cross-firm prediction of our model can be summarized as:

P.2 (Cross): A greater degree of contractibility of upstream inputs should decrease a firm’s propensity to integrate upstream (as opposed to downstream) inputs when the firm is in a final-good industry with low ρ_j (alternatively, $\rho_j - \alpha_j$). Conversely, it should increase that propensity when the firm is in a final-good industry with a high ρ_j (alternatively, $\rho_j - \alpha_j$).

The corresponding prediction at the firm-input pair level can be stated as:

P.2 (Within): A greater degree of contractibility of inputs upstream of a given input (relative to the inputs downstream of it) should decrease the propensity of a firm to integrate that input when the firm faces a low ρ_j (alternatively, $\rho_j - \alpha_j$). Conversely, it should increase the propensity to integrate that input when the firm faces a high ρ_j (alternatively, $\rho_j - \alpha_j$).

4.1 Cross-firm Results

We first exploit variation in integration choices across firms to assess the validity of our model’s predictions. To examine prediction P.1 (Cross), we estimate the following regression:

$$\log R_{jpc} = \beta_0 + \beta_1 \mathbf{1}(\rho_j > \rho_{med}) + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \quad (21)$$

The dependent variable is the log ratio-upstreamness measure, defined in equation (18), which captures the propensity of each parent p with primary SIC industry j to integrate relatively upstream inputs. Note that the subscript c is introduced to index the country where the parent is located, as we will include a full set of country fixed effects, D_c , among the controls. We report standard errors clustered at the level of the SIC output industry j .

The key regressor of interest is the dummy variable $\mathbf{1}(\rho_j > \rho_{med})$, which identifies whether the elasticity of demand ρ_j is above the median value of ρ across industries. This variable is meant to pick out industries that fall under the sequential complements case. Prediction P.1 suggests that β_1 should be negative: As we transition to industries that fall under the complements case, the propensity to integrate upstream relative to downstream inputs should fall. For all the specifications that we describe below, we will also run regressions in which the demand elasticity ρ_j is instead replaced by our proxy for $\rho_j - \alpha_j$; in (21), this means that we will test P.1 using an indicator variable for whether $\rho_j - \alpha_j$ exceeds its median value across industries j .

We include a list of industry and firm controls in the above specification. The vector X_j includes measures of factor intensity, R&D intensity, and a value-added to shipments ratio (see Appendix A-2 for a more detailed description, as well as Table A-3 for basic summary statistics). The vector W_p contains parent firm characteristics obtained from WorldBase. This includes several variables that reflect the size of the parent, namely the number of establishments, whether it is a multinational, as well as log total employment and log total sales.²² We also account for the age of the parent by including the year of its establishment (or in which current ownership took control). We view X_j and W_p strictly as auxiliary controls, in the sense that the model does not deliver direct predictions that would lead us to clearly sign their effects on the ratio-upstreamness measure.

Table 1 reports the results of estimating (21). Column (1) presents a stripped-down specification in which only $\mathbf{1}(\rho_j > \rho_{med})$ and parent country fixed effects are included. The estimated coefficient on our proxy for the complements case is negative and significant at the 10% level, already confirming that the propensity to integrate upstream stages is lower in industries that face a high demand elasticity, consistent with prediction P1 (Cross). This result becomes even more significant (at the 1% level) as we successively add the output industry variables X_j in column (2), and the parent controls W_p in column (3). Looking at these auxiliary variables, the estimates indicate that there is a tendency towards upstream integration in more equipment capital-intensive industries, as well as in firms with more establishments, younger firms, and in multinationals.

²²For employment and sales, we also include dummy variables for whether the respective variables were based on actual data or were otherwise estimated/approximated by WorldBase.

The remaining columns in Table 1 explore alternative elasticity measures to capture industries in the complements case. Column (4) restricts the construction of ρ_j to the use of product-level elasticities classified by the UN BEC as either consumption or capital goods (dropping the intermediate-use products), while column (5) further limits this to just consumption goods elasticities. These refinements would in principle yield elasticities that pertain more directly to final-goods demand. Reassuringly though, this does not change the key finding of a negative and highly significant coefficient on the high-elasticity dummy, even though the number of observations falls as SIC industries that are composed entirely of intermediate-use goods are dropped from the sample. Finally, column (6) brings in information related to the demand elasticity for intermediate inputs, through the proxy for α_j . The key right-hand side variable is now an indicator for whether $\rho_j - \alpha_j$ is larger or smaller than its median value, where ρ_j is the demand elasticity from column (5) based on consumption goods only and the construction of α_j was described earlier (in Section 3). We continue to find that the propensity to integrate upstream stages is lower for industries that more likely correspond to the complements case on the basis of $\rho_j - \alpha_j$.

[Table 1 about here]

We also test prediction P.1 (Cross) by using specifications based upon a finer cut by quintiles of our proxy for ρ_j (alternatively, $\rho_j - \alpha_j$):

$$\log R_{jpc} = \beta_0 + \sum_{n=2}^5 \beta_n \mathbf{1}(\rho_j \in \text{Quint}_n(\rho)) + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \quad (22)$$

Here, $\mathbf{1}(\rho_j \in \text{Quint}_n(\rho))$ is an indicator variable for whether the demand elasticity for industry j belongs in the n -th quintile of that variable; the first quintile dummy is the omitted category. This approach has the advantage of allowing for more flexibility in the relationship between our empirical proxy for ρ_j and our ratio-upstreamness dependent variable.

Table 2 repeats the exercise in Table 1 using the above quintile specification. In line with our model's predictions, the magnitude of the estimated negative coefficient increases steadily as we move from the second to the fifth elasticity quintile throughout columns (2)-(6). As in Table 1, the regression based on the most stringent refinement of the ρ_j proxy – that in column (5) using consumption goods elasticities alone – yields the largest point estimates for the coefficients of interest. The implied magnitudes of these effects is fairly sizeable: Looking at columns (5) and (6), the fifth-quintile point estimates of -0.1849 and -0.1026 correspond to a range of between a half to a full standard deviation decrease (relative to the first quintile) in the propensity to integrate upstream inputs.

[Table 2 about here]

To assess prediction P.2 (Cross), we augment the specifications in (21) and (22) in order to

uncover the effects of upstream contractibility on integration:

$$\begin{aligned} \log R_{jpc} = & \beta_0 + \beta_1 \mathbf{1}(\rho_j > \rho_{med}) + \beta_{U1} \mathbf{1}(\rho_j < \rho_{med}) \times UpstCont_j + \beta_{U2} \mathbf{1}(\rho_j > \rho_{med}) \times UpstCont_j \\ & + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}, \text{ and} \end{aligned} \quad (23)$$

$$\begin{aligned} \log R_{jpc} = & \beta_0 + \sum_{n=2}^5 \beta_n \mathbf{1}(\rho_j \in Quint_n(\rho)) + \sum_{n=1}^5 \beta_{Un} \mathbf{1}(\rho_j \in Quint_n(\rho)) \times UpstCont_j \\ & + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \end{aligned} \quad (24)$$

In the median cutoff specification in (23), we interact the dummy variables $\mathbf{1}(\rho_j < \rho_{med})$ and $\mathbf{1}(\rho_j > \rho_{med})$ with the “upstream contractibility” measure, $UpstCont_j$. Based on the second prediction of our model, we would expect $\beta_{U1} < 0$ and $\beta_{U2} > 0$ in (23). Likewise in (24), we interact each of the quintile dummies with $UpstCont_j$, where the theory would lead us to expect that $\beta_{U1} < 0$ and $\beta_{U5} > 0$.²³

In Table 3, we report the results of (23). Notice that the estimated coefficient on the proxy for the complements case, $\mathbf{1}(\rho_j > \rho_{med})$, is negative and significant, as in the previous regressions in Table 1.²⁴ Turning to the interactions with $UpstCont_j$, the estimated coefficient in the complements case is positive and statistically significant, while that in the substitutes case is negative and also highly significant. This is entirely in line with the predictions of the model: Firms that fall under the complements case would have a lower propensity to integrate upstream stages, but this tendency is weakened among those industries whose production processes inherently exhibit a greater degree of upstream contractibility. The converse holds for the substitutes case, with $UpstCont_j$ instead lowering the propensity to integrate upstream stages when $\rho_j < \rho_{med}$. Note that these results hold when restricting the elasticity measure to HS codes classified as consumption or capital goods in column (2), when further limiting this to consumption goods elasticities only in column (3), and when using the proxy for $\rho_j - \alpha_j$ to distinguish between the two cases in column (4).

[Table 3 about here]

Table 4 confirms that the predictions related to upstream contractibility continue to hold with the more flexible quintile elasticity specification in (24). The main effects of the quintile elasticity dummies exhibit a pattern similar to that in the more parsimonious regressions in Table 3, with negative and significant coefficients especially as we transition to the higher quintiles. We perform a test for whether the effect of being in the fifth quintile, evaluated at the median in-sample value of $UpstCont_j$, is in fact significantly different from zero. The p-values reported in each column confirm that this is indeed the case, so that the propensity to integrate upstream inputs is lower in the fifth relative to the first elasticity quintile. This holds true regardless of the variant of the elasticity

²³The correlation between $UpstCont_j$ and the ρ_j proxy is small and never exceeds 0.06 in absolute value when we look across the various versions of the demand elasticity measure that we have constructed. The interaction effect is thus unlikely to be picking up a non-linear effect of the demand elasticity.

²⁴We have also verified that the overall effect of the $\mathbf{1}(\rho_j > \rho_{med})$ variable – taking into account its main effect and that through the interaction term with upstream contractibility – is indeed negative when evaluated at the median in-sample value of $UpstCont_j$. The p-value for this coefficient test is reported for each column in Table 3.

proxy used in each respective column. Of note, we find that in the complements case, a higher degree of upstream contractibility does counteract the above tendency to integrate upstream inputs, as the estimated coefficient on the fifth elasticity quintile interacted with $UpstCont_j$ is positive and statistically significant (at the 1% level) across all columns. Conversely, the interaction term between the first elasticity quintile dummy and $UpstCont_j$ bears the opposite sign, indicating that upstream contractibility instead acts to raise the propensity to integrate *downstream* inputs in this latter case. This last pattern appears most strongly in columns (1)-(3), where a demand elasticity associated with the output industry ρ_j is used to separate the complements from the substitutes cases. In column (4), where $\rho_j - \alpha_j$ is used instead, the largest negative effect appears to be concentrated in the second elasticity quintile. The overall message we obtain is in line with prediction P.2 (Cross), which relates integration decisions to the sequencing of high- versus low-contractibility inputs.

[Table 4 about here]

We have subjected the cross-firm regressions to an extensive series of robustness checks, a number of which are reported in Tables 5 and 6, as well as in the Appendix Tables. In the interest of space, we present results based on the specification from column (3) of Table 4, which uses the ρ_j measure constructed from consumption goods elasticities only; the results with the alternative elasticity measures are qualitatively similar and available on request.

