

NBER WORKING PAPER SERIES

INTERNALIZING GLOBAL VALUE CHAINS:
A FIRM-LEVEL ANALYSIS

Laura Alfaro
Pol Antràs
Davin Chor
Paola Conconi

Working Paper 21582
<http://www.nber.org/papers/w21582>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2015

We thank participants at the following conferences: ERWIT, Barcelona GSE Summer Forum, Princeton IES Summer Workshop, NBER ITI and OE meetings, ETSG, Asia Pacific Trade Seminars, the AEA meetings, the EEA meetings, the Global Fragmentation of Production and Trade Policy conference (ECARES), the World Bank GVC conference, the World Bank Kuala Lumpur conference, and the Trade and Macro Interdependence in the Age of GVCs conference (Lithuania). In addition, we thank seminar audiences at LSE, the Paris Trade Workshop, MIT Sloan, Boston College, Warwick, Ferrara, Munich, Sapienza, Bologna, Nottingham, Bank of Italy, HKUST, HKU, NUS, Singapore Management University, UIBE, Nottingham Ningbo, and University of Tokyo. We are particularly grateful to Kamran Bilir, Arnaud Costinot, Thibault Fally, Ali Horta csu, Thierry Mayer, Peter Morrow, Claudia Steinwender, David Weinstein, and the anonymous referees for their detailed comments on a previous draft. Chor thanks colleagues at the Global Production Networks Centre (GPN@NUS) for their engaging discussions. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2015 by Laura Alfaro, Pol Antràs, Davin Chor, and Paola Conconi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Internalizing Global Value Chains: A Firm-Level Analysis
Laura Alfaro, Pol Antràs, Davin Chor, and Paola Conconi
NBER Working Paper No. 21582
September 2015, Revised December 2016
JEL No. D23,F14,F23,L20

ABSTRACT

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

Laura Alfaro
Harvard Business School
Morgan Hall 263
Soldiers Field
Boston, MA 02163
and NBER
lalfaro@hbs.edu

Pol Antràs
Department of Economics
Harvard University
1805 Cambridge Street
Littauer Center 207
Cambridge, MA 02138
and NBER
pantras@fas.harvard.edu

Davin Chor
Department of Economics
National University of Singapore
1 Arts Link
Singapore 117570
davinchor@nus.edu.sg

Paola Conconi
ECARES, Université Libre de Bruxelles
Avenue F.D. Roosevelt, 42
1050 Bruxelles - Belgium
and CEPR
pconconi@ulb.ac.be

1 Introduction

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

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

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

Although the global fragmentation of production has featured prominently in the trade literature (e.g., Johnson and Noguera, 2012), much less attention has been placed on how the position of a given production stage in the value chain affects firm boundary choices, and firm organizational decisions more broadly. Furthermore, most studies on this topic have been mainly theoretical in nature.¹ To a large extent, this theoretical bias is explained by the challenges one faces when taking models of global value chains to the data. Ideally, researchers would like to access comprehensive datasets that would enable them to track the flow of goods within value chains across borders and organizational forms. Trade statistics are useful in capturing the flows of goods when they cross a particular border, and some countries' customs offices also record whether goods flow in and out of a country within or across firm boundaries. Nevertheless, once a good leaves a country, it is

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

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

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

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

Perhaps surprisingly, we show that the key predictions of Antràs and Chor (2013) continue to hold in this richer environment with input asymmetries. In particular, a firm’s decision to integrate upstream or downstream suppliers depends crucially on the relative size of the elasticity of demand for its final good and the elasticity of substitution across production stages. When demand is elastic or inputs are not particularly substitutable, inputs are sequential complements, in the sense that the

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

marginal incentive of a supplier to undertake relationship-specific investments is higher, the larger are the investments by upstream suppliers. In this case, the firm finds it optimal to integrate only the most downstream stages, while contracting at arm's length with upstream suppliers in order to incentivize their investment effort. When instead demand is inelastic or inputs are sufficiently substitutable, inputs are sequential substitutes, i.e., investments by upstream suppliers lower the investment incentives of downstream suppliers. When this is the case, the firm would choose to integrate relatively upstream stages, while engaging in outsourcing to downstream suppliers. While the profile of marginal productivities and costs along the value chain does not detract from this core prediction, it does shape the measure of stages (i.e., how many inputs) the firm ends up finding optimal to integrate in both the complements and the substitutes cases.

We develop several extensions of the benchmark model that are relevant for our empirical analysis. First, we map the asymmetries across inputs to differences in their inherent degree of contractibility. We show that the propensity of a firm to integrate a given stage is shaped in subtle ways by the contractibility of upstream and downstream stages. Intuitively, in production processes that feature a high degree of contractibility among upstream relative to downstream inputs, firms need to rely less on the organizational mode to counteract the distortions associated with inefficient investments upstream. Hence, high levels of upstream contractibility tend to reduce the set of outsourced stages when inputs are sequential complements, while reducing the set of integrated stages when inputs are sequential substitutes.

Second, we incorporate heterogeneity across final good producers in their core productivity, while introducing fixed costs of integrating suppliers, as in Antràs and Helpman (2004). With these features, more productive firms would (*ceteris paribus*) integrate a larger number of inputs, in both the complements and substitutes cases. This is because more productive firms find it easier to amortize the fixed costs associated with integrating suppliers, and thus find it optimal to integrate stages that smaller firms can only profitably outsource. This extension also suggests that productivity differences within an industry should have a distinct effect on integration choices: more productive firms should have a higher propensity to integrate downstream (relative to upstream) suppliers when inputs are sequential substitutes, but a higher propensity to integrate upstream (relative to downstream) suppliers when inputs are sequential complements.

Finally, we consider a scenario in which integration is infeasible for certain segments of the value chain, for example, due to exogenous technological or regulatory factors. We show that even when integration is sparse (as is the case in our data), the model's predictions continue to describe firm boundary choices for those inputs along the value chain over which integration is feasible.

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

In our empirical analysis, we study the determinants of a firm’s propensity to integrate upstream versus downstream inputs. To distinguish between integrated and non-integrated inputs, we rely on the methodology of Fan and Lang (2000), combining information on firms’ reported activities with Input-Output tables (see also Acemoglu *et al.*, 2009; and Alfaro *et al.*, 2016). To capture the position of different inputs along the value chain, we compute a measure of the upstreamness of each input i in the production of output j using U.S. Input-Output Tables. This extends the measure of the upstreamness of an industry with respect to final demand from Fally (2012) and Antràs *et al.* (2012) to the bilateral industry-pair level. To provide a test of the model, we exploit information from WorldBase on the primary activity of each firm, and use estimates of demand elasticities from Broda and Weinstein (2006), as well as measures of contractibility from Nunn (2007).

We first examine how firms’ organizational choices depend on the elasticity of demand for their final good. In line with the first prediction of the model, we find that the higher the elasticity of demand faced by the parent firm, the lower the average upstreamness of its integrated inputs relative to the upstreamness of its non-integrated inputs. This result is illustrated in a simple (unconditional) form in Figure 1, based on different quintiles of the parent firm’s elasticity of demand. As seen in the left panel of the figure, the average upstreamness of integrated inputs is much higher when the parent company belongs to an industry with a low demand elasticity than when it belongs to one associated with a high demand elasticity. Conversely, the right panel shows that the average upstreamness of non-integrated stages is greater the higher the elasticity of demand faced by the parent’s final good.³




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

The above pattern is robust in the regression analysis, even when controlling for a comprehensive

³Figure 1 is plotted using only inputs i that rank within the top 100 manufacturing inputs in terms of total requirements coefficients of the parent’s output industry j . The average for each firm is computed weighting each input by its total requirements coefficient tr_{ij} , while excluding integrated stages belonging to the same industry j as the parent; a simple unweighted average across firms in the elasticity quintile is then illustrated. The figures obtained when considering all manufacturing inputs, when computing unweighted averages over inputs, and when considering the output industry j as an input are all qualitatively similar.

list of firm characteristics (e.g., size, age, employment, sales), using different measures of the demand elasticity, as well as in different subsamples of firms (e.g., restricting to domestic firms, or to multinationals). We also show that our results hold in specifications where the elasticity of demand is replaced by the difference between this same elasticity and a proxy for the degree of input substitutability associated with the firm’s production process. We reach a similar conclusion when we exploit within-firm variation in integration patterns. In these specifications, we find that a firm’s propensity to integrate is generally larger for downstream inputs, but disproportionately so for firms facing high demand elasticities.

We report two further empirical regularities that are strongly consistent with the model’s implications. First, we find that firms’ ownership decisions are shaped by the contractibility of upstream versus downstream inputs: a greater degree of “upstream contractibility” increases the likelihood that a firm integrates upstream inputs, when the firm faces a high elasticity of demand (i.e., in the complements case); conversely, it increases the propensity to outsource upstream inputs, when the firm’s demand elasticity is low (i.e., in the substitutes case). This is in line with the intuition that greater upstream contractibility lowers a firm’s need to rely on decisions over organizational mode to elicit the right incentives from suppliers positioned at early stages in the value chain. Second, we find that more productive firms integrate more inputs in industries across all the demand elasticity quintiles. Moreover, consistent with the theory, more productive firms exhibit a higher propensity to integrate relatively downstream (respectively, upstream) inputs when the elasticity of demand for their final product is low (respectively, high).

This body of findings suggests that contractual frictions play a crucial role in shaping the integration choices of firms around the world, with the patterns being consistent with a view of firm boundary choices that is rooted in the property-rights approach to the theory of the firm. As we shall discuss later in the paper, it would be much harder to rationalize our empirical results invoking the transaction-cost theory of Coase and Williamson, which views integration as a means to circumvent contractual frictions vis-à-vis independent suppliers. More generally, the rich differential effects predicted and observed in the complements and substitutes cases are not straightforward to rationalize by invoking alternative theories of firm boundaries.

It is useful to further discuss how our analysis relates to other recent work on vertical linkages at the firm level. In an influential study, Atalay *et al.* (2014) find little evidence of intrafirm shipments between related plants within the United States; they instead present complementary evidence which indicates that firm boundaries are more influenced by the transfer of intangible inputs, than by the transfer of physical goods. Our theoretical model is abstract enough to allow one to interpret the sequential investments as resulting in either tangible or intangible transfers across establishments; and our empirical analysis takes into account both manufacturing and non-manufacturing inputs (including services). That said, due to the inherent difficulties of recording and measuring intangible inputs, we believe that our empirical results speak more to the optimal provision of incentives along sequential value chains involving tangible inputs. It is important to stress, however, that our findings should not be interpreted as invalidating the intangibles hypothe-

sis. Relatedly, our analysis suggests that intrafirm trade flows are an imperfect proxy for the extent to which firms react to contractual insecurity by internalizing particular stages of their global value chains. As the “sparse integration” extension of our model shows, internalization decisions along value chains are consistent with an arbitrarily low level of intrafirm trade relative to the overall transaction volume in these chains. This helps reconcile our findings with those of Ramondo *et al.* (2016), who find that intrafirm trade between U.S. multinationals and their affiliates abroad is highly concentrated among a small number of large affiliates.

By conducting our analysis at the firm level, we are able to greatly improve upon the empirical evidence provided in Antràs and Chor (2013), which was based on industry-level data on U.S. intrafirm import shares and lacked direct information on the U.S. entity internalizing these foreign purchases. Moreover, by extending the theory in several directions, we have generated a richer set of predictions about firms’ boundary choices that we can bring to the data. Our work is closely related to two contemporaneous papers with similar goals. Del Prete and Rungi (2015) employ a dataset of about 4,000 multinational business groups to explore the correlation between the average “downstreamness” of integrated affiliates (relative to final demand) and that of the parent firm itself (also relative to final demand). They find that this correlation varies depending on the size of the demand elasticity faced by the parent firm, in a manner reminiscent of the predictions in Antràs and Chor (2013). Their work is however silent on the production line position of non-integrated inputs and does not incorporate an industry-pair measure of the upstreamness of affiliates relative to their parents. Luck (2014) reports corroborating evidence based on city-level evidence on the export-import activities of processing firms in China, though his work adopts a value-added notion of production line position (rather than one rooted in actual production staging). As insightful as these contributions are, we view the empirical strategy developed in this paper as a more direct firm-level test of the propositions of the theory. More generally, our paper is related to a recent empirical literature testing various aspects of the property-rights theory of multinational firm boundaries. This includes Antràs (2003), Yeaple (2006), Nunn and Treffer (2008, 2013), Corcos *et al.* (2013), Defever and Toubal (2013), Díez (2014), and Antràs (2015), among others.⁴

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 the data. Section 4 outlines our empirical methodology and presents our findings in detail. Section 5 concludes. The appendices contain additional material related to both the theory and the empirical analysis.

2 Theoretical Framework

In this section, we develop our model of sequential production. We first describe a generalized version of the model in Antràs and Chor (2013) that incorporates heterogeneity across inputs beyond their position along the value chain. We then consider several extensions to derive additional theoretical results and enrich the set of predictions that can be brought to the data.

⁴Even more broadly, our work is related to the extensive empirical literature on firm boundaries, which is nicely overviewed in Lafontaine and Slade (2007), and Bresnahan and Levin (2012).

2.1 Benchmark Model with Heterogeneous Inputs

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:

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

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), 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, where we rely on sectoral estimates of demand elasticities. Given that we largely focus on the problem of a representative firm, we abstain from indexing variables by firm or sector to keep the notation tidy.

Obtaining the finished product requires the completion of a unit measure of production stages. These stages are indexed by $i \in [0, 1]$, with a larger i corresponding to stages further downstream and thus closer to the finished product. Denote 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:

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

where θ is a productivity parameter, $\alpha \in (0, 1)$ is a parameter that captures the (symmetric) degree of substitutability among the stage inputs, the shifters $\psi(i)$ reflect asymmetries in the marginal product of different inputs' investments, 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.⁵

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, the supplier of input i must undertake a relationship-specific investment entailing a marginal cost of $c(i)$ 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

⁵In fact, one can 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 $\psi(i) x(i)$.

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

This concludes the description of the assumptions of the model. Figure 2 outlines the timing of events of the game played by the firm and the unit measure of suppliers. Antràs and Chor (2013) provide an extensive discussion of the robustness of their key results should ex-ante transfers between the firm and the suppliers be allowed, and under alternative bargaining protocols that allow suppliers to lay claim over part of the revenues that are realized downstream of i . Although a similar robustness analysis could be carried out in this current richer framework, we will abstain from doing so due to space constraints.

Despite the presence of additional sources of input asymmetries, captured by the functions $\psi(i)$ and $c(i)$, the subgame perfect equilibrium of the above game can be derived in a manner similar to Antràs and Chor (2013). 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

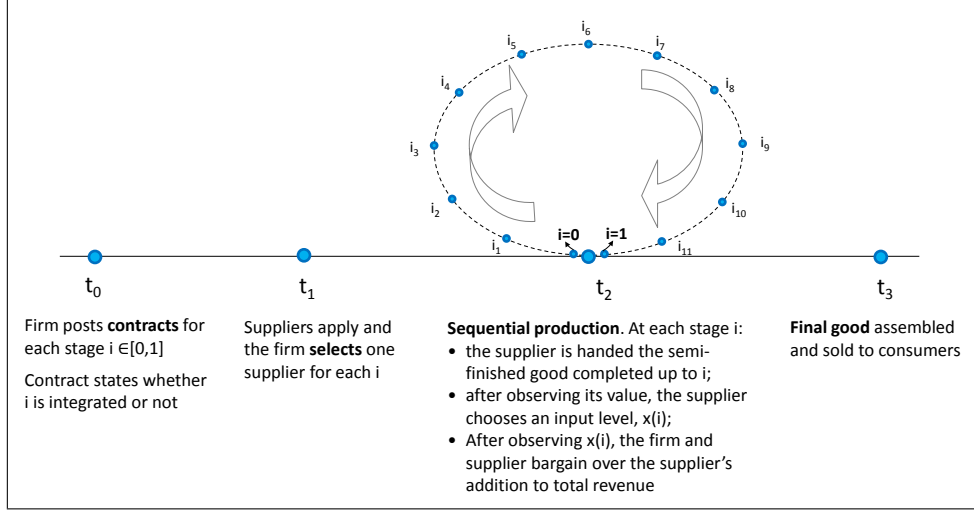


Figure 2: Timing of Events

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

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

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

Now consider the bargaining between the firm and the supplier at stage m . Because inputs are customized to the needs of the firm, the supplier's outside option at the bargaining stage is 0 and the quasi-rents over which the firm and the supplier negotiate are given by the incremental contribution to total revenue generated by supplier m at that stage. Applying Leibniz' rule to (3), this is given by:

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

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

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

Notice that the marginal return to investing in $x(m)$ is increasing in the demand level A , while it decreases in the marginal cost c . Furthermore, this marginal return is increasing in supplier m 's bargaining share $1 - \beta(m)$, and thus, other things equal, outsourcing provides higher-powered incentives for the supplier to invest. This is a standard feature of property-rights models. The more novel property of program (5) is that a supplier's marginal return to invest at stage m is shaped by all investment decisions in prior stages, i.e., $\{x(i)\}_{i=0}^m$, as captured by the value of production secured up to stage m , i.e., $r(m)$. The nature of such dependence is in turn crucially shaped by the

relative size of the demand elasticity parameter ρ and the input substitutability parameter α . When $\rho > \alpha$, investment choices are *sequential complements* in the sense that higher investment levels by upstream suppliers increase the marginal return of supplier m 's own investment. Conversely, when $\rho < \alpha$, investment choices are *sequential substitutes* because high values of upstream investments reduce the marginal return to investing in $x(m)$. We shall refer to $\rho > \alpha$ as the *complements* case and to $\rho < \alpha$ as the *substitutes* case, as in Antràs and Chor (2013).

