Internalizing Global Value Chains: A Firm-Level Analysis *

Laura Alfaro
HBS

Pol Antràs
Harvard

Davin Chor
NUS

Paola Conconi
ECARES

April 2015

Preliminary Version.

Abstract

In recent decades, technological progress and falling trade barriers have allowed firms to slice up their value chains, retaining within their boundaries and in their domestic economies only a subset of their production stages. A key question facing firms worldwide is figuring out which segments of the value chain are more profitably offshored, outsourced, or both. Building on Antràs and Chor (2013), we describe a property-rights model in which the organization of a firm's manufacturing process is shaped by characteristics of the different stages of production and their position in the value chain. To assess the evidence, we use the WorldBase dataset, which contains plant-level information on the production activities of firms located in a large set of countries. We combine this information with Input-Output data to construct firm-level measures of the upstreamness of integrated and non-integrated stages. In line with the model's predictions, we find that whether a firm integrates suppliers located upstream or downstream in the value chain depends crucially on the size of the elasticity of demand faced by the firm. Moreover, we show that the propensity of a firm to internalize a given stage in the value chain is shaped by the contractibility of stages of production upstream from that stage relative to the contractibility of stages of production downstream from that stage.

JEL classifications: F14, F23, D23, L20.

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

*We thank seminar participants at the "Global Fragmentation of Production and Trade Policy" conference in Brussels, HKUST, HKU, NUS, and Singapore Management University for useful comments. All errors are our own.

1 Introduction

Responding to improvements in communication technologies, competition pressures, and reduced trade barriers, firms have increasingly sliced up their value chains, retaining within firm boundaries and in their domestic economies only a small subset of the stages in these value chains. A key question facing firms worldwide is figuring out which segments of the value chain are more profitably offshored, outsourced, or both.

The evolution of the semiconductors in the last fifty years exemplifies these trends. The first semiconductor chips were manufactured in the United States by vertically integrated firms such as IBM and Texas Instruments. These firms kept within firm boundaries the design, fabrication, assembly, and testing of integrated circuits. Starting in the 1960s, these firms began offshoring the assembly of microchips to wholly-owned producers in emerging Asian economies, which would then ship the assembled chips back to the United States for testing (see Brown and Linden, 2005). In subsequent years, assembly was largely outsourced to independent contractors in Asia, who also began performing the final testing stage before distribution. Beginning in the 1970s, the semiconductors industry witnessed a new wave of offshoring in Asia, this time affecting the more upstream wafer fabrication stage. A significant share of fabrication plants offshore are independently-owned 'pure-play' foundries. In the last two decades, the industry has undergone a third wave of offshoring, with the most upstream design stage now being performed in several Asian locations, though to a large extent in wholly-owned subsidiaries.

Although the disintegration of production has featured prominently in the trade literature, much less attention has been placed on the importance of the position of a given production stage in the value chain for firm boundary choices. Only in recent years have theoretical frameworks been developed that highlight the role of the sequentiality of production for the global sourcing decisions of firms. Recent papers in this literature 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).[1] Although some of these papers provide suggestive empirical evidence, these contributions study firms' organizational decisions mostly from a theoretical perspective. To a large extent, this theoretical bias in the literature is explained by the challenges one faces when taking these models to the data. Ideally, one would like to access comprehensive datasets that would allow the researcher to track the flow of goods within value chains across borders and organizational forms. Trade statistics are useful in tracking the flows of goods when they cross a particular border, and some countries' customs offices, such as the U.S. Bureau of Customs and Border Protection, also record whether goods flow in or out of a country within or across firm boundaries. Nevertheless, once a good leaves a country, it is virtually impossible with 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 production activities of

[1] 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).

firms operating in many countries and industries can be combined with information from standard Input-Output tables to study firm boundaries along value chains. A key advantage of our firm-level approach is that it allows us to study how the internalization of the different stages of a *specific* firm's manufacturing process are shaped by the characteristics of these different stages, and most notably, by their position in the value chain. In more precise terms, the richness of our data allows us to run specifications that exploit variation in organizational features *within* firms and across their various manufacturing stages. Available theoretical frameworks of sequential production are highly stylized and do not feature asymmetries across production stages other than in their position in the value chain, so a second contribution of this paper is to develop a richer framework of firm behavior that can guide an empirical analysis using firm-level data.

On the theoretical side, we extend the property-rights theory of the organization of production in Antràs and Chor (2013).[2] We focus on the problem of a firm controlling the production process associated with a final-good manufacturing variety that consumers demand according to a constant price elasticity demand schedule. The manufacture of the final good entails a large number of production stages that need to be performed in a predetermined order. The different stages 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 physical output of the final good is determined by a function that is isomorphic to a constant elasticity of substitution (CES) production 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 it. A key organizational decision faced by the final-good producer is which (if any) suppliers along the value chain to own and which to outsource components from. As in Grossman and Hart (1986), the internalization 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.

Initially, we follow Antràs and Chor (2013) and develop a baseline model that isolates the role of "downstreamness" on organizational choices by treating all production stages as symmetric on the technology, contracting and cost sides. In this setting, organizational decisions have spillovers along the value chain, because relationship-specific investments made by upstream suppliers affect the incentives of suppliers in downstream stages. This symmetric model delivers the result that whether the incentives of the firm to integrate suppliers are higher or lower for relatively upstream suppliers depends crucially on the relative size of the elasticity of demand faced by the firm and the elasticity of substitution across production stages. When demand is elastic or inputs are not particularly substitutable, the firm finds it optimal to integrate only the most downstream stages,

[2] The 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 (2014) for a comprehensive overview of this literature.

while when demand is inelastic or inputs are sufficiently substitutable, only the most upstream stages can be profitably integrated.

We next consider a richer environment with additional sources of variation across inputs. More specifically, we allow for a non-symmetric CES technology aggregating inputs (reflecting, for instance, that some inputs might be more value-enhancing than others), as well as heterogeneity 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 in 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 internalize a given stage in the value chain is shaped by the contractibility of stages of production upstream from that stage relative to the contractibility of stages of production downstream from that stage. Intuitively, in production processes that feature a high degree of contractibility among upstream relative to downstream inputs, firms have less need to rely upon 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 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.

To assess the validity of the model's predictions, we employ the WorldBase dataset of Dun and Bradstreet (D&B), which provides detailed plant-level information for public and private companies in many countries and territories. For each plant, the dataset includes information about its primary production activity and its secondary activities. Plants belonging to the same firm can be linked via information on domestic and global parents using their DUNS numbers, this being a unique identifier.[3] Our main sample consists of around 120,000 manufacturing parent firms in 89 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 combine information on firms' production activities with input-output tables (see also Acemoglu *et al.*, 2009; and Alfaro *et al.*, 2013). To capture the position of different industries along the value chain, we compute measures of the upstreamness for each input i and output j industry pair using U.S. Input-Output Tables. This extends to the bilateral industry pair level the measure for the upstreamness of an industry with respect to final demand from Fally (2012) and Antràs *et al.* (2012). To test the validity of these predictions, we exploit information from WorldBase on the primary activity of each firm and use estimates of demand elasticities from Broda and Weinstein (2006), as

[3] D&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.

well as measures of production contractibility from Nunn (2007).

We first examine how firms' organizational choices depend on the elasticity of demand for their final goods.[4] In line with the first prediction of the model, we find that the higher is the elasticity of demand faced by a (parent) firm, the lower is the average upstreamness of its integrated suppliers versus the upstreamness of its non-integrated suppliers. This result is robust to controlling for different firm characteristics of parent firms (e.g., size, age, location, multinational status), using different elasticity and upstreamness measures, and different samples of firms. It also continues to hold when we exploit only within-firm variation in firm boundaries. In these within-firm specifications, we find that a firm's propensity to outsource is generally larger for upstream inputs, but disproportionately so for firms facing high demand elasticities. Our results are very robust and can be illustrated in a simple (unconditional) form in Figure 1. As seen in the left-panel of the figure, the average upstreamness of integrated stages is much higher when the parent company belongs to a sector with a low demand elasticity than when it belongs to a sector associated with a high demand elasticity. This is true regardless of whether, when computing the average, one includes or not integrated stages belonging to the same sector as the parent, which tend to have a low upstreamness measure. Conversely, the right-panel of the figure shows that the average upstreamness of non-integrated stages appears to be higher, the higher is the elasticity of demand faced by the parent company.[5]

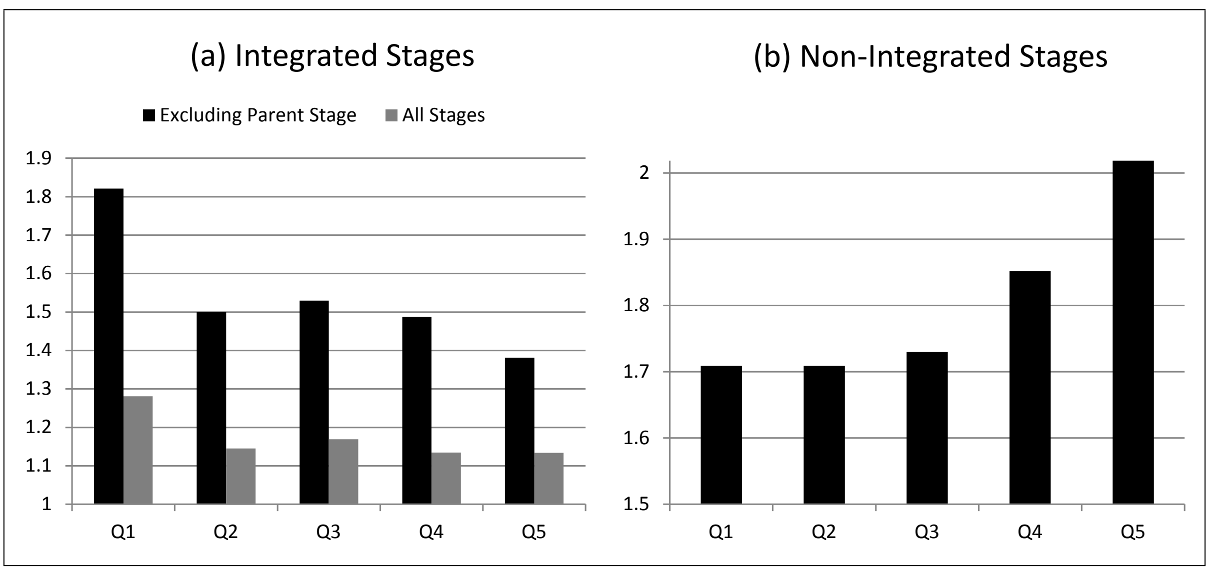

Figure 1: Average Upstreamness of Production Stages, by Quintile of Parent's Demand Elasticity

After documenting this fact from a variety of angles, we turn to studying how firms' ownership

[4] Although the theory highlights the importance of the size of the elasticity of demand faced by the firm *relative* to the elasticity of substitution across its production stages, we lack estimates of the latter elasticity.

[5] To be more precise, Figure 1 is plotted using only input stages i that rank within the top hundred manufacturing inputs in terms of total requirements coefficients in the production of the parent output industry j. Furthermore, the plotted averages are computed weighting each input by its total requirement coefficient in the production of j. The figures obtained when considering all manufacturing inputs or when computing unweighted averages are qualitatively very similar.

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, and this effect is larger when the firm faces a higher elasticity of demand. This results is in line with the second prediction of our theoretical model, according to which a greater level of upstream contractibility reduces a firm's need to rely on decisions over organizational mode to elicit the right incentives from suppliers. 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. intra-firm import shares and lacked direct information on the U.S. entity internalizing these foreign purchases. In that respect, our paper is closely related to two recent (and contemporaneous) papers with similar goals in mind. 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 and that of the parent firm itself. They find that this correlation varies depending on the size of the demand elasticity faced by the parent firm, in a manner consistent with the predictions of the theory in Antràs and Chor (2013). Their work is however silent on the production line position of non-integrated inputs. 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 measures). 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.[6]

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. A recent influential paper by Atalay *et al.* (2014) finds little evidence of significant commodity shipments across plants in U.S. domestic firms.[7] This finding has been interpreted as indicating that firm boundaries are shaped by issues related to the transfer of intangible inputs, rather than of physical goods. As Antràs (2015) has stressed, however, contractual frictions in the exchange of physical inputs are much more likely to be relevant in international transactions than in domestic ones, particularly when those domestic transactions involve U.S. producers (as in Atalay *et al.*'s work). It is thus plausible that firm boundaries might be shaped by different factors in our sample of firms from around the world. Moreover, even if some of the internalizations we identify in the D&B data are in fact unrelated to 'vertical' motives, it is not obvious that this would bias our estimates in any particular direction, since this would constitute measurement error in our key endogenous variable. Of course, the existence of alternative theories of internalization has some incidence on the overall power of the tests we perform, but the patterns we unveil are rather subtle and not easily explained

[6] Naturally, our paper is related more broadly to the empirical literature on firm boundaries, which is nicely overviewed in Lafontaine and Slade (2007) and Bresnahan and Levin (2012).