In Table 5, we show that our results are robust to: (i) examining different subsamples of firms; and (ii) controlling for additional firm and industry variables. As discussed above, a recent study by Atalay *et al.* (2014) finds little evidence of significant domestic commodity shipments across plants owned by the same U.S. firm. In related work, Ramondo *et al.* (2015) document that the bulk of intrafirm trade involving a U.S. multinational parent tends to be concentrated among a small number of its large foreign affiliates. If only some of the listed SIC input activities of a firm’s affiliates are actually used in the production of the parent’s final good, our dependent variable will then suffer from measurement error. Crucially, our results continue to hold in column (1) of Table 5, where we restrict the analysis to the subset of single-establishment firms, for which such concerns related to cross-plant input shipments should be muted.²⁵ For these self-contained firms, we continue to find significant effects on the quintile elasticity dummies, as well as similar patterns on the interaction terms with $UpstCont_j$, i.e., a negative and significant coefficient for the first-quintile interaction, but the opposite sign for the fifth-quintile interaction. (We have reproduced the specifications in Table 4 in full using the single-establishment firm subsample in Appendix Table A-5 for the interested reader.) In column (2), we restrict the analysis to observations that have establishments in more than one country, i.e., multinational firms. The empirical findings remain largely intact, despite the fact that the number of observations decreases substantially with this cut of the dataset.²⁶

²⁵In the case of single-plant firms, it is unlikely that a parent would not use inputs produced in its own establishment.

²⁶The results are also unaffected if we expand our sample by lowering the employment threshold to a minimum of 10 employees, or if we restrict the sample to parents labeled as “global ultimates” (results available upon request).

In the remaining columns of Table 5, we consider several variables that have appeared elsewhere in the literature on firm-level vertical integration. Column (3) adds the share of direct input use in the production of j that could be obtained from within firm boundaries; for each parent, this is the sum of the direct coefficients of the inputs in $I(p)$ (see Acemoglu *et al.* 2009, and Alfaro *et al.* 2013). Next, column (4) controls for the share of total requirements value that each parent could in principle source from an overseas affiliate, together with a set of country fixed effects that indicate whether the parent has an establishment located in the country in question. Column (5) tests for whether the results might be driven by double marginalization motives, wherein parent firms would have an incentive to integrate inputs that exhibit a low demand elasticity, for which the markups charged by arm’s length suppliers would be higher. We control here for the (log) tr_{ij} -weighted average of the demand elasticity of inputs used by industry j . In addition, we include a tr_{ij} -weighted covariance of the input demand elasticity and $upst_{ij}$, to see if the correlation between these elasticities and production line position might matter. (Here, the demand elasticity associated with each input is computed using only those constituent HS10 products classified as intermediates by the UN BEC.) Our results remain robust to the inclusion of these variables, even when they are jointly entered into the regression (column (6)).²⁷

[Table 5 about here]

In Table 6, we report several checks based on alternative constructions of the ratio-upstreamness dependent variable. The version of R_{jpc} in column (1) is based on $upst_{ij}$ values obtained from a random pick when the mapping from I-O to SIC codes yielded multiple matches for a non-manufacturing input i . In column (2), we limit the set $S(j)$ in the construction of R_{jpc} to those inputs for which we observe at least one parent firm in j in our sample integrating the input in question (“ever-integrated” inputs). We alternatively restrict $S(j)$ to manufacturing inputs in column (3), and further drop the parent SIC from the set of inputs in column (4). Our findings are broadly consistent, with the main exception being the final column of Table 6. There, $UpstCont_j$ does reduce the propensity to integrate upstream in the first quintile (the substitutes case), but the point estimates for the fifth-quintile interactions (the complements case) are not significantly different from zero. Note though that the overall effect of being in quintile-5 (when evaluated at the median in-sample value of $UpstCont_j$) remains negative and significant, with the p-value from this coefficient test being 0.0013; in other words, the results in column (4) are still very much consistent with the earlier prediction P.1 (Cross). There is moreover a large decrease in the number of available observations in column (4), since this variant of the ratio-upstreamness measure can only be computed for those parent firms that have integrated at least one other manufacturing input apart from the parent’s primary SIC code.

[Table 6 about here]

²⁷Interestingly, the weighted covariance between the input elasticity and upstreamness has a coefficient with the expected sign (negative and significant), consistent with the interpretation that the presence of demand-*inelastic* inputs upstream in the production process would be associated with more upstream integration.

We briefly describe here a number of other notable checks that we performed, which are reported in the Appendix. In Table A-4, we present several alternative treatments for parent firms that could be active as output producers in multiple industries. We first check whether the patterns are similar when limiting the sample to parents that have only one manufacturing SIC code, i.e., that do not report any secondary SIC manufacturing activities (columns (1) and (2)). Alternatively, we can designate the output industry j to be the SIC code of the parent (among the up to six codes reported) that is most proximate to final demand, on the basis of the Fally (2012) and Antràs *et al.* (2012) measure of upstreamness (columns (3) and (4)). Last but not least, we have constructed R_{jpc} taking in turn each secondary manufacturing SIC code as the parent’s output industry j . The regression in (24) is then run, pooling across the multiple R_{jpc} values per parent (columns (5) and (6)); two-way clustered standard errors by SIC output industry and by parent firm are reported (Cameron *et al.* 2011). Overall, our regression findings remain stable under each of these approaches to account for multi-product firms.²⁸

Summing up, the cross-firm regressions provide strong evidence that the propensity for a firm to integrate relatively upstream inputs is weakest when the demand elasticity faced by that industry is largest, in line with prediction P.1 (Cross). Prediction P.2 (Cross) concerning the impact of upstream contractibility along the sequence of inputs also finds strong support in the data.

4.2 Within-firm Results

We exploit our data further, by exploring whether the patterns of integration within firms are consistent with the model’s predictions. To study within-firm integration decisions, we restructure the data so that an observation is now a SIC input i by parent p pair. To assess the validity of prediction P.1 (Within), we run the following two types of specifications:

$$D_INT_{ijp} = \gamma_0 + \gamma_1 \mathbf{1}(\rho_j < \rho_{med}) \times upst_{ij} + \gamma_2 \mathbf{1}(\rho_j > \rho_{med}) \times upst_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp} \quad (25)$$

$$D_INT_{ijp} = \gamma_0 + \sum_{n=1}^5 \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times upst_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (26)$$

The dependent variable, D_INT_{ijp} , is a 0-1 indicator for whether the firm p with primary output j has integrated the input i within firm boundaries. The key explanatory variables are the terms involving $upst_{ij}$ and its interactions with the elasticity variables. D_p denotes a full set of parent fixed effects. These specifications therefore allow us to study the integration decisions of individual firms, how they are affected by the upstreamness of the inputs, and whether these effects vary across the complements and substitutes cases. In particular, our theory would suggest that $\gamma_1 > 0$ and $\gamma_5 < 0$ (in the quintile specification), although we shall see below that the empirical results

²⁸We have also explored the robustness of our results to the inclusion of several controls related to various dimensions of input contractibility, such as (i) the contractibility of the output industry j itself, (ii) a tr_{ij} -weighted average of the contractibility of the inputs used by j , and (iii) a set of interactions between each quintile dummy and a tr_{ij} -weighted variance of the contractibility of the inputs used by j . The results are available upon request.

will be consistent with a weaker form of this prediction.

We include two additional sets of controls. The first is a dummy variable $\mathbf{1}(i = j)$ that is equal to 1 if and only if input i has the same SIC code as the output industry j . In such instances, D_INT_{ijp} always takes on a value of 1, as $j \in I(j)$ by definition. Including this dummy allows us to focus on the effects of $upst_{ij}$ for manufacturing inputs other than the “self-SIC”. Second, in our most stringent specifications, we use a full set of dummies for the input SIC code, D_i , which allows us to control for any input characteristics that might affect a firm’s propensity to integrate it. When these input fixed effects are used, only covariates that vary at the input-output (i - j) pair level can be identified in the estimation.

We estimate (25) and (27) as a linear probability model, with standard errors clustered by i - j pair. To keep the analysis tractable, we limit the sample to the top 100 manufacturing inputs i used by j , as ranked by the total requirements coefficient tr_{ij} . This covers between 88-98% of the total requirements value for each output industry. We focus on the subsample of parent firms that have integrated at least one manufacturing input other than the parent’s self-SIC code (the set of firms that appear in column (4) of Table 5), in order to avoid including firms for whom occurrences of integration are sparse.

To assess prediction P.2 (Within), we extend (25) and (26) by adding the interactions between the ρ_j indicator variables and $ContUpToi_{ij}$, where the latter measure captures the contractibility of all inputs up to i in the production of j :

$$D_INT_{ijp} = \gamma_0 + \gamma_1 \mathbf{1}(\rho_j < \rho_{med}) \times upst_{ij} + \gamma_2 \mathbf{1}(\rho_j > \rho_{med}) \times upst_{ij} + \delta_1 \mathbf{1}(\rho_j < \rho_{med}) \times ContUpToi_{ij} + \delta_2 \mathbf{1}(\rho_j > \rho_{med}) \times ContUpToi_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (27)$$

$$D_INT_{ijp} = \gamma_0 + \sum_{n=1}^5 \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times upst_{ij} + \sum_{n=1}^5 \delta_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times ContUpToi_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (28)$$

Recall that $ContUpToi_{ij}$ was constructed in (20) as an empirical proxy for $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ from the model. Looking back at the expression for the optimal bargaining share, $\beta^*(m)$, in equation (13), one would then expect that “contractibility up to i ” would raise the propensity to integrate input i if industry j came under the complements case ($\delta_5 > 0$), while having the opposite effect in the substitutes case ($\delta_1 < 0$). There is a further implication from (13), namely that having controlled for the role of $ContUpToi_{ij}$, one should no longer expect to see that $upst_{ij}$ would have a significant effect on integration decisions. This is because the effect of m on $\beta^*(m)$ is captured entirely by the $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ term.

The findings from the within-firm estimation are reported in Table 7. Following (25) and (27), we first adopt the median elasticity cutoff dummies (constructed from consumption-goods elasticities only) to differentiate between output industries in the complements and substitutes cases. Column (1) reveals a negative and significant effect of $upst_{ij}$ in industries that feature an above-median

demand elasticity. This dovetails with prediction P.1 (Within) in that the propensity to integrate declines the more upstream the input in question for firms that fall under the complements case. The coefficient obtained for the interaction between $\mathbf{1}(\rho_j < \rho_{med})$ and $upst_{ij}$ is also negative, albeit of a smaller magnitude. While the sign is at odds with a strict statement of the theory’s prediction in the substitutes case, it is consistent with the weaker conclusion that the effect of upstreamness in lowering the propensity to integrate is stronger in the complements case. Note that the “self-SIC” dummy emerges with a positive and highly significant effect, with a point estimate close to 1, as it does in all remaining columns.