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

Plugging the first-order condition from (5) into (4), and solving the resulting separable differential equation, we show in Appendix A-1 that one can express the equilibrium volume of input m services $x^*(m)$ as a function of the whole path of bargaining shares $\{\beta(i)\}_{i \in [0, m]}$ up to stage m :

$$x^*(m) = A\theta^{\frac{\rho}{1-\rho}} \left(\frac{1-\rho}{1-\alpha}\right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{1}{1-\rho}} \left(\frac{1-\beta(m)}{c(m)}\right)^{\frac{1}{1-\alpha}} \psi(m)^{\frac{\alpha}{1-\alpha}} \left[\int_0^m \left(\frac{(1-\beta(i))\psi(i)}{c(i)}\right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}}. \quad (6)$$

It is then straightforward to see that $x^*(m) > 0$ for all m as long as $\beta(m) < 1$. This in turn implies that the firm has every incentive to abide by the proper (or technological) sequencing of production, so that $I^*(m) = 1$ for all m (consistent with our expressions above).

To complete the description of the equilibrium, we roll back to the initial period prior to any production taking place, in which the firm decides whether the contract associated with a given input m is associated with integration or outsourcing. This amounts to choosing $\{\beta(i)\}_{i \in [0, 1]}$ to maximize $\pi_F = \int_0^1 \beta(i)r'(i)di$, with $r'(m)$ given in equation (4), $x^*(m)$ in equation (6), and $\beta(i) \in \{\beta_V, \beta_O\}$. After several manipulations, the problem of choosing the optimal organizational structure can be reduced to the program:

$$\begin{aligned} \max_{\beta(i)} \quad & \pi_F = \Theta \int_0^1 \beta(i) \left(\frac{(1-\beta(i))\psi(i)}{c(i)}\right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di \\ \text{s.t.} \quad & \beta(i) \in \{\beta_V, \beta_O\}, \end{aligned} \quad (7)$$

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

It will prove useful to consider a relaxed version of program (7) in which rather than constraining $\beta(i)$ to equal β_V or β_O , we allow the firm to freely choose the function $\beta(i)$ from the whole set of piecewise continuously differentiable real-valued functions. Defining:

$$v(i) \equiv \int_0^i \left(\frac{(1-\beta(k))\psi(k)}{c(k)}\right)^{\frac{\alpha}{1-\alpha}} dk, \quad (8)$$

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

$$\pi_F(v) = \Theta \int_0^1 \left(1 - v'(i)^{\frac{1-\alpha}{\alpha}} \frac{c(i)}{\psi(i)} \right) v'(i) v(i)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di. \quad (9)$$

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

$$\beta^*(m) = 1 - \alpha \left[\frac{\int_0^m (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk} \right]^{\frac{\alpha-\rho}{\alpha}}. \quad (10)$$

Notice that because the term inside the square brackets is a monotonically increasing function of m , expression (10) 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 .⁶ In the complements case ($\rho > \alpha$), the incentive to integrate suppliers increases as we move downstream in the value chain. Intuitively, given sequential complementarity, the firm is particularly concerned about incentivizing upstream suppliers to raise their investment effort, in order to generate positive spillovers on the investment levels of downstream suppliers. Instead, in the substitutes case ($\rho < \alpha$), the firm is less concerned with underinvestment by upstream suppliers, while capturing rents upstream is particularly appealing when marginal revenue falls quickly with output.

A remarkable feature of equation (10) is that the slope of $\partial\beta^*(m)/\partial m$ is governed by the sign of $\rho - \alpha$ *regardless* of the paths of $\psi(k)$ and $c(k)$. It is worth pausing to explain why this result is not straightforward. Notice that a disproportionately high value of $\psi(m)$ at a given stage m can be interpreted as that stage being relatively important in the production process. Indeed, in a model with complete contracts, the share of m in the total input purchases of the firm would be a monotonically increasing function of $\psi(m)$. According to one of the canonical results of the property-rights literature, one would then expect the incentive to outsource such a stage to be particularly large (see, in particular, Proposition 1 in Antràs, 2014). Intuitively, outsourcing provides higher-powered incentives to suppliers, and minimizing underinvestment inefficiencies is particularly beneficial for inputs that are relatively important in production. In terms of the notation of the model, one might have thus expected the optimal division of surplus $\beta^*(m)$ to be decreasing in stage m 's importance $\psi(m)$. For the same reason, and given that input shares are monotonically decreasing in the marginal cost $c(m)$, one might have also expected the share $\beta^*(m)$ to be increasing in $c(m)$. As intuitive as this reasoning might appear, one would then be led to conclude that if the path of $\psi(m)$ were sufficiently increasing in m – or the path of $c(m)$ were

⁶Although Antràs and Chor (2013) considered a variant of their model with heterogeneity in $\psi(i)$ and $c(i)$, they 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).

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 (10) 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 result bears some resemblance to the classic result in consumption theory that an agent’s dynamic utility-maximizing level of consumption should be growing or declining over time according to whether the real interest rate is greater or less than the rate of time preference, regardless of the agent’s income path. It is important to stress, however, that the paths of $\psi(m)$ and $c(m)$ are not irrelevant for the incentive to integrate suppliers along the value chain (in the same manner that the path of income is not irrelevant in the dynamic consumption problem). Equation (10) illustrates that the incentives to integrate a particular input will be notably shaped by the size of the ratio $\psi(k)/c(k)$ for inputs upstream from input m relative to the average size of this ratio along the whole value chain.

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

So far, we have focused on a characterization of the optimal bargaining share $\beta^*(m)$, but the above results can easily be turned into statements regarding the propensity of firms to integrate ($\beta^*(m) = \beta_V$) or outsource ($\beta^*(m) = \beta_O$) the different stages of the value chain. In particular, in Appendix A-1 we show that:

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. 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.

Figure 3 illustrates that the main result in Proposition 1 concerning the optimal pattern of ownership along the value chain depends critically on whether the stage inputs are sequential complements or substitutes. When the demand faced by the final-good producer is sufficiently




Figure 3: Firm Boundary Choices along the Value Chain

elastic, then there exists a unique cutoff stage such that all inputs prior to that cutoff are outsourced, and all inputs (if any) downstream of it are integrated. The converse prediction holds when demand is sufficiently inelastic (i.e., in the sequential substitutes case): the firm would instead integrate relatively upstream inputs, while outsourcing would take place relatively downstream.

Although the last statement in Proposition 1 follows pretty immediately from our discussion of the properties of the solution $\beta^*(m)$ to the relaxed problem, it can also be shown more directly by explicitly characterizing the thresholds m_C^* and m_S^* . For the sequential complements case, we show in Appendix A-1 that, provided that integration and outsourcing coexist along the value chain, the threshold m_C^* is given by:

$$\frac{\int_0^{m_C^*} (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk} = \left\{ 1 + \left(\frac{1-\beta_O}{1-\beta_V} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{1-\frac{\beta_O}{\beta_V}}{1-\left(\frac{1-\beta_O}{1-\beta_V}\right)^{-\frac{\alpha}{1-\alpha}}} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}. \quad (11)$$

Notice then that the larger the value of $\psi(k)/c(k)$ in upstream production stages (in the numerator of the left-hand side) relative to downstream production stages, the lower the value of m_C^* will be; the set of integrated stages will thus be larger.⁷ (The analogous expression for m_S^* in the substitutes case is reported in Appendix A-1.)

2.2 Extensions

In this section, we develop three extensions of our framework to further guide the firm-level empirical analysis conducted later in this paper.

A. Heterogeneous Contractibility of Inputs

So far, we have been agnostic about the underlying drivers of input heterogeneity in the model. In order to develop empirical tests of Proposition 1 – and especially its last statement – it is important to map variation in the ratio $\psi(m)/c(m)$ along the value chain to certain observables. With that in mind, in this section we explore the link between $\psi(m)$ and the degree of contractibility of

⁷In the complements case, integration and outsourcing coexist along the value chain when $\beta_V(1-\beta_V)^{\frac{\alpha}{1-\alpha}} > \beta_O(1-\beta_O)^{\frac{\alpha}{1-\alpha}}$, which ensures $m_C^* < 1$. When instead $\beta_V(1-\beta_V)^{\frac{\alpha}{1-\alpha}} < \beta_O(1-\beta_O)^{\frac{\alpha}{1-\alpha}}$, the firm finds it optimal to outsource *all* stages, i.e., $m_C^* = 1$.

different stage inputs. In Appendix A-1, we shall also briefly relate marginal cost variation in $c(m)$ along the value chain to the sourcing location decisions of the firm.⁸

Remember that in our benchmark model, $x(m)$ captures the services related to the non-contractible aspects of input production, in the sense that the volume $x(m)$ cannot be disciplined via an initial contract and is chosen unilaterally by suppliers. Conversely, we shall now assume that $\psi(m)$ encapsulates investments and other aspects of production that are specified in the initial contract in a way that precludes any deviation from that agreed level. In light of equation (2), 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 .⁹ 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 non-contractible investments be constant along the value chain, i.e., $c(m) = c$ for all m .

In terms of the timing of events summarized in Figure 2, notice that nothing has changed except for the fact that the initial contract also specifies the profit-maximizing choice of $\psi(m)$ along the value chain. Furthermore, once the levels of $\psi(m)$ have been set at stage t_0 , the subgame perfect equilibrium is identical to that in our previous model in which $\psi(m)$ was assumed exogenous. This implies that the firm's optimal ownership structure along the value chain will seek to maximize the program in (7), and the solution of this problem will be characterized by Proposition 1.

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

$$\tilde{\pi}_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} c^{\frac{-\rho}{1-\rho}} \Gamma(\beta_O, \beta_V) \left[\int_0^1 \psi(i)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - \int_0^1 \frac{(\psi(i))^\phi}{\mu(i)} di, \quad (12)$$

where remember that $\Theta = A\theta^{\frac{\rho}{1-\rho}} \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}} > 0$, and where $\Gamma(\beta_O, \beta_V) > 0$ is a function of β_O and β_V , as well as of α and ρ (see Appendix A-1 for the full expression). The choice of the profit-maximizing path of $\psi(m)$ will thus seek to maximize π_F in (12).

A notable feature of equation (12) 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

⁸In the absence of contractual frictions, $\psi(m)/c(m)$ would be 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.

⁹Acemoglu *et al.* (2007) also model input production as involving a Cobb-Douglas function of contractible and non-contractible 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.

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 non-contractible 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 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 (12) imply that for any two inputs at stages m and m' , we have:

$$\left(\frac{\psi(m)}{\psi(m')}\right)^{\phi - \frac{\alpha}{1-\alpha}} = \frac{\mu(m)}{\mu(m')}. \quad (13)$$

For the second-order conditions of problem (12) 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 $1/\mu(m)$.¹⁰ In light of our discussion in the last section, this implies:

Proposition 2. 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 the contractibility $\mu(m)$ for upstream inputs relative to downstream inputs.

Figure 4 illustrates the key result of Proposition 2. Intuitively, the higher the contractibility of the upstream inputs, the less firms need to rely on upstream organizational decisions as a way to counteract the distortions associated with inefficient investments by upstream suppliers. As a consequence, high levels of upstream contractibility tend to reduce the set of outsourced stages whenever final-good demand is elastic or inputs are not too substitutable, while they tend to reduce the set of integrated stages whenever final-good demand is inelastic or inputs are highly substitutable.




Figure 4: The Effect of an Increase in Upstream Contractibility

By mapping variation in $\psi(m)$ to the degree of input contractibility, Proposition 2 helps operationalize our previous Proposition 1. More specifically, in our empirical analysis, we will employ empirical proxies for input contractibility to develop a sector-level measure of the extent to which non-contractibilities feature disproportionately in upstream versus downstream stages in the pro-

¹⁰The inequality $\phi > \alpha/(1-\alpha)$ is necessary but not sufficient for the second-order conditions to be satisfied (see Appendix A-1).

duction 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.

B. Heterogeneous Productivity of Final Good Producers

Our model incorporates heterogeneity across final good producers in terms of their demand level A and their core productivity θ . In this section, we show how such heterogeneity shapes firm boundary choices along the value chain, in the presence of fixed organizational costs associated with vertically integrating production stages. More specifically, we shall now assume that if a firm wants to integrate a given stage $i \in [0, 1]$, it needs to pay a fixed cost equal to $f_V > 0$.¹¹

In order to facilitate a swifter transition to the empirical analysis, we shall revert back to our benchmark model with exogenous paths of $\psi(i)$ and $c(i)$. We will relegate most mathematical details to Appendix A-1, in which we show that Proposition 1 continues to apply in an environment with fixed costs of integration. More precisely, there continue to exist thresholds $m_C^* \in (0, 1]$ and $m_S^* \in (0, 1]$ such that all production stages $m \in [0, m_C^*)$ are outsourced and all stages $m \in [m_C^*, 1]$ are integrated in the complements case, while all production stages $m \in [0, m_S^*)$ are integrated and all stages $m \in [m_S^*, 1]$ are outsourced in the substitutes case. Furthermore, one can still show that both m_C^* and m_S^* are lower, the higher is the ratio $\psi(m)/c(m)$ for upstream inputs relative to downstream inputs.

These characterization results can be obtained even though the equations determining the cut-offs m_C^* and m_S^* are now significantly more involved. For instance, m_C^* is now the solution to the following implicit function:

$$\left(\frac{\psi(m_C^*)}{c(m_C^*)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^{m_C^*} \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \times \left\{ \left(1 - \frac{\beta_O}{\beta_V} \right) - \left(1 - \left(\frac{1-\beta_V}{1-\beta_O} \right)^{\frac{\alpha}{1-\alpha}} \right) \left[1 + \left(\frac{1-\beta_V}{1-\beta_O} \right)^{\frac{\alpha}{1-\alpha}} \frac{\int_{m_C^*}^1 \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk}{\int_0^{m_C^*} \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk} \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \right\} = \frac{f_V}{\Psi A \theta^{\frac{\rho}{1-\rho}}}, \quad (14)$$

where $\Psi = (1 - \beta_O)^{\frac{\rho}{1-\rho}} \beta_V \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}}$. Invoking the second-order conditions necessary for the solution m_C^* to be unique, we can establish that the left-hand side of (14) is increasing in m_C^* , and thus this threshold is necessarily a decreasing function of the level of firm demand A or firm productivity θ . Following analogous steps, in Appendix A-1 we show that, in the substitutes case, m_S^* is instead increasing in both A and θ . In words, this result indicates that regardless of the sign of $\rho - \alpha$, relatively more productive firms will tend to integrate a larger interval of production stages. The intuition behind this is simple: more productive firms will find it easier to amortize

¹¹Our results below would continue to hold in the presence of fixed costs f_O associated with outsourcing stages, as long as those fixed costs are lower than f_V .

the fixed cost associated with integrating more stages.

In our empirical analysis, we will explore whether the observed intra-industry heterogeneity in integration choices is in accordance with these predictions, which we summarize as:

Proposition 3. In the presence of fixed costs of integration, the statements in Proposition 1 continue to hold. Furthermore, the cutoff m_C^* is decreasing in firm-level demand A and firm-level productivity θ , while m_S^* is increasing in A and θ .




Figure 5: The Effect of an Increase in Productivity of the Final Good Producer

Figure 5 illustrates how an increase in the productivity θ of the final good producer (or an increase in firm-level demand A) affects integration choices along the value chain. The interval of integrated stages expands in both cases, but in a manner that would lead us to observe relatively more internalization of upstream stages when inputs are sequential complements, and conversely relatively more internalization of downstream stages in the substitutes case.