[7] A related study by Ramondo *et al.* (2014) documents that the bulk of intra-firm trade between foreign affiliates and U.S. parent is concentrated among a small number of large foreign affiliates.

(to the best of our knowledge) by alternative theories.

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. Section 5 presents our baseline empirical results, and includes several additional robustness checks. Section 6 concludes.

2 Theoretical Framework

In this section, we develop our model of sequential production. We will 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: Isolating the Role of Downstreamness

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-level 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.

As in Antràs and Chor (2013), 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:

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

where θ is a firm-specific productivity parameter, $\alpha \in (0, 1)$ 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.[8]

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 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 internalization affect the game played between the firm and the unit measure of suppliers? As in Antràs and Chor (2013), and following the property-rights theory of firm boundaries, we let the *effective* bargaining power of the firm vis-à-vis a particular supplier depend on whether the firm owns this supplier or not. 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

[8] In 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 $x(i)$.

that tilts the division of surplus in its favor. As in Antràs and Chor (2013), 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 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 their analysis by exploring a richer framework with several sources of asymmetry across production stages, most notably incorporating heterogeneity in the degree of contractibility of inputs along the value chain.

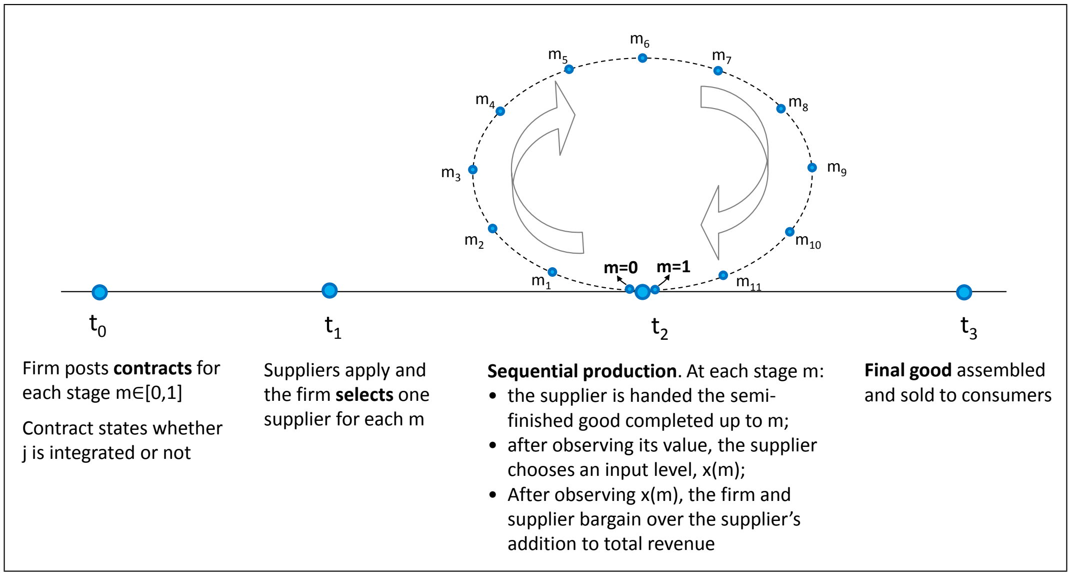

Figure 2: Timing of Events

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

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

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 sale 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, which applying Leibniz' rule to (3) is given by

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

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

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

Notice that the marginal return to investing in $x(m)$ is increasing in the demand level, A, firm productivity θ, 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 remember is what the firm and supplier m bargain over – might 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 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 (consistently 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, Antràs and Chor (2013) show that the problem of choosing the optimal organizational structure can be reduced to the program:

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

As mentioned above, we will not provide a full characterization 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). Suffices to say, at this point, that the program in (6) can be written as a fairly standard calculus of variations problem, from which 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.

Proposition 1 indicates that the optimal pattern of ownership along the value chain depends critically on whether production stages are sequential complements or substitutes. When the demand faced by the final-good producer is sufficiently elastic, then there exists a unique cutoff production stage such that all stages prior to this cutoff are outsourced, while all stages (if any) after that threshold are integrated. Intuitively, when inputs are sequential complements, the firm chooses to forgo control rights over upstream suppliers in order to incentivize their investment effort, since this generates positive spillovers on the investment decisions to be made by downstream suppliers. When demand is instead sufficiently inelastic, the converse prediction holds: it is optimal to integrate relatively upstream stages, and if outsourcing is observed along the value chain, it necessarily occurs relatively downstream.[9]

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 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.

[9] Antràs and Chor (2013) further show that the thresholds m_C^* and m_S^* can be solved for in closed form as a function of the bargaining shares β_V and β_O and the key parameters ρ and α, and they show that a decrease in ρ necessarily expands the range of stages that are vertically integrated (i.e., m_C^* increases and m_S^* decreases with ρ). The intuition for this result is that when the firm has a relatively high market power (low ρ), it will tend to place a relatively high weight on the rent-extraction motive for integration and will thus be less concerned with the investment inefficiencies caused by such integration.

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 and to the location of input 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

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

where $\psi\left(i\right)$ 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 $\psi\left(i\right)$ 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 $\psi\left(i\right)^{\alpha}$. The program related to the optimal investment choice of supplier at stage m is similarly modified to:

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

The first-order condition associated with this program, relates investment $x(m)$ at stage m with the bargaining share as well as the revenue $r\left(m\right)$ secured by suppliers upstream from m. When plugging this expression for $x(m)$ into $r'(m) = \frac{\rho}{\alpha} \left(A^{1-\rho}\theta^{\rho}\right)^{\frac{\alpha}{\rho}} r(m)^{\frac{\rho-\alpha}{\rho}} \psi\left(m\right)^{\alpha} x(m)^{\alpha}$, we show in the Appendix that one obtains a separable differential equation for $r\left(m\right)$ with solution:

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

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 $\beta\left(i\right)$ – along the value chain. As shown in the Appendix, this reduces to a simple generalization of program (6):

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

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

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. For instance, we will analyze 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

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

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

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

In the Appendix, 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:

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

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.[10] 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. As emphasized by Antràs and Chor (2013) this is true regardless of the path of $\psi(i)$ and $c(i)$.

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 input m in the total input purchases of the firm is a monotonically increasing function of $\psi(m)$. According to one of the canonical results of the property-rights literature (cf., Grossman

[10] Antràs 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).

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 to be, one would then be led to conclude that if the path of $\psi(m)$ were to be sufficiently increasing in m – or the path of $c(m)$ were to be 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.[11] Equation (13) illustrates, in particular, that the incentives to integrate a particular input will be notably shaped by the size of the ratio $\psi(m)/c(m)$ 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 $\psi(m)/c(m)$ for inputs upstream from m relative to its value for inputs downstream from m, the higher will the incentive of the firm to integrate stage m.

The intuition behind this result is as follows. Remember that when inputs are sequential 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 $\psi(i)/c(i)$ is common for all stages. Conversely, whenever $\rho < \alpha$, investments are sequentially 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 the Appendix 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 com-

[11] This is 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 at that period.

plements 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 the Appendix, integration and outsourcing coexist along the value chain provided that $\beta_V (1-\beta_V)^{\frac{\alpha}{1-\alpha}} > \beta_O (1-\beta_O)^{\frac{\alpha}{1-\alpha}}$.[12] In such a case, the threshold m_C^* is given by

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

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.

Another way to illustrate this result is by considering the case in which the ratio $\psi(m)/c(m)$ increases or decreases at constant rate $\lambda \in \mathbb{R}$ along the value chain. Note that λ is allowed to be negative and that a lower λ implies that upstream investments are relatively more important than downstream investments. In such a case, we have that $\psi(m)/c(m) = \varkappa e^{\lambda m}$ for some $\varkappa > 0$, and (14) reduces to

$$\frac{e^{\frac{\lambda\alpha}{1-\alpha} m_C^*} - 1}{e^{\frac{\lambda\alpha}{1-\alpha}} - 1} = \left\{ 1 + \left(\frac{1-\beta_O}{1-\beta_V}\right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1-\frac{\beta_O}{\beta_V}}{1-\left(\frac{1-\beta_O}{1-\beta_V}\right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}.$$

The left-hand-side of this expression is clearly increasing in m_C^*, and it can also be shown to be decreasing in λ. As a result, m_C^* is lower and integration is more prevalent, the lower is λ, consistently with the last statement of Proposition 2. In the Appendix, we follow similar steps to show that the threshold m_S^* in the substitutes case is also decreasing in λ, implying that for $\rho < \alpha$, integration is *less* prevalent whenever λ is high.

2.3 Contractibility and the Integration Decision

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 attempt 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

[12] When 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 find it optimal to outsource *all* production 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 comove 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 at different stages of the value chain. At the end of this section, 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.[13] 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 of the model is identical to the one of our previous model in which $\psi(m)$ was assumed exogenous. This implies that the firm's optimal ownership structure along the value chain will

[13] Acemoglu, 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.

seek to maximize the program in (10), and the solution of this problem will be characterized by Proposition 2.

As shown in the Appendix, after some cumbersome manipulations, one can further show, after solving for the optimal choice of $\beta(m) \in \{\beta_V, \beta_O\}$, that firm profits net of contracting costs can be expressed as

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

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

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

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 in turn implies that:

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.

By mapping variation in $\psi(m)$ to the degree of input contractibility, Proposition 3 helps operationalize our previous, more abstract Proposition 2. More specifically, in the next section we will employ empirical proxies for input contractibility to develop a sector-level measure of the extent to which non-contractibilities features 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.

When introducing heterogeneity in contractibility, we have made the strong assumption that the marginal cost faced by firms is common for all stages. How important is this restrictive assumption for the predictions in Proposition 3? In the Appendix, we show that in the presence of variation in marginal costs along the value chain, equations (15) and (16) are only slightly modified and whether the ratio $\psi(m)/c(m)$ is higher for upstream or downstream stages will continue to be monotonically related to the degree of contractibility $\mu(m)$ for upstream inputs relative to downstream inputs except in cases in which marginal costs are positively correlated with contractibility along the value chain. We also show in the Appendix, however, that it is much more plausible to expect a *negative* correlation between stage contractibility and stage marginal costs. In particular, when thinking about the global sourcing decisions of firms as shaping the path of marginal costs along the value chain, we find that the marginal benefit from marginal cost reductions is particularly high for stages with high contractibility levels, regardless of the position of those stages in the value chain. In sum, the ratio $\psi(m)/c(m)$ will tend to be large for stages with high levels of contractibility $\mu(m)$, thus justifying our approach in the empirical analysis.[14]

3 Dataset and Key Variables

We turn now to our empirical evidence. The key object we aim to measure in the data is the relative propensity of firms to integrate or outsource inputs at different positions in the value chain. For that purpose, we need firm-specific measures of input integration and outsourcing, as well as an index of the 'upstreamness' of these various inputs. To assess the validity of our model, we also need proxies for whether a final-good industry is characterized by the complements or substitutes case, as well as a measure of input contractibility. In this section, we discuss the dataset that we employ to identify integrated inputs and the construction of our key variables.[15]

3.1 The WorldBase Dataset

In order to construct measures of internalization in value chains, we use data from Dun & Bradstreet's (D&B) WorldBase, a dataset covering public and private companies across more than 100 countries and territories. The WorldBase dataset has been used extensively in the literature, in particular to explore research questions related to the organizational practices of firms around the

[14] Due to data limitations, we will not able to empirically study how the integration decision of firms is affected by whether firms face relatively higher production costs upstream or downstream. More specifically, although we could construct marginal cost proxies based on the location of the firm's integrated plants, we lack information on where outsourced components are sourced from.

[15] We postpone the discussion of some auxiliary control variables until we detail our empirical specifications in the next section.

world.[16] Cross-country empirical investigations at the firm level are challenging due to data limitations, as there are few high-quality datasets that are comparable across countries; when such data are available, these tend to be limited mostly to advanced countries.[17] One of the advantages of WorldBase is its coverage of a wide sample of countries at different development levels. Another advantage is that the unit of observation is the establishment/plant, namely a single physical location where industrial operations or services are performed, or business is conducted. Each establishment in WorldBase comes with a unique identifier, called a DUNS number, as well as a name and address. The data also reports ownership linkages across establishments, allowing us to link each plant to a domestic or global ultimate owner.