[Table 7 about here]

In column (2), we present results when interaction terms involving the median cutoff dummies and the “contractibility up to i ” measure are introduced to the regression, as in (27). The effect of $ContUpToi_{ij}$ is indeed positive and significant when $\rho_j > \rho_{med}$, which is in line with prediction P.2 (Within) for the complements case, in that a greater degree of contractibility upstream of input i raises the likelihood that we observe i being integrated. Once again however, we find that the point estimate for the interaction term when $\rho_j < \rho_{med}$ is consistent with a weaker form of the model’s prediction, being positive though smaller in magnitude compared to the corresponding coefficient for the above-median interaction term. (The p-value reported in column (2) confirms that we can reject the null hypothesis that the estimated coefficients in the above- and below-median elasticity cases are equal.) These patterns persist even when we include SIC input fixed effects to control for any characteristics specific to inputs i (column (3)), or use the log total requirements coefficient as an additional control (column (4)). Note that this latter tr_{ij} variable enters with a positive and significant coefficient, suggesting that firms are more likely to integrate when an input is more important in production. Last but not least, column (5) presents the results when using the $\rho_j - \alpha_j$ proxy as the elasticity measure of interest. The findings are retained, although the difference in the effect of $ContUpToi_{ij}$ in the below- and above-median elasticity cases is now marginally insignificant.

In Table 8, we repeat the above using a more extensive set of quintile elasticity dummies instead for the elasticity measure, following (26) and (28). The patterns we find here are qualitatively very similar. When the quintile dummies are interacted with $upst_{ij}$ in column (1), notice that the coefficients become successively more negative in the higher elasticity quintiles, consistent once again with the integration of upstream inputs being less likely in the complements relative to the substitutes cases. This neat pattern for the effect of $upst_{ij}$ however disappears when the analogous interactions involving $ContUpToi_{ij}$ are added (column (2)); instead, it is the effect of “contractibility” up to i that increases monotonically across the quintiles.²⁹ This confirms that it is the effect of the contractibility of upstream inputs, rather than upstreamness *per se*, that matters for integration patterns. This is precisely what our model predicts, since what matters for the

²⁹A formal test for the equality of the first and fifth quintile coefficients is rejected at conventional significance levels (see the p-value reported in each column).

optimal organizational decision at stage i is $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$, and that having controlled for this, the upstreamness of i in the production of j should have no further effect. These patterns remain even with the inclusion of SIC input fixed effects and the log total requirements coefficient, as well as when variation in α_j is brought to bear on our elasticity measure (columns (3)-(5)).

[Table 8 about here]

We conclude the empirical discussion with several robustness checks on our within-firm regressions. The specifications that we report in Table 9 are based on that in column (4) of the preceding Table 8. Column (1) demonstrates what happens when we restrict the sample once again to single establishment firms. In columns (2) and (3), we focus on subsets of firms that feature more interesting variation in their integration patterns. Column (2) drops firms that do not have an integrated manufacturing input (apart from the self-SIC) that ranks in the top 100 inputs used as determined by the total requirements coefficient. Alternatively, column (3) retains only those parents that have integrated at least three of their top-100 manufacturing inputs. The findings we obtain from these different subsample cuts of the data turn out to be very similar to those presented already in Table 8. Column (4) adopts a different treatment of the self-SIC code, which is classified rather mechanically as an integrated input in our regressions. Here, the self-SIC is instead dropped altogether from the estimation. While we lose some statistical significance on the effect of “contractibility up to i ” in the lowest elasticity quintiles, the positive and significant coefficients in the highest quintiles that map to the complements case are preserved. Last but not least, in column (5), we include the full set of quintile elasticity dummies interacted with “contractibility at i ”, where this latter variable is given by: $\frac{tr_{ij} cont_i}{\sum_{k \in S^m(j)} tr_{kj} cont_k}$. In words, this is the component of $ContUpToi_{ij}$ that is accrued at stage i itself. The results indicate that it is indeed the profile of contractibility prior to input i , rather than that at stage i , that matters for explaining integration patterns.

[Table 9 about here]

5 Conclusion

Revolutionary advances in information and communication technology combined with major reductions in trade barriers have led to the fragmentation of production processes across firms and countries. The emergence of global value chains has attracted much attention from policymakers and academics alike. However, there are still few systematic empirical studies attempting to shed light on the determinants of firms’ decision to control different segments of their production processes. In this paper, we show how detailed data on the activities of firms around the world can be combined with information from standard Input-Output tables to study integration choices along value chains.

Building on Antràs and Chor (2013), we describe a property-rights model in which a firm’s boundaries are shaped by characteristics of the different stages of production and their position in

the value chain. As available theoretical frameworks of sequential production are highly stylized, a key contribution of this paper is to develop a richer framework of firm behavior that can guide an empirical analysis using firm-level data.

To assess the evidence, we use the WorldBase dataset, which contains establishment-level information on the activities of firms located in a large set of countries. We combine this information with Input-Output tables to construct firm-level measures of the upstreamness of integrated and non-integrated stages. The richness of our data allows us to run specifications that exploit variation in organizational features across firms, as well as within firms and across their various manufacturing stages.

In line with the model's predictions, we find that whether a firm integrates suppliers located upstream or downstream depends crucially on the size of the elasticity of demand faced by the firm. Moreover, the relative propensity to integrate upstream (as opposed to downstream) inputs depends also on the extent to which contractible inputs tend to be located in the early or late stages of the production process. The firm-level patterns that we uncover provide strong evidence that considerations driven by contractual frictions critically shape firms' ownership decisions along their value chains.

References

- Acemoglu, Daron, Pol Antràs, and Elhanan Helpman (2007), "Contracts and Technology Adoption," *American Economic Review* 97(3): 916-943.
- Acemoglu, Daron, Simon Johnson, and Todd Mitton (2009), "Determinants of Vertical Integration: Financial Development and Contracting Costs," *Journal of Finance* 63(3): 1251-1290.
- Alfaro, Laura, and Andrew Charlton (2009), "Intra-Industry Foreign Direct Investment," *American Economic Review* 99(5): 2096-2119.
- Alfaro, Laura, and Maggie Xiaoyang Chen (2014), "The Global Agglomeration of Multinational Firms," *Journal of International Economics* 94(2): 263-276.
- Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew Newman (2013), "Do Price Determine Vertical Integration?" *NBER Working Paper* 16118.
- Antràs, Pol (2003), "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics* 118(4): 1375-1418.
- Antràs, Pol (2015), *Global Production: Firms, Contracts and Trade Structure*, Princeton University Press, forthcoming.
- Antràs, Pol, and Davin Chor (2013), "Organizing the Global Value Chain," *Econometrica* 81(6): 2127-2204.
- 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.
- Atalay, Enghin, Ali Hortaçsu, and Chad Syverson (2014), "Vertical Integration and Input Flows," *American Economic Review* 104(4): 1120-1148.

- Baldwin, Richard, and Anthony Venables (2013), "Spiders and Snakes: Offshoring and Agglomeration in the Global Economy," *Journal of International Economics* 90(2): 245-254.
- Becker, Randy A., and Wayne B. Gray (2009), "NBER-CES Manufacturing Industry Database (1958-2005)".
- Bresnahan, Timothy, and Jonathan Levin (2012), "Vertical Integration and Market Structure," in *The Handbook of Organizational Economics*, Princeton University Press.
- Broda, Christian, and David Weinstein (2006), "Globalization and the Gains from Variety," *Quarterly Journal of Economics* 121(2): 541-585.
- Brown, Clair, and Greg Linden (2005), "Offshoring in the Semiconductor Industry: A Historical Perspective," *Brookings Trade Forum* (Offshoring White-Collar Work), pp.279-333.
- Cameron, Colin, Jonah Gelbach, and Douglas Miller (2011), "Robust Inference with Multi-way Clustering," *Journal of Business and Economic Statistics* 29(2): 238-249.
- Caves, Richard (1975), "Diversification, Foreign Investment and Scale in North American Manufacturing Industries," *Canadian Public Policy* 2: 274-276.
- Costinot, Arnaud, Jonathan Vogel, and Su Wang (2013), "An Elementary Theory of Global Supply Chains," *Review of Economic Studies* 80(1): 109-144.
- Corcos, Gregory, Delphine Irac, Delphine M., Giordano Mion, and Thierry Verdier (2013), "The Determinants of Intrafirm Trade: Evidence from French Firms," *Review of Economics and Statistics* 95(3): 825-838.
- Defever, Fabrice, and Farid Toubal (2013), "Productivity, Relationship-Specific Inputs and the Sourcing Modes of Multinationals," *Journal of Economic Behavior and Organization* 94: 345-357.
- Del Prete, Davide, and Armando Rungi (2015), "Organizing the Global Value Chain: A Firm Level Test," mimeo.
- Dietzenbacher, Erik, Isidoro Romero Luna, and Niels S. Bosma (2005), "Using Average Propagation Lengths to Identify Production Chains in the Andalusian Economy," *Estudios de Economía Aplicada* 23(2): 405-422.
- Díez, Federico (2010), "The Asymmetric Effects of Tariffs on Intra-firm Trade and Offshoring Decisions," *Journal of International Economics* 93(1): 76-91.
- Dixit, Avinash, and Gene Grossman (1982), "Trade and Protection with Multistage Production," *Review of Economic Studies* 49(4): 583-594.
- Fajgelbaum, Pablo, Gene Grossman, and Elhanan Helpman (2015), "A Linder Hypothesis for Foreign Direct Investment," *Review of Economic Studies* 82(1): 83-121.
- Fally, Thibault (2012), "On the Fragmentation of Production in the U.S.," mimeo.
- Fally, Thibault, and Russell Hillbery (2014), "A Coasian Model of International Production Chains," mimeo.
- Fan, Joseph P. H., and Larry H. P. Lang (2000), "The Measurement of Relatedness: An Application to Corporate Diversification," *Journal of Business* 73(4): 629-660.
- Feenstra, Robert C., John Romalis, and Peter K. Schott, (2002), "U.S. Imports, Exports and Tariff Data, 1989-2001," *NBER Working Paper* 9387.

- 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.
- Harrison, Ann E., Inessa Love, and Margaret S. McMillan (2004), "Global Capital Flows and Financing Constraints," *Journal of Development Economics* 75(1): 269-301.
- Robert C. Johnson, and Guillermo Noguera (2012), "Accounting for Intermediates: Production Sharing and Trade in Value Added," *Journal of International Economics* 86(2): 224-236.
- Kikuchi, Tomoo, Kazuo Nishimura, and John Stachurski (2014), "Transaction Costs, Span of Control and Competitive Equilibrium," mimeo.
- 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.
- Lafontaine, Francine, and Margaret Slade (2007), "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature* 45(3): 629-685.
- Luck, Philip (2014), "Global Supply Chains and Vertical Integration: Evidence from China," mimeo.
- 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 (2013), "Incomplete Contracts and the Boundaries of the Multinational Firm," *Journal of Economic Behavior and Organization* 94: 330-344.
- Ramondo, Natalia, Veronica Rappoport, and Kim Ruhl (2015), "Horizontal versus Vertical Foreign Direct Investment: Evidence from U.S. Multinationals," *Journal of International Economics*, forthcoming.
- Rauch, James E. (1999), "Networks versus Markets in International Trade," *Journal of International Economics* 48(1): 7-35.
- Sanyal, Kalyan K., and Ronald W. Jones, (1982), "The Theory of Trade in Middle Products," *American Economic Review* 72(1): 16-31.
- Xing, Yuqing (2011), "How the iPhone widens the US trade deficit with China," VoxEU.org, 10 April.
- Yeaple, Stephen R. (2006), "Offshoring, Foreign Direct Investment, and the Structure of U.S. Trade," *Journal of the European Economic Association* 4(2-3): 602-611.
- Yi, Kei-Mu (2003), "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111(1): 52-102.