C. Sparse Integration and Intrafirm Trade

Our framework has the strong implication that the sets of integrated and outsourced stages are both connected and jointly constitute a partition of $[0, 1]$. As might have been expected, this strong prediction of the model is not borne out in the data. In fact, integrated stages are very sparse in our dataset and the overwhelming majority of them “border” with outsourced stages immediately upstream and downstream from them.¹² This paucity of integration might be due to technological or regulatory factors that make vertical integration infeasible for certain production stages. We next briefly outline a third extension of our model that accommodates such sparsity, and we demonstrate that it does not undermine the validity of the key predictions of the model that we will take to the data.

A simple way to render integration infeasible for certain segments of the value chain is to assume that the fixed costs of integrating those segments is arbitrarily large. In terms of our second extension above, we thus now have that the fixed cost of integration is stage-specific and takes a value of $f_V(m) = +\infty$ for any $m \in \Upsilon$, where Υ is the set of stages that cannot possibly

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

be integrated. For simplicity, we assume that the fixed costs of integration are finite and identical for all remaining stages, so $f_V(m) = f_V$ for $m \in \Omega$, where Ω is the set of integrable stages (i.e., $\Omega = [0, 1] \setminus \Upsilon$). Clearly, by making the set Υ larger and larger, one can make integration decisions arbitrarily sparse in our model. As we show in Appendix A-1, despite the presence of the exogenously non-integrable stages, we can establish that:

Proposition 4. If $\rho > \alpha$, the firm cannot possibly find it optimal to integrate a positive measure of stages located upstream from a positive measure of outsourced stages $(\tilde{m}, \tilde{m} + \varepsilon) \in \Omega$ that could have been integrated. If $\rho < \alpha$, the firm cannot possibly find it optimal to integrate a positive measure of stages located downstream from a positive measure of outsourced stages $(\tilde{m}, \tilde{m} + \varepsilon) \in \Omega$ that could have been integrated.

Naturally, Proposition 4 provides a much weaker characterization of the integration decisions of firms along their value chain than our previous Propositions 1-3. Yet, a corollary of Proposition 4 is that, holding constant the set Υ of stages that cannot possibly be integrated, the average upstreamness of integrated stages relative to the average upstreamness of outsourced stages should be lower when $\rho > \alpha$ than when $\rho < \alpha$. This relative upstreamness of integrated and non-integrated stages is what we refer to as “ratio-upstreamness” in our regression analysis, and will be one of the key metrics employed to assess the empirical validity of the model.

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

3 Dataset and Key Variables

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

3.1 The WorldBase Dataset

Our core firm-level dataset is Dun & Bradstreet’s (D&B) WorldBase, which provides comprehensive coverage of public and private companies across more than 100 countries and territories. WorldBase

has been used extensively in the literature, in particular to explore research questions related to the organizational practices of firms around the world.¹³ Cross-country studies at the firm level are challenging, as there are few high-quality datasets that are comparable across countries; when such data are available, these tend to be limited to advanced countries. One of the advantages of WorldBase is thus the inclusion of a wide set of countries at different levels of development.¹⁴

Another key advantage is that the unit of observation is the establishment, namely a single physical location where industrial operations or services are performed, or business is conducted. Each establishment in WorldBase is assigned a unique identifier, called a DUNS number.¹⁵ Where applicable, the DUNS number of the global ultimate owner is also reported, which allows us to keep track of ownership linkages within the dataset. In addition, WorldBase provides information on: (i) the location (address) of each establishment; (ii) the 4-digit SIC code (1987 version) of its primary industry, and the SIC codes of up to five secondary industries; (iii) the year it was started or in which current ownership took control; and (iv) basic data on employment and sales.

Note that each firm in the data is either: (i) a single-establishment firm; or (ii) is identified in WorldBase as a “global ultimate”. The former refers to a business entity whose entire activity is in one location, and which does not report ownership links with other establishments in WorldBase. For the latter, D&B WorldBase defines a “global ultimate” to be the top, most important, responsible entity within a corporate family tree, that has more than 50% ownership of other establishments. We link each global ultimate to all its identified majority-owned subsidiaries, both in manufacturing and non-manufacturing, by using the DUNS number of the global ultimate that is reported for establishments. The set of integrated SIC activities for a single-establishment plant is simply the list of up to six SIC codes associated with it. The set of integrated SIC codes for a global ultimate is the complete list of SIC activities that is performed either in its headquarters or by one of its subsidiaries. Moving forward, we will refer simply to each observation as a “parent” firm, indexed by p .

For our analysis, we use the 2004/2005 WorldBase vintage and focus on parent firms in the manufacturing sector – i.e., whose primary SIC code lies between 2000 and 3999 – with a minimum total employment (across all establishments) of 20. To be clear, while each parent firm in our sample has a primary SIC code in manufacturing, we nevertheless include all the parent firm’s integrated SIC activities (both in manufacturing and non-manufacturing) in the exercise that follows. In all, our sample contains 320,254 parent firms from 116 countries; 259,312 of these are

¹³An early example was Caves’ (1975) analysis of size and diversification patterns between Canadian and U.S. plants. More recent uses include Harrison *et al.* (2004), Acemoglu *et al.* (2009), Alfaro and Charlton (2009), Alfaro and Chen (2014), Fajgelbaum *et al.* (2015), and Alfaro *et al.* (2016).

¹⁴The data in WorldBase is compiled from a large number of underlying sources, including partner firms, business registers, telephone directory records, company websites, and even self-registration. See Alfaro and Charlton (2009) for a more detailed discussion, and comparisons with other data sources such as the Bureau of Economic Analysis (BEA) data on U.S. multinational activity.

¹⁵D&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.

single-establishment firms, while 60,942 are global ultimates. Among the global ultimates, 6,370 observations have subsidiaries in more than one country, and can thus be labeled as multinational firms. Panel A of Appendix Table A-1 provides some descriptive statistics for our full sample, as well as for the subset of multinationals. Not surprisingly, multinationals are on average larger in terms of employment, sales and number of integrated SIC codes, as compared to the typical firm in our data. We will show nevertheless that our core findings concerning the relationship between “upstreamness” and integration patterns are stable when we look at different subsamples.

3.2 Key Variables

Integrated and Outsourced Inputs

For each parent, WorldBase provides us with information on the inputs that are integrated within the firm’s ownership boundaries. In order to further identify which inputs are outsourced, we combine the above with information from U.S. Input-Output (I-O) Tables, following the methodology of Fan and Lang (2000). (See also Acemoglu *et al.*, 2009, and Alfaro *et al.*, 2016.)

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

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

To implement the above, we turn to the 1992 U.S. Benchmark I-O Tables from the Bureau of Economic Analysis (BEA).¹⁶ The U.S. Tables are one of the few publicly-available I-O accounts that provide a level of industry detail close to the 4-digit SIC codes used in WorldBase, while the 1992 vintage is the most recent year for which the BEA provides a concordance from its I-O industry classification to the 1987 SIC system.¹⁷ Readers familiar with these tables will be aware that the concordance is not a one-to-one key. This is not a major problem given our focus on parents whose primary output j is in manufacturing, as the key assigns a unique 6-digit I-O industry to each 4-digit SIC code between 2000 and 3999. Outside these sectors, in those inputs i whose 6-digit I-O industry code maps to multiple 4-digit SIC codes, we split the total requirements value tr_{ij} equally across the multiple SIC codes that i maps to.

Panel A of Table A-1 shows that the mean tr_{ij} value associated with the inputs integrated by firms in WorldBase is 0.019241 (or 0.006774 when the I-O diagonal entries are dropped); this is larger than the average tr_{ij} value across the 416,349 (i, j) pairs in the I-O Tables that are relevant to our study (0.001311). In other words, firms tend to integrate stages that are more important in terms of total requirements usage.¹⁸ We can further report that 98.0% of the (i, j) pairs in our WorldBase sample, namely inputs i that are integrated by a parent firm with output industry j , are relevant for production in the sense that $tr_{ij} > 0$.¹⁹ As mentioned before, firms tend to integrate very few of the inputs necessary to produce their final good. The median number of integrated stages is 2, compared to a median number of non-integrated stages equal to 906. There is considerable skewness, however, as the corresponding 90th, 95th, and 99th percentiles of the number of integrated stages are 3, 4, and 6, while the maximum number is 254.²⁰ As discussed below, however, integrated inputs tend to be “bunched” together along the value chain, consistent with our model.

Upstreamness

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

¹⁶The BEA draws on records of movements across establishments when constructing the U.S. I-O Tables. See Chapters 3 and 6 of the BEA’s “Concepts and Methods of the U.S. Input-Output Accounts”, available at http://www.bea.gov/papers/pdf/IOmanual_092906.pdf.

¹⁷This concordance is available from: <http://www.bea.gov/industry/exe/ndn0017.exe>.

¹⁸In the 1992 U.S. I-O Tables, there are altogether 416,349 I-O pairs that are relevant to our study, namely that involve an SIC manufacturing output j and an SIC input i (either in manufacturing or non-manufacturing), with $tr_{ij} > 0$. Of these, 57,057 or 13.7% can be found in our WorldBase firm sample of integrated input by parent primary industry pairs. The share is very similar if the input-output pairs along the diagonal are excluded from consideration (13.6% = 56,612/415,904).

¹⁹85.6% of these pairs actually exceed the median tr_{ij} value of 0.000163 (where this median is taken over the same 416,349 I-O pairs from the preceding footnote). We obtain similarly high relevance rates when restricting the count to manufacturing inputs only, or if we drop the self-SIC of the parent firm (i.e., pairs where $i = j$).

²⁰The median number of integrated inputs is very similar when computed industry-by-industry, varying between 1 and 3. On the other hand, the maximum number of integrated inputs exhibits more variation across industries, ranging from 3 to 254 (with a median value of 26).

$$upst_{ij} = \frac{dr_{ij} + 2 \sum_{k=1}^N dr_{ik} dr_{kj} + 3 \sum_{k=1}^N \sum_{l=1}^N dr_{ik} dr_{kl} dr_{lj} + \dots}{dr_{ij} + \sum_{k=1}^N dr_{ik} dr_{kj} + \sum_{k=1}^N \sum_{l=1}^N dr_{ik} dr_{kl} dr_{lj} + \dots}. \quad (15)$$

Observe that dr_{ij} is the value of i that enters exactly one stage prior to the production of j , that $\sum_{k=1}^N dr_{ik} dr_{kj}$ is the value of i that enters two stages prior to production of j , and so on and so forth. The denominator in (15) is therefore equal to tr_{ij} , written as an infinite sum over the value of i 's use that enters exactly n stages removed from the production of j (where $n = 1, 2, \dots, \infty$). The numerator is similarly an infinite sum, but there each input use term is multiplied by an integer equal to the number of stages upstream at which the input value enters the production process. Looking then at (15), $upst_{ij}$ is a weighted average of the number of stages it takes for i to enter in j 's production, where the weights correspond to the share of tr_{ij} that enters at that corresponding upstream stage. In particular, a larger $upst_{ij}$ means that a greater share of the total input use value of i is accrued further upstream in the production process for j . We thus refer to $upst_{ij}$ simply as the ‘‘upstreamness’’ of i in the production of j .

Note that $upst_{ij} \geq 1$ by construction, with equality if and only if $tr_{ij} = dr_{ij}$, namely when the entirety of the input use of i goes directly into the production of j via one stage. With some matrix algebra, one can see that the numerator of (15) is equal to the (i, j) -th entry of $[I - D]^{-2}D$. Together with the formula for tr_{ij} noted earlier (i.e., the (i, j) -th entry of $[I - D]^{-1}D$), one can then calculate $upst_{ij}$ when provided with the direct requirements matrix, D .

Two additional remarks are in order. First, we should stress the distinction between $upst_{ij}$ and the upstreamness measure put forward previously in Fally (2012) and Antràs *et al.* (2012). The measure in this earlier work served to capture the average production line position of each industry i with respect to final demand (i.e., consumption and investment), whereas our current $upst_{ij}$ instead reflects the position of input i with respect to output industry j . This is therefore a measure of production staging specific to each input-output industry pair, which we can directly map to the firm-level observations in our dataset to assess the validity of the model's predictions. Second, $upst_{ij}$ also has the interpretation of an ‘‘average propagation length’’, a concept introduced in Dietzenbacher *et al.* (2005) to capture the average number of stages taken by a shock in i to spread to industry j . Dietzenbacher *et al.* (2005) in fact show that this average propagation length has the appealing property that it is invariant to whether one adopts a forward or backward linkage perspective when computing the average number of stages between a pair of industries.

We use the direct requirements matrix derived from the 1992 U.S. I-O Tables to calculate $upst_{ij}$.²¹ We first obtain $upst_{ij}$ for each 6-digit I-O industry pair, before mapping these to 4-digit SIC codes. As mentioned earlier, each 4-digit manufacturing SIC code is mapped to a single 6-digit I-O code; this means that we can uniquely assign an $upst_{ij}$ value to SIC code pairs where both the input i and output j are in manufacturing. The complications arise only when we have a non-

²¹We apply an open-economy and net-inventories correction to the direct requirements matrix D , before calculating tr_{ij} and $upst_{ij}$. This involves a simple adjustment to each dr_{ij} to take into account input flows across borders, as well as into and out of inventories, on the assumption that these flows occur in proportion to what is observed in domestic input-output transactions; see Antràs *et al.* (2012) for details.

manufacturing input i which maps to multiple 6-digit I-O codes. We adopt a range of approaches in such cases, by taking either: (i) the simple mean of $upst_{ij}$ over constituent I-O codes of the SIC input industry; (ii) the median value; (iii) a random pick; or (iv) the tr_{ij} -weighted average value. Reassuringly, the pairwise correlation of the upstreamness measures obtained under these different treatments is very high (> 0.98), and our regression results will not depend on which specific approach we adopt, so we will focus on the version that uses a simple mean as our baseline. To be clear, what this yields is a measure of the average number of production stages based on the I-O classification system that are traversed between a given pair of SIC industries.

Panel B of Appendix Table A-1 presents some basic information on the total requirements and upstreamness variables after the mapping to SIC codes. Figure 6 provides an illustration of the variation contained in the $upst_{ij}$ measure, even when focusing on one particular input industry, in this case Tires and Inner Tubes (SIC 3011). Notice that $upst_{ij}$ is indeed smaller for j sectors such as Mobile Homes (2451), Lawn and Garden Equipment (3524), Industrial Trucks and Tractors (3537), Motorcycles, Bicycles, and Parts (3751), and Transportation Equipment, n.e.c. (3799), these being industries that use tires almost exclusively as a direct input. For comparison and to illustrate the difference, Figure 6 also depicts the upstreamness of Tires with respect to final demand (from Antràs *et al.* 2012); this is the horizontal line with value 2.0954.




Figure 6: Upstreamness of Tires (SIC 3011) in the Production of all Other Manufacturing Industries

As noted before, firms tend to integrate few inputs. This is a key feature of the data that our model can accommodate, as explained in the discussion on “sparse integration” in Section 2.2.C. Our upstreamness measure allows us to examine the extent to which – though sparse – integrated inputs nevertheless tend to be “bunched” together, consistent with the environment described in

this earlier extension on “sparse integration”. To do so, we focus on firms that report at least two secondary manufacturing SIC codes (on top of their primary output industry j). Table A-2 computes the probability that a pair of randomly drawn integrated manufacturing SICs of a given firm would belong to any two quintiles of $upst_{ij}$, where j is the output industry of the firm and the quintiles are taken over all SIC manufacturing inputs i in the value chain for producing j ; the reported probability is an average across all firms under consideration. From Table A-2, one can see that firms are clearly more likely to integrate inputs in the first quintile of upstreamness than in the other quintiles. Leaving aside this first quintile, note that the probability that a firm integrates an input is significantly higher when it already owns an input in the same quintile, and furthermore these probabilities fall for quintiles that are further apart. These patterns are suggestive of the existence of “bunching” along the value chain in the integration decisions of firms.