To be more specific, we use several different categories of information recorded in WorldBase:

1. Industry: the 4-digit SIC code of the primary industry in which each establishment operates, and for most countries, the SIC codes of as many as five secondary industries, listed in descending order of importance.[18]

2. Ownership: the identity of the establishment's global parent, if any.

3. Location: country, state, city, and street address of each plant.

4. Additional information: sales, employment and age of each plant.

Our main sample is drawn from the 2004/2005 WorldBase dataset vintage. We focus on manufacturing firms, namely observations that are identified as global parents (or in the terminology of WorldBase, "global ultimates"), whose primary SIC code lies between 2000 and 3999, since these are in principle the parent firms that would best fit our theory of sequential production processes. This yields a sample of 116,843 global parents from 89 countries. Of these, 6,983 observations can be classified as multinational firms, these being those global parents that have establishments in more than one country. Note that parents are defined in the data as entities that have legal and financial responsibility for another firm. We link these global parents to their identified majority-owned subsidiaries, both in manufacturing and in non-manufacturing, using the reported DUNS number of the global parent reported by establishments.[19]

[16] Early examples include Caves' (1975) analysis of size and diversification patterns between Canadian and US plants. More recent uses include Harrison *et al.* (2004), Acemoglu *et al.* (2009), Alfaro and Charlton (2009), Alfaro *et al.* (2013), Alfaro and Chen (2014), and Fajgelbaum *et al.* (2015).

[17] Economic censuses of firms are infrequently collected due to high costs and institutional restrictions, especially in poor countries. No single institution has the capacity or resources to undertake the primary collection of such census data for a wide range of countries and periods. Researchers thus have turned to other sources, such as business "compilations" (registries, tax sources) or surveys. Compared to other international datasets, the data in WorldBase are compiled from a large number of sources (e.g., partner firms, telephone directory records, websites, self-registration). See Alfaro and Charlton (2009) for a more detailed discussion of the WorldBase data and comparisons with other data sources.

[18] D&B uses the United States Government Department of Commerce, Office of Management and Budget, Standard Industrial Classification Manual 1987 edition to classify business establishments.

[19] It turns out that not all establishments in WorldBase report a global parent. A large number of observations are actually single establishment firms that are self-contained entities. We exclude these from the analysis as our model likely serves as a better description of large firms with complex production processes.

The upper panel of Table 1 provides some descriptive statistics for the set of global parents. A reasonably large share of these observations that are labeled as "global ultimates" in WorldBase are actually linked to only one establishment, corresponding to the parent firm itself. The median number of SIC codes under the umbrella of a global parent and all its subsidiaries is also small (only two). We will nevertheless show that our core results are robust in subsamples of the data comprising larger global parents, as defined by the number of employees, subsidiaries or SIC codes.

[Table 1 about here]

3.2 Key Variables

Integrated and Outsourced Inputs

We begin by describing how we combine the above information from the WorldBase dataset with information on industry production linkages contained in Input-Output (I-O) tables to identify which inputs are integrated and outsourced by each firm (see also Acemoglu *et al.*, 2009; and Alfaro *et al.*, 2013).

It helps at this juncture to introduce some notation to fix ideas. Consider a setting with $N > 1$ industries. We denote *output* industries with j and *input* industries with i. For each pair of industries $1 \leq i, j \leq N$, the I-O accounts report the dollar value of i used directly as an input in the production of a dollar's worth 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, however, each input i is used not just directly, but could also enter further upstream, more than one stage removed from the actual production of j. The total dollar value of i used either directly or indirectly to produce a dollar's worth of j is called the total requirements coefficient, tr_{ij}, and we view this as reflecting the overall importance of the said input for the production of j. As is well known, tr_{ij} is given by the (i, j)-th entry of $[I - D]^{-1}D$; here, I is the identity matrix and $[I - D]^{-1}$ is the Leontief inverse matrix.

In our empirical analysis, we use the primary SIC code reported in WorldBase by each parent firm p to identify its relevant final good j. We then use the information from I-O Tables to deduce the set of 4-digit SIC inputs $S(j)$ that are used (either directly or indirectly) in the production of j, namely: $S(j) = \{i : tr_{ij} > 0\}$.

We can then identify which inputs are integrated and which inputs are outsourced by each firm. We define $I(p) \subseteq S(j)$ to be the set of integrated inputs of parent firm p. These are identified in WorldBase by all the primary and secondary SIC codes of the parent firm p and all its subsidiaries. These represent the set of inputs that the firm can in principle obtain within its ownership boundaries. Finally, we define the complement set $NI(p) = S(j) \setminus I(p)$ to be the set of non-integrated SICs for a parent p whose primary output industry is j. (Note that under the above construction, the primary SIC activity of the parent firm j is automatically classified as an element of $I(p)$; we will later show our results to be robust to excluding this "self-SIC" code from the set of integrated inputs.)

Upstreamness

We next turn to discuss how we employ I-O tables to calculate a measure of the (average) 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 compute the expression:

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

Focusing first on the denominator in (17), recall that dr_{ij} is the value of i that enters directly in (i.e., one stage immediately 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 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 actual production of j (where $n = 1, 2, \ldots, \infty$). 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 simply a weighted-average of the number of stages it takes for input i to enter in the production of j, where the weights correspond to the share of tr_{ij} that enters at that corresponding upstream stage of production. 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.

Several remarks are in order. First, 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. Second, while $upst_{ij}$ is built on a logic analogous to the upstreamness measure introduced in Fally (2012) and Antràs *et al.* (2012), the two measures are distinct. This earlier measure sought to capture the average production line position of each industry i with respect to final demand (i.e., consumption and investment), whereas $upst_{ij}$ in this paper instead captures the position of input i with respect to the specific output industry j. What we now construct is therefore a measure of production staging specific to each input-output industry pair.[20] Third, one can show with some matrix algebra 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} provided earlier, it is then straightforward to calculate $upst_{ij}$ given information from I-O Tables on the direct requirements matrix, D.

To perform this calculation, we turn to the 1992 U.S. Benchmark I-O Tables from the Bureau of Economic Analysis, specifically the table on the Use of Commodities by Industries after Redefinitions (in producer prices) for information on input requirements. We utilize the U.S. Tables, as these are one of the few publicly-available national I-O accounts that provide a level of industry detail close to the SIC 4-digit level codes found in WorldBase. 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

[20] The $upst_{ij}$ measure is also related to the concept of "average propagation length" in Dietzenbacher *et al.* (2005), which captures the average number of stages taken by a shock to output in industry i to spread to industry j.

SIC system.[21] We first compute tr_{ij} and $upst_{ij}$ for each 6-digit I-O industry pair, before mapping these to 4-digit SIC codes.[22] Readers familiar with the U.S. I-O Tables will be aware that this concordance is not a one-to-one key. This multiple matching problem is however not a major problem for us given our focus on global parents whose primary output j is in manufacturing, since the key assigns a unique I-O industry to each 4-digit SIC code between 2000 and 3999. The mapping concerns apply in our setting only to non-manufacturing inputs i, as multiple 6-digit I-O codes can be assigned to a single SIC 4-digit code. We use a variety of approaches to deal with such cases, by taking: (i) the simple mean of $upst_{ij}$ over constituent I-O codes for the SIC input industry; (ii) the median value; (iii) a random pick; and (iv) the tr_{ij}-weighted average value.[23] Reassuringly, the pairwise correlation of the upstreamness measures obtained under these different treatments is very high (> 0.98), so we will focus on the version that uses the simple mean as our baseline.

To be clear, what the above construction and mapping yields is a measure of the average number of production stages based on the I-O classification system that are traversed between each pair of SIC input and manufacturing output industries. Table 2 presents some basic information on the total requirements coefficients and upstreamness variables after the mapping to SIC codes. A natural first-pass question to ask is how relevant the integrated SIC codes in WorldBase actually are for the production of the primary output j of their global parent. We can report that 98.3% of the observed (i, j) pairs in our sample are relevant in the sense that $tr_{ij} > 0$, while 82.9% of these pairs exceed the median positive tr_{ij} value (of 0.000163, across all input-output pairs in which the output industry is in manufacturing).[24] Panel B of Table 2 further suggests that the relationships we pick up in WorldBase do in fact reflect actual production linkages: Among the most commonly observed integrated (i, j) pairs, we have for example Commercial Printing, Lithographic (SIC 2752) as an input to Newspapers (SIC 2711), and Ready-Mixed Concrete (SIC 3272) as an input to Concrete Products, n.e.c. (SIC 3272).

[Table 2 about here]

Ratio-Upstreamness

In order to see whether the variation *across* global parents in integration decisions is consistent with our theory, we will 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.

[21] This concordance is available upon request. The BEA matches its 6-digit industry codes to 1987 U.S. SIC codes; see: http://www.bea.gov/industry/exe/ndn0017.exe.

[22] We apply the 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; see Antràs *et al.* (2012) for details.

[23] We also account for these mapping issues when converting the tr_{ij} values from I-O to SIC codes. For a given manufacturing output industry j, suppose a particular input i has a 4-digit I-O code that maps to multiple 4-digit SIC codes. We assign the total requirements value of i's use in the production of j equally across the multiple destination SIC codes that i maps to.

[24] We obtain similarly high relevance rates if we restrict our count to manufacturing inputs only, or if we drop repeated integrated SIC codes within each global parent. Likewise, our conclusions are unchanged if we drop pairs where $i = j$, to ensure that the large diagonal entries in the I-O tables is not unduly driving these statistics.

More specifically, we construct the following R_{jp} measure for each global parent, by taking the ratio of a weighted-average upstreamness of its integrated inputs relative to that of its non-integrated inputs:

$$R_{jp} = \frac{\sum_{i \in I(p)} \ \theta^{I}_{ijp} upst_{ij}}{\sum_{i \in NI(p)} \ \theta^{NI}_{ijp} upst_{ij}}. \tag{18}$$

The weights here are given by: $\theta^{I}_{ijp} = tr_{ij} / \sum_{i \in I(p)} tr_{ij}$ and $\theta^{NI}_{ijp} = tr_{ij} / \sum_{i \in NI(p)} tr_{ij}$, these being proportional to the total requirements coefficients to capture the relative importance of each input in the production of j. For convenience (and for want of a better name), we will refer to R_{jp} as the "ratio-upstreamness" of parent firm p. Note that by design, R_{jp} increases the greater is the propensity of p to integrate relatively upstream inputs, while outsourcing its more downstream inputs. In what follows, we will also consider several variants of "ratio-upstreamness" to assess it robustness under different constructions. These include: (i) restricting $S(j)$ to the set of "ever-integrated" inputs, namely the set of inputs i for which we actually observe 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, i.e., $i = j$, from $S(j)$. The last variant in particular will help in excluding subsidiaries that are purely horizontal affiliates in the nature of their activity.

The lower panel of Table 1 presents summary statistics for the different "ratio-upstreamness" measures. Note that R_{jp} tends to take on values less than one for most variants. This is due to the fact that the parent industry itself automatically belongs in the set $I(p)$ for these versions of R_{jp}, and the upstreamness of an industry j's use of itself as an input ($upst_{jj}$) tends to be relatively small in the data. From (18), this acts to lower the numerator of R_{jp}. When we exclude the self-SIC from the construction of "ratio-upstreamness", we see that R_{jp} is now more centered around median value of close to 1. The pairwise correlation between the different versions of R_{jp} is high (typically > 0.8), except when the self-SIC is excluded (the correlation with the baseline measures drops to about 0.15).

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

Demand Elasticity

As highlighted by our theory, the incentives to internalize upstream or downstream suppliers are crucially affected by whether the elasticity of demand faced by the firm is higher or lower than the elasticity of technological substitutability across its inputs. Unfortunately, estimates of cross-input substitutability are not readily available for a broad sets of industries, nor it is clear how one could construct them based on available data sources. Our approach will then rely on variation in the parent's demand elasticity, and associate the sequential complements case with high values of ρ_j and the substitutes case with low values of ρ_j. This amounts to assuming that any cross-sectoral variation in α_j is not systematically correlated with the demand elasticity parameter ρ_j.

To proxy for ρ_j, we use data on U.S. import demand elasticities from Broda and Weinstein (2006), which we aggregate up to the SIC industry level from their original estimates at the detailed HS10 product level using US import trade values as weights (see the Data Appendix for further details). We will also pursue several refinements of the proxy for ρ_j, by using only those HS product-level elasticities that correspond to consumption and capital goods in the United Nations' Classification by Broad Economic Categories (BEC); elasticities corresponding to products classified as intermediates are thus excluded. The idea here is that our model should better apply to final goods-producing industries, so that an elasticity constructed based such products would be 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 from the regression sample, as these are composed entirely of intermediate goods in the eyes of the BEC classification.

Input Contractibility

Our model suggests that the relative propensity to integrate upstream (as opposed to downstream) inputs depends also on the extent to which contractible inputs tend to be "front-loaded" or located in the early stages of the production process. To explore the role of input-contractibility, we construct the measure $UpstCont_j$ for each SIC output industry j, which captures the tendency of its high-contractibility inputs to be located at relatively upstream production stages.