Appendices

A-1 Theoretical Appendix

A-1.1 Derivation of Program (10)

In this Appendix, we provide more details on firm behavior conditional on the path of ownership structure along the value chain. Notice first that solving program (8), we obtain the following optimal choice of investment by the supplier at stage m :

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

Plugging this express into the marginal contribution function $r'(m) = \frac{\rho}{\alpha} (A^{1-\rho})^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} \psi(m)^\alpha x(m)^\alpha$ delivers the following separable differential equation:

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

It is straightforward to verify that the solution to this differential equation is given by equation (9) in the main text, from which can also conclude that

$$x(m) = A \left(\frac{1 - \rho}{1 - \alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{\rho}{c} \right)^{\frac{1}{1-\rho}} \left(\frac{(1 - \beta(m)) \psi(m)}{c(m)} \right)^{\frac{1}{1-\alpha}} \left[\int_0^m \left(\frac{(1 - \beta(i)) \psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}, \quad (\text{A-1})$$

which is clearly positive as long as $\beta(m) < 1$. This implies that the firm's profits are enhanced at each stage m by having suppliers provide compatible inputs, which in turn implies that it is optimal to abide by the proper sequencing of production, or $I^*(m) = 1$ for all m (as implicitly assumed in the previous expressions).

The firm thus chooses the path of $\beta(i)$ that maximizes its profits $\pi_F = \int_0^1 \beta(i) r'(i) di$. Differentiating (9), we can express this profit function as

$$\pi_F = A \frac{\rho}{\alpha} \left(\frac{1 - \rho}{1 - \alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}} \int_0^1 \beta(i) \left(\frac{(1 - \beta(i)) \psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^i \left(\frac{(1 - \beta(k)) \psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di, \quad (\text{A-2})$$

which coincides with the expression in program (10) in the main text, and which in turn constitutes a generalization of program (6).

A-1.2 Derivation of Equation (13)

As pointed out in the main text, we can express program (10) as a standard calculus of variation problem where the firm chooses the real-value function v that maximizes the functional

$$\pi_F(v) = \Theta \int_0^1 \left(1 - v'(i)^{\frac{1-\alpha}{\alpha}} \frac{c(i)}{\psi(i)} \right) v'(i) v(i)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di,$$

where

$$v(i) \equiv \int_0^i \left(\frac{(1 - \beta(k)) \psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk. \quad (\text{A-3})$$

The Euler-Lagrange equation associated with this problem is given by:

$$\frac{\rho - \alpha}{\alpha(1 - \rho)} \left[1 - v'(i)^{\frac{1-\alpha}{\alpha}} \frac{c(i)}{\psi(i)} \right] v'(i) [v(i)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}-1} = \frac{d}{di} \left[[v(i)]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(1 - \frac{1}{\alpha} v'(i)^{\frac{1-\alpha}{\alpha}} \frac{c(i)}{\psi(i)} \right) \right],$$

which after a couple of manipulations can be reduced to the following differential equation:

$$\frac{\rho - \alpha}{1 - \rho} \frac{[v'(i)]^2}{v(i)} + \frac{v''(i)}{v'(i)} = -\frac{\alpha}{1 - \alpha} \frac{d(c(i)/\psi(i))/di}{c(i)/\psi(i)} \quad (\text{A-4})$$

To solve (A-4), integrate both sides with respect to i , and exponentiate to get:

$$v'(i) v(i)^{\frac{\rho-\alpha}{1-\rho}} = C_1 (\psi(i)/c(i))^{\frac{\alpha}{1-\alpha}}, \quad (\text{A-5})$$

where $C_1 > 0$ is a constant of integration. Given the definition of $v(i)$ in (A-3), we can express this equation as:

$$(1 - \beta(i))^{\frac{\alpha}{1-\alpha}} = C_1 \left(\int_0^i \left(\frac{(1 - \beta(k)) \psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right)^{\frac{\alpha-\rho}{1-\rho}}, \quad (\text{A-6})$$

Denoting $z(i) \equiv (1 - \beta(i))^{\frac{\alpha}{1-\alpha}}$, we can express (A-6) as

$$\left(\frac{z(i)}{C_1} \right)^{\frac{1-\rho}{\alpha-\rho}} = \int_0^i z(k) \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk,$$

which after differentiation delivers

$$\frac{1 - \rho}{\alpha - \rho} \left(\frac{z(i)}{C_1} \right)^{\frac{1-\rho}{\alpha-\rho}} \frac{z'(i)}{z(i)} = z(i) \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}}.$$

This change of variable has thus allowed us to arrive at a separable differential equation in $z(i)$, which has solution:

$$z(i) = (C_1)^{\frac{1-\rho}{1-\alpha}} \left(\frac{1 - \alpha}{1 - \rho} \right)^{\frac{\alpha-\rho}{1-\alpha}} \left[\int_0^i \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\alpha-\rho}{1-\alpha}}.$$

Given the definition of $z(i)$ and imposing the transversality condition

$$1 - \frac{1}{\alpha} v'(1)^{\frac{1-\alpha}{\alpha}} \frac{c(1)}{\psi(1)} = 0 \implies 1 - \beta(1) = \alpha,$$

we finally obtain equation (13).

A-1.3 Proof of Proposition 2

The proof is a generalization of the proof of Proposition 2 in Antràs and Chor (2013). It is straightforward to see from equation (13), that when $\rho > \alpha$, $\lim_{m \rightarrow 0} \beta^*(m) \rightarrow -\infty$, and it is thus optimal for the firm to choose β_O (namely outsourcing) for the most upstream stages in the neighborhood of $m = 0$. Conversely, whenever $\rho < \alpha$ case, $\lim_{m \rightarrow 0} \beta^*(m) = 1$, it is optimal for the firm to choose β_V (namely integration) for those upstream stages in the 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, contrary to the claim in Proposition 2, there were to exist a stage $\tilde{m} \in (0, 1)$ such that a measurable set of stages upstream from \tilde{m} are integrated, while a measurable set of stages downstream from \tilde{m} are outsourced. Now consider two positive constants ε_L and ε_R such that

$$\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di = \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di. \quad (\text{A-7})$$

These constants can always be chosen small enough such that not only do they satisfy (A-7), but they are also such that the set of stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ are integrated, while stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$ are outsourced. Denote by Π_1 firm profits under this suggested ownership structure. We shall consider an alternative organizational mode such that the firm instead chooses to outsource the stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ and integrates the stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$, while retaining the same organizational decision for all other stages. Denote the profits of this alternative organizational form by Π_2 . If we can show that this reorganization necessarily increases firm profits, i.e., $\Pi_1 < \Pi_2$, then we will have shown that our posited deviation from the optimal pattern in Proposition 2 is inconsistent with profit maximization.

Note that we can rewrite firm profits in (10) as

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \int_0^1 \beta(i) \frac{\partial \left(\left[\int_0^i ((1-\beta(k))\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right)}{\partial i} di, \quad (\text{A-8})$$

It is useful to distinguish four regions in the set of stages: (i) all stages upstream from $\tilde{m} - \varepsilon_L$, (ii) those in $(\tilde{m} - \varepsilon_L, \tilde{m})$; (iii) those in $(\tilde{m}, \tilde{m} + \varepsilon_R)$, and (iv) all stages downstream from $\tilde{m} + \varepsilon_R$. Note that the profits generated by all stages in the first region are common for the profits functions Π_1 and Π_2 , so we can ignore them hereafter. Less trivially, the profits generated in the last region are also common in the profit functions Π_1 and Π_2 . To see this, and to keep the notation manageable, define

$$\begin{aligned} \gamma(i) &= (\psi(i)/c(i))^{\frac{\alpha}{1-\alpha}}, \\ \mathcal{A} &= \int_0^{\tilde{m}-\varepsilon_L} ((1-\beta(k))\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk, \text{ and} \\ \mathcal{D} &= \int_{\tilde{m}+\varepsilon_R}^i ((1-\beta(k))\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk. \end{aligned}$$

Notice next that in light of equation (A-8), the part of profits Π_1 associated with stages $m > \tilde{m} + \varepsilon_R$ is

$$\Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \int_{\tilde{m}+\varepsilon_R}^1 \beta(i) \frac{\partial}{\partial i} \left(\mathcal{A} + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(k) dk + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(k) dk + \mathcal{D} \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} di,$$

while for profits Π_2 , these same profits are given by

$$\Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \int_{\tilde{m}+\varepsilon_R}^1 \beta(i) \frac{\partial}{\partial i} \left(\mathcal{A} + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(k) dk + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(k) dk + \mathcal{D} \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} di,$$

But given (A-7), we have that $\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(k) dk = \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(k) dk$ and so these two expressions are equal.

In order to compare the relative size of Π_1 and Π_2 , it thus suffices to compare profits associated only with the intervals $(\tilde{m} - \varepsilon_L, \tilde{m})$ and $(\tilde{m}, \tilde{m} + \varepsilon_R)$. Again invoking equation (A-8), and after some manipulations, we find that

$$\begin{aligned} \Pi_1 - \Pi_2 \propto (\beta_V - \beta_O) & \left[\left(\mathcal{A} + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} + \left(\mathcal{A} + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right. \\ & \left. - \left(\mathcal{A} + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(i) di \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - \mathcal{A}^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right]. \end{aligned}$$

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 $y = \mathcal{A}$, $a = (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di$ and $b = (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di$, 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 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.

A-1.4 Derivation of Thresholds m_C^* and m_S^* in Section 2.2

Consider first the complements case $\rho > \alpha$, in which all stages upstream from m_C^* are outsourced, while all stages downstream from m_C^* are integrated. We can then use (A-8) to express profits as

$$\begin{aligned} \pi_F = & \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \beta_O (1-\beta_O)^{\frac{\rho}{1-\rho}} \left(\int_0^{m_C} \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \\ & + \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \beta_V \left[\left((1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_0^{m_C} \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{m_C}^1 \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right. \\ & \left. - \left((1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_0^{m_C} \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right]. \end{aligned} \quad (\text{A-9})$$

Taking the first-order-condition with respect to the threshold m_C and rearranging, we then find

$$(\beta_V - \beta_O) (1-\beta_O)^{\frac{\rho}{1-\rho}} = \beta_V \left((1-\beta_O)^{\frac{\alpha}{1-\alpha}} - (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \right) \left[(1-\beta_O)^{\frac{\alpha}{1-\alpha}} + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \frac{\int_{m_C}^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^{m_C} (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk} \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}},$$

from which equation (14) can easily be obtained. Notice that $m_C^* \in (0, 1)$ requires the right-hand-side of (14) to be lower than one, which in turn requires

$$\left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\frac{\alpha}{1-\alpha}} > \frac{\beta_O}{\beta_V},$$

or simply $\beta_V (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} > \beta_O (1 - \beta_O)^{\frac{\alpha}{1-\alpha}}$, as claimed in the main text.