Ratio-Upstreamness

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

$$R_{jp} = \frac{\sum_{i \in I(p)} \theta_{ijp}^I upst_{ij}}{\sum_{i \in NI(p)} \theta_{ijp}^{NI} upst_{ij}}, \quad (16)$$

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

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

Our first set of regression specifications will use “ratio-upstreamness” as the dependent variable,

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

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

Demand Elasticity

As highlighted in our theory, the incentives to integrate upstream or downstream suppliers are crucially affected by whether the elasticity of demand faced by the firm (ρ_j) is higher or lower than the elasticity of technological substitution across its inputs (α_j). For practical reasons, we focus on variation in the former in most of the regressions, since detailed estimates of demand elasticities are available from standard sources. To capture ρ_j , we use the U.S. import demand elasticities from Broda and Weinstein (2006). The original estimates are for HS10 products, and we average these up to the SIC industry level using U.S. import trade values as weights (see Appendix A-2 for further details). Since the HS10 codes are highly disaggregated, this should in principle provide a good proxy for ρ_j in the model, short of having actual firm-level elasticities. We will also pursue several refinements of ρ_j , by using only elasticities for those HS10 codes deemed as consumption and capital goods in the United Nations’ Classification by Broad Economic Categories (BEC). (The omitted category is goods classified as intermediates.) As the model arguably applies better to final goods, a demand elasticity constructed based on such products should yield a cleaner proxy for ρ_j . Note that when refining the construction in this manner, about half of the 459 SIC manufacturing industries are dropped, namely those industries composed entirely of intermediate goods.

The UN BEC classification also provides a basis for constructing a proxy for α_j . From the model, α_j is closely related to the elasticity of demand for each intermediate input by firms in industry j . We therefore begin by computing the average demand elasticity for each 4-digit SIC code using now only those HS10 elasticities that correspond to products classified as *intermediates*, in an analogous fashion to the construction of the ρ_j refinements above. We construct our proxy for α_j as the weighted average of the intermediate-good demand elasticities across inputs i used in j ’s production, with weights proportional to the total requirements coefficients, tr_{ij} .

In principle, the value of $\rho_j - \alpha_j$ could then be used to distinguish whether a given industry j falls in the complements or substitutes case. Nevertheless, since our proxies of ρ_j and especially α_j are imperfect, in our baseline regressions we will associate the sequential complements case with high values of ρ_j and the substitutes case with low values of ρ_j . This approach is valid insofar

as the demand elasticity and input substitutability parameters are relatively uncorrelated across industries.²² For corroboration, we will also report specifications in which the complements and substitutes cases are related to the size of the difference $\rho_j - \alpha_j$.

Input Contractibility

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

We start by following Nunn (2007) in constructing a measure of input contractibility for each SIC industry. The basis for this measure is the Rauch (1999) classification of products into whether they are: (i) homogeneous; (ii) reference-priced; or (iii) differentiated in nature. The “contract-intensity” of an industry is then the share of the constituent HS product codes in the composition of the industry’s input use that is classified as differentiated (i.e., neither homogeneous nor reference-priced), on the premise that it is inherently more difficult to specify and enforce the terms of contractual agreements for such products. As our interest is in the converse concept of contractibility, we use instead one minus the Nunn measure of contract-intensity.²³ Denote this metric of input contractibility for industry i by $cont_i$. Then, for each output industry j , we calculate $UpstCont_j$ as a weighted covariance between the upstreamness and the contractibility of its manufacturing inputs, namely:

$$UpstCont_j = \sum_{i \in S^m(j)} \theta_{ij}^m (upst_{ij} - \overline{upst}_{ij}) (cont_i - \overline{cont}_i), \quad (17)$$

where $S^m(j)$ is the set of all manufacturing inputs used in the production of j (i.e., with $tr_{ij} > 0$). The weights are given by: $\theta_{ij}^m = tr_{ij} / \sum_{k \in S^m(j)} tr_{kj}$, while $\overline{upst}_{ij} = \sum_{i \in S^m(j)} \theta_{ij}^m upst_{ij}$ and $\overline{cont}_i = \sum_{i \in S^m(j)} \theta_{ij}^m cont_i$ are total requirements weighted averages of the upstreamness and contractibility variables respectively. Therefore, if high-contractibility inputs tend to be located at earlier production stages, this will lead to a larger (more positive) covariance and hence a higher $UpstCont_j$; we refer to such an industry as exhibiting a greater degree of “upstream contractibility”.²⁴

In the within-firm regressions, we can perform a more detailed test of the role of contractibility

²²Indeed, the pairwise correlation between the constructed proxy for α_j and the measures of ρ_j (both the baseline measure and its refinements) is low, ranging between -0.026 and 0.083 . As reported in Table A-3, our proxies for α_j are on average higher than those for ρ_j .

²³In Nunn’s (2007) notation, the measure of input contractibility that we use is equal to $1 - z^{rs1}$. The results reported are based on the “conservative” Rauch (1999) classification, but are robust when using the alternative “liberal” classification instead.

²⁴We have also experimented with alternative measures of $UpstCont_j$, by taking a ratio of the tr_{ij} -weighted upstreamness of inputs classified as being of high contractibility relative to those classified as low contractibility, in a manner analogous to the construction of the ratio-upstreamness measure in (16). To distinguish high- versus low-contractibility inputs, we have adopted either the first tercile, median, or second tercile values of $cont_i$ across the 459 SIC manufacturing industries as cutoffs. The results with these alternative versions of $UpstCont_j$ all continue to lend strong support to the model (results available on request).

in explaining the propensity to integrate particular inputs. Motivated by the theory, we construct the variable $ContUpTol_{ij}$, which is an input-output industry pair-specific measure of the “contractibility up to i in the production of j ”. This is computed as:

$$ContUpTol_{ij} = \frac{\sum_{k \in S_i^m(j)} tr_{kj} cont_k}{\sum_{k \in S^m(j)} tr_{kj} cont_k}, \quad (18)$$

where the relevant set of inputs, $S_i^m(j)$, for the sum in the numerator is those manufacturing inputs that are located upstream of and including i itself in the production of j , i.e., $S_i^m(j) = \{k \in S^m(j) : upst_{kj} \geq upst_{ij}\}$. The denominator thus sums up the product of the total requirements and contractibility values across all manufacturing inputs used in the production of j , with the numerator being the partial sum excluding all inputs downstream of i . The construction of (18) is intended to approximate the $\frac{\int_0^i (\psi(k))^{\frac{1-\alpha}{\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{1-\alpha}{\alpha}} dk}$ term, which appears in equation (10) in the theory. There, it was shown that “contractibility up to i ” plays a central role in the expression for the optimal β^* , and hence the propensity towards integration of each input in the production of j .

4 Empirical Methodology and Results

We translate the propositions from the model into a series of empirical predictions that can be taken to the data. According to Proposition 1, integration patterns along the value chain should vary systematically for industries that fall under the sequential complements versus substitutes cases. Our approach to distinguish between these two cases focuses on variation in ρ_j (or alternatively, $\rho_j - \alpha_j$). However, the limitations inherent in how ρ_j and α_j are constructed mean that we cannot use them to precisely delineate where the cutoff between the complements and substitutes cases lies. Consequently, what we test in our regressions is a milder version of Proposition 1, that examines whether the propensity to integrate upstream stages falls as ρ_j (alternatively, $\rho_j - \alpha_j$) increases, and we more confidently move towards the complements case. We thus formulate the first cross-firm prediction of our model as follows:

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

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

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

The first extension of the model developed in Section 2.2.A provides us with further predictions that emerge from considering heterogeneity in the contractibility of inputs. In particular, Propo-

sition 2 suggests that the relative propensity to integrate upstream inputs depends on the extent to which contractible inputs tend to be located in the early stages of production. Moreover, the effect of “upstream contractibility” varies subtly across the sequential complements and substitutes cases. The second cross-firm prediction of our model can be summarized as:

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

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

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

From the second extension of the model developed in Section 2.2.B, we can derive predictions concerning the role of the productivity of final good producers. The results in Proposition 3 can be stated in testable form as follows:

P.3: More productive firms should integrate more inputs, irrespective of ρ_j (or $\rho_j - \alpha_j$). Relative to less productive firms, they should have a higher propensity to integrate downstream (relative to upstream inputs) when ρ_j (alternatively, $\rho_j - \alpha_j$) is low, and a higher propensity to integrate upstream (relative to downstream inputs) when ρ_j (alternatively, $\rho_j - \alpha_j$) is high.

4.1 Cross-firm Results

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

$$\log R_{jpc} = \beta_0 + \beta_1 \mathbf{1}(\rho_j > \rho_{med}) + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \quad (19)$$

The dependent variable is the log ratio-upstreamness measure, defined in equation (16), which captures the propensity of each parent p with primary SIC industry j to integrate relatively upstream inputs. Note that the subscript c is introduced to index the country where the parent is located, as we will include a full set of country fixed effects, D_c , among the controls. We report standard errors clustered at the level of the SIC output industry j .

The key regressor of interest is the dummy variable $\mathbf{1}(\rho_j > \rho_{med})$, which identifies whether the elasticity of demand ρ_j is above the median value of ρ across industries. This variable is meant to pick out industries that fall under the sequential complements case. Prediction P.1 (Cross)

suggests that β_1 should be negative: As we transition to industries that fall under the complements case, the propensity to integrate upstream relative to downstream inputs should fall. For all the specifications that we describe below, we will also run regressions in which the demand elasticity ρ_j is instead replaced by our proxy for $\rho_j - \alpha_j$; in (19), this means that we will test P.1 (Cross) using an indicator variable for whether $\rho_j - \alpha_j$ exceeds its median value across industries j .

We include a list of industry and firm controls in the above specification. The vector X_j includes measures of factor intensity, R&D intensity, and a value-added to shipments ratio (see Appendix A-2 for a more detailed description, as well as Table A-3 for basic summary statistics). The vector W_p contains parent firm characteristics obtained from WorldBase. This includes several variables that reflect the size of the parent, namely the number of establishments, whether it is a multinational, as well as log total employment and log total sales.²⁵ We also account for the age of the parent by including the year of its establishment (or in which current ownership took control). We view X_j and W_p strictly as auxiliary controls, in the sense that the model does not deliver direct predictions that would lead us to clearly sign their effects on the ratio-upstreamness measure.

Table 1 reports the results of estimating (19). Column (1) presents a stripped-down specification in which only $\mathbf{1}(\rho_j > \rho_{med})$ and parent country fixed effects are included. The estimated coefficient on our proxy for the complements case is negative and significant at the 10% level, already confirming that the propensity to integrate upstream stages is lower in industries that face a high demand elasticity, consistent with prediction P.1 (Cross). This result becomes even more significant (at the 1% level) as we successively add the output industry variables X_j in column (2), and the parent controls W_p in column (3). Looking at these auxiliary variables, the estimates indicate that there is a tendency towards upstream integration in more equipment capital-intensive industries, as well as in firms with more establishments, younger firms, and in multinationals.

The remaining columns in Table 1 explore alternative elasticity measures to capture industries in the complements case. Column (4) restricts the construction of ρ_j to the use of product-level elasticities classified by the UN BEC as either consumption or capital goods (dropping the intermediate-use products), while column (5) further limits this to just consumption goods elasticities. These refinements would in principle yield elasticities that pertain more directly to final-goods demand. Reassuringly, this does not change the key finding of a negative and highly significant coefficient on the high-elasticity dummy, even though the number of observations falls as SIC industries that are composed entirely of intermediate-use goods are dropped from the sample. Finally, column (6) brings in information related to the demand elasticity for intermediate inputs, through the proxy for α_j . The key right-hand side variable is now an indicator for whether $\rho_j - \alpha_j$ is larger or smaller than its median value, where ρ_j is the demand elasticity from column (5) based on consumption goods only and the construction of α_j was described earlier (in Section 3). We continue to find that the propensity to integrate upstream stages is lower for industries that more likely correspond to the complements case on the basis of $\rho_j - \alpha_j$.²⁶

²⁵For employment and sales, we also include dummy variables for whether the respective variables were based on actual data or were otherwise estimated/approximated by WorldBase.

²⁶Figures A-1 and A-2 in the Appendix illustrate these patterns of integration decisions using examples from our

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

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

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

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

$$\log R_{jpc} = \beta_0 + \sum_{n=2}^5 \beta_n \mathbf{1}(\rho_j \in \text{Quint}_n(\rho)) + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \quad (20)$$

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

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

Although we do not deny that our empirical results might be consistent with other theoretical frameworks, it is worth highlighting that they are hard to rationalize with the transaction-cost approach of Coase and Williamson. This latter approach is based on the idea that, when dealing with integrated suppliers, firms can settle issues by fiat, authority, or disciplinary action, thereby circumventing the contractual frictions that plague market transactions with independent suppliers. In Appendix A-1, we develop a ‘transaction-cost’ variant of our model, and show that it delivers a set of implications that are diametrically opposite to those of our property-rights model, and that are thus inconsistent with our empirical results.²⁷

To assess the validity of prediction P.2 (Cross), we augment the specifications in (19) and (20)

sample of firms. A Danish firm active in Boat Building and Repairing (SIC 3732) illustrates the complements case, as this sector exhibits an above-median ρ_j and $\rho_j - \alpha_j$ value regardless of the variant of the demand elasticity proxy considered. The firm has integrated only one SIC activity other than its primary SIC, namely Internal Combustion Engines (SIC 3519), which is clearly one of the most downstream among the top 100 manufacturing inputs by total requirements value used by SIC 3732 (see Figure A-1). Conversely, a Swedish producer of Household Furniture (SIC 2519) illustrates the substitutes case, as this sector is consistently classified with below-median ρ_j and $\rho_j - \alpha_j$ values. This firm has integrated only one SIC activity other than its primary SIC, namely Fabricated Metal Products (SIC 3499), which is among the most upstream of its top 100 manufacturing inputs (see Figure A-2).

²⁷To see the reason for this contrary result, take as an example the complements case, where underinvestment by upstream suppliers is particularly costly. Under the transaction-cost approach, this would call for the integration of upstream stages, in order to avert the underinvestment encountered when contracting with arm’s length suppliers.

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

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

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

in order to uncover the effects of upstream contractibility on integration:

$$\begin{aligned} \log R_{jpc} = & \beta_0 + \beta_1 \mathbf{1}(\rho_j > \rho_{med}) + \beta_{U1} \mathbf{1}(\rho_j < \rho_{med}) \times UpstCont_j + \beta_{U2} \mathbf{1}(\rho_j > \rho_{med}) \times UpstCont_j \\ & + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}, \text{ and} \end{aligned} \quad (21)$$

$$\begin{aligned} \log R_{jpc} = & \beta_0 + \sum_{n=2}^5 \beta_n \mathbf{1}(\rho_j \in Quint_n(\rho)) + \sum_{n=1}^5 \beta_{Un} \mathbf{1}(\rho_j \in Quint_n(\rho)) \times UpstCont_j \\ & + \beta_X X_j + \beta_W W_p + D_c + \epsilon_{jpc}. \end{aligned} \quad (22)$$

In the median cutoff specification in (21), we interact the dummy variables $\mathbf{1}(\rho_j < \rho_{med})$ and $\mathbf{1}(\rho_j > \rho_{med})$ with the “upstream contractibility” measure, $UpstCont_j$. Based on the second prediction of our model, we would expect $\beta_{U1} < 0$ and $\beta_{U2} > 0$ in (21). Likewise in (22), we interact each of the quintile dummies with $UpstCont_j$, where the theory would lead us to expect that $\beta_{U1} < 0$ and $\beta_{U5} > 0$.²⁸

The results of (21) are reported in Table 3. Notice that the estimated coefficient on the proxy for the complements case, $\mathbf{1}(\rho_j > \rho_{med})$, is negative and significant, as in the previous regressions in Table 1.²⁹ Turning to the interactions with $UpstCont_j$, the estimated coefficient in the complements case is positive and statistically significant, while that in the substitutes case is negative and also highly significant. This is entirely in line with the predictions of the model: Firms that fall under the complements case would have a lower propensity to integrate upstream stages, but this tendency is weakened among those industries whose production processes inherently exhibit a greater degree of upstream contractibility. The converse holds for the substitutes case, with $UpstCont_j$ instead lowering the propensity to integrate upstream stages when $\rho_j < \rho_{med}$. Note that these results hold when restricting the elasticity measure to HS codes classified as consumption or capital goods in column (2), when further limiting this to consumption goods elasticities only in column (3), and when using the proxy for $\rho_j - \alpha_j$ to distinguish between the two cases in column (4).

Table 4 confirms that the predictions related to upstream contractibility continue to hold with the more flexible quintile elasticity specification in (22). The main effects of the quintile elasticity dummies exhibit a pattern similar to that in the more parsimonious regressions in Table 2, with negative and significant coefficients especially as we transition to the higher quintiles. We perform a test for whether the effect of being in the fifth quintile, evaluated at the median in-sample value of $UpstCont_j$ in that fifth demand elasticity quintile, is in fact significantly different from zero. The p-values reported in each column confirm that this is indeed the case, so that the propensity to integrate upstream inputs is lower in the fifth relative to the first elasticity quintile; this holds true regardless of the variant of the elasticity proxy used across the columns. Of note, we find

²⁸The correlation between $UpstCont_j$ and the ρ_j proxy is small and never exceeds 0.06 in absolute value when we look across the various versions of the demand elasticity measure that we have constructed. The interaction terms are thus unlikely to be picking up a non-linear effect of the demand elasticity.