The $UpstCont_j$ variable is constructed as follows. First, we follow the methodology of Nunn (2007) to construct measures of contractibility for each SIC input industry. Nunn (2007) developed a measure of contract-intensity, based on the Rauch (1999) classification of products into: (i) homogeneous products (that can be bought and sold on an open exchange); (ii) reference-priced products; and (iii) differentiated products. He then computed a measure of "relationship-specificity" or the "contract-intensity" of an industry as the share of constituent HS6 product codes in the composition of the industry's input use that are classified as being differentiated, on the premise that it is inherently more difficult to spell out and enforce the terms of contractual agreements for such products. As we are interested in the converse concept of contractibility, we therefore take one minus this Nunn measure of contract-intensity, namely the share of constituent product codes that are either homogeneous goods or reference-priced.[25] This give us a metric of the contractibility for each SIC manufacturing industry. (We restrict our attention here to manufacturing inputs as the Rauch (1999) measure is available only for tradable products.) Second, we need to take a stand on what constitutes a high-contractibility versus a low-contractibility input. As a baseline, we use the median contractibility value across the 459 SIC manufacturing industries to make this distinction. The results are very similar when using the upper tercile or lower tercile contractibility value as

[25] Specifically, our measure of the contractibility of an input is $1 - z^{rs1}$ in Nunn's (2007) notation. The findings are qualitatively similar if we were to exclude reference-priced products from this measure of contractibility (available on request). Likewise, the results are robust to using the "conservative" or "liberal" classification of products from Rauch (1999), although the "conservative" classification yields a measure of industry contractibility that exhibits more variation.

the cutoff instead (available on request). As the final step, for each output industry j, we calculate $UpstCont_j$ as the following ratio, in a manner analogous to the earlier ratio-upstreamness measure:

$$UpstCont_j = \frac{\sum_{i \in \mathcal{H}} \theta_{ij}^H upst_{ij}}{\sum_{i \in \mathcal{L}} \theta_{ij}^L upst_{ij}}. \tag{19}$$

Here, $\mathcal{H}$ denotes the set of SIC manufacturing industries with above-median contractibility values, with the converse set being $\mathcal{L}$; the weights are given by: $\theta_{ij}^H = tr_{ij}/\sum_{i \in \mathcal{H}} tr_{ij}$ and $\theta_{ij}^L = tr_{ij}/\sum_{i \in \mathcal{L}} tr_{ij}$. (19) thus takes a total-requirements weighted-average of the upstreamness of high-contractibility inputs, and divides it by the corresponding weighted-average upstreamness of low-contractibility inputs. If high-contractibility inputs tend to be located at earlier production stages, this would raise the numerator in (19) and reduce its denominator, yielding a higher value for $UpstCont_j$.

As described below in our exposition of the empirical results, we have also experimented with a *smoother* measure of upstream contractibility, computed for each industry j as the total-requirements weighted-covariance of $upst_{ij}$ and the contractibility of the inputs i (see, in particular, Appendix Table 2).

We also construct the variable $ContUpToi_{ij}$, which is our input-output pair-specific measure of the "contractibility up to i in the production of j"to be used later in the within-firm regression specifications. This is constructed as:

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

where the set of inputs, $S^m(j)$, over which the sum in the denominator is taken is the set of all manufacturing inputs used in the production of j (i.e., with $tr_{kj} > 0$), while 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 (i.e., $S_i^m(j) = \{k \in S^m(j) : upst_{kj} \geq upst_{ij}\}$). The denominator sums up the product of the total requirements coefficient and contractibility values across all inputs, with the numerator being the partial sum involving all inputs that are upstream of and including the input i. The construction of (20) is thus intended to capture the effect of the $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ terms from the theory, in particular the model with stage asymmetries in which inputs differ for exogenous reasons in their degree of contractibility.

4 Empirical Methodology

In our empirical analysis, we will assess the validity of the predictions of our theoretical model concerning patterns of ownership along the value chain. In doing so, we will exploit variation in integration choices both across and within firms to assess the above predictions.

According to Proposition 1, integration choices along the value chain should vary systematically between industries that fall under the sequential complements versus substitutes cases. While the model indicates that the relevant conditions delineating these cases are given by $\rho > \alpha$ and $\rho < \alpha$, as mentioned above, we are forced to rely purely on variation in ρ_j to identify these two cases.

A direct test of Proposition 1 would be to verify that $R_{jpc} < 1$ in the complements case (when ρ_j is high) and $R_{jpc} > 1$ in the substitutes case (when ρ_j is high). Unfortunately, we cannot test this stronger prediction given that, as discussed above, the construction of the "ratio-upstreamness" dependent variable yields measures that tend to be uniformly smaller than one. We can, however, test a weaker version of Proposition 1, examining whether the propensity to integrate upstream stages rises as ρ_j increases and we move towards the complements case. We can 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 the elasticity of demand for its final product.

When studying internalization decisions within firms across inputs, we can restate this first prediction as:

P.1 (Within): The effect of the upstreamness of an input on the propensity of a firm to integrate that input is lower (i.e., less positive, or more negative), the larger is the firm's elasticity of demand for its final product.

The extension of the model developed in Section 2 provides us with further testable predictions that emerge from considering heterogeneous contractibility levels across the inputs required in production. In particular, Proposition 3 suggests that the relative propensity to integrate upstream (as opposed to downstream) inputs depends also on the extent to which contractible inputs tend to be "front-loaded" or located in the early stages of the production process. Moreover, the effect of "upstream contractibility" varies subtly across the sequential complements and substitutes cases. The second cross-firm testable prediction of our model can be stated as follows:

P.2 (Cross): A greater degree of contractibility of its upstream inputs should decrease a firm's propensity to integrate upstream versus downstream inputs when the firm faces a low demand elasticity, but it should increase that propensity when the firm faces a high demand elasticity.

When turning it into an firm-input pair level prediction, this second prediction 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 lower the propensity of a firm to integrate that input when the firm faces a low demand elasticity. Conversely, it should raise the propensity to integrate that input when the firm faces a high demand elasticity.

Let us now turn to a more precise description of how we tests these predictions using both cross-firm as well as within-firm variation.

4.1 Cross-firm Specifications

To verify the validity of prediction P.1, we will first estimate the following regression model:

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

The dependent variable is the log ratio-upstreamness measure, which captures the propensity of each parent firm p, with primary SIC activity j to integrate relatively upstream inputs. Note that we introduce the subscript c to index the country where the parent firm is located, and 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.

The key regressor of interest is the dummy variable $\mathbf{1}(\rho_j > \rho_{med})$, which identifies industries where the elasticity of demand ρ_j is above the industry median. 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.

The regression in (21) includes on its right-hand side several additional controls. We add a vector X_j of variables corresponding to the primary output industry of the parent firm.[26] These include several measures of factor intensity – non-production employment, equipment capital, plant capital, and materials, all in per worker terms – as well as the value-added to shipments ratio. These were computed from the NBER-CES dataset using averages over the years 2001-2005. We further control for R&D intensity, specifically log(0.001 + R&D expenditures / Sales), as computed by Nunn and Trefler (2013) from firm-level data in Orbis. These have been carefully mapped to SIC industry categories from their original HS6 codes. (Please see the Data Appendix for a more detailed description, as well as Appendix Table 1 for basic summary statistics.)

Given the firm-level nature of the empirical analysis, we also add a vector of parent firm characteristics, W_p, obtained from WorldBase. Several of these variables reflect the size of the global parent, namely the number of establishments, log total employment, and log total sales. As the latter two variables are often estimated or missing from the original data, we also put in two 0-1 indicators on the right-hand side to control for whether the employment or sales figures (respectively) are based on actual data (coded as 0) or estimated/missing (coded as 1).[27] We also control for the age of the parent firm by using the year of its establishment (or in which the current ownership took control). Last but not least, an indicator variable for whether the parent is a multinational is included.

[26] The ability to use directly-observed characteristics of the output industry is an added advantage that comes from the firm-level dimension of the dataset. In the previous work in Antràs and Chor (2013), the implications of the model were tested using data on US intra-firm import shares, in which the identity of the buying industry could not be ascertained. There, what were used as controls were instead computed variables for the characteristics of the "average-buyer" industry.

[27] Our results are robust to further including a full set of dummy variables for trade status, namely whether the firm is an importer, an exporter, both an importer and exporter, a trade intermediary. As this information is not available for a relatively large subset of the data, we do not report the detailed results from using these controls (available on request).

We view the above variables in X_j and W_p strictly as auxiliary controls, in the sense that our model does not deliver direct predictions that would lead us to clearly sign the effects of these variables on the "ratio-upstreamness" measure. For example, several of these industry variables, such as the equipment capital intensity and R&D intensity, have been commonly used in the empirical literature as proxies for the headquarter intensity of the industry in question. Although the property-rights approach to the theory of the firm would predict that integration would be more likely in high headquarter-intensity industries, the theory is actually silent as to how this should affect the relative propensity to integrate upstream versus downstream inputs.

An alternative way to test our prediction P.1 is to explore specifications based upon a finer cut by quintiles of our proxy for ρ_j:

$$\log R_{jpc} = \beta_0 + \beta_k \sum_{k=2}^{5} \mathbf{1}(\rho_j \in Quint_k(\rho)) + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}, \tag{22}$$

where $\mathbf{1}(\rho_j \in Quint_n(\rho))$ is an indicator variable for whether the demand elasticity proxy for industry j belongs in the n-th quintile. This approach has the advantage of relaxing the degree of uncorrelatedness between α_j and ρ_j necessary for our empirical strategy to be sound. In light of the model, we expect the coefficients β_k to be lower (more negative), the higher is k.

To assess the validity of prediction P.2, we adopt the more flexible specification based on quintile elasticity dummies and estimate the following cross-firm regression:

$$\begin{aligned} \log R_{jpc} \quad = \quad & \beta_0 + \beta_k \sum_{k=2}^{5} \mathbf{1}(\rho_j \in Quint_k(\rho)) + \gamma_k \sum_{k=1}^{5} \mathbf{1}(\rho_j \in Quint_k(\rho)) \times \log UpstCont_j \\ & + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \end{aligned} \tag{23}$$

We interact each of the five quintile dummies for ρ_j with the measure $UpstCont_j$, which captures the tendency for high-contractibility inputs to be located at relatively upstream production stages. Based on the second prediction of our model, we would then expect the coefficients γ_k to increase in k, with $\gamma_1 < 0$ and $\gamma_5 > 0$.

4.2 Within-firm Specifications

In this second part of the empirical exercise, we exploit variation in the patterns of integration within firms across different inputs. To this end, we restructure the dataset so that an observation is now an SIC input (i) by parent firm (p) pair. The analysis pursues regressions of the following form:

$$\begin{aligned} D\_INT_{ijp} \quad = \quad & \gamma_0 + \sum_{n=1}^{5} \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times ContUpToi_{ij} \\ & + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp} \end{aligned} \tag{24}$$

The dependent variable, D_INT_{ijp}, is an indicator variable for whether the firm p with primary output industry j has integrated input i within firm boundaries. In practice, we will estimate (24) as a linear probability model, with standard errors clustered at the input-output pair (i-j) level. To keep the analysis tractable, we limit the sample here only 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.[28] The relevant variation in D_INT_{ijp} for the above regression is that at the within-firm level, since the right-hand side includes a full set of parent firm fixed effects, D_p. These will also absorb any systematic differences arising from the output industry or country of the parent firm.

The key explanatory variables in (24) are the terms involving $ContUpToi_{ij}$, which is our input-output pair-specific measure of the "contractibility up to i in the production of j". By interacting this measure of $ContUpToi_{ij}$ with each of the quintile elasticity dummies in (24), we seek to tease out the differential effect of "contractibility up to i" across the substitutes versus complements cases as suggested in the theory.

Two further remarks are in order. First, our empirical setting has the feature that inputs from the same SIC 4-digit code as the parent output industry are viewed as integrated within firm boundaries, i.e., $j \in I(j)$ by definition. This is an aspect of the integration patterns of the firm that is not particularly interesting, since it stems from a mechanical treatment of these "self-SIC" inputs. We therefore add a dummy variable $\mathbf{1}(i = j)$ to the right-hand side of (24), which equals one if and only if the input i belongs to the same SIC output industry as the parent firm. The effects of "contractibility up to i" are therefore estimated off the integration patterns observed from manufacturing inputs other than the "self-SIC". Related to this, we use for the within-firm regressions the sample of parent firms from that have integrated at least one other manufacturing input apart from the parent's primary SIC code. This avoids including in the regression sample parent firms for which the dependent variable is equal to 0, except for the 'special' case of the "self-SIC". Second, in some specifications, we will further add a full set of dummies to control for the identity of the input i, D_i. When these are used, only covariates that vary at the input-output pair level can be identified in the estimation.