The threshold in the substitutes case can be derived in an analogous way. In fact, it is straightforward to see that m_S will be chosen to maximize a profit function identical to that in (A-9) with β_O replacing β_V throughout and vice versa. As a result, m_S^* is given by

$$\frac{\int_0^{m_S^*} (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di} = \left\{ 1 + \left(\frac{1 - \beta_V}{1 - \beta_O} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{\frac{\beta_V}{\beta_O} - 1}{\left(\frac{1 - \beta_V}{1 - \beta_O} \right)^{-\frac{\alpha}{1-\alpha}} - 1} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}.$$

A-1.5 Derivation of Equation (15)

Let us ignore contracting costs for now and focus on the first term of equation (15). Consider first the complements case. We begin by plugging the equilibrium equation (14) determining the thresholds m_C^* in into the profit function (A-9). After a few simplifications, this delivers

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} (1 - \beta_O)^{\frac{\rho}{1-\rho}} (\mathcal{H}_C)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \left[(\beta_O - \beta_V) + \beta_V \left(\frac{1 - \frac{\beta_O}{\beta_V}}{1 - \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\rho(1-\alpha)}{\rho-\alpha}} \right]$$

where

$$\mathcal{H}_C = \left\{ 1 + \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1 - \frac{\beta_O}{\beta_V}}{1 - \left(\frac{1 - \beta_O}{1 - \beta_V} \right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}.$$

Hence, we can write profits as

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \Gamma_C(\beta_V, \beta_O, \rho, \alpha).$$

In the substitutes case, we obtain

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} (1 - \beta_V)^{\frac{\rho}{1-\rho}} (\mathcal{H}_S)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \left[(\beta_V - \beta_O) + \beta_O \left(\frac{1 - \frac{\beta_V}{\beta_O}}{1 - \left(\frac{1 - \beta_V}{1 - \beta_O} \right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\rho(1-\alpha)}{\rho-\alpha}} \right]$$

where

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

or

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \Gamma_S(\beta_V, \beta_O, \rho, \alpha).$$

Overall, we then see that profits can be expressed compactly as

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \Gamma(\beta_V, \beta_O), \quad (\text{A-10})$$

where

$$\Gamma(\beta_V, \beta_O) = \begin{cases} \Gamma_C(\beta_V, \beta_O, \rho, \alpha) & \text{if } \rho \geq \alpha \\ \Gamma_S(\beta_V, \beta_O, \rho, \alpha) & \text{if } \rho < \alpha \end{cases}.$$

It is straightforward to verify that equation $\Gamma_S(\beta_V, \beta_O, \rho, \alpha)$ is identical to $\Gamma_S(\beta_V, \beta_O, \rho, \alpha)$ except for the fact that β_V is replaced by β_O and β_O is replaced by β_V .

Obtaining equations (15) from the more general equation (A-10) is then trivial. Notice, however, that when studying the optimal choice of $\psi(m)$, the first-order condition of this generalized version of program (15) now delivers that, for two inputs at stages m and m' , we have

$$\frac{\psi(m)/c(m)}{\psi(m')/c(m')} = \left(\frac{\mu(m)/c(m)}{\mu(m')/c(m')} \right)^{\phi - \frac{\alpha}{1-\alpha}}, \quad (\text{A-11})$$

which generalizes equation (16) in the main text, and where again the second-order conditions impose $\phi > \alpha/(1-\alpha)$. This equation illustrates that the ratio $\psi(m)/c(m)$ will tend to comove with contractibility along the value chain as long as contractibility and marginal costs are not positively correlated. But notice that plugging (A-11) into (A-10), we have that the effect of a reduction in the marginal cost of a given stage i will be increasing in the level of contractibility $\mu(i)$. As a result, if we were to interpret the path of marginal costs as being the outcome of an optimal global sourcing model, then we would expect, other things equal, that the firm would be particularly willing to achieve marginal cost reductions for highly contractible stages, thus resulting in a negative correlation between $c(i)$ and $\mu(i)$.

A-2 Data Appendix: Construction of the Industry Controls

Import demand elasticities: Based on the U.S. HS10 product import demand elasticities estimated by Broda and Weinstein (2006). These are mapped into SIC categories using concordance weights based on total U.S. imports between 1989-2006 from Feenstra *et al.* (2002). 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 shares the same first nine digits. This was done successively up to codes that share the same first two digits, to assign as many HS10 codes with elasticities as possible. The corresponding elasticity for each SIC 4-digit code is then taken as the trade-weighted average over its constituent HS10 elasticities. A number of 4-digit SIC codes remain without elasticities, as these codes are not used in the U.S. import trade data. This arises because customs is unable to distinguish the source industry of certain goods on the basis of their physical specimen; for example, it cannot distinguish SIC2011 (Meat Packing Plants) from SIC2013 (Sausages and Other Prepared Meats). In such instances, U.S. customs assigns all the goods value to one of the possible SIC codes, and excludes the others. Table 3 in Feenstra *et al.* (2002) provides a list of such excluded codes and their corresponding destination codes, allowing us to compute a trade-weighted elasticity value of the respective destination codes to obtain an elasticity for each excluded code. There were 51 4-digit SIC codes that were successfully assigned in this way. After these steps, there were still 10 4-digit SIC codes missing an elasticity. A trade-weighted average elasticity over all 4-digit SIC categories that share the same first three digits, and if necessary those that share the same first two digits, was computed for these remaining 10 codes.

Contractibility: Following the methodology in Nunn (2007). This in turn uses the Rauch (1999) classification of goods as either homogeneous, reference-priced, or differentiated. Rauch’s original classification is in SITC Rev 2. We use Feenstra *et al.* (2002) to obtain a master-list of HS by SITC Rev 2 by SIC triplets. We then associate the Rauch codings from each SITC Rev 2 to all the HS10 products that fall under it. For each SIC 4-digit category, we calculated the fraction of HS10 constituent codes classified as neither reference-priced nor traded on an organized exchange; we term this the specificity of the industry in question. We used the procedure described above for import demand elasticities to assign the specificity values for missing 4-digit SIC codes. In the next step, we take a direct requirements weighted average over the specificities of all inputs purchased, to obtain Nunn’s (2007) contract-intensity measure for each 4-digit SIC code. We take one minus the contract-intensity to get a measure of contractibility.

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. 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. These are computed as averages over 2001-2005, using the annual data for 4-digit SIC industries. For a small number of industries without 2001-2005 data, we used an average over a late in-sample five-year window. One further variable – value-added over total shipments – was constructed in the same way.

R&D intensity: From Nunn and Treffer (2013), who calculated R&D expenditures to total sales on an annual basis for HS6 products, using the U.S. firms in the Orbis dataset. We use an analogous procedure to that described above for the import demand elasticities, to first assign R&D intensity values using the trade-weighted average over HS codes that share the same first five digits, and so on successively until the same first two digits. These are then converted to 4-digit SIC codes using a trade-weighted average intensity of constituent HS6 codes. All concordance weights are from Feenstra *et al.* (2002). Missing codes were assigned following the same procedure as for the import demand elasticities.

Table 1: Upstreamness of Integrated vs Non-Integrated Inputs: Median Elasticity Cutoff

Dependent variable:	Log Ratio-Upstreamness					
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Elas _j > Median)	-0.0354* [0.0204]	-0.0612*** [0.0188]	-0.0604*** [0.0185]	-0.0593*** [0.0215]	-0.1138*** [0.0261]	-0.1073*** [0.0275]
Log (Skilled Emp./Workers) _j		0.0100 [0.0243]	0.0091 [0.0245]	0.0111 [0.0278]	-0.0219 [0.0360]	-0.0082 [0.0364]
Log (Equip. Capital/Workers) _j		0.1139*** [0.0206]	0.1120*** [0.0202]	0.0808*** [0.0207]	0.0835*** [0.0254]	0.0960*** [0.0262]
Log (Plant Capital/Workers) _j		-0.0405* [0.0229]	-0.0397* [0.0225]	-0.0174 [0.0274]	-0.0320 [0.0322]	-0.0417 [0.0317]
Log (Materials/Workers) _j		-0.0279 [0.0222]	-0.0289 [0.0222]	-0.0393* [0.0229]	-0.0059 [0.0296]	-0.0129 [0.0294]
R&D intensity _j		0.0049 [0.0058]	0.0039 [0.0058]	0.0103 [0.0074]	0.0058 [0.0085]	0.0024 [0.0091]
(Value-added/Shipments) _j		-0.1050 [0.1278]	-0.1141 [0.1286]	-0.0705 [0.1294]	0.1683 [0.1587]	0.1600 [0.1573]
Log (No. of Establishments) _p			0.0574*** [0.0032]	0.0614*** [0.0037]	0.0661*** [0.0049]	0.0652*** [0.0048]
Year Started _p			0.0001 [0.0001]	0.0001 [0.0001]	0.0002* [0.0001]	0.0002** [0.0001]
Dummy: Multinational _p			0.0102** [0.0050]	0.0147** [0.0065]	0.0259*** [0.0081]	0.0286*** [0.0083]
Log (Total Employment) _p			-0.0010 [0.0016]	-0.0002 [0.0017]	-0.0007 [0.0019]	-0.0006 [0.0020]
Log (Total USD Sales) _p			0.0006 [0.0008]	0.0000 [0.0010]	0.0001 [0.0013]	0.0005 [0.0013]
Elasticity based on:	All goods	All goods	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Parent country dummies	Y	Y	Y	Y	Y	Y
Observations	316,977	316,977	286,072	206,490	144,107	144,107
No. of industries	459	459	459	305	219	219
R ²	0.0334	0.1372	0.1447	0.1511	0.2051	0.2027

Notes: The sample comprises all firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. A median cutoff dummy is used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Columns (1)-(3) use a measure based on all available HS10 elasticities from Broda and Weinstein (2006); column (4) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; column (5) further restricts this to consumption goods; column (6) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include parent country fixed effects. Columns (3)-(6) also include indicator variables for whether the reported employment and sales data respectively are estimated/missing/from the low end of a range, as opposed to being from actual data (coefficients not reported).