²⁹We have verified that the overall effect of the $\mathbf{1}(\rho_j > \rho_{med})$ variable – taking into account its main effect and that through the interaction term with upstream contractibility – is indeed negative when evaluated at the median in-sample value of $UpstCont_j$ for industries that exhibit an above-median ρ_j . The p-value for this coefficient test on the overall effect of $\mathbf{1}(\rho_j > \rho_{med})$ in the complements case is reported for each column in Table 3.

Table 3: Effect of Upstream Contractibility: Median Elasticity Cutoff

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Elas _j > Median)	-0.0910*** [0.0210]	-0.1306*** [0.0256]	-0.1432*** [0.0263]	-0.1372*** [0.0249]
Upstream Contractibility _j				
× Ind.(Elas _j < Median)	-0.8943*** [0.2869]	-1.1148*** [0.3838]	-1.2395*** [0.4345]	-1.2195*** [0.4363]
× Ind.(Elas _j > Median)	0.5044*** [0.1717]	1.0224*** [0.1571]	0.8871*** [0.1505]	0.9451*** [0.1415]
p-value: Q5 at median <i>UpstCont_j</i>	[0.0004]	[0.0054]	[0.0000]	[0.0000]
Elasticity based on:	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	286,072	206,490	144,107	144,107
No. of industries	459	305	219	219
R ²	0.1882	0.2609	0.2910	0.2888

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

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

We have subjected the cross-firm regressions to an extensive series of robustness checks, which

Table 4: Effect of Upstream Contractibility: Elasticity Quintiles

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 $Elas_j$)	-0.0350 [0.0300]	-0.0611 [0.0396]	-0.0490 [0.0429]	0.0763** [0.0323]
Ind.(Quintile 3 $Elas_j$)	-0.1104*** [0.0288]	-0.0566 [0.0405]	-0.0683** [0.0328]	-0.0476** [0.0223]
Ind.(Quintile 4 $Elas_j$)	-0.1207*** [0.0304]	-0.1605*** [0.0292]	-0.1611*** [0.0277]	-0.1185*** [0.0236]
Ind.(Quintile 5 $Elas_j$)	-0.1409*** [0.0297]	-0.1760*** [0.0306]	-0.1643*** [0.0292]	-0.1108*** [0.0260]
Upstream Contractibility $_j$				
× Ind.(Quintile 1 $Elas_j$)	-1.5540*** [0.4934]	-1.5492*** [0.4177]	-1.8562*** [0.4446]	-0.8114 [0.5369]
× Ind.(Quintile 2 $Elas_j$)	-0.9810*** [0.3165]	-0.5723 [0.5973]	-0.6886 [0.7621]	-2.0195*** [0.6896]
× Ind.(Quintile 3 $Elas_j$)	0.3271 [0.2408]	-0.3234 [0.3742]	-0.4171 [0.3855]	0.1796 [0.1727]
× Ind.(Quintile 4 $Elas_j$)	0.3849 [0.2867]	1.0662*** [0.2319]	0.6855*** [0.2106]	0.9811*** [0.2565]
× Ind.(Quintile 5 $Elas_j$)	0.7106*** [0.2148]	1.0530*** [0.2149]	1.1171*** [0.2273]	1.0419*** [0.2275]
p-value: Q5 at median $UpstCont_j$	[0.0002]	[0.0001]	[0.0000]	[0.0000]
Elasticity based on:	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	286,072	206,490	144,107	144,107
No. of industries	459	305	219	219
R ²	0.2204	0.2792	0.3064	0.3191

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

are reported in Tables A-4 to A-8 in the Appendix. We present there results based on our preferred specification in column (3) of Table 4, which uses the ρ_j measure constructed from consumption goods elasticities only; the results with the alternative elasticity proxies are qualitatively similar and available on request. We briefly discuss below the nature of these sensitivity results, while leaving the details to Appendix A-3.

In Table A-4, we show that the patterns are robust when examining different subsamples of firms. More specifically, we obtain similar results when restricting the sample respectively to single-establishment firms, to domestic firms, or to multinationals (i.e., firms with establishments in more than one country).³⁰ In Table A-5, we experiment with specifications that control for additional firm and industry variables that relate to alternative motives for the vertical integration decisions of firms (see Appendix A-3 for more details). These controls have an immaterial impact on our estimates, even when these variables are jointly entered into the regression. Turning to Table A-6, we demonstrate that our results are robust to alternative treatments of the identity of the primary output industry for multi-product firms. In particular, our results hold when we re-designate the output industry of each firm to be the SIC manufacturing code that is most downstream with respect to final demand, on the basis of the Antràs *et al.* (2012) measure. We also show that the patterns are similar when limiting the sample to firms whose primary SIC code is its only manufacturing SIC activity. In Tables A-7 and A-8, we report several checks based on alternative constructions of the ratio-upstreamness dependent variable. These include restricting the set $S(j)$ to manufacturing inputs, to “ever-integrated” inputs, and to more relevant inputs with larger tr_{ij} values. Our findings are broadly robust, with the main exception being the results pertaining to the interaction with $UpstCont_j$ when only manufacturing inputs excluding the parent SIC are considered; there, the coefficients do not turn positive even for the highest quintile interactions.

We next move to test prediction P.3 of our model concerning the role of heterogeneity in firm productivity. We are limited here to using a simple measure of log sales per worker, computed using total sales and employment across all establishments of the parent, to proxy for the firm-level parameter θ in the model, as WorldBase contains little information on the operations of firms beyond this.³¹ Moreover, when available, these variables are often based on estimates rather than administrative data. To reduce the possible influence of such measurement error, we use a dummy variable, $\mathbf{1}(\theta_p > \theta_{j,med})$, which identifies highly productive firms as those with above-median productivity within each output industry j .

According to the first part of prediction P.3, more productive firms should integrate more inputs,

³⁰The check related to single-establishment firms is reassuring in light of the findings in Atalay *et al.* (2014), documenting small volumes of domestic shipments across plants owned by the same U.S. parent, and Ramondo *et al.* (2016), indicating that the bulk of intrafirm trade involving a U.S. multinational parent tends to be concentrated among a small number of its large foreign affiliates. In the case of single-establishment firms, it is unlikely that a parent would not use the inputs produced in its own establishment.

³¹Because log sales per worker is a measure of revenue-based productivity, it captures variation in both θ and A . Proposition 3 shows, however, that our comparative statics results hold regardless of whether one varies θ or A .

in both the complements and substitutes cases. To verify this, we estimate the following:

$$\log(\text{No. of integrated inputs})_{jpc} = \beta_0 + \sum_{n=2}^5 \beta_n \mathbf{1}(\theta_p > \theta_{j,med}) \times \mathbf{1}(\rho_j \in \text{Quint}_n(\rho)) + \beta_W W_p + D_{jc} + \epsilon_{jpc}. \quad (23)$$

Note that the appropriate source of variation that we focus on here is that across firms within a given industry, so the above regression is estimated with country-industry fixed effects, D_{jc} . We accordingly can include the same set of firm-level variables, W_p , used in the previous specifications in Tables 1-4, but not the vector of industry controls, X_j .

Table 5: Within-Sector, Cross-Firm Heterogeneity in Effects

Dependent variable:	Log No. of Integrated SICs		Log Ratio-Upstreamness	
	(1)	(2)	(3)	(4)
Ind.(Log(Sales/Emp) _p > Median)				
× Ind.(Quintile 1 Elas _j)	0.0195*** [0.0066]	0.0123 [0.0081]	-0.0026** [0.0013]	-0.0023** [0.0010]
× Ind.(Quintile 2 Elas _j)	0.0190 [0.0117]	0.0216*** [0.0066]	-0.0002 [0.0018]	-0.0035* [0.0020]
× Ind.(Quintile 3 Elas _j)	0.0342*** [0.0120]	0.0373** [0.0171]	0.0039 [0.0033]	0.0064** [0.0027]
× Ind.(Quintile 4 Elas _j)	0.0334*** [0.0095]	0.0286*** [0.0092]	0.0061*** [0.0014]	0.0060*** [0.0014]
× Ind.(Quintile 5 Elas _j)	0.0212* [0.0109]	0.0204* [0.0106]	0.0082*** [0.0024]	0.0078*** [0.0024]
Elasticity based on:	BEC cons.	BEC cons. & α proxy	BEC cons.	BEC cons. & α proxy
Firm controls	Y	Y	Y	Y
Parent country-industry pair dummies	Y	Y	Y	Y
Observations	142,135	142,135	142,135	142,135
No. of industries	219	219	219	219
R ²	0.3809	0.3809	0.7665	0.7666

Notes: The sample comprises all firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable in columns (1)-(2) is the log number of 4-digit SIC codes integrated by the firm, while that in columns (3)-(4) is the baseline log ratio-upstreamness measure described in Section 3. Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries; columns (1) and (3) use the elasticity measure based only on HS10 codes classified as consumption goods in the UN BEC, while columns (2) and (4) use the consumption-goods-only demand elasticity minus the proxy for α to distinguish between the complements and substitutes cases. All columns include parent country by parent primary SIC industry pair dummies, as well as the full list of firm-level variables used in the earlier specifications of Table 2, columns (3)-(6).

The results from running (23) are reported in the first two columns of Table 5. We use a measure of ρ_j constructed using consumption-goods only demand elasticities in column (1), before using the alternative proxy for $\rho_j - \alpha_j$ in column (2). The estimated coefficients of the $\mathbf{1}(\theta_p >$

$\theta_{j,med}) \times \mathbf{1}(\rho_j \in Quint_n(\rho))$ interaction terms are almost all positive and significant, confirming that more productive firms tend to integrate more SIC activities; this is regardless of whether the industry in question falls closer to being in the substitutes or complements case.

According to the second part of prediction P.3, more productive firms should exhibit a higher ratio-upstreamness in the complements case. The underlying intuition is that more productive firms would be able to bear the higher fixed costs of integrating a larger set of stages within firm boundaries, and so would engage in integrating some upstream inputs when compared against smaller, less productive firms in the same industry. Conversely, the opposite would hold in the substitutes case, with more productive firms instead featuring a lower ratio-upstreamness. To test for such a pattern in the data, we therefore replace the dependent variable in (23) with $\log R_{jpc}$ and re-estimate the regression. The results are reported in columns (3)-(4) of Table 5, these being based respectively on the ρ_j and $\rho_j - \alpha_j$ proxies used in the first two columns of the same table. The patterns that emerge are entirely in line with prediction P.3: the estimated coefficient of the interaction term between $\mathbf{1}(\theta_p > \theta_{j,med})$ and the first quintile of ρ_j is negative and significant, while the corresponding coefficient for the fifth quintile of ρ_j is positive and significant. Thus, more productive firms have a lower (respectively, higher) propensity to integrate upstream inputs when the elasticity of demand for their final product is low (respectively, high).³²

Summing up, the cross-firm regressions provide strong evidence that the propensity for a firm to integrate relatively upstream inputs is weakest when the demand elasticity faced by that industry is largest, in line with prediction P.1 (Cross). Predictions P.2 (Cross) and P.3 concerning the role of the contractibility of inputs along the value chain and the productivity of final good producers also find strong support in the data.

4.2 Within-firm Results

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

$$D_INT_{ijp} = \gamma_0 + \gamma_1 \mathbf{1}(\rho_j < \rho_{med}) \times upst_{ij} + \gamma_2 \mathbf{1}(\rho_j > \rho_{med}) \times upst_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp} \quad (24)$$

$$D_INT_{ijp} = \gamma_0 + \sum_{n=1}^5 \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times upst_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (25)$$

The dependent variable, D_INT_{ijp} , is a 0-1 indicator for whether the firm p with primary output j has integrated the input i within firm boundaries. The key explanatory variables are the terms involving $upst_{ij}$ and its interactions with the elasticity variables. D_p denotes a full set of parent

³²The alert reader may wonder why we do not further explore triple interaction specifications involving the demand elasticity quintile, the high-productivity dummy, and the upstream contractibility variable. However, as discussed in the proof of Proposition 3 in Appendix A-1, it is in general not possible to sign this effect in the theory.

fixed effects. These specifications therefore allow us to study the integration decisions of individual firms, how they are affected by the upstreamness of the inputs, and whether these effects vary across the complements and substitutes cases. In particular, our theory would suggest that $\gamma_1 > 0$ and $\gamma_5 < 0$ (in the quintile specification), although we shall see below that the empirical results are consistent with a weaker form of this prediction.

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

We estimate (24) and (25) as a linear probability model, with standard errors clustered by i - j pair. To keep the analysis tractable, we limit the sample to the top 100 manufacturing inputs i used by j , as ranked by the total requirements coefficient tr_{ij} . This covers between 88-98% of the total requirements value for each output industry. We focus on the subsample of parent firms that have integrated at least one manufacturing input other than the parent’s self-SIC code, in order to avoid including firms for whom occurrences of integration are exceedingly sparse.

To assess prediction P.2 (Within), we extend (24) and (25) by adding the interactions between the ρ_j indicator variables and $ContUpToi_{ij}$, where the latter measure captures the contractibility of all inputs up to i in the production of j :

$$D_INT_{ijp} = \gamma_0 + \gamma_1 \mathbf{1}(\rho_j < \rho_{med}) \times upst_{ij} + \gamma_2 \mathbf{1}(\rho_j > \rho_{med}) \times upst_{ij} + \delta_1 \mathbf{1}(\rho_j < \rho_{med}) \times ContUpToi_{ij} + \delta_2 \mathbf{1}(\rho_j > \rho_{med}) \times ContUpToi_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (26)$$

$$D_INT_{ijp} = \gamma_0 + \sum_{n=1}^5 \gamma_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times upst_{ij} + \sum_{n=1}^5 \delta_n \mathbf{1}(\rho_j \in Quint_n(\rho)) \times ContUpToi_{ij} + \gamma_S \mathbf{1}(i = j) + D_i + D_p + \epsilon_{ijp}. \quad (27)$$

Recall that $ContUpToi_{ij}$ was constructed in (18) as an empirical proxy for $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ from the model. Looking back at the expression for the optimal bargaining share, $\beta^*(m)$, in equation (10), one would then expect that “contractibility up to i ” would raise the propensity to integrate input i if industry j came under the complements case ($\delta_5 > 0$), while having the opposite effect in the substitutes case ($\delta_1 < 0$). There is a further implication from (10), namely that having controlled for $ContUpToi_{ij}$, one should no longer expect to see that $upst_{ij}$ would have a significant effect on integration decisions, since the effect of m on $\beta^*(m)$ is captured entirely by the $\frac{\int_0^i (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}{\int_0^1 (\psi(k))^{\frac{\alpha}{1-\alpha}} dk}$ term.

The findings from the within-firm estimation are reported in Table 6. Following (24) and (26), we first adopt the median elasticity cutoff dummies (constructed from consumption-goods elasticities

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

In column (2), we present results when interaction terms involving the median cutoff dummies and the “contractibility up to i ” measure are introduced to the regression, as in (26). The effect of $ContUpToi_{ij}$ is indeed positive and significant when $\rho_j > \rho_{med}$, which is in line with prediction P.2 (Within) for the complements case, in that a greater degree of contractibility upstream of input i raises the likelihood that we observe i being integrated. Once again however, the point estimate for the interaction term when $\rho_j < \rho_{med}$ is consistent with a weaker form of the model’s prediction, being positive though smaller in magnitude compared to the corresponding coefficient for the above-median interaction term. (The p-value reported in column (2) confirms that we can reject the null hypothesis that the estimated coefficients in the above- and below-median elasticity cases are equal.) These patterns persist even when we include SIC input dummies to control for any characteristics specific to inputs i (column (3)), or use the log total requirements coefficient as an additional control (column (4)). Note that this latter tr_{ij} variable enters with a positive and significant coefficient, suggesting that firms are more likely to integrate inputs that are more important in production. Last but not least, column (5) presents the results when using the $\rho_j - \alpha_j$ proxy as the elasticity measure of interest. The findings are retained, although the difference in the effect of $ContUpToi_{ij}$ in the below- and above-median elasticity cases is marginally insignificant.