5 Empirical Results

5.1 Cross-firm Specifications

The regression in (21) serves as a means to validate prediction P.1 of our model, based on the information on firm SIC activities contained in WorldBase. Table 3 reports these first results using the median elasticity cutoff for our explanatory variable of interest. Column (1) presents a stripped-down specification in which only $\mathbf{1}(\rho_j > \rho_{med})$ and parent country fixed effects are included. The

[28] The focus on manufacturing inputs is also motivated by a practical consideration, namely that the measures of contractibility based on the Nunn (2007) methodology are only available for manufacturing and a small number of additional tradable inputs. In particular, the methodology does not yield contractibility values for services.

estimated coefficient on our proxy for the complements case is negative and strongly significant, confirming that the propensity to integrate upstream stages is lower when $\rho_j > \rho_{med}$, consistent with the prediction of Proposition 1. Columns (2) and (3) show that this result is very robust, even as we successively add the parent industry variables, X_j, and the parent firm controls, W_p, respectively. Looking at these auxiliary variables, the estimates indicate that there is a tendency for more upstream integration in more equipment capital-intensive industries, as well as in firms that have more establishments, younger firms, and in MNCs.

The remaining columns in Table 3 adopt the refinements in the ρ proxy, to allow it to better reflect elasticities that pertain to final-goods demand. Column (4) restrict its construction 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. 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 industry categories that are composed entirely of intermediate-use goods are dropped from the sample.

[Table 3 about here]

Table 4 repeats the preceding exercise, but uses instead the set of quintile elasticity dummies, $\mathbf{1}(\rho_j \in Quint_n(\rho))$ for $n = 2, \ldots, 5$, in place of the median cutoff dummy. The results are similarly supportive of the model, with the magnitude of the estimated negative coefficient increasing steadily as we move from the second to the fifth elasticity quintile (the first quintile dummy is the omitted category) throughout Columns (2)-(5). As in Table 3, the regression based on the most stringent refinement of the ρ proxy – that in Column (5) using consumption goods elasticities alone – yields the largest point estimates for the coefficients of interest; in terms of economic magnitudes, these suggest that the difference between the lowest and highest elasticity quintiles corresponds to a one standard-deviation decrease in the propensity to integrate upstream inputs as captured by the log ratio-upstreamness variable.

[Table 4 about here]

In Table 5, we explore whether the predictions related to the degree of contractibility of upstream inputs are borne out in the WorldBase data. As described in the specification in (23), this is done by adding further interaction terms between the quintile elasticity dummies and the $UpstCont_j$ measure given by equation (19). Looking first at the main effects of the quintile elasticity dummies, these continue to exhibit similar patterns to the more parsimonious specification in Table 4, with negative and significant coefficients particularly 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 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 ρ proxy used in each column. Of note, we find that in the complements case, a higher degree of upstream contractibility does in fact counteract the above tendency to integrate upstream inputs, as the estimated coefficient on $\mathbf{1}(\rho_j \in Quint_5(\rho)) \times UpstCont_j$ consistently emerges as positive and statistically significant (at the 1% level). Conversely, in the substitutes case, the interaction between $UpstCont_j$ and the first elasticity quintile dummy has the opposite sign, indicating that the degree of upstream contractibility instead acts to raise the propensity to integrate *downstream* inputs. (The above results are very similar when using a more continuous measure of upstream contractibility, computed for each industry j as the total-requirements weighted-covariance of $upst_{ij}$ and the contractibility of the inputs i; these are presented in Appendix Table 2.)

The above findings are in line with prediction P.2 of our model, which related integration decisions to the sequencing of high- versus low-contractibility inputs. As emphasized in the theory section, a greater level of upstream contractibility enables firms to spell out investment levels by suppliers that are closer to the first-best levels for the early stages of production. This reduces the need to rely on decisions over organizational mode to elicit the right incentives from these suppliers. In the complements case, this reduces the propensity to outsource upstream, and invites the firm to push the cutoff stage between outsourcing and integration further upstream. On the other hand, in the substitutes case, the firm has more leeway to integrate fewer stages upstream, as the negative spillover effects from upstream investment on downstream supplier effort are ameliorated by being better able to contract explicitly with upstream suppliers. It bears noting that these findings are more consistent with a view of integration decisions that is rooted in the property-rights approach to the theory of the firm. The transactions-cost approach, which views integration as a means to circumvent contractual frictions vis-à-vis independent suppliers, would intuitively deliver a different set of implications related to upstream contractibility. For example, in the complements case, greater upstream contractibility would instead lead the parent firm to outsource more upstream stages.

[Table 5 about here]

Table 6 restricts the regression sample to larger firms, to ensure that the above effects are not driven by smaller, less economically-relevant entities. In the interest of space, we report results based on the ρ proxy constructed using consumption goods elasticities only (other results available on request). We continue to find significant effects on the quintile elasticity dummies, as well as their respective interaction terms with $UpstCont_j$. This holds even as we narrow the sample successively down to parent firms with at least 20 employees (Column (1)), with at least two distinct establishments (Column (2)), and with establishments in more than one country, i.e., multinational firms (Column (3)). Finally, in Column (4), we take out those multinationals from Column (3) that report only one SIC activity code within firm boundaries. This in principle removes multinational firms that undertake FDI that is purely horizontal in nature (i.e., for the purpose of gaining access to the foreign market), rather than to facilitate vertical specialization in production. Of note, the previous findings from Table 5 still hold even as the number of observations decreases with each subsample cut of the dataset.

[Table 6 about here]

We also take steps to address concerns over the treatment of parent firms that are actively engaged in multiple output industries, bearing in mind that WorldBase provides up to five additional secondary 4-digit SIC codes for each establishment. We can nevertheless report very similar results when limiting the sample to those parent firms that report only a primary (and no secondary) SIC activities. Alternatively, we have also constructed the ratio-upstreamness measure taking in turn each secondary manufacturing SIC activity as the parent firm's output industry j. We then run regression (23) pooling across the multiple R_{jpc} values per parent firm, and reporting multi-way clustered standard errors by SIC output industry and by parent firm (Cameron *et al.* 2011). These results are reported in full in Appendix Table 3, confirming that our conclusions are unchanged under either treatment of parent firms with multiple SIC codes.

Table 7 undertakes a final set of robustness checks, that build on the specification run in Column (3) of Table 5. (Recall that this uses the ρ proxy based on consumption goods elasticities only.) In Column (1), we directly include the contractibility of the output industry j, to explore if that has an effect on integration decisions over the inputs used by the parent firm. Also, to verify that it is the production line position of high- versus low-contractibility inputs that matters, and not just the raw variation across inputs in contractibility, we also include here the interaction between each quintile dummy $\mathbf{1}(\rho_j \in Quint_n(\rho))$ and a total-requirements weighted standard deviation across the contractibility of the inputs used by j. The estimates in Column (1) continue to point to a negative and significant effect on upstream integration in the complements case, with this effect being moderated in industries that exhibit a greater degree of upstream contractibility (see the quintile 5 interactions). Conversely, the estimates of the effect of upstream contractibility in the substitutes case remain negative, though now marginally insignificant.

In Columns (2)-(5), we revert to the baseline specification in (23), and instead examine the use of alternative constructions of the ratio-upstreamness dependent variable. The version of R_{jpc} in Column (2) is based on $upst_{ij}$ values from a random pick, when the mapping from I-O to SIC codes yielded multiple matches for a non-manufacturing input i and manufacturing output industry j. In Column (3), we limit the set $S(j)$ in the construction of R_{jpc} to those inputs for which we observe a parent firm in j integrating the input in question ("ever-integrated" inputs). We alternatively restrict $S(j)$ to the set of manufacturing inputs in Column (4), and further drop the parent SIC from the set of inputs in Column (5). Our findings are broadly consistent, with the main exception being the final column of Table 7. There, $UpstCont_j$ does reduce the propensity to integrate upstream in quintile 1 (the substitutes case), but the point estimates for the quintile 5 interactions (the complements case) are not significantly different from zero. Note however that there is a large decrease in the number of available observations in Column (5), 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 7 about here]

In sum, the cross-firm regressions provide strong evidence that the propensity for a firm to integrates relatively upstream inputs is strongest when the demand elasticity faced by that industry is largest, in line with prediction P.1. Moreover, prediction P.2 regarding how upstream contractibility would affect integration patterns along the sequence of production inputs also finds support in the data.

5.2 Within-firm Specifications

Table 8 reports the estimates from regression 24. Before examining the effect of "contractibility up to i", we first explore what effect our more basic upstreamness measure has on the propensity to integrate particular inputs, as well as how any such effect would differ for parent firms in high vs low demand elasticity industries. For this, we include the following covariates on the right-hand side of the regression: $\sum_{n=1}^{5} \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times upst_{ij}$. The results in Column (1) confirm a negative and significant effect of $upst_{ij}$ in the fifth elasticity quintile. This dovetails with the predictions of our model, as the propensity to integrate declines the more upstream the input in question is in the complements case. That said, the estimated coefficient for the first elasticity quintile is also negative, albeit of a smaller magnitude; a formal test for the quality of the first and fifth quintile coefficients is rejected at conventional significance levels. As expected, the "self-SIC" dummy emerges with a positive and highly significant effect, with a point estimate approximately equal to 1, as it does in all remaining columns.

To the above, Column (2) includes the set of quintile elasticity dummies interacted with the "contractibility up to i" terms. Of note, the effect of $ContUpToi_{ij}$ is positive and statistically significant in the fifth elasticity quintile. This is entirely consistent with the prediction from Proposition 2 that a greater degree of contractibility upstream of input i would raise the propensity to integrate i itself. On the other hand, however, the point estimate for the interaction term with the first elasticity quintile also turns up as positive. Although this is at variance with what the model predicts for the substitutes case, observe nevertheless that this positive effect of $ContUpToi_{ij}$ is indeed the weakest here before increasing steadily as we progress to higher and higher elasticity quintiles. (The p-value reported in the table confirms that we can reject the null hypothesis that the coefficients corresponding to quintiles 1 and 5 are equal.) Thus, the data at least exhibits the property that the effect of "contractibility up to i" in raising the likelihood of integration is highest in industries that would correspond to the complements case. Moreover, observe that the prior coefficients of the quintile elasticity and $upst_{ij}$ interactions are now for the most part close to zero, once we have controlled for the effect of $ContUpToi_{ij}$. This confirms that it is the effect of the contractibility of upstream inputs, rather than the upstreamness of i itself, that matters for integration patterns. It is useful to point out that this is precisely what the model would predict, since what matters for the optimal organizational decision 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.

[Table 8 about here]

In Column (3), we add the contractibility of input i itself interacted with the full set of elasticity quintile dummies. While these terms do appear to have significant effects in Column (3), these dissipate when we further add the SIC input fixed effects (the D_i's) in Column (4). Most importantly, the effects of "contractibility up to i" and how these vary across elasticity quintiles remains very stable throughout the different regressions.

Table 9 provides corroborating evidence from regressions that look at the subsets of firms by output elasticity quintile. For each quintile, we control for both parent firm and SIC input fixed effects (D_p and D_i respectively), while also including the "self-SIC" dummy. Even within each quintile, the propensity to integrate rises with the contractibility of inputs up to i, $ContUpToi$, but this effect is clearly the strongest in quintile 5, which maps most closely to the complements case. Once again, the upstreamness of i's use in the production of j *per se* has no significant effect on integration patterns, once the contractibility of upstream inputs has been controlled for.

[Table 9 about here]

6 Conclusion

This paper shows how available data on the production activities of firms operating in many countries and industries can be combined with information from standard Input-Output tables to study the organization of firms along global 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 plant-level information on the production activities of firms located in a large set of countries. We combine this information with Input-Output data 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 in the value chain 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 empirical patterns that we uncover here thus provide an additional source of evidence that considerations driven by contractual frictions are indeed critical for the decisions that firms make over the organization of their value-chain activities.

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 (2014), "Grossman-Hart (1986) Goes Global: Incomplete Contracts, Property Rights, and the International Organization of Production," *Journal of Law, Economics and Organization* 30(suppl 1): 118-175.

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)".

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012), "Americans Do IT Better: US Multinationals and the Productivity Miracle," *American Economic Review* 102(1): 167-201.

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.

Del Prete, Davide, and Armando Rungi (2014), "Organizing the Global Value Chain: A Firm Level Test," mimeo.

Dietzenbacher, Erik, Isidoro Romero Luna, and Niels S. Bosma (2005), "Using Average Propogration Lengths to Identify Production Chains in the Andalusian Economy," *Estudios de Economia Aplicada* 23(2): 405-422.

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

Dixit, Avinash, and Gene Grossman (1982), "Trade and Protection with Multistage Production," *Review of Economic Studies* 49(4): 583-594.

Fally, Thibault (2012), "On the Fragmentation of Production in the U.S.," mimeo.

Fajgelbaum, Pablo, Gene Grossman, and Elhanan Helpman (2015), "A Linder Hypothesis for Foreign Direct Investment," *Review of Economic Studies* 82, 83-121.