Table 2: Upstreamness of Integrated vs Non-Integrated Inputs: Elasticity Quintiles

Dependent variable:	Log Ratio-Upstreamness					
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Quintile 2 $Elas_j$)	-0.0209 [0.0345]	-0.0290 [0.0319]	-0.0278 [0.0314]	-0.0590 [0.0447]	-0.0802* [0.0474]	0.0634 [0.0550]
Ind.(Quintile 3 $Elas_j$)	-0.0742** [0.0336]	-0.0802** [0.0316]	-0.0782** [0.0309]	-0.0569 [0.0454]	-0.0982** [0.0429]	-0.0379* [0.0224]
Ind.(Quintile 4 $Elas_j$)	-0.0480 [0.0365]	-0.0893*** [0.0337]	-0.0881*** [0.0331]	-0.1068** [0.0459]	-0.1685*** [0.0457]	-0.0942*** [0.0259]
Ind.(Quintile 5 $Elas_j$)	-0.0588 [0.0377]	-0.0955*** [0.0325]	-0.0947*** [0.0318]	-0.1156*** [0.0420]	-0.1849*** [0.0459]	-0.1026*** [0.0317]
Log (Skilled Emp./Workers) $_j$		0.0080 [0.0238]	0.0069 [0.0239]	0.0073 [0.0290]	-0.0290 [0.0379]	-0.0215 [0.0386]
Log (Equip. Capital/Workers) $_j$		0.1127*** [0.0195]	0.1112*** [0.0192]	0.0731*** [0.0183]	0.0768*** [0.0205]	0.0949*** [0.0257]
Log (Plant Capital/Workers) $_j$		-0.0331 [0.0210]	-0.0325 [0.0207]	-0.0087 [0.0228]	-0.0240 [0.0276]	-0.0316 [0.0290]
Log (Materials/Workers) $_j$		-0.0311 [0.0222]	-0.0322 [0.0222]	-0.0397* [0.0237]	-0.0099 [0.0290]	-0.0190 [0.0317]
R&D intensity $_j$		0.0053 [0.0058]	0.0044 [0.0057]	0.0113 [0.0070]	0.0048 [0.0086]	0.0017 [0.0103]
(Value-added/Shipments) $_j$		-0.1270 [0.1295]	-0.1356 [0.1301]	-0.0840 [0.1323]	0.1725 [0.1699]	0.1453 [0.1665]
Log (No. of Establishments) $_p$			0.0570*** [0.0031]	0.0612*** [0.0037]	0.0661*** [0.0047]	0.0640*** [0.0052]
Year Started $_p$			0.0001 [0.0001]	0.0001* [0.0001]	0.0002** [0.0001]	0.0003*** [0.0001]
Dummy: Multinational $_p$			0.0105** [0.0048]	0.0125** [0.0060]	0.0192** [0.0079]	0.0304*** [0.0085]
Log (Total Employment) $_p$			-0.0003 [0.0016]	0.0004 [0.0017]	0.0005 [0.0019]	-0.0005 [0.0019]
Log (Total USD Sales) $_p$			0.0003 [0.0008]	-0.0004 [0.0009]	-0.0003 [0.0011]	-0.0001 [0.0012]
Elasticity based on:	All goods	All goods	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Parent country dummies	Y	Y	Y	Y	Y	Y
Observations	316,977	316,977	286,072	206,490	144,107	144,107
No. of industries	459	459	459	305	219	219
R ²	0.0449	0.1504	0.1580	0.1770	0.2333	0.2268

Notes: The sample comprises all firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Columns (1)-(3) use a measure based on all available HS10 elasticities from Broda and Weinstein (2006); column (4) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; column (5) further restricts this to consumption goods; column (6) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitute cases. All columns include parent country fixed effects. Columns (3)-(6) also include indicator variables for whether the reported employment and sales data respectively are estimated/missing/from the low end of a range, as opposed to being from actual data (coefficients not reported).

Table 3: Effect of Upstream Contractibility: Median Elasticity Cutoff

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Elas _j > Median)	-0.0910*** [0.0210]	-0.1306*** [0.0256]	-0.1432*** [0.0263]	-0.1372*** [0.0249]
Upstream Contractibility _j				
× Ind.(Elas _j < Median)	-0.8943*** [0.2869]	-1.1148*** [0.3838]	-1.2395*** [0.4345]	-1.2195*** [0.4363]
× Ind.(Elas _j > Median)	0.5044*** [0.1717]	1.0224*** [0.1571]	0.8871*** [0.1505]	0.9451*** [0.1415]
p-value: Q5 at median <i>UpstCont_j</i>	[0.0001]	[0.0004]	[0.0000]	[0.0000]
Elasticity based on:	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	286,072	206,490	144,107	144,107
No. of industries	459	305	219	219
R ²	0.1882	0.2609	0.2910	0.2888

Notes: The sample comprises all firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . A median cutoff dummy is used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Column (1) uses a measure based on all available HS10 elasticities from Broda and Weinstein (2006); column (2) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; column (3) further restricts this to consumption goods; column (4) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table 4: Effect of Upstream Contractibility: Elasticity Quintiles

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 $Elas_j$)	-0.0350 [0.0300]	-0.0611 [0.0396]	-0.0490 [0.0429]	0.0763** [0.0323]
Ind.(Quintile 3 $Elas_j$)	-0.1104*** [0.0288]	-0.0566 [0.0405]	-0.0683** [0.0328]	-0.0476** [0.0223]
Ind.(Quintile 4 $Elas_j$)	-0.1207*** [0.0304]	-0.1605*** [0.0292]	-0.1611*** [0.0277]	-0.1185*** [0.0236]
Ind.(Quintile 5 $Elas_j$)	-0.1409*** [0.0297]	-0.1760*** [0.0306]	-0.1643*** [0.0292]	-0.1108*** [0.0260]
Upstream Contractibility $_j$				
× Ind.(Quintile 1 $Elas_j$)	-1.5540*** [0.4934]	-1.5492*** [0.4177]	-1.8562*** [0.4446]	-0.8114 [0.5369]
× Ind.(Quintile 2 $Elas_j$)	-0.9810*** [0.3165]	-0.5723 [0.5973]	-0.6886 [0.7621]	-2.0195*** [0.6896]
× Ind.(Quintile 3 $Elas_j$)	0.3271 [0.2408]	-0.3234 [0.3742]	-0.4171 [0.3855]	0.1796 [0.1727]
× Ind.(Quintile 4 $Elas_j$)	0.3849 [0.2867]	1.0662*** [0.2319]	0.6855*** [0.2106]	0.9811*** [0.2565]
× Ind.(Quintile 5 $Elas_j$)	0.7106*** [0.2148]	1.0530*** [0.2149]	1.1171*** [0.2273]	1.0419*** [0.2275]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.0000]	[0.0005]
Elasticity based on:	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	286,072	206,490	144,107	144,107
No. of industries	459	305	219	219
R ²	0.2204	0.2792	0.3064	0.3191

Notes: The sample comprises all firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Column (1) uses a measure based on all available HS10 elasticities from Broda and Weinstein (2006); column (2) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; column (3) further restricts this to consumption goods; column (4) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table 5: Different Subsamples and Additional Controls

Dependent variable:	Log Ratio-Upstreamness					
	Single estab. (1)	MNC (2)	Robustness to additional controls			
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Quintile 2 $Elas_j$)	-0.0461 [0.0445]	-0.0870*** [0.0288]	-0.0429 [0.0414]	-0.0491 [0.0430]	-0.0492 [0.0403]	-0.0418 [0.0386]
Ind.(Quintile 3 $Elas_j$)	-0.0630* [0.0338]	-0.0787*** [0.0279]	-0.0549* [0.0305]	-0.0683** [0.0328]	-0.0532* [0.0308]	-0.0384 [0.0293]
Ind.(Quintile 4 $Elas_j$)	-0.1625*** [0.0284]	-0.1103*** [0.0268]	-0.1601*** [0.0253]	-0.1613*** [0.0277]	-0.1437*** [0.0230]	-0.1444*** [0.0213]
Ind.(Quintile 5 $Elas_j$)	-0.1638*** [0.0299]	-0.1206*** [0.0330]	-0.1546*** [0.0269]	-0.1642*** [0.0292]	-0.1666*** [0.0258]	-0.1565*** [0.0233]
Upstream Contractibility $_j$						
× Ind.(Quintile 1 $Elas_j$)	-1.8620*** [0.4612]	-1.5014*** [0.3691]	-1.6826*** [0.4083]	-1.8554*** [0.4451]	-1.6147*** [0.3643]	-1.4820*** [0.3275]
× Ind.(Quintile 2 $Elas_j$)	-0.7401 [0.8055]	0.2330 [0.3979]	-0.6775 [0.7338]	-0.6876 [0.7626]	-0.5599 [0.7994]	-0.6227 [0.7701]
× Ind.(Quintile 3 $Elas_j$)	-0.4965 [0.3919]	0.2476 [0.2838]	-0.5875 [0.3681]	-0.4186 [0.3854]	-0.4597 [0.4041]	-0.6614* [0.3966]
× Ind.(Quintile 4 $Elas_j$)	0.6749*** [0.2162]	0.5686** [0.2484]	0.5891*** [0.1714]	0.6850*** [0.2105]	0.6457*** [0.2157]	0.5434*** [0.1890]
× Ind.(Quintile 5 $Elas_j$)	1.1025*** [0.2321]	0.9941*** [0.2949]	0.9582*** [0.2165]	1.1183*** [0.2272]	1.1302*** [0.2518]	0.9516*** [0.2393]
Vertical Integration Index $_p$			-1.1296*** [0.2065]			-1.1144*** [0.2044]
Foreign integrated tr. share $_p$				-1.0690*** [0.1330]		-0.2034* [0.1214]
Log (Input Elasticity) $_j$					-0.2999*** [0.1099]	-0.2853*** [0.1024]
Wtd. Cov. of Input Elasticity $_j$ and upstreamness $_{ij}$					-0.4963*** [0.1718]	-0.4330*** [0.1555]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0030]	[0.0000]	[0.0000]	[0.0000]	[0.0000]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y	Y	Y
Subsidiary country dummies	N	N	N	Y	N	Y
Observations	117,956	2,490	144,107	144,107	144,107	144,107
No. of industries	219	199	219	219	219	219
R ²	0.2990	0.2467	0.3526	0.3079	0.3204	0.3655

Notes: Columns (1)-(2) restrict to different subsets of firms from the 2004/2005 vintage of D&B WorldBase, as described in each column heading. Columns (3)-(6) examine robustness to additional control variables. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries; the elasticity measure used is that whose construction is restricted to only the HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table 6: Alternative Constructions of Ratio-Upstreamness

Dependent variable:	Random pick	“Ever-Integrated” inputs	Mfg. inputs only	Mfg. inputs only, drop parent SIC
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 $Elas_j$)	-0.0481 [0.0428]	-0.0240 [0.0413]	-0.0385 [0.0497]	-0.0262 [0.0926]
Ind.(Quintile 3 $Elas_j$)	-0.0687** [0.0329]	-0.0402 [0.0341]	-0.0786** [0.0394]	-0.0642 [0.0514]
Ind.(Quintile 4 $Elas_j$)	-0.1574*** [0.0277]	-0.1293*** [0.0307]	-0.1825*** [0.0320]	-0.1388** [0.0661]
Ind.(Quintile 5 $Elas_j$)	-0.1652*** [0.0303]	-0.1313*** [0.0261]	-0.1762*** [0.0396]	-0.2958*** [0.0934]
Upstream Contractibility _{<i>j</i>}				
× Ind.(Quintile 1 $Elas_j$)	-1.8583*** [0.4454]	-0.8338*** [0.3137]	-2.1696*** [0.4819]	-1.1117* [0.5749]
× Ind.(Quintile 2 $Elas_j$)	-0.6960 [0.7602]	-0.8880 [0.7960]	-0.9343 [0.9046]	0.0021 [0.8379]
× Ind.(Quintile 3 $Elas_j$)	-0.4193 [0.3873]	0.0377 [0.4977]	-0.2726 [0.4890]	-1.8093* [0.9849]
× Ind.(Quintile 4 $Elas_j$)	0.6473*** [0.2126]	0.9039*** [0.3313]	0.8981*** [0.2504]	-2.5374*** [0.7379]
× Ind.(Quintile 5 $Elas_j$)	1.1816*** [0.2803]	1.3664*** [0.2992]	1.1370*** [0.3822]	-0.0754 [1.1158]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.0000]	[0.0013]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	144,107	144,107	143,846	46,992
No. of industries	219	219	219	218
R ²	0.3059	0.1950	0.3311	0.1216