In Table 7, we repeat the above using a more extensive set of quintile elasticity dummies instead for the elasticity measure, following (25) and (27). The patterns we find here are qualitatively very similar. When the quintile dummies are interacted with $upst_{ij}$ in column (1), notice that the coefficients become successively more negative in the higher elasticity quintiles, consistent once again with the integration of upstream inputs being less likely in the complements relative to the substitutes cases. This neat pattern for the effect of $upst_{ij}$ however disappears when the analogous interactions involving $ContUpToi_{ij}$ are added (column (2)); instead, it is the effect of “contractibility up to i ” that increases monotonically across the quintiles.³³ This confirms that it is the effect of the contractibility of upstream inputs, rather than upstreamness *per se*, that matters for integration patterns. This is precisely what our model predicts, since what matters for the optimal organizational decision at stage i is $\frac{\int_0^i(\psi(k))^{\frac{1-\alpha}{\alpha}} dk}{\int_0^1(\psi(k))^{\frac{1-\alpha}{\alpha}} dk}$, and that having controlled for this, the

³³A formal test for the equality of the first and fifth quintile coefficients is rejected at conventional significance levels (see the p-value reported in each column).

Table 6: Integration Decisions within Firms (Top 100 Inputs): Median Elasticity Cutoff

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

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

upstreamness of *i* in the production of *j* should have no further effect. These patterns remain even with the inclusion of SIC input fixed effects and the log total requirements coefficient, as well as when variation in α_j is brought to bear on our elasticity measure (columns (3)-(5)).

We conclude the empirical discussion with several checks on the within-firm regressions; these are reported in more detail in Appendix A-3. Table A-10 shows that the patterns are qualitatively similar when we in turn restrict the sample to single-establishment firms, to domestic firms, and to multinationals, with the only exception being in the quintile elasticity regressions for the subsample

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

Dependent variable:	Indicator variable: Input Integrated?				
	(1)	(2)	(3)	(4)	(5)
Upstreamness _{<i>ij</i>}					
× Ind.(Quintile 1 Elas _{<i>j</i>})	-0.0056*** [0.0009]	0.0005 [0.0014]	-0.0034** [0.0016]	0.0011 [0.0015]	0.0030*** [0.0011]
× Ind.(Quintile 2 Elas _{<i>j</i>})	-0.0085*** [0.0019]	-0.0001 [0.0035]	-0.0038 [0.0035]	0.0002 [0.0033]	-0.0010 [0.0027]
× Ind.(Quintile 3 Elas _{<i>j</i>})	-0.0100*** [0.0012]	-0.0001 [0.0027]	-0.0018 [0.0026]	0.0019 [0.0025]	-0.0008 [0.0046]
× Ind.(Quintile 4 Elas _{<i>j</i>})	-0.0098*** [0.0021]	0.0084*** [0.0024]	0.0024 [0.0016]	0.0064*** [0.0017]	0.0070*** [0.0019]
× Ind.(Quintile 5 Elas _{<i>j</i>})	-0.0113*** [0.0021]	0.0054* [0.0028]	0.0024 [0.0019]	0.0059*** [0.0020]	0.0060*** [0.0020]
Contractibility up to <i>i</i> (in prod. of <i>j</i>)					
× Ind.(Quintile 1 Elas _{<i>j</i>})		0.0234*** [0.0052]	0.0217*** [0.0048]	0.0108** [0.0049]	0.0157*** [0.0049]
× Ind.(Quintile 2 Elas _{<i>j</i>})		0.0339*** [0.0128]	0.0261*** [0.0093]	0.0117 [0.0100]	0.0047 [0.0073]
× Ind.(Quintile 3 Elas _{<i>j</i>})		0.0365*** [0.0082]	0.0304*** [0.0080]	0.0146* [0.0082]	0.0132 [0.0141]
× Ind.(Quintile 4 Elas _{<i>j</i>})		0.0669*** [0.0157]	0.0398*** [0.0086]	0.0239*** [0.0086]	0.0254*** [0.0088]
× Ind.(Quintile 5 Elas _{<i>j</i>})		0.0685*** [0.0134]	0.0456*** [0.0095]	0.0304*** [0.0093]	0.0322*** [0.0090]
Dummy: Self-SIC	0.9794*** [0.0018]	0.9699*** [0.0028]	0.9340*** [0.0085]	0.9313*** [0.0085]	0.9313*** [0.0085]
Log (Total Requirements _{<i>ij</i>})				0.0055*** [0.0009]	0.0054*** [0.0008]
p-value: Contractibility up to <i>i</i> , Quintile 1 minus Quintile 5	—	[0.0015]	[0.0217]	[0.0559]	[0.1001]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons. & α proxy
Firm fixed effect	Y	Y	Y	Y	Y
Input industry <i>i</i> fixed effects	N	N	Y	Y	Y
Observations	4,707,722	4,707,722	4,707,722	4,707,722	4,707,722
No. of parent firms	46,992	46,992	46,992	46,992	46,992
No. of <i>i-j</i> pairs	21,836	21,836	21,836	21,836	21,836
R ²	0.5342	0.5359	0.5594	0.5598	0.5599

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

of multinational firms. Table A-11 further reports on a series of tests that confirm the robustness of the results to: (i) dropping parent firms that do not have an integrated manufacturing input (apart from the self-SIC) among the top 100 inputs as ranked by the total requirements value; (ii) focusing on parents that have integrated at least three of their top-100 manufacturing inputs; (iii) dropping the self-SIC from the estimation; and (iv) including the full set of quintile elasticity dummies interacted with a measure of the “contractibility at i ” (as opposed to “up to i ”).

5 Conclusion

The emergence of global value chains in recent decades has attracted much attention from policymakers and academics alike. However, there are currently few systematic empirical studies attempting to shed light on the determinants of firms’ decisions to control different segments of their production processes. In this paper, we show how detailed data on the activities of firms around the world can be combined with information from standard Input-Output tables to study such integration choices along value chains.

Building on Antràs and Chor (2013), we describe a property-rights model in which a firm’s boundaries are shaped by characteristics of the different stages of production and their position in the value chain. As available theoretical frameworks of sequential production are highly stylized, a key contribution of this paper is to develop a richer theoretical framework of firm behavior that can guide an empirical analysis using firm-level data. Our model delivers several testable predictions, suggesting that the propensity to integrate upstream versus downstream inputs should depend crucially on the elasticity of demand for the final product, the degree of contractibility of the inputs, and the productivity of the parent firm. One of our model’s extensions also helps to rationalize the sparsity of integrated inputs, showing that low levels of intrafirm trade can be consistent with the property-rights approach to firm boundaries, in which final good producers choose to own suppliers to better discipline their behavior.

To assess the evidence, we use the WorldBase dataset, which contains establishment-level information on the activities of firms located in a large set of countries. We combine this information with Input-Output tables to construct firm-level measures of the upstreamness of integrated and non-integrated stages. The richness of our data allows us to run specifications that exploit variation in organizational features across firms, as well as within firms and across their various manufacturing stages. In line with our model’s predictions, we find that whether a firm integrates suppliers located upstream or downstream depends crucially on the elasticity of demand faced by the firm. Moreover, the relative propensity to integrate upstream (as opposed to downstream) inputs depends on the extent to which contractible inputs tend to be located in the early or late stages of the production process, as well as on the productivity of final good producers. The firm-level patterns that we uncover provide strong evidence that considerations driven by contractual frictions critically shape firms’ ownership decisions along their value chains.

References

- Acemoglu, Daron, Pol Antràs, and Elhanan Helpman (2007), “Contracts and Technology Adoption,” *American Economic Review* 97(3): 916-943.
- Acemoglu, Daron, Simon Johnson, and Todd Mitton (2009), “Determinants of Vertical Integration: Financial Development and Contracting Costs,” *Journal of Finance* 63(3): 1251-1290.
- Alfaro, Laura, and Andrew Charlton (2009), “Intra-Industry Foreign Direct Investment,” *American Economic Review* 99(5): 2096-2119.
- Alfaro, Laura, and Maggie Xiaoyang Chen (2014), “The Global Agglomeration of Multinational Firms,” *Journal of International Economics* 94(2): 263-276.
- Alfaro, Laura, Paola Conconi, Harald Fadinger, and Andrew Newman (2016), “Do Price Determine Vertical Integration?” *Review of Economic Studies* 83(3): 855-888.
- 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.
- 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.
- Antràs, Pol, and Elhanan Helpman (2004), “Global Sourcing,” *Journal of Political Economy* 112(3): 552-580.
- Atalay, Engin, Ali Hortaçsu, and Chad Syverson (2014), “Vertical Integration and Input Flows,” *American Economic Review* 104(4): 1120-1148.
- Baldwin, Richard, and Anthony Venables (2013), “Spiders and Snakes: Offshoring and Agglomeration in the Global Economy,” *Journal of International Economics* 90(2): 245-254.
- Becker, Randy A., and Wayne B. Gray (2009), “NBER-CES Manufacturing Industry Database (1958-2005)”.
- Bresnahan, Timothy, and Jonathan Levin (2012), “Vertical Integration and Market Structure,” in *The Handbook of Organizational Economics*, Princeton University Press.
- Broda, Christian, and David Weinstein (2006), “Globalization and the Gains from Variety,” *Quarterly Journal of Economics* 121(2): 541-585.
- Brown, Clair, and Greg Linden (2005), “Offshoring in the Semiconductor Industry: A Historical Perspective,” *Brookings Trade Forum* (Offshoring White-Collar Work), pp.279-333.

- Cameron, Colin, Jonah Gelbach, and Douglas Miller (2011), “Robust Inference with Multi-way Clustering,” *Journal of Business and Economic Statistics* 29(2): 238-249.
- Caves, Richard (1975), “Diversification, Foreign Investment and Scale in North American Manufacturing Industries,” *Canadian Public Policy* 2: 274-276.
- Costinot, Arnaud, Jonathan Vogel, and Su Wang (2013), “An Elementary Theory of Global Supply Chains,” *Review of Economic Studies* 80(1): 109-144.
- Corcos, Gregory, Delphine Irac, Delphine M., Giordano Mion, and Thierry Verdier (2013), “The Determinants of Intrafirm Trade: Evidence from French Firms,” *Review of Economics and Statistics* 95(3): 825-838.
- Defever, Fabrice, and Farid Toubal (2013), “Productivity, Relationship-Specific Inputs and the Sourcing Modes of Multinationals,” *Journal of Economic Behavior and Organization* 94: 345-357.
- Del Prete, Davide, and Armando Rungi (2015), “Organizing the Global Value Chain: A Firm Level Test,” mimeo.
- Dietzenbacher, Erik, Isidoro Romero Luna, and Niels S. Bosma (2005), “Using Average Propagation Lengths to Identify Production Chains in the Andalusian Economy,” *Estudios de Economía Aplicada* 23(2): 405-422.
- Díez, Federico (2014), “The Asymmetric Effects of Tariffs on Intra-firm Trade and Offshoring Decisions,” *Journal of International Economics* 93(1): 76-91.
- Dixit, Avinash, and Gene Grossman (1982), “Trade and Protection with Multistage Production,” *Review of Economic Studies* 49(4): 583-594.
- Fajgelbaum, Pablo, Gene Grossman, and Elhanan Helpman (2015), “A Linder Hypothesis for Foreign Direct Investment,” *Review of Economic Studies* 82(1): 83-121.
- Fally, Thibault (2012), “On the Fragmentation of Production in the U.S.,” mimeo.
- Fally, Thibault, and Russell Hillbery (2014), “A Coasian Model of International Production Chains,” mimeo.
- Fan, Joseph P. H., and Larry H. P. Lang (2000), “The Measurement of Relatedness: An Application to Corporate Diversification,” *Journal of Business* 73(4): 629-660.
- Feenstra, Robert C., John Romalis, and Peter K. Schott, (2002), “U.S. Imports, Exports and Tariff Data, 1989-2001,” *NBER Working Paper* 9387.
- Grossman, Sanford J., and Hart, Oliver D. (1986), “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy* 94(4): 691-719.
- Harms, Philipp, Oliver Lorz, and Dieter Urban (2012), “Offshoring along the Production Chain,” *Canadian Journal of Economics* 45(1): 93-106.
- Harrison, Ann E., Inessa Love, and Margaret S. McMillan (2004), “Global Capital Flows and Financing Constraints,” *Journal of Development Economics* 75(1): 269-301.
- Johnson, Robert C., and Guillermo Noguera (2012), “Accounting for Intermediates: Production Sharing and Trade in Value Added,” *Journal of International Economics* 86(2): 224-236.

- Kikuchi, Tomoo, Kazuo Nishimura, and John Stachurski (2014), "Transaction Costs, Span of Control and Competitive Equilibrium," mimeo.
- Kohler, Wilhelm (2004), "International Outsourcing and Factor Prices with Multistage Production," *Economic Journal* 114(494): C166-C185.
- Kremer, Michael (1993), "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics* 108(3): 551-575.
- Lafontaine, Francine, and Margaret Slade (2007), "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature* 45(3): 629-685.
- Luck, Philip (2014), "Global Supply Chains and Vertical Integration: Evidence from China," mimeo.
- Nunn, Nathan (2007), "Relationship-Specificity, Incomplete Contracts and the Pattern of Trade," *Quarterly Journal of Economics* 122(2): 569-600.
- Nunn, Nathan, and Daniel Treffer (2008), "The Boundaries of the Multinational Firm: An Empirical Analysis," in E. Helpman, D. Marin, and T. Verdier (eds.), *The Organization of Firms in a Global Economy*, Harvard University Press.
- Nunn, Nathan, and Daniel Treffer (2013), "Incomplete Contracts and the Boundaries of the Multinational Firm," *Journal of Economic Behavior and Organization* 94: 330-344.
- Ramondo, Natalia, Veronica Rappoport, and Kim Ruhl (2016), "Horizontal versus Vertical Foreign Direct Investment: Evidence from U.S. Multinationals," *Journal of International Economics*, 98(1): 51-59.
- Rauch, James E. (1999), "Networks versus Markets in International Trade," *Journal of International Economics* 48(1): 7-35.
- Sanyal, Kalyan K., and Ronald W. Jones, (1982), "The Theory of Trade in Middle Products," *American Economic Review* 72(1): 16-31.
- Xing, Yuqing (2011), "How the iPhone widens the US trade deficit with China," VoxEU.org, 10 April.
- Yeaple, Stephen R. (2006), "Offshoring, Foreign Direct Investment, and the Structure of U.S. Trade," *Journal of the European Economic Association* 4(2-3): 602-611.
- Yi, Kei-Mu (2003), "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111(1): 52-102.

Appendices

A-1 Theoretical Appendix

A-1.1 Derivation of Program (7)

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 (5), we obtain the following optimal choice of investment by the supplier at stage m :

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

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

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

It is straightforward to verify that the solution to this differential equation (with the initial condition $r(0) = 0$) is given by:

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

from which we can obtain the expression for $x^*(m)$ in equation (6).

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

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

which coincides with the expression in program (7) in the main text.

A-1.2 Derivation of Equation (10)

As pointed out in the main text, we can express program (7) 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 $\Theta = A \theta^{\frac{\rho}{1-\rho}} \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}}$, and:

$$v(i) \equiv \int_0^i \left(\frac{(1-\beta(k)) \psi(k)}{c(k)} \right)^{\frac{1-\alpha}{1-\alpha}} dk. \quad (\text{A-2})$$

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

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

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

$$\frac{\rho - \alpha}{1 - \rho} \frac{v'(i)}{v(i)} + \frac{v''(i)}{v'(i)} = -\frac{\alpha}{1 - \alpha} \frac{d(c(i)/\psi(i))/di}{c(i)/\psi(i)}. \quad (\text{A-3})$$

To solve (A-3), integrate both sides with respect to i , and exponentiate to get:

$$v'(i) v(i)^{\frac{\rho-\alpha}{1-\rho}} = C_1 (\psi(i)/c(i))^{\frac{\alpha}{1-\alpha}}, \quad (\text{A-4})$$

where $C_1 > 0$ is a constant of integration. Given the definition of $v(i)$ in (A-2), equation (A-4) can be rewritten as:

$$(1 - \beta(i))^{\frac{\alpha}{1-\alpha}} = C_1 \left(\int_0^i \left(\frac{(1 - \beta(k)) \psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right)^{\frac{\alpha-\rho}{1-\rho}}. \quad (\text{A-5})$$

Denoting $z(i) \equiv (1 - \beta(i))^{\frac{\alpha}{1-\alpha}}$, we can express (A-5) as:

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

which after differentiation delivers:

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

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

$$z(m)^{\frac{1-\alpha}{\alpha-\rho}} - z(0)^{\frac{1-\alpha}{\alpha-\rho}} = (C_1)^{\frac{1-\rho}{\alpha-\rho}} \left(\frac{1 - \alpha}{1 - \rho} \right) \left[\int_0^m \left(\frac{\psi(k)}{c(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right].$$

To simplify the above, note that (A-6) implies $z(0)^{\frac{1-\alpha}{\alpha-\rho}} = 0$. Recalling the definition $z(m) \equiv (1 - \beta(m))^{\frac{\alpha}{1-\alpha}}$, 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 the full solution as spelled out in equation (10) in the main text.