Fally, Thibault and Russell Hillbery (2014), "A Coasian Model of International Production Chains," mimeo UC Berkeley.

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. McMillian (2004), "Global Capital Flows and Financing Constraints," *Journal of Development Economics* 75, 269-301.

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," UC Davis 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 Trefler (2013), "Incomplete Contracts and the Boundaries of the Multinational Firm," *Journal of Economic Behavior and Organization* 94: 330-344.

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

Ramondo, Natalia, Veronica Rappoport, and Kim Ruhl (2014), "Horizontal versus Vertical Foreign Direct Investment: Evidence from U.S. Multinationals," mimeo.

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.

Yi, Kei-Mu (2003), "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111(1): 52-102.

A Theoretical Appendix

A.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:

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

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

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

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

$$x(m) = A\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(\frac{\rho}{c}\right)^{\frac{1}{1-\rho}} \theta^{\frac{\rho}{1-\rho}} \left(\frac{(1-\beta(m))\psi(m)}{c(m)}\right)^{\frac{1}{1-\alpha}} \left[\int_0^m \left(\frac{(1-\beta(i))\psi(i)}{c(i)}\right)^{\frac{\alpha}{1-\alpha}} di\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}, \tag{25}$$

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

$$\pi_F = A\frac{\rho}{\alpha}\left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (\rho\theta)^{\frac{\rho}{1-\rho}} \int_0^1 \beta(i) \left(\frac{(1-\beta(i))\psi(i)}{c(i)}\right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}} dk\right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di, \tag{26}$$

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

A.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

$$v(i) \equiv \int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}} dk. \tag{27}$$

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

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

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

$$\frac{\rho-\alpha}{1-\rho}\frac{\left[v'(i)\right]^2}{v(i)}+\frac{v''(i)}{v'(i)}=-\frac{\alpha}{1-\alpha}\frac{d\left(c(i)/\psi(i)\right)/di}{c(i)/\psi(i)} \tag{28}$$

To solve (28), integrate both sides with respect to i, and exponentiate to get:

$$v'(i)\,v(i)^{\frac{\rho-\alpha}{1-\rho}}=C_1\left(\psi(i)/c(i)\right)^{\frac{\alpha}{1-\alpha}}, \tag{29}$$

where $C_1>0$ is a constant of integration. Given the definition of $v(i)$ in (27), we can express this equation as:

$$(1-\beta(i))^{\frac{\alpha}{1-\alpha}}=C_1\left(\int_0^i\left(\frac{(1-\beta(k))\,\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}}dk\right)^{\frac{\alpha-\rho}{1-\rho}}, \tag{30}$$

Denoting $z(i)\equiv(1-\beta(i))^{\frac{\alpha}{1-\alpha}}$, we can express (30) 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 to a separable differential equation, with 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.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 $\rho>\alpha$, $\lim_{m\to 0}\beta^*(m)\to-\infty$, 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, $\lim_{m\to 0}\beta^*(m)=1$, and instead 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

$$\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di = \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di. \tag{31}$$

These constants can always be chosen small enough such that not only do they satisfy (31), 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

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \int_0^1 \beta(i) \frac{\partial \left(\left[\int_0^i ((1-\beta(k))\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right)}{\partial i} di, \tag{32}$$

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 thereafter. 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

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

Notice next that in light of equation (32), the part of profits Π_1 associated with stages $m > \tilde{m}+\varepsilon_R$ is given by

$$\begin{aligned} &\Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \beta_V \left(\mathcal{A} + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(i) di + \mathcal{D} \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \\ &-\Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \beta_V \left(\mathcal{A} + (1-\beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(i) di + (1-\beta_O)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(i) di \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}, \end{aligned}$$

while for profits Π_2, these same profits are given by

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

But given (31), we have that $\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}}\gamma(i)\,di=\int_{\tilde{m}}^{\tilde{m}+\varepsilon_R}\gamma(i)\,di$ and thus these two expressions are identical.

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 (32), and after some manipulations, we find that

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

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 $y=\mathcal{A}$, $a=(1-\beta_O)^{\frac{\rho}{1-\rho}}\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}}\gamma(i)\,di$ and $b=(1-\beta_V)^{\frac{\rho}{1-\rho}}\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}}\gamma(i)\,di$, 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.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 (32) to express profits as

$$\pi_F = \Theta\frac{\alpha(1-\rho)}{\rho(1-\alpha)}\beta_O(1-\beta_O)^{\frac{\rho}{1-\rho}}\left(\int_0^{m_C}\left(\frac{\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}}dk\right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \tag{33}$$
$$+\Theta\frac{\alpha(1-\rho)}{\rho(1-\alpha)}\beta_V\left[\begin{array}{c}\left((1-\beta_O)^{\frac{\alpha}{1-\alpha}}\int_0^{m_C}\left(\frac{\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}}dk+(1-\beta_V)^{\frac{\alpha}{1-\alpha}}\int_{m_C}^1\left(\frac{\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}}dk\right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}\\-\left((1-\beta_O)^{\frac{\alpha}{1-\alpha}}\int_0^{m_C}\left(\frac{\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}}dk\right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}\end{array}\right].$$

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

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

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 $\beta_V\left(1-\beta_V\right)^{\frac{\alpha}{1-\alpha}}>\beta_O\left(1-\beta_O\right)^{\frac{\alpha}{1-\alpha}}$, 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 (33) with β_O replacing β_V throughout and vice versa. As a result, m_S^* is given by

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

and in our example in which $\psi\left(k\right)/c\left(k\right)$ grows (or declines) at a constant rate λ, we can again conclude that m_S^* is also decreasing in λ, as stated in the main text.

A.5 Derivation of Equation (15)

Let us ignore contracting costs for now and focus on the first term of equation (15). Also, we will derive an expression that applies even when there is variation in both $\psi\left(i\right)$ and $c\left(i\right)$. This will prove useful when our claims towards the end of section 2.3. Consider first the complements case. We begin by plugging the equilibrium equation (14) determining the thresholds m_C^* in into the profit function (33). After a few simplifications, this delivers

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

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

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

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

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \Gamma(\beta_V, \beta_O), \tag{34}$$

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 (34) 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

$$\frac{\psi(m)/c(m)}{\psi(m')/c(m')} = \left(\frac{\mu(m)/c(m)}{\mu(m')/c(m')} \right)^{\phi - \frac{\alpha}{1-\alpha}}, \tag{35}$$

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 (35) into (34), 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 $\mu(i)$.

B Data Appendix

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 category 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 assign 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.

Two refinements of the above baseline elasticity were also constructed. The first uses only those underlying HS10 product codes that are classified by the UN Broad Economic Categories (BEC) as being either consumption or capital goods; this excludes products that are classified as being for intermediate use. The second further restricts to those HS10 product codes that are classified as consumption goods.

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 Trefler (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 HS5 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.

Contractibility: Following the methodology in Nunn (2007). This in turn uses the Rauch (1999) classification of goods into 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 obtainNunn's (2007) contract-intensity measure for each 4-digit SIC code. We take one minus the contract-intensity to get a measure of contractibility.

Firm-Level Data: From WorldBase, a database compiled by Dun & Bradstreet (D&B). The data covers public and private companies across more than 100 countries and territories. We use the 2004/2005 vintage. The unit of observation in WorldBase is the establishment. Establishments, which we also refer to as plants, like firms have their own addresses, business names, and managers, but might be partly or wholly owned by other firms. We use the following categories of data that WorldBase records for each establishment:

i. Industry information including the four-digit SIC code of the primary industry in which each establishment operates and, for most countries, the SIC codes of up to five secondary industries, listed in descending order of importance. D&B uses the United States Government Department of Commerce, Office of Management and Budget, Standard Industrial Classification Manual 1987 edition to classify business establishments. In 1963, the firm introduced the Data Universal Numbering System—The D&B D-U-N-S® Number—used to identify businesses numerically for data-processing purposes. The system supports the linking of plants and firms across countries.

ii. Ownership information including information about firms' family members (number of family members, domestic parent, and global parent), status (joint venture, corporation, partnership), and position in the hierarchy (branch, division, headquarters).

iii. Detailed location information including the country, state, city, and street address of each family member.

We focus on manufacturing firms, namely observations that are identified as global parents (or in the terminology of WorldBase, "global ultimates"), whose primary SIC code lies between 2000 and 3999. Parents are defined in the data as entities that have legal and financial responsibility for another firm. We link these global parents to their identified majority-owned subsidiaries in WorldBase, using the reported DUNS number of the global parent reported by establishments.

Input-Output Data: Taken from the Bureau of Economic Analysis (BEA), Benchmark IO

Tables. We use the Use of Commodities by Industries after Redefinitions 1992 (Producers' Prices) tables. While the BEA employs six-digit input-output industry codes, WorldBase uses the SIC industry classification. We use the BEA provided concordance guide. The key is not one-to-one key, although the multiple matching problem is not particularly relevant when analyzing plants in the manufacturing sector (for which the key is almost one-to-one). For codes for which the match is not one-to-one, we use different approaches: i) simple average of distij over constituent IO1992 input categories; ii) simple median; iii) random pick; and iv) tr_{ij} weighted-average.

Table 1
Summary Statistics: Global Parent Firms

	10th	Median	90th	Mean	Std Dev
A: <u>Global parent firm variables</u>					
All global parents:					
Number of Establishments (incl. self)	1	1	2	1.77	5.81
Number of countries (incl. self)	1	1	1	1.14	1.03
Number of integrated SIC codes	1	2	4	2.35	3.41
Year started	1948	1985	2000	1977	26.17
Log (Total employment), 107656 obs	1.099	3.219	5.704	3.322	1.856
Log (Sales in USD), 87675 obs	12.795	15.305	17.844	15.325	2.055
MNCs only, 6983 obs:					
Number of Establishments (incl. self)	2	3	15	8.05	22.32
Number of countries (incl. self)	2	2	6	3.36	3.51
Number of integrated SIC codes	2	4	16	7.73	11.45
B: <u>Ratio-Upstreamness measures</u>					
Baseline (mean)	0.490	0.558	0.698	0.586	0.136
Baseline (random pick)	0.494	0.557	0.698	0.586	0.136
Manufacturing inputs only	0.547	0.620	0.779	0.645	0.161
Ever-integrated inputs only (mean)	0.564	0.659	0.821	0.693	0.178
Exclude parent sic (mean)	0.586	0.953	1.607	1.049	0.401
Exclude parent sic, manufacturing only	0.589	1.065	2.110	1.257	0.625

Note: Based on the sample of 116,483 global ultimates in the Dun & Bradstreet database (year=2005) whose primary SIC activity is in manufacturing. 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; as this mapping is unambiguous for manufacturing inputs, there is no need to distinguish between these treatments when restricting to manufacturing inputs only.

Table 2
Upstreamness: Summary Statistics and Some Examples

	10th	Median	90th	Mean	Std Dev
A: From Input-Output Tables (*i* =input; *j* =output)					
(for *j* in manufacturing only: 416,349 obs.)					
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

B: Top ten most commonly observed SIC input-output pairs (in D&B)
(for *i and j* in manufacturing only)

SIC input, *i*	**SIC output, *j***	**No. such pairs**	**Upst_ij**
Cookies and Crackers (2052)	Bread, Cake and Related Products (2051)	497	3.135
Commercial Printing, Lithographic (2752)	Commercial Printing, n.e.c. (2759)	439	1.186
Periodicals (2721)	Newspapers (2711)	391	1.409
Commercial Printing, n.e.c. (2759)	Commercial Printing, Lithographic (2752)	319	1.186
Commercial Printing, Lithographic (2752)	Newspapers (2711)	299	1.348
Women's and Misses' Outerwear, n.e.c. (2339)	Men's and Boys' Clothing, n.e.c. (2329)	287	1.106
Typesetting (2791)	Commercial Printing, Lithographic (2752)	280	1.151
Bookbinding and Related Work (2789)	Commercial Printing, Lithographic (2752)	273	2.192
Sausages and Other Prepared Meats (2013)	Meat Packing Plants (2011)	272	1.329
Ready-Mixed Concrete (3273)	Concrete Products, n.e.c. (3272)	190	1.074

Notes: Based on the U.S. 1992 Benchmark Input-Output Tables. Summary statistics in Panel A are for the respective variables after mapping from the original IO1992 to SIC industry codes.