Notes: The sample comprises firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The four columns use variants of the log ratio-upstreamness measure as the dependent variable, as described in the column headings and the main text. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries; the elasticity measure used is that whose construction is restricted to only the HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table 7: Integration Decisions within Firms (Top 100 Inputs): Median Elasticity Cutoff

Dependent variable:	Indicator variable: Input Integrated?				
	(1)	(2)	(3)	(4)	(5)
Upstreamness _{ij}					
× Ind.(Elas _j < Median)	-0.0080*** [0.0010]	-0.0005 [0.0019]	-0.0039* [0.0023]	0.0002 [0.0021]	0.0009 [0.0016]
× Ind.(Elas _j > Median)	-0.0103*** [0.0013]	0.0065*** [0.0016]	0.0026** [0.0012]	0.0064*** [0.0014]	0.0050** [0.0025]
Contractibility up to <i>i</i> (in prod. of <i>j</i>)					
× Ind.(Elas _j < Median)		0.0294*** [0.0067]	0.0240*** [0.0053]	0.0106* [0.0061]	0.0102** [0.0047]
× Ind.(Elas _j > Median)		0.0633*** [0.0092]	0.0426*** [0.0058]	0.0274*** [0.0055]	0.0268*** [0.0097]
Dummy: Self-SIC	0.9794*** [0.0017]	0.9700*** [0.0027]	0.9339*** [0.0084]	0.9312*** [0.0085]	0.9318*** [0.0085]
Log (Total Requirements _{ij})				0.0055*** [0.0009]	0.0055*** [0.0009]
p-value: Contractibility up to <i>i</i> , high vs low Elas _j	—	[0.0023]	[0.0127]	[0.0216]	[0.1040]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons. & α proxy
Firm fixed effects	Y	Y	Y	Y	Y
Input industry <i>i</i> fixed effects	N	N	Y	Y	Y
Observations	4,707,722	4,707,722	4,707,722	4,707,722	4,707,722
No. of parent firms	46,992	46,992	46,992	46,992	46,992
No. of <i>i</i> - <i>j</i> pairs	21,836	21,836	21,836	21,836	21,836
R ²	0.5342	0.5357	0.5594	0.5598	0.5598

Notes: Each observation is a SIC input by parent firm pair, where the set of parent firms is that from the Table 6, Column (4) regression, namely firms with primary SIC industry in manufacturing and employment of at least 20, which have integrated at least one manufacturing input apart from the output self-SIC. Manufacturing inputs ranked in the top 100 by total requirements coefficients of the SIC output industry are included. Standard errors are clustered by input-output industry pair; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is a 0-1 indicator for whether the SIC input is integrated. The “Contractibility up to *i*” measure is the share of the total-requirements weighted contractibility of inputs that has been accrued in production upstream of and including input *i* in the production of output *j*. The median cutoff dummies in columns (1)-(4) are based on the elasticity measure constructed using only those HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC; column (5) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include parent firm fixed effects, while columns (3)-(5) also include SIC input industry fixed effects.

Table 8: Integration Decisions within Firms (Top 100 Inputs): Elasticity Quintiles

Dependent variable:	Indicator variable: Input Integrated?				
	(1)	(2)	(3)	(4)	(5)
Upstreamness _{<i>ij</i>}					
× Ind.(Quintile 1 Elas _{<i>j</i>})	-0.0056*** [0.0009]	0.0005 [0.0014]	-0.0034** [0.0016]	0.0011 [0.0015]	0.0030*** [0.0011]
× Ind.(Quintile 2 Elas _{<i>j</i>})	-0.0085*** [0.0019]	-0.0001 [0.0035]	-0.0038 [0.0035]	0.0002 [0.0033]	-0.0010 [0.0027]
× Ind.(Quintile 3 Elas _{<i>j</i>})	-0.0100*** [0.0012]	-0.0001 [0.0027]	-0.0018 [0.0026]	0.0019 [0.0025]	-0.0008 [0.0046]
× Ind.(Quintile 4 Elas _{<i>j</i>})	-0.0098*** [0.0021]	0.0084*** [0.0024]	0.0024 [0.0016]	0.0064*** [0.0017]	0.0070*** [0.0019]
× Ind.(Quintile 5 Elas _{<i>j</i>})	-0.0113*** [0.0021]	0.0054* [0.0028]	0.0024 [0.0019]	0.0059*** [0.0020]	0.0060*** [0.0020]
Contractibility up to <i>i</i> (in prod. of <i>j</i>)					
× Ind.(Quintile 1 Elas _{<i>j</i>})		0.0234*** [0.0052]	0.0217*** [0.0048]	0.0108** [0.0049]	0.0157*** [0.0049]
× Ind.(Quintile 2 Elas _{<i>j</i>})		0.0339*** [0.0128]	0.0261*** [0.0093]	0.0117 [0.0100]	0.0047 [0.0073]
× Ind.(Quintile 3 Elas _{<i>j</i>})		0.0365*** [0.0082]	0.0304*** [0.0080]	0.0146* [0.0082]	0.0132 [0.0141]
× Ind.(Quintile 4 Elas _{<i>j</i>})		0.0669*** [0.0157]	0.0398*** [0.0086]	0.0239*** [0.0086]	0.0254*** [0.0088]
× Ind.(Quintile 5 Elas _{<i>j</i>})		0.0685*** [0.0134]	0.0456*** [0.0095]	0.0304*** [0.0093]	0.0322*** [0.0090]
Dummy: Self-SIC	0.9794*** [0.0018]	0.9699*** [0.0028]	0.9340*** [0.0085]	0.9313*** [0.0085]	0.9313*** [0.0085]
Log (Total Requirements _{<i>ij</i>})				0.0055*** [0.0009]	0.0054*** [0.0008]
p-value: Contractibility up to <i>i</i> , Quintile 1 minus Quintile 5	—	[0.0015]	[0.0217]	[0.0559]	[0.1000]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons. & α proxy
Firm fixed effect	Y	Y	Y	Y	Y
Input industry <i>i</i> fixed effects	N	N	Y	Y	Y
Observations	4,707,722	4,707,722	4,707,722	4,707,722	4,707,722
No. of parent firms	46,992	46,992	46,992	46,992	46,992
No. of <i>i-j</i> pairs	21,836	21,836	21,836	21,836	21,836
R ²	0.5342	0.5359	0.5594	0.5598	0.5599

Notes: Each observation is a SIC input by parent firm pair, where the set of parent firms is that from the Table 6, column (4) regression, namely firms with primary SIC industry in manufacturing and employment of at least 20, which have integrated at least one manufacturing input apart from the output self-SIC. Manufacturing inputs ranked in the top 100 by total requirements coefficients of the SIC output industry are included. Standard errors are clustered by input-output industry pair; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is a 0-1 indicator for whether the SIC input is integrated. The “Contractibility up to *i*” measure is the share of the total-requirements weighted contractibility of inputs that has been accrued in production upstream of and including input *i* in the production of output *j*. The quintile dummies in columns (1)-(4) are based on the elasticity measure constructed using only those HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC; column (5) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include parent firm fixed effects, while columns (3)-(5) also include SIC input industry fixed effects.

Table 9: Integration Decisions within Firms (Top 100 Inputs): Robustness

Dependent variable:	Indicator variable: Input Integrated?				
	Single estab.	# non-self-SIC integ. inputs ≥ 1	# integ. inputs ≥ 3	Drop self-SIC	Contractibility at i
	(1)	(2)	(3)	(4)	(5)
Upstreamness $_{ij}$					
× Ind.(Quintile 1 Elas $_j$)	0.0008 [0.0014]	0.0019 [0.0022]	0.0020 [0.0039]	0.0006 [0.0015]	0.0001 [0.0014]
× Ind.(Quintile 2 Elas $_j$)	0.0001 [0.0033]	0.0006 [0.0055]	-0.0038 [0.0098]	0.0001 [0.0035]	0.0001 [0.0033]
× Ind.(Quintile 3 Elas $_j$)	0.0027 [0.0021]	0.0041 [0.0035]	0.0047 [0.0059]	0.0009 [0.0026]	0.0017 [0.0025]
× Ind.(Quintile 4 Elas $_j$)	0.0063*** [0.0018]	0.0108*** [0.0026]	0.0115*** [0.0038]	0.0068*** [0.0017]	0.0055*** [0.0018]
× Ind.(Quintile 5 Elas $_j$)	0.0053*** [0.0018]	0.0102*** [0.0031]	0.0096** [0.0044]	0.0050** [0.0023]	0.0043** [0.0021]
Contractibility up to i (in prod. of j)					
× Ind.(Quintile 1 Elas $_j$)	0.0090* [0.0046]	0.0158** [0.0074]	0.0338*** [0.0128]	0.0071 [0.0049]	0.0094** [0.0046]
× Ind.(Quintile 2 Elas $_j$)	0.0114 [0.0098]	0.0228 [0.0166]	0.0271 [0.0292]	0.0094 [0.0105]	0.0163 [0.0108]
× Ind.(Quintile 3 Elas $_j$)	0.0151** [0.0076]	0.0236** [0.0114]	0.0426** [0.0189]	0.0082 [0.0085]	0.0176** [0.0084]
× Ind.(Quintile 4 Elas $_j$)	0.0231*** [0.0088]	0.0393*** [0.0127]	0.0527*** [0.0188]	0.0253*** [0.0086]	0.0221** [0.0093]
× Ind.(Quintile 5 Elas $_j$)	0.0270*** [0.0087]	0.0501*** [0.0141]	0.0599*** [0.0200]	0.0256*** [0.0093]	0.0206** [0.0095]
Dummy: Self-SIC	0.9333*** [0.0092]	0.9031*** [0.0114]	0.8409*** [0.0165]		0.9312*** [0.0087]
Log (Total Requirements $_{ij}$)	0.0051*** [0.0008]	0.0084*** [0.0013]	0.0136*** [0.0021]	0.0056*** [0.0008]	0.0046*** [0.0011]
p-value: Contractibility up to i , Quintile 1 minus Quintile 5	[0.0581]	[0.0282]	[0.2540]	[0.0798]	[0.2860]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Firm fixed effects	Y	Y	Y	Y	Y
Input industry i fixed effects	Y	Y	Y	Y	Y
Contractibility at i interactions	N	N	N	N	Y
Observations	3,608,516	3,001,343	700,443	4,662,172	4,707,722
No. of parent firms	36,019	29,967	6,995	46,992	46,992
No. of i - j pairs	21,836	21,835	20,223	21,633	21,836
R ²	0.5799	0.4668	0.3750	0.0791	0.5601

Notes: Each observation is a SIC input by parent firm pair, where the set of parent firms is that from the Table 6, column (4) regression, namely firms with primary SIC industry in manufacturing and employment of at least 20, which have integrated at least one manufacturing input apart from the output self-SIC. Manufacturing inputs ranked in the top 100 by total requirements coefficients of the SIC output industry are included. Standard errors are clustered by input-output industry pair; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is a 0-1 indicator for whether the SIC input is integrated. The “Contractibility up to i ” measure is the share of the total-requirements weighted contractibility of inputs that has been accrued in production upstream of and including input i in the production of output j . The quintile dummies are based on the elasticity measure constructed using only those HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent firm fixed effects and SIC input industry fixed effects. Column (5) further controls for the full set of quintile elasticity dummies interacted with the “Contractibility at i ” measure, namely the share of the total-requirements weighted contractibility of inputs accrued at stage i itself (coefficients not reported).