A-1.3 Proof of Proposition 1

The proof is a generalization of that for Proposition 2 in Antràs and Chor (2013). It is straightforward to see from equation (10), that when $\rho > \alpha$, $\lim_{m \rightarrow 0} \beta^*(m) \rightarrow -\infty$, and it is thus optimal for the firm to choose β_O (namely outsourcing) for the most upstream stages in the neighborhood of $m = 0$. Conversely, when $\rho < \alpha$, $\lim_{m \rightarrow 0} \beta^*(m) = 1$, and 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 1 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 by contradiction. Suppose that, contrary to the claim in Proposition 1, there were to exist a stage $\tilde{m} \in (0, 1)$ such that a measurable set of stages immediately upstream from \tilde{m} are integrated, while a measurable set of stages immediately downstream from \tilde{m} are outsourced. Now consider two positive constants ε_L and ε_R such that:

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

These constants can always be chosen to be small enough such that they satisfy (A-7), and moreover are such that the set of stages $(\tilde{m} - \varepsilon_L, \tilde{m})$ is 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 in which the firm instead chooses to outsource the stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ and to integrate the stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$, while retaining the same organizational decision for all other stages in the unit interval. Denote the profits of this alternative organizational form by Π_2 . We will now show that this reorganization necessarily increases firm profits, i.e., $\Pi_1 < \Pi_2$, so that the posited deviation from the optimal pattern in Proposition 1 is inconsistent with profit maximization.

Note that we can rewrite firm profits in (7) as:

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

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

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

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

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

while for profits Π_2 , these same profits are given by:

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

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

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

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

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

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

A-1.4 Derivation of m_C^* and m_S^* Thresholds

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

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

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

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

from which equation (11) can easily be obtained. Notice that for a strictly interior solution, i.e., $m_C^* \in (0, 1)$, the right-hand side of (11) would need to be smaller than one, which in turn requires:

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

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

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

$$\frac{\int_0^{m_S^*} (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di}{\int_0^1 (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} di} = \left\{ 1 + \left(\frac{1-\beta_V}{1-\beta_O} \right)^{\frac{\alpha}{1-\alpha}} \left[\left(\frac{\frac{\beta_V}{\beta_O} - 1}{\left(\frac{1-\beta_V}{1-\beta_O} \right)^{-\frac{\alpha}{1-\alpha}} - 1} \right)^{\frac{\alpha(1-\rho)}{\rho-\alpha}} - 1 \right] \right\}^{-1}. \quad (\text{A-10})$$

A-1.5 Derivation of Equation (12) and Proposition 2

In the extension in Section 2.2.A, recall that the profits of the firm are given by: $\tilde{\pi}_F = \pi_F - \int_0^1 \frac{(\psi(i))^\phi}{\mu(i)} di$, where the second term captures the contracting costs. Focus first on the π_F term.

Consider the complements case. We begin by plugging equation (11), which pins down the m_C^* threshold, into the profit function (A-9). After a few simplifications, this delivers:

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

where:

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

Hence, we can write profits as:

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

In the substitutes case, we have an analogous expression:

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

where:

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

so that:

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

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

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \left[\int_0^1 \left(\frac{\psi(i)}{c(i)} \right)^{\frac{\alpha}{1-\alpha}} di \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \Gamma(\beta_V, \beta_O), \quad (\text{A-11})$$

where:

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

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

Obtaining equation (12) from the more general equation (A-11) is then trivial. Notice, however, that when studying the optimal choice of $\psi(m)$ that maximizes $\tilde{\pi}_F$, the first-order condition with respect to $\psi(m)$ now delivers that, for two inputs at stages m and m' , we have:

$$\left(\frac{\psi(m)/c(m)}{\psi(m')/c(m')} \right)^{\phi - \frac{\alpha}{1-\alpha}} = \frac{\mu(m)}{\mu(m')} \left(\frac{c(m)}{c(m')} \right)^{-\phi}, \quad (\text{A-12})$$

which generalizes equation (13) in the main text. Moreover, one can show that the second-order condition with respect to $\psi(m)$, when evaluated at the optimal $\psi(m)$, simplifies to:

$$\frac{\rho - \alpha}{(1-\alpha)(1-\rho)} \frac{(\psi(m)/c(m))^{\frac{\alpha}{1-\alpha}}}{\int_0^1 (\psi(i)/c(i))^{\frac{\alpha}{1-\alpha}} di} + \frac{\alpha}{1-\alpha} - \phi < 0.$$

In particular, the restriction: $\phi > \alpha/(1-\alpha)$ is necessary to ensure that the second-order condition holds in the complements case. Equation (A-12) thus illustrates that the ratio $\psi(m)/c(m)$ will tend to comove with contractibility along the value chain as long as contractibility and marginal costs are not positively correlated. But notice that plugging (A-12) into (A-11), we have that the effect of a reduction in the marginal cost of a given stage m will be increasing in the level of contractibility $\mu(m)$. 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 $c(m)$ and $\mu(m)$.

Turning now specifically to Proposition 2, we have developed the argument in the main text that the characterization of the optimal organizational mode from Proposition 1, in particular how this hinges on whether ρ is greater or less than α , continues to hold. This is because these predictions hold taking the profile of the $\psi(m)$'s as given; the mapping of these $\psi(m)$'s to heterogeneous contractibility across stages does not detract from this conclusion. Assuming that marginal costs of production are constant ($c(m) = c$) across all stages m , the optimal level of the $\psi(m)$'s that will be specified in the initial contract varies inversely with the exogenous contracting cost $1/\mu(m)$ pertaining to that stage. We thus associate a larger $\psi(m)$ with a higher degree of contractibility, in the sense that it is less costly to contract upon $\psi(m)$.

The second part of Proposition 2 speaks to how an increase in the contractibility of upstream relative to downstream inputs affects the m_C^* and m_S^* thresholds. To ease notation, define: $\tilde{\psi}(m) \equiv \psi(m)^{\frac{\alpha}{1-\alpha}}$. In the complements case, in light of equation (11), the natural notion of what constitutes a greater degree of “upstream contractibility” is an increase in the integral of the $\tilde{\psi}(m)$'s over all $m \in [0, m_C^*)$, that nevertheless holds the overall contractibility of the production process constant. In differential calculus notation, this

translates to: $\int_0^{m_C^*} d\tilde{\psi}(m)dm > 0$ and $\int_0^1 d\tilde{\psi}(m)dm = 0$. Taking the total derivative of (11), one obtains:

$$\tilde{\psi}(m_C^*)dm_C^* + \int_0^{m_C^*} d\tilde{\psi}(m)dm = 0,$$

from which it follows that $dm_C^* < 0$ in response to an increase in upstream contractibility, as claimed in the proposition. For the substitutes case, a similar argument can be applied to (A-10) to establish that $dm_S^* < 0$ in response to a differential change in the profile of the $\tilde{\psi}(m)$'s that satisfies: $\int_0^{m_S^*} d\tilde{\psi}(m)dm > 0$ and $\int_0^1 d\tilde{\psi}(m)dm = 0$.

A-1.6 Proof of Proposition 3

With firm heterogeneity in core productivity and incorporating fixed costs of integration, the firm's profits are now given by: $\pi_F - \int_0^1 f_V \mathbf{1}(\beta(i) = \beta_V) di$, where $\mathbf{1}(\beta_i = \beta_V)$ is an indicator function equal to 1 if and only if stage i is integrated by the firm. The proof is presented below for the complements case; that for the substitutes case follows in analogous fashion.

Suppose that $\rho > \alpha$. We first show that despite the introduction of fixed costs of integration, the optimal organizational mode continues to feature outsourcing of stages $[0, m_C^*)$ up to a cutoff stage m_C^* , and integration for all stages $[m_C^*, 1]$ further downstream. This is established through a proof by contradiction. Suppose there exists a $\tilde{m} \in (0, 1)$, such that there is a non-zero measure of integrated stages immediately upstream of \tilde{m} and a non-zero measure of outsourced stages immediately downstream of it. Pick two positive constants ε_L and ε_R satisfying equation (A-7); these two constants can always be chosen to be sufficiently small so that $(\tilde{m} - \varepsilon_L, \tilde{m})$ lies within the subset of integrated stages immediately upstream of \tilde{m} , and $(\tilde{m}, \tilde{m} + \varepsilon_R)$ is within the subset of outsourced stages immediately downstream of \tilde{m} .

If $\varepsilon_L \geq \varepsilon_R$, compare profits under the organizational mode where stages $(\tilde{m} - \varepsilon_L, \tilde{m})$ are integrated and stages $(\tilde{m}, \tilde{m} + \varepsilon_R)$ are outsourced, against an alternative where $(\tilde{m} - \varepsilon_L, \tilde{m})$ is outsourced and $(\tilde{m}, \tilde{m} + \varepsilon_R)$ is integrated, holding the organizational decisions over all other stages constant. The proof in Section A-1.3 showed that π_F is strictly higher under the latter organizational mode. The fixed costs of integration that are incurred would also be (weakly) lower under the latter option, since a (weakly) smaller measure of stages is integrated. This alternative organizational mode is thus more profitable, and yields the desired contradiction.

If instead $\varepsilon_L < \varepsilon_R$, a more involved argument is needed. Compare now profits under the organizational mode where stages $(\tilde{m} - \varepsilon_L, \tilde{m})$ are integrated and stages $(\tilde{m}, \tilde{m} + \varepsilon_L)$ are outsourced, versus an alternative where $(\tilde{m} - \varepsilon_L, \tilde{m})$ is outsourced and $(\tilde{m}, \tilde{m} + \varepsilon_L)$ is integrated, holding the organizational decisions over all other stages constant. Let the profits associated with the former set of organizational decisions be Π_1^f , while that for the latter be Π_2^f . By construction, the incurred fixed costs of integration are exactly equal under both organizational modes, so one can focus solely on π_F . Bearing in mind the expression for π_F from (A-8), consider the respective contribution to profits of: (i) stages in $[0, \tilde{m} - \varepsilon_L]$; (ii) those in $(\tilde{m} - \varepsilon_L, \tilde{m})$; (iii) those in $(\tilde{m}, \tilde{m} + \varepsilon_L)$; and (iv) stages in $[\tilde{m} + \varepsilon_L, 1]$.

It is straightforward to see that the contribution of stages in the first region is identical in both Π_1^f and Π_2^f . As for the fourth region, the contribution of these stages to Π_1^f is:

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

where we define: $\gamma(i) \equiv (\psi(i)/c(i))^{\frac{\alpha}{1-\alpha}}$, $\mathcal{A}^f \equiv \int_0^{\tilde{m}-\varepsilon_L} (1 - \beta(k))^{\frac{\alpha}{1-\alpha}} \gamma(k) dk$, $\mathcal{B}^f \equiv \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(k) dk$, $\mathcal{C}^f \equiv$

$\int_{\tilde{m}}^{\tilde{m}+\varepsilon_L} \gamma(k) dk$, and $\mathcal{D}^f \equiv \int_{\tilde{m}+\varepsilon_L}^i (1 - \beta(k))^{\frac{\alpha}{1-\alpha}} \gamma(k) dk$. On the other hand, the contribution from these stages to Π_2^f is equal to:

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

Since $\varepsilon_L < \varepsilon_R$, we have: $\mathcal{B}^f = \int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} \gamma(k) dk = \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} \gamma(k) dk > \int_{\tilde{m}}^{\tilde{m}+\varepsilon_L} \gamma(k) dk = \mathcal{C}^f$. Bear in mind also that: $(1 - \beta_O)^{\frac{\alpha}{1-\alpha}} > (1 - \beta_V)^{\frac{\alpha}{1-\alpha}}$. Comparing the last two equations above, it follows that when $\rho > \alpha$, the stages in $[\tilde{m} + \varepsilon_L, 1]$ contribute more to profits in Π_2^f than in Π_1^f .

It remains to compare the relative contributions due to the middle two sets of stages, i.e., $(\tilde{m} - \varepsilon_L, \tilde{m})$ and $(\tilde{m}, \tilde{m} + \varepsilon_L)$. Let $\tilde{\Pi}_1^f$ refer to the profits under Π_1^f that accrue from these subsets of stages, and likewise define $\tilde{\Pi}_2^f$ analogously for Π_2^f . Using (A-8) and after some algebra, it can be shown that:

$$\begin{aligned} \tilde{\Pi}_1^f - \tilde{\Pi}_2^f &\propto (\beta_V - \beta_O) \left[\left(\mathcal{A}^f + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} + \left(\mathcal{A}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - \left(\mathcal{A}^f \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right] \\ &\quad - \beta_V \left(\mathcal{A}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \mathcal{C}^f \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \\ &\quad + \beta_O \left(\mathcal{A}^f + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{C}^f \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \end{aligned}$$

We now proceed to show that if the value of ε_L that was initially chosen was sufficiently small, then $\tilde{\Pi}_1^f - \tilde{\Pi}_2^f < 0$. Observe that $\tilde{\Pi}_1^f - \tilde{\Pi}_2^f = 0$ at $\varepsilon_L = 0$. Given this, it then suffices to show that $\frac{\partial}{\partial \varepsilon_L} (\tilde{\Pi}_1^f - \tilde{\Pi}_2^f) < 0$ at $\varepsilon_L = 0$. Differentiating the above expression with Leibniz's rule yields:

$$\begin{aligned} \frac{\partial}{\partial \varepsilon_L} (\tilde{\Pi}_1^f - \tilde{\Pi}_2^f) &\propto (\beta_V - \beta_O) \left(\mathcal{A}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \left(-(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} - \varepsilon_L) + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} - \varepsilon_L) \right) \\ &\quad - \beta_V \left(\mathcal{A}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \mathcal{C}^f \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \\ &\quad \quad \times \left(-(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} - \varepsilon_L) + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} - \varepsilon_L) + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} + \varepsilon_L) \right) \\ &\quad + \beta_O \left(\mathcal{A}^f + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \mathcal{B}^f + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \mathcal{C}^f \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} + \varepsilon_L). \end{aligned}$$

The above steps use the fact that: (i) $\frac{\partial}{\partial \varepsilon_L} \mathcal{A}^f = -(1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \gamma(\tilde{m} - \varepsilon_L)$, since for ε_L sufficiently small, $\tilde{m} - \varepsilon_L$ is within the positive measure of stages immediately upstream of \tilde{m} that is initially integrated; (ii) that $\frac{\partial}{\partial \varepsilon_L} \mathcal{B}^f = \gamma(\tilde{m} - \varepsilon_L)$; and (iii) that $\frac{\partial}{\partial \varepsilon_L} \mathcal{C}^f = \gamma(\tilde{m} + \varepsilon_L)$. As $\varepsilon_L \rightarrow 0$, the above simplifies to: $\frac{\partial}{\partial \varepsilon_L} (\tilde{\Pi}_1^f - \tilde{\Pi}_2^f) \propto -(\beta_V - \beta_O) (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} (\mathcal{A})^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \gamma(\tilde{m}) < 0$. It follows that $\tilde{\Pi}_1^f - \tilde{\Pi}_2^f < 0$ when ε_L is positive but sufficiently small. Summarizing the comparison of profits across all four subsets of stages under Π_1^f and Π_2^f , the alternative organizational mode that generates Π_2^f delivers higher profits than Π_1^f , which yields the desired contradiction once again. This concludes the proof that the optimal organizational mode remains as described in Proposition 2, even though fixed costs of integration have been introduced.

To solve for the cutoff stage m_C^* in the complements case, we appeal to the expression for π_F in (A-9). Taking the first-order condition with respect to m_C in the profit function: $\pi_F - \int_0^1 f_V \mathbf{1}(\beta(i) = \beta_V) di = \pi_F - (1 - m_C) f_V$ and rearranging, this delivers the implicit function that pins down m_C^* as reported in equation (14) in Section 2.2.B. Note that in the special case of $f_V = 0$, (14) simplifies to the expression for m_C^* in the benchmark model in (11).