Table 3
Upstreamness of Integrated vs Non-Integrated Inputs: Median Elasticity Cutoff

Dependent variable:	**Log Ratio-Upstreamnes**				
	(1)	(2)	(3)	(4)	(5)
Ind.(Elas > Median)	-0.0417**	-0.0681***	-0.0677***	-0.0667***	-0.1096***
	[0.0207]	[0.0186]	[0.0181]	[0.0214]	[0.0248]
Log (Skilled Emp./Worker)		0.0004	0.0034	0.0000	-0.0310
		[0.0231]	[0.0224]	[0.0259]	[0.0322]
Log (Equip. Capital / Worker)		0.1094***	0.1067***	0.0798***	0.0846***
		[0.0219]	[0.0211]	[0.0226]	[0.0265]
Log (Plant Capital / Worker)		-0.0217	-0.0237	0.0026	-0.0038
		[0.0227]	[0.0223]	[0.0281]	[0.0328]
Log (Materials / Worker)		-0.0527**	-0.0487**	-0.0651**	-0.0471
		[0.0247]	[0.0228]	[0.0257]	[0.0325]
R&D intensity		0.0082	0.0059	0.0113	0.0067
		[0.0055]	[0.0053]	[0.0068]	[0.0076]
Value-added / Shipments		-0.1580	-0.1427	-0.1299	0.0673
		[0.1148]	[0.1108]	[0.1178]	[0.1527]
Log (No. of Establishments)			0.0519***	0.0564***	0.0579***
			[0.0031]	[0.0039]	[0.0048]
Year Started			0.0001*	0.0002**	0.0002*
			[0.0001]	[0.0001]	[0.0001]
Dummy: Multinational			0.0100**	0.0146**	0.0281***
			[0.0049]	[0.0062]	[0.0076]
Elasticity based on:	All codes	All codes	All codes	BEC cons. & cap. goods	BEC cons. only
Parent country dummies?	Y	Y	Y	Y	Y
Controls for emp. & sales?	N	N	Y	Y	Y
Observations	115,800	115,800	84,171	62,377	44,895
No. of industries	459	459	459	305	219
R^2	0.0671	0.1674	0.1896	0.2053	0.2393

Notes: Each observation is a global ultimate parent firm in 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 ratio-upstreamness measure described in Section 3.2 of the main text. A median cutoff dummy is used to distinguish parents with primary SIC output that are in high vs low import demand elasticity industries. Columns (1)-(3) use an elasticity measure constructed from all available HS10 elasticities from Broda-Weinstein (2006); Column (4) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC, while Column (5) restricts this to consumption goods only. All columns include parent country fixed effects. Columns (3)-(5) also control for log employment and log sales of the parent firm, as well as indicators for whether each of these respective variables was based on actual data or was estimated/missing in Worldbase (coefficients not reported).

Table 4
Upstreamness of Integrated vs Non-Integrated Inputs: By Elasticity Quintiles

Dependent variable:	**Log Ratio-Upstreamness**				
	(1)	(2)	(3)	(4)	(5)
Ind.(Quintile 2 Elas)	-0.0205	-0.0304	-0.0313	-0.0629	-0.0805*
	[0.0307]	[0.0277]	[0.0282]	[0.0426]	[0.0453]
Ind.(Quintile 3 Elas)	-0.0677**	-0.0784***	-0.0797***	-0.0713*	-0.1026**
	[0.0308]	[0.0293]	[0.0295]	[0.0424]	[0.0415]
Ind.(Quintile 4 Elas)	-0.0334	-0.0832***	-0.0845***	-0.1035**	-0.1506***
	[0.0336]	[0.0312]	[0.0311]	[0.0432]	[0.0449]
Ind.(Quintile 5 Elas)	-0.0715*	-0.1021***	-0.1043***	-0.1287***	-0.1890***
	[0.0375]	[0.0315]	[0.0312]	[0.0418]	[0.0448]
Log (Skilled Emp./Worker)		0.0001	0.0022	-0.0042	-0.0370
		[0.0225]	[0.0219]	[0.0274]	[0.0335]
Log (Equip. Capital / Worker)		0.1084***	0.1058***	0.0750***	0.0800***
		[0.0207]	[0.0198]	[0.0199]	[0.0214]
Log (Plant Capital / Worker)		-0.0154	-0.0167	0.0134	0.0053
		[0.0211]	[0.0206]	[0.0235]	[0.0287]
Log (Materials / Worker)		-0.0561**	-0.0520**	-0.0707***	-0.0541*
		[0.0243]	[0.0223]	[0.0257]	[0.0314]
R&D intensity		0.0078	0.0058	0.0112*	0.0039
		[0.0053]	[0.0052]	[0.0063]	[0.0079]
Value-added / Shipments		-0.1732	-0.1572	-0.1454	0.0707
		[0.1159]	[0.1113]	[0.1188]	[0.1617]
Log (No. of Establishments)			0.0521***	0.0575***	0.0600***
			[0.0030]	[0.0035]	[0.0044]
Year Started			0.0001**	0.0002**	0.0002***
			[0.0001]	[0.0001]	[0.0001]
Dummy: Multinational			0.0100**	0.0131**	0.0226***
			[0.0047]	[0.0057]	[0.0073]
Elasticity based on:	All codes	All codes	All codes	BEC cons. & cap. goods	BEC cons. only
Parent country dummies?	Y	Y	Y	Y	Y
Controls for emp. & sales?	N	N	Y	Y	Y
Observations	115,800	115,800	84,171	62,377	44,895
No. of industries	459	459	459	305	219
R^2	0.0777	0.1773	0.2005	0.2300	0.2707

Notes: Each observation is a global ultimate parent firm in 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 ratio-upstreamness measure described in Section 3.2 of the main text. Quintile dummies are used to distinguish parents with primary SIC output that are in high vs low import demand elasticity industries. Columns (1)-(3) use an elasticity measure constructed from all available HS10 elasticities from Broda-Weinstein (2006); Column (4) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC, while Column (5) restricts this to consumption goods only. All columns include parent country fixed effects. Columns (3)-(5) also control for log employment and log sales of the parent firm, as well as indicators for whether each of these respective variables was based on actual data or was estimated/missing in Worldbase (coefficients not reported).

Table 5
Effect of Upstream Contractibility

Dependent variable:	**Log Ratio-Upstreamness**		
	(1)	(2)	(3)
Ind.(Quintile 2 Elas)	-0.0290	-0.0441*	-0.0405
	[0.0186]	[0.0238]	[0.0286]
Ind.(Quintile 3 Elas)	-0.0639***	-0.0538**	-0.0617**
	[0.0205]	[0.0246]	[0.0251]
Ind.(Quintile 4 Elas)	-0.0617***	-0.0753***	-0.0914***
	[0.0223]	[0.0247]	[0.0278]
Ind.(Quintile 5 Elas)	-0.0835***	-0.1041***	-0.0876***
	[0.0207]	[0.0233]	[0.0292]
"Upstream Contractibility"			
X Ind.(Quintile 1 Elas)	-0.1685**	-0.2170***	-0.2270***
	[0.0684]	[0.0635]	[0.0640]
X Ind.(Quintile 2 Elas)	-0.0966**	-0.0673	-0.0834
	[0.0436]	[0.0721]	[0.0802]
X Ind.(Quintile 3 Elas)	0.0533	0.0616*	0.1049***
	[0.0443]	[0.0362]	[0.0382]
X Ind.(Quintile 4 Elas)	0.0476	0.1650***	0.1105***
	[0.0443]	[0.0398]	[0.0373]
X Ind.(Quintile 5 Elas)	0.1204***	0.1962***	0.2434***
	[0.0390]	[0.0352]	[0.0329]
p-value: Q5 at median Upst. Cont.	[0.0000]	[0.0001]	[0.0001]
Elasticity based on:	All codes	BEC cons. & cap. goods	BEC cons. only
Industry controls?	Y	Y	Y
Firm controls?	Y	Y	Y
Parent country fixed effects?	Y	Y	Y
Observations	84,171	62,377	44,895
No. of industries	459	305	219
R^2	0.2399	0.3174	0.3470

Notes: Each observation is a global ultimate parent firm in 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 ratio-upstreamness measure described in Section 3.2 of the main text. "Upstream Contractibility" is the log ratio of the total-requirements weighted-average upstreamness of high-contractibility inputs relative to that of low-contractibility inputs, where a median cutoff is used to distinguish between high- and low-contractibility inputs. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies in Column (1) are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006); Column (2) restricts this to HS codes classified as consumption or capital goods in the UN BEC, while Column (3) further restricts this to consumption goods only. All columns include parent country fixed effects, as well as the set of parent industry characteristics and parent firm controls used in Tables 3 and 4.

Table 6
Robustness: Across Different Subsamples

Dependent variable:	Log Ratio-Upstreamness Measure			
	Emp.>=20	**Emp.>=20 & Subs.>=2**	**Emp.>=20 & MNC**	**Emp.>=20 & MNC & SICs>=2**
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 Elas)	-0.0450	-0.0467	-0.0516*	-0.0511*
	[0.0290]	[0.0304]	[0.0297]	[0.0298]
Ind.(Quintile 3 Elas)	-0.0603**	-0.0627**	-0.0468	-0.0455
	[0.0255]	[0.0280]	[0.0302]	[0.0304]
Ind.(Quintile 4 Elas)	-0.0931***	-0.0778***	-0.0616**	-0.0605**
	[0.0278]	[0.0295]	[0.0278]	[0.0282]
Ind.(Quintile 5 Elas)	-0.0987***	-0.0806**	-0.0667*	-0.0633*
	[0.0290]	[0.0323]	[0.0343]	[0.0353]
"Upstream Contractibility"				
X Ind.(Quintile 1 Elas)	-0.2208***	-0.2056***	-0.1858***	-0.1870***
	[0.0633]	[0.0652]	[0.0595]	[0.0604]
X Ind.(Quintile 2 Elas)	-0.0686	-0.0591	-0.0025	-0.0035
	[0.0803]	[0.0803]	[0.0576]	[0.0576]
X Ind.(Quintile 3 Elas)	0.0988**	0.1060*	0.0834	0.0853
	[0.0398]	[0.0568]	[0.0689]	[0.0693]
X Ind.(Quintile 4 Elas)	0.1173***	0.1052**	0.0854*	0.0832*
	[0.0393]	[0.0490]	[0.0435]	[0.0449]
X Ind.(Quintile 5 Elas)	0.2364***	0.2575***	0.2123***	0.2016***
	[0.0345]	[0.0369]	[0.0516]	[0.0531]
p-value: Q5 at median Upst. Cont.	[0.0000]	[0.0009]	[0.0631]	[0.0906]
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 fixed effects?	Y	Y	Y	Y
Observations	26,151	7,805	2,490	2,419
No. of industries	219	216	199	197
R^2	0.3307	0.3086	0.2403	0.2292

Notes: Each observation is a global ultimate parent firm in 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 ratio-upstreamness measure described in Section 3.2 of the main text. "Upstream Contractibility" is the log ratio of the total-requirements weighted-average upstreamness of high-contractibility inputs relative to that of low-contractibility inputs, where a median cutoff is used to distinguish between high- and low-contractibility inputs. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies in all columns are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent country fixed effects, as well as the set of parent industry characteristics and parent firm controls used in Tables 3 and 4.

Table 7
Robustness: Alternative constructions of Ratio-Upstreamness

Dependent variable:	Log Ratio-Upstreamness Measure				
	More cont. controls	**Random pick**	**"Ever-Integrated" Inputs**	**Mfg. Inputs only**	**Mfg. Inputs and Drop parent SIC**
	(1)	(2)	(3)	(4)	(5)
Ind.(Quintile 2 Elas)	-0.2932	-0.0396	-0.0494*	-0.0274	0.0237
	[0.2978]	[0.0285]	[0.0257]	[0.0318]	[0.0902]
Ind.(Quintile 3 Elas)	-1.0567***	-0.0633**	-0.0369	-0.0538*	-0.0915
	[0.3082]	[0.0253]	[0.0254]	[0.0293]	[0.0630]
Ind.(Quintile 4 Elas)	-0.7486**	-0.0886***	-0.0608**	-0.0884***	-0.1930**
	[0.3089]	[0.0278]	[0.0277]	[0.0307]	[0.0764]
Ind.(Quintile 5 Elas)	-0.6888**	-0.0819***	-0.0987***	-0.0923**	-0.2491**
	[0.2790]	[0.0295]	[0.0289]	[0.0359]	[0.0997]
"Upstream Contractibility"					
X Ind.(Quintile 1 Elas)	-0.1493	-0.2286***	-0.0705	-0.3133***	-0.2565***
	[0.1101]	[0.0635]	[0.0607]	[0.0695]	[0.0954]
X Ind.(Quintile 2 Elas)	-0.0862	-0.0807	-0.1097	-0.1058	0.1134
	[0.0838]	[0.0804]	[0.0943]	[0.0923]	[0.1278]
X Ind.(Quintile 3 Elas)	-0.1848*	0.1098***	0.1398***	0.1030	-0.2827
	[0.0972]	[0.0401]	[0.0534]	[0.0655]	[0.2202]
X Ind.(Quintile 4 Elas)	-0.0195	0.1044***	0.1246**	0.1204***	-0.3512**
	[0.0782]	[0.0388]	[0.0580]	[0.0396]	[0.1395]
X Ind.(Quintile 5 Elas)	0.1282**	0.2758***	0.2823***	0.1410**	-0.0239
	[0.0551]	[0.0410]	[0.0384]	[0.0582]	[0.2007]
p-value: Q5 at median Upst. Cont.	[0.0123]	[0.0002]	[0.0000]	[0.0026]	[0.0134]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls?	Y	Y	Y	Y	Y
Firm controls?	Y	Y	Y	Y	Y
Parent country fixed effects?	Y	Y	Y	Y	Y
Observations	44,895	44,895	44,895	44,780	14,503
No. of industries	219	219	219	218	216
R^2	0.3706	0.3558	0.2578	0.3339	0.1116

Notes: Each observation is a global ultimate parent firm in 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 log ratio-upstreamness measure used in Tables 3 and 4. "Upstream Contractibility" is constructed as the log ratio of the total-requirements weighted-average upstreamness of high contractibility inputs relative to that of low contractibility inputs, where a median cutoff is used to distinguish between high and low contractibility inputs. Results tabled here are based on the '1c' version of the contractibility measure. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies in all columns are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006) classified as consumption goods in the UN BEC. The additional contractibility controls included in Column 1 are the contractibility of the parent primary SIC, and the elasticity quintile dummies interacted with a total-requirements weighted average standard deviation of the contractibility of the inputs used. The remaining columns use different constructions of the ratio-upstreamness dependent variable, as described in the main text. All columns include parent country fixed effects, as well as the set of parent industry characteristics and parent firm controls used in Tables 3 and 4.