Table A-1: Firm Characteristics

	10th	Median	90th	Mean	Std Dev
A: Firm variables					
Number of Establishments (incl. self)	1	1	1	1.22	3.44
Number of countries (incl. self)	1	1	1	1.05	0.62
Number of integrated SIC codes	1	2	3	1.95	2.21
Year started	1948	1984	1999	1976.84	24.68
Log (Total employment)	3.045	3.807	5.557	4.088	1.080
Log (Sales in USD) (288,627 obs.)	12.522	15.202	17.059	14.803	2.573
MNCs only (6,370 obs.)					
Number of Establishments (incl. self)	2	3	17	8.48	22.74
Number of countries (incl. self)	2	2	6	3.47	3.64
Number of integrated SIC codes	2	5	17	8.10	11.88
Year started	1917	1968	1995	1960.29	33.88
Log (Total employment)	3.912	5.737	8.522	6.031	1.788
Log (Sales in USD) (5,891 obs.)	15.895	17.997	20.934	18.208	1.978
B: From Input-Output Tables					
Total Requirements coefficient	0.000006	0.000163	0.002322	0.001311	0.008026
Baseline Upstreamness measure (mean)	1.838	3.094	4.285	3.097	0.955
C: Ratio-Upstreamness measures					
Baseline (mean)	0.494	0.561	0.691	0.590	0.141
Baseline (random pick)	0.495	0.561	0.692	0.590	0.141
Ever-integrated inputs only	0.583	0.656	0.803	0.692	0.179
Manufacturing inputs only	0.548	0.633	0.798	0.657	0.174
Exclude parent sic, manufacturing only	0.590	1.100	2.128	1.269	0.625

Notes: Panels A and C are tabulated for the sample of 320,254 firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. For the Ratio-Upstreamness measures, “mean” and “random pick” refer to the treatment adopted for non-manufacturing inputs when mapping from the original IO1992 to SIC codes. Panel B is computed for all input (i) and output (j) SIC industry pairs for which j is in manufacturing (416,349 observations).

Table A-2: “Bunching” of Integrated Inputs by Quintiles of Upstreamness

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Quintile 1	0.409	0.056	0.045	0.040	0.026
Quintile 2	0.056	0.034	0.016	0.012	0.009
Quintile 3	0.045	0.016	0.024	0.014	0.010
Quintile 4	0.040	0.012	0.014	0.023	0.017
Quintile 5	0.026	0.009	0.010	0.017	0.018

Notes: Probability matrix constructed using the subset of 15,179 firms that have integrated at least two manufacturing inputs other than the parent industry self-SIC. For the a -th row and b -th column, we compute the probability that any two randomly drawn integrated manufacturing input SICs of the firm in question come from the a -th and b -th quintiles of $upst_{ij}$ values, where j is the SIC output industry of the firm and the quintiles are taken over all SIC manufacturing inputs i . A simple average of the probabilities across all 15,179 firms is reported.

Table A-3: Industry Characteristics

	10th	Median	90th	Mean	Std Dev
Import demand elasticity (all codes)	2.300	4.820	20.032	8.569	10.181
Import demand elasticity (BEC cons. & cap.)	1.983	4.500	20.289	8.819	11.722
Import demand elasticity (BEC cons. only)	2.000	4.639	15.992	8.366	11.881
BEC cons. import demand elasticity minus α proxy	-9.086	-4.266	7.783	-1.294	12.314
Log (Skilled Emp./Workers)	-1.750	-1.363	-0.778	-1.308	0.377
Log (Equip. Capital/Workers)	2.869	4.043	5.163	4.039	0.867
Log (Plant Capital/Workers)	2.517	3.302	4.524	3.426	0.755
Log (Materials/Workers)	3.898	4.596	5.681	4.702	0.726
R&D intensity: Log (0.001+ R&D/Sales)	-6.908	-6.097	-3.426	-5.506	1.463
Value-added/Shipments	0.357	0.518	0.660	0.514	0.119
Contractibility	0.091	0.362	0.816	0.410	0.265
Upstream Contractibility	-0.069	0.018	0.101	0.015	0.069

Notes: Summary statistics taken over the 459 SIC manufacturing industries, except for: (i) the “BEC cons. & cap.” elasticity, which is available for only 305 industries; and (ii) the “BEC cons. only” elasticity, which is available for 219 industries. The “contractibility” and “upstream contractibility” measures are based on the Rauch (1999) “conservative” classification; both homogeneous and reference-priced products are considered to be contractible.

Table A-4: Parent Firms with Multiple SIC Output Activities

Dependent variable:	Log Ratio-Upstreamness					
	Single mfg. output SIC		Most downstream mfg. output SIC		Firm by mfg. output SIC (two-way cluster)	
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Quintile 2 $Elas_j$)	-0.0419 [0.0464]	0.0937*** [0.0325]	-0.0385 [0.0413]	0.0637* [0.0343]	-0.0476 [0.0428]	0.0687** [0.0331]
Ind.(Quintile 3 $Elas_j$)	-0.1021*** [0.0292]	-0.0279 [0.0235]	-0.0218 [0.0458]	-0.0604** [0.0242]	-0.0362 [0.0398]	-0.0629*** [0.0228]
Ind.(Quintile 4 $Elas_j$)	-0.1521*** [0.0305]	-0.1046*** [0.0237]	-0.1453*** [0.0292]	-0.1180*** [0.0261]	-0.1642*** [0.0256]	-0.1280*** [0.0247]
Ind.(Quintile 5 $Elas_j$)	-0.1521*** [0.0306]	-0.0902*** [0.0270]	-0.1479*** [0.0315]	-0.1120*** [0.0287]	-0.1680*** [0.0286]	-0.1230*** [0.0264]
Upstream Contractibility $_j$						
× Ind.(Quintile 1 $Elas_j$)	-1.9121*** [0.4691]	-0.3997 [0.4808]	-1.5441*** [0.4565]	-1.2602* [0.6518]	-1.7766*** [0.4150]	-1.2353** [0.6099]
× Ind.(Quintile 2 $Elas_j$)	-0.7892 [0.7723]	-2.1371*** [0.6992]	-0.4465 [0.6290]	-1.6280** [0.6543]	-0.5588 [0.7887]	-1.8332*** [0.7036]
× Ind.(Quintile 3 $Elas_j$)	0.1059 [0.2068]	0.0528 [0.1843]	-0.8717 [0.6070]	0.3641** [0.1795]	-0.8416 [0.5438]	0.3168* [0.1813]
× Ind.(Quintile 4 $Elas_j$)	0.6619*** [0.2346]	1.0140*** [0.2784]	0.6961*** [0.2113]	0.9758*** [0.2560]	0.6808*** [0.2039]	0.9299*** [0.2512]
× Ind.(Quintile 5 $Elas_j$)	1.1166*** [0.2104]	1.0096*** [0.2307]	1.2292*** [0.2641]	1.1572*** [0.2593]	1.1637*** [0.2544]	1.1215*** [0.2371]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0067]	[0.0001]	[0.0017]	[0.0000]	[0.0001]
Elasticity based on:	BEC cons.	BEC cons. & α proxy	BEC cons.	BEC cons. & α proxy	BEC cons.	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y	Y	Y
Observations	97,174	97,174	146,844	146,844	211,232	211,232
No. of industries	219	219	219	219	—	—
R ²	0.3308	0.3471	0.2647	0.2807	0.2881	0.3023

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Columns (1) and (2) restrict the sample to those firms with at least 20 employees that report only one SIC manufacturing output activity, this being their primary SIC industry; robust standard errors clustered by output industry are reported. For columns (3) and (4), we designate as the output industry the SIC manufacturing activity of the firm that has the smallest upstreamness value with respect to final demand, for the purposes of constructing the ratio-upstreamness measure; robust standard errors clustered by this output industry are reported. For columns (5) and (6), each observation is a firm by SIC output activity pair. The ratio-upstreamness variable is constructed treating in turn each SIC manufacturing activity as the output industry for the firm in question; two-way clustered standard errors – by firm and by SIC output activity – are reported. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries; the elasticity measure used in even-numbered columns is that whose construction is restricted to only the HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC; odd-numbered columns use the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table A-5: Single-Plant Firms Only

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 $Elas_j$)	-0.0349 [0.0306]	-0.0568 [0.0414]	-0.0461 [0.0445]	0.0811** [0.0327]
Ind.(Quintile 3 $Elas_j$)	-0.1103*** [0.0290]	-0.0521 [0.0414]	-0.0630* [0.0338]	-0.0454** [0.0221]
Ind.(Quintile 4 $Elas_j$)	-0.1224*** [0.0307]	-0.1606*** [0.0299]	-0.1625*** [0.0284]	-0.1214*** [0.0235]
Ind.(Quintile 5 $Elas_j$)	-0.1419*** [0.0301]	-0.1744*** [0.0313]	-0.1638*** [0.0299]	-0.1104*** [0.0259]
Upstream Contractibility $_j$				
× Ind.(Quintile 1 $Elas_j$)	-1.6348*** [0.5101]	-1.5595*** [0.4312]	-1.8620*** [0.4612]	-0.7765 [0.5300]
× Ind.(Quintile 2 $Elas_j$)	-1.0054*** [0.3304]	-0.6377 [0.6395]	-0.7401 [0.8055]	-2.0780*** [0.7329]
× Ind.(Quintile 3 $Elas_j$)	0.3175 [0.2436]	-0.3535 [0.3828]	-0.4965 [0.3919]	0.1444 [0.1743]
× Ind.(Quintile 4 $Elas_j$)	0.3935 [0.2919]	1.0661*** [0.2392]	0.6749*** [0.2162]	0.9702*** [0.2619]
× Ind.(Quintile 5 $Elas_j$)	0.7122*** [0.2182]	1.0384*** [0.2208]	1.1025*** [0.2321]	1.0298*** [0.2299]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.0000]	[0.0005]
Elasticity based on:	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	233,529	168,617	117,956	117,956
No. of industries	459	305	219	219
R ²	0.2163	0.2699	0.2990	0.3166

Notes: The sample comprises only single establishment firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the baseline log ratio-upstreamness measure described in Section 3.2. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used to produce good j . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Column (1) uses a measure based on all available HS10 elasticities from Broda and Weinstein (2006); column (2) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; column (3) further restricts this to consumption goods only; column (4) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).