Table 8
Integration Decisions within Firms (Top 100 Inputs)

Dependent variable:	**Indicator variable: Input Integrated?**			
	(1)	(2)	(3)	(4)
Upstreamness_ij				
X Ind.(Quintile 1 Elas_j)	-0.0068***	0.0016	0.0021	-0.0037*
	[0.0009]	[0.0017]	[0.0017]	[0.0019]
X Ind.(Quintile 2 Elas_j)	-0.0093***	-0.0000	0.0002	-0.0045
	[0.0020]	[0.0036]	[0.0036]	[0.0037]
X Ind.(Quintile 3 Elas_j)	-0.0123***	-0.0022	-0.0016	-0.0040
	[0.0018]	[0.0042]	[0.0042]	[0.0038]
X Ind.(Quintile 4 Elas_j)	-0.0107***	0.0080***	0.0076***	0.0015
	[0.0016]	[0.0021]	[0.0020]	[0.0017]
X Ind.(Quintile 5 Elas_j)	-0.0127***	0.0061*	0.0059*	0.0027
	[0.0022]	[0.0033]	[0.0032]	[0.0025]
"Contractibility up to i" (in prod. of j)				
X Ind.(Quintile 1 Elas_j)		0.0323***	0.0356***	0.0278***
		[0.0070]	[0.0071]	[0.0062]
X Ind.(Quintile 2 Elas_j)		0.0375***	0.0378***	0.0295***
		[0.0134]	[0.0133]	[0.0107]
X Ind.(Quintile 3 Elas_j)		0.0378***	0.0360***	0.0324***
		[0.0114]	[0.0113]	[0.0097]
X Ind.(Quintile 4 Elas_j)		0.0699***	0.0668***	0.0446***
		[0.0123]	[0.0116]	[0.0083]
X Ind.(Quintile 5 Elas_j)		0.0761***	0.0750***	0.0521***
		[0.0153]	[0.0149]	[0.0120]
Contractibility of input i				
X Ind.(Quintile 1 Elas_j)			-0.0190***	-0.0079
X Ind.(Quintile 2 Elas_j)			-0.0106***	0.0019
X Ind.(Quintile 3 Elas_j)			-0.0193***	-0.0040
X Ind.(Quintile 4 Elas_j)			-0.0123***	0.0039
X Ind.(Quintile 5 Elas_j)			-0.0098*	0.0068
Dummy: Self-SIC	0.9760***	0.9651***	0.9636***	0.9275***
	[0.0018]	[0.0029]	[0.0030]	[0.0074]
p-value: Quintile 5 - Quintile 1 effect of "Contractibility up to i"	---	[0.0087]	[0.0157]	[0.0671]
Observations	1,452,817	1,452,817	1,452,817	1,452,817
No. of parent firms	14,503	14,503	14,503	14,503
No. of i-j pairs	21,635	21,635	21,635	21,635
R^2	0.4990	0.5008	0.5015	0.5253

Notes: Each observation is a SIC input by global ultimate parent pair; inputs ranked in the top 100 by total requirements coefficients of the primary SIC output industry of the parent 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 in question is integrated within the boundaries of the global parent firm. The "contractibility up to i" measure is the share of the total-requirements weighted-average contractibility of inputs that has been accrued in production upstream of and including input i in the production of output j. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent firm fixed effects, while Column 4 further includes SIC input industry fixed effects. Standard errors for auxiliary control variables are suppressed to save space (full details available on request).

Table 9
Integration Decisions over the Top 100 Inputs: By Elasticity Quintile

Dependent variable:	**Indicator variable: Input Integrated?**				
BEC cons. Elas_j:	Quintile 1 (1)	Quintile 2 (2)	Quintile 3 (3)	Quintile 4 (4)	Quintile 5 (5)
Contractibility up to i (in prod. of j)	0.0338***	0.0264***	0.0321***	0.0312***	0.0532***
	[0.0063]	[0.0077]	[0.0094]	[0.0098]	[0.0150]
Upstreamness_ij	0.0001	-0.0072*	-0.0030	0.0008	0.0001
	[0.0018]	[0.0043]	[0.0044]	[0.0021]	[0.0031]
Dummy: Self-SIC	0.9217***	0.9247***	0.9401***	0.8226***	0.8767***
	[0.0128]	[0.0266]	[0.0135]	[0.0448]	[0.0378]
Firm fixed effects?	Y	Y	Y	Y	Y
Input industry (i) fixed effects?	Y	Y	Y	Y	Y
Observations	332,351	408,227	271,730	222,704	217,805
No. of parent firms	3317	4074	2710	2227	2175
No. of input-output (ij) industry pairs	4206	4411	4304	4401	4313
R^2	0.5158	0.5565	0.4957	0.5636	0.5661

Notes: Each observation is a SIC input by global ultimate parent pair; inputs ranked in the top 100 by total requirements coefficients of the primary SIC output activity of the parent 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 in question is integrated within the boundaries of the global parent firm. The "contractibility up to i" measure is the share of the total-requirements weighted-average contractibility of inputs that has been accrued in production upstream of and including input i in the production of output j. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent firm fixed effects and SIC input industry fixed effects.

Appendix Table 1
Summary Statistics: Industry Characteristics

	10th	Median	90th	Mean	Std Dev
SIC characteristics (459 industries)					
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
Log (Skilled Emp./Worker)	-1.750	-1.363	-0.778	-1.308	0.377
Log (Capital/Worker)	3.493	4.428	5.591	4.495	0.794
Log (Equip. Capital / Worker)	2.869	4.043	5.163	4.039	0.867
Log (Plant Capital / Worker)	2.517	3.302	4.524	3.426	0.755
Log (Materials / Worker)	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 (Rauch cons., homog. only)	0.006	0.021	0.183	0.073	0.132
Contractibility (Rauch cons., homog.+ref.priced)	0.091	0.362	0.816	0.410	0.265
Upst. contractibility (Rauch cons., homog. only)	0.659	1.011	1.498	1.054	0.333
Upst. contractibility (Rauch cons., homog.+ref.priced)	0.549	0.914	1.438	0.966	0.352

Notes: For the 459 SIC manufacturing industries.

Appendix Table 2
Effect of Upstream Contractibility: Alternative Measure

Dependent variable:	Log Ratio-Upstreamness		
	(1)	(2)	(3)
Ind.(Quintile 2 Elas)	-0.0407	-0.0740**	-0.0572
	[0.0282]	[0.0337]	[0.0363]
Ind.(Quintile 3 Elas)	-0.1150***	-0.0871**	-0.0998***
	[0.0295]	[0.0362]	[0.0297]
Ind.(Quintile 4 Elas)	-0.1126***	-0.1576***	-0.1528***
	[0.0312]	[0.0271]	[0.0262]
Ind.(Quintile 5 Elas)	-0.1417***	-0.1748***	-0.1592***
	[0.0289]	[0.0275]	[0.0269]
"Upstream Contractibility"			
X Ind.(Quintile 1 Elas)	-1.2784***	-1.5249***	-1.8220***
	[0.4564]	[0.3683]	[0.3826]
X Ind.(Quintile 2 Elas)	-0.8160***	-0.3932	-0.6059
	[0.2640]	[0.4604]	[0.5864]
X Ind.(Quintile 3 Elas)	0.4082*	-0.0452	0.0563
	[0.2361]	[0.3314]	[0.3535]
X Ind.(Quintile 4 Elas)	0.3364	1.0129***	0.6766***
	[0.2762]	[0.2170]	[0.1989]
X Ind.(Quintile 5 Elas)	0.7606***	1.0618***	1.2564***
	[0.1941]	[0.1913]	[0.2188]
p-value: Q5 at median Upst. Cont.	[0.0000]	[0.0000]	[0.0000]
Elasticity based on:	All codes	BEC cons. & cap. goods	BEC cons. only
Industry controls?	Y	Y	Y
Firm controls?	Y	Y	Y
Parent country fixed effects?	Y	Y	Y
Observations	84,171	62,377	44,895
No. of industries	459	305	219
R^2	0.2568	0.3286	0.3531

Notes: Each observation is a global ultimate parent firm in 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 log ratio-upstreamness measure used in Tables 3 and 4. "Upstream Contractibility" is constructed as a total-requirements weighted-covariance between the upstreamness and contractibility of the input. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies in Column (1) are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006); Column (2) restricts this to HS codes classified as consumption or capital goods in the UN BEC, while Column (3) further restricts this to consumption goods only. All columns include parent country fixed effects, as well as the set of parent industry characteristics and parent firm controls used in Tables 3 and 4.

Appendix Table 3
Addressing parent firms with multiple SIC output activities

Dependent variable:	**Log Ratio-Upstreamness Measure**			
	Restrict to single SIC code parents		**Parent firm by SIC output (two-way cluster)**	
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 Elas)	-0.0782	-0.0375	-0.0769*	-0.0379
	[0.0490]	[0.0301]	[0.0410]	[0.0280]
Ind.(Quintile 3 Elas)	-0.1140**	-0.0721***	-0.0901**	-0.0505*
	[0.0448]	[0.0261]	[0.0390]	[0.0263]
Ind.(Quintile 4 Elas)	-0.1489***	-0.0893***	-0.1504***	-0.0938***
	[0.0485]	[0.0297]	[0.0407]	[0.0269]
Ind.(Quintile 5 Elas)	-0.1886***	-0.0805***	-0.1871***	-0.0876***
	[0.0476]	[0.0305]	[0.0424]	[0.0297]
"Upstream Contractibility"				
X Ind.(Quintile 1 Elas)		-0.2353***		-0.2159***
		[0.0638]		[0.0612]
X Ind.(Quintile 2 Elas)		-0.0965		-0.0588
		[0.0857]		[0.0782]
X Ind.(Quintile 3 Elas)		0.1330***		0.0826*
		[0.0367]		[0.0429]
X Ind.(Quintile 4 Elas)		0.1063**		0.1058***
		[0.0413]		[0.0369]
X Ind.(Quintile 5 Elas)		0.2466***		0.2527***
		[0.0349]		[0.0370]
p-value: Q5 at median Upst. Cont.		[0.0004]		[0.0017]
Elasticity based on:	BEC cons. only	BEC cons. only	BEC cons. only	BEC cons. only
Industry controls?	Y	Y	Y	Y
Firm controls?	Y	Y	N	Y
Parent country fixed effects?	Y	Y	Y	Y
Observations	32,126	32,126	64,281	64,281
No. of industries	218	218	---	---
R^2	0.2764	0.3673	0.2633	0.3270

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Columns 1 and 2 restrict the regression sample to those parent firms that report only one primary SIC activity. Robust standard errors clustered by SIC industry are reported. For Columns (3) and (4), each observation is a parent firm by SIC output activity pair; the ratio-upstreamness dependent variable is constructed for each separate SIC output activity of each parent firm. Two-way clustered standard errors -- by parent firm and by SIC output activity -- are reported. "Upstream Contractibility" is constructed as the log ratio of the total-requirements weighted-average upstreamness of high contractibility inputs relative to that of low contractibility inputs, where a median cutoff is used to distinguish between high and low contractibility inputs. Results tabled here are based on the '1c' version of the contractibility measure. The quintile dummies in all columns are based on an elasticity measure constructed from all available HS10 elasticities in Broda-Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent country fixed effects, as well as the set of parent industry characteristics and parent firm controls used in Tables 3 and 4.