It remains to show that the predictions related to how ‘‘upstream contractibility’’ affects the cutoff stage

carry through even in the presence of fixed costs of integration. As in Section A-1.5, suppose $c(m) = c$ for all stages and define $\tilde{\psi}(m) \equiv \psi(m)^{\frac{\alpha}{1-\alpha}}$. Consider the effects of an increase in upstream contractibility, wherein: $\int_0^{m_C^*} d\tilde{\psi}(m)dm > 0$ while $\int_{m_C^*}^1 d\tilde{\psi}(m)dm = -\int_0^{m_C^*} d\tilde{\psi}(m)dm < 0$, so that $\int_0^1 d\tilde{\psi}(m)dm = 0$. Denote the left-hand side of equation (14) by $F(m_C^*)$. Taking the total derivative of (14), one obtains:

$$0 = F'(m_C^*)dm_C^* + \tilde{\psi}(m_C^*)\frac{\rho - \alpha}{\alpha(1 - \rho)} \left[\int_0^{m_C^*} \tilde{\psi}(k) dk \right]^{\frac{\rho - \alpha}{\alpha(1 - \rho)} - 1} G(m_C^*) \int_0^{m_C^*} d\tilde{\psi}(k) dk,$$

where:

$$\begin{aligned} G(m_C^*) &= \left(1 - \frac{\beta_O}{\beta_V}\right) - \left(1 - \left(\frac{1 - \beta_V}{1 - \beta_O}\right)^{\frac{\alpha}{1-\alpha}}\right)^2 \left[1 + \left(\frac{1 - \beta_V}{1 - \beta_O}\right)^{\frac{\alpha}{1-\alpha}} \frac{\int_{m_C^*}^1 \tilde{\psi}(k) dk}{\int_0^{m_C^*} \tilde{\psi}(k) dk}\right]^{\frac{\rho - \alpha}{\alpha(1 - \rho)} - 1} \\ &> \left(1 - \frac{\beta_O}{\beta_V}\right) - \left(1 - \left(\frac{1 - \beta_V}{1 - \beta_O}\right)^{\frac{\alpha}{1-\alpha}}\right) \left[1 + \left(\frac{1 - \beta_V}{1 - \beta_O}\right)^{\frac{\alpha}{1-\alpha}} \frac{\int_{m_C^*}^1 \tilde{\psi}(k) dk}{\int_0^{m_C^*} \tilde{\psi}(k) dk}\right]^{\frac{\rho - \alpha}{\alpha(1 - \rho)}} \\ &= \frac{F(m_C^*)}{\tilde{\psi}(m_C^*)} \left[\int_0^{m_C^*} \tilde{\psi}(k) dk \right]^{-\frac{\rho - \alpha}{\alpha(1 - \rho)}}. \end{aligned}$$

This last step follows from a substitution that uses the first-order condition for m_C^* . Clearly, the coefficient of the $\int_0^{m_C^*} d\tilde{\psi}(k) dk$ term in the total derivative is positive in the complements case. Recall also that the second-order condition for m_C^* implies that $F'(m^*) > 0$. A quick inspection of the total derivative then shows that $dm_C^* < 0$ when $\int_0^{m_C^*} d\tilde{\psi}(k) dk > 0$. Thus, the response of the cutoff stage to a greater degree of upstream contractibility continues to be characterized by the statement in Proposition 2, even in this extension of the model.

As a further consequence of this argument, it is in general not straightforward to sign the cross-partial effect of θ and upstream contractibility on the cutoff stage, m_C^* , as this would require making non-standard assumptions regarding the third-derivative of the firm's profit function, i.e., on the behavior of $F''(m_C^*)$. This is why we do not pursue specifications in the empirics that involve triple interactions between the ρ quintiles, upstream contractibility, and firm productivity.

A-1.7 Proof of Proposition 4

We illustrate this for the complements case ($\rho > \alpha$); the mechanics of the proof carry over to the substitutes case. Suppose to the contrary that $I_0 \equiv (\tilde{m}, \tilde{m} + \varepsilon) \in \Omega$ is a positive measure of discretionarily outsourced stages, located downstream of a positive measure of integrated stages. Denote this latter positive measure of integrated stages by I_1 , where $I_1 \in \Omega$ by definition. There are two cases to consider: (i) I_1 is immediately upstream of \tilde{m} , i.e., $I_1 = (\tilde{m} - \tilde{\varepsilon}, \tilde{m})$, for some $\tilde{\varepsilon} > 0$; and (ii) I_1 is not immediately upstream of \tilde{m} , i.e., $I_1 = (m_1 - \tilde{\varepsilon}, m_1)$, where $m_1 < \tilde{m}$ and $[m_1, \tilde{m}] \in \Upsilon$, i.e., the intervals I_0 and I_1 are separated by a positive measure of exogenously outsourced stages.

Consider first case (i). Without loss of generality, we can select two positive constants ε_L and ε_R such that $\varepsilon_L, \varepsilon_R < \min\{\varepsilon, \tilde{\varepsilon}\}$, which moreover satisfy equation (A-7). The same argument from the proof of Proposition 1 in Section A-1.3 can then be applied: If we were to interchange the organizational mode, to instead outsource the stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ and integrate the stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$, this necessarily results in a strict increase in profits. This yields the desired contradiction, as it cannot then be optimal to have the

stages in $(\tilde{m} - \tilde{\varepsilon}, \tilde{m})$ integrated, while those in $(\tilde{m}, \tilde{m} + \varepsilon)$ are discretionarily outsourced.

Consider next case (ii). Denote by $\tilde{\Pi}_1$ the configuration of organizational modes in which the stages in $I_1 = (m_1 - \tilde{\varepsilon}, m_1)$ are integrated, while those in $I_0 = (\tilde{m}, \tilde{m} + \varepsilon)$ are discretionarily outsourced. We will compare this against $\tilde{\Pi}_2$, which is the profits from the configuration where I_1 is instead outsourced and I_0 is integrated, holding the organizational mode of all other stages in $[0, m_1 - \varepsilon_L]$, $[m_1, \tilde{m}]$, and $[\tilde{m} + \varepsilon_R, 1]$ constant. Note in particular that the stages in $[m_1, \tilde{m}]$ are all exogenously outsourced. We now select ε_L and ε_R , so that $0 < \varepsilon_L, \varepsilon_R < \min\{\varepsilon, \tilde{\varepsilon}\}$ and:

$$\int_{m_1 - \varepsilon_L}^{m_1} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di = \int_{\tilde{m}}^{\tilde{m} + \varepsilon_R} (\psi(i)/c(i))^{\alpha/(1-\alpha)} di.$$

We now use the expression for firm profits from (A-8), and distinguish between five sets of stages: (i) all stages upstream of $m_1 - \varepsilon_L$; (ii) stages in $(m_1 - \varepsilon_L, m_1)$; (iii) in $[m_1, \tilde{m}]$; (iv) in $(\tilde{m}, \tilde{m} + \varepsilon_R)$; and (v) in $[\tilde{m} + \varepsilon_R, 1]$. Using the same arguments as in Section A-1.3, the profits associated with both the first and fifth sets of stages can be shown to cancel out exactly when comparing their contributions to $\tilde{\Pi}_1$ and $\tilde{\Pi}_2$. As for the remaining three sets of stages, one can show after some algebra that:

$$\begin{aligned} \tilde{\Pi}_1 - \tilde{\Pi}_2 \propto (\beta_V - \beta_O) & \left[-\tilde{\mathcal{A}} \frac{\rho(1-\alpha)}{\alpha(1-\rho)} + \left(\tilde{\mathcal{A}} + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_{m_1 - \varepsilon_L}^{m_1} \gamma(k) dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right. \\ & + \left(\tilde{\mathcal{A}} + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1 - \varepsilon_L}^{m_1} \gamma(k) dk + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1}^{\tilde{m}} \gamma(k) dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \\ & \left. - \left(\tilde{\mathcal{A}} + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1 - \varepsilon_L}^{m_1} \gamma(k) dk + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1}^{\tilde{m}} \gamma(k) dk + (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m} + \varepsilon_R} \gamma(k) dk \right)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right], \end{aligned}$$

where: $\gamma(k) = (\psi(k)/c(k))^{\frac{\alpha}{1-\alpha}}$, and: $\tilde{\mathcal{A}} = \int_0^{m_1 - \varepsilon_L} ((1 - \beta(k)) \psi(k)/c(k))^{\frac{\alpha}{1-\alpha}} dk$. Again, we can show that $\tilde{\Pi}_1 - \tilde{\Pi}_2 < 0$ by invoking the inequality: $(y + a + b)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - (y + b)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} > (y + a)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} - (y)^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}}$, which holds when $\rho > \alpha$, and substituting in: $y = \tilde{\mathcal{A}}$, $a = (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1 - \varepsilon_L}^{m_1} \gamma(k) dk + (1 - \beta_O)^{\frac{\alpha}{1-\alpha}} \int_{m_1}^{\tilde{m}} \gamma(k) dk$ and $b = (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_{\tilde{m}}^{\tilde{m} + \varepsilon_R} \gamma(k) dk = (1 - \beta_V)^{\frac{\alpha}{1-\alpha}} \int_{m_1 - \varepsilon_L}^{m_1} \gamma(k) dk$.

A-1.8 A Transaction-Cost Variant of the Model

In this Appendix, we consider a transaction-cost variant of our model and compare its predictions to those derived in our property-rights framework. The transaction-cost approach of Coase and Williamson is based on the notion that integration is a means to circumvent the contractual frictions that plague market transactions with independent suppliers. The idea is that, when dealing with integrated suppliers, firms can settle issues by fiat, authority, or disciplinary action, although possibly at the cost of incurring higher *governance* costs. A natural way to capture this tradeoff in our framework is to allow the firm to determine the investments made by integrated suppliers (as well as their compensation) without worrying about negotiations and holdup problems. More specifically, assume that when a supplier is owned by the final-good producer, the firm has the authority to force the supplier to choose a level of investment at stage m that maximizes the supplier's incremental contribution to revenue net of the cost of investment provision. More formally, we assume that, under integration, $x(m)$ is set to maximize $r'(m) - c(m)x(m)$ rather than $(1 - \beta_O)r'(m) - c(m)x(m)$, where $r'(m)$ is given in equation (4) in the main text. Thus, integration resolves the hold-up problem at stage m . Although it would be feasible to study the case in which the firm pre-determines these investments

in an initial contract in a full profit-maximizing manner (as in our extension with contractible investments), it is simplest to consider this case in which the hold-up problem is resolved sequentially.

We also assume that the firm can set the supplier's compensation to exactly meet the supplier's participation constraint. If the supplier's outside option is 0, then the firm captures the full surplus $r'(m) - c(m)x(m)$ at each stage m . It should be clear that without any costs of integration, the firm would have every incentive to integrate all suppliers. To generate a tradeoff, we introduce governance costs in the form of an increase in the marginal cost of production by a factor $\lambda > 1$, so that the marginal cost of an integrated supplier is $c(m) = \lambda\tilde{c}(m)$ if m is integrated and $c(m) = \tilde{c}(m)$ if m is outsourced. In order for integration to still enhance the efficiency of supplier investments, we assume that governance costs are low enough so that $\lambda < 1/(1 - \beta_O)$.

The assumed sequential nature of contracting (even with integration) allows us to solve this version of the model in a manner analogous to that in our benchmark model. In particular, the path of investments continues to be determined by equation (6) in the main text, with $\beta(m) = \beta_O$ if supplier m is outsourced, and with $\beta(m) = 0$ if supplier m is integrated (and with the path of marginal costs potentially affected by integration decisions). The expression for firm profits π_F is slightly different due to the fact that the firm does not capture $\beta(m)r'(m)$ but rather $r'(m) - c(m)x(m)$ at stage m . Using the fact that, under integration, $x(m)$ will satisfy $c(m)x(m) = \alpha r'(m)$, we can express firm profits as:

$$\pi_F = \Theta \int_0^1 \beta_F(i) \left(\frac{\beta_S(i)\psi(i)}{\tilde{c}(i)} \right)^{\frac{\alpha}{1-\alpha}} \left[\int_0^i \left(\frac{\beta_S(i)\psi(k)}{\tilde{c}(k)} \right)^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho-\alpha}{\alpha(1-\rho)}} di, \quad (\text{A-13})$$

where for the case of outsourcing, we have $\beta_F(i) = \beta_O$ and $\beta_S(i) = 1 - \beta_O$, while for the case of integration, we instead have $\beta_F(i) = 1 - \alpha$ and $\beta_S(i) = 1/\lambda$. The constant Θ is defined as $\Theta = A\theta^{\frac{\rho}{1-\rho}} \frac{\rho}{\alpha} \left(\frac{1-\rho}{1-\alpha} \right)^{\frac{\rho-\alpha}{\alpha(1-\rho)}} \rho^{\frac{\rho}{1-\rho}}$. It is possible to show that if $\beta_O < 1 - \alpha$, integration would be the optimal organizational form for all stages $m \in [0, 1]$. We thus focus on the more interesting case $\beta_O > 1 - \alpha$, in which integration and outsourcing may coexist along the value chain.

The key question is then: in which type of stages will the firm find it particularly valuable to resolve the hold-up problem? Our results on the optimal bargaining shares $\beta^*(m)$ in the benchmark model in the main paper suggest that resolving the hold-up problem (i.e., reducing supplier underinvestment) is particularly beneficial in upstream stages for $\rho > \alpha$, and in downstream stages for $\rho < \alpha$. In other words, and *conversely* to our results in the property-rights model, downstreamness should have a negative effect on integration whenever inputs are sequential complements ($\rho > \alpha$), while it should have a positive effect on integration when inputs are sequential substitutes ($\rho < \alpha$).

We next formalize this result along the lines of the proof of Proposition 1. Take the case $\rho > \alpha$, and suppose there exists a stage $\tilde{m} \in (0, 1)$ and a positive constant $\varepsilon > 0$ such that stages in $(\tilde{m} - \varepsilon, \tilde{m})$ are outsourced, while stages in $(\tilde{m}, \tilde{m} + \varepsilon)$ are integrated. This situation would provide a counterexample of our conjecture that, in this transaction-cost model, the firm will only find it profitable to integrate the most upstream stages. We shall then show that this counterexample leads to a contradiction.

Let the firm profits associated with this scenario be denoted by Π_1 . As in the proof of Proposition 1, consider two positive constants ε_L and ε_R such that:

$$\int_{\tilde{m}-\varepsilon_L}^{\tilde{m}} (\psi(i)/\tilde{c}(i))^{\alpha/(1-\alpha)} di = \int_{\tilde{m}}^{\tilde{m}+\varepsilon_R} (\psi(i)/\tilde{c}(i))^{\alpha/(1-\alpha)} di. \quad (\text{A-14})$$

These constants can always be chosen small enough such that the set of stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ are outsourced,

while stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$ are integrated. We shall consider an alternative organizational mode such that the firm instead chooses to integrate the stages in $(\tilde{m} - \varepsilon_L, \tilde{m})$ and outsources the stages in $(\tilde{m}, \tilde{m} + \varepsilon_R)$, while retaining the same organizational decisions for all other stages. Denote the profits of this alternative organizational form by Π_2 . Our goal is then to show that $\Pi_1 < \Pi_2$, which would constitute a contradiction to our claim.

Note that we can rewrite firm profits in (A-13) as:

$$\pi_F = \Theta \frac{\alpha(1-\rho)}{\rho(1-\alpha)} \int_0^1 \beta_F(i) \frac{\partial \left(\left[\int_0^i (\beta_S(k)\psi(k)/\tilde{c}(k))^{\frac{\alpha}{1-\alpha}} dk \right]^{\frac{\rho(1-\alpha)}{\alpha(1-\rho)}} \right)}{\partial i} di. \quad (\text{A-15})$$

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 profit functions Π_1 and Π_2 , so we can ignore them hereafter. Less trivially, the profits generated in the last region are also common in the profit functions Π_1 and Π_2 ; this is easily verified following the steps in Proposition 1).

Defining: $\gamma(i) = (\psi(i)/\tilde{c}(i))^{\frac{\alpha}{1-\alpha}}$ and $\mathcal{A} = \int_0^{\tilde{m}-\varepsilon_L} (\beta_S(k)\psi(k)/\tilde{c}(k))^{\frac{\alpha}{1-\alpha}} dk$, and invoking equation (A-8), we find after some manipulations that:

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

Following the same arguments as in the proof of Proposition 1, it is easy to see that the term in large square brackets is negative. Because $\beta_O > 1 - \alpha$, we can then conclude that $\Pi_1 < \Pi_2$. In sum, in an equilibrium in which integration and outsourcing coexist, integrated stages cannot possibly be downstream from outsourced stages, which is diametrically different from the conclusion we obtained in our property-rights framework.

Carrying out an analogous analysis for the sequential substitutes case, we can then conclude that in this transaction-cost variant of the model:

Proposition 1 (Transaction-Cost Approach). 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 integrated; and (ii) all stages $m \in [m_C^*, 1]$ are outsourced 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 outsourced within firm boundaries; and (ii) all stages $m \in [m_S^*, 1]$ are integrated.

As should be clear, the results we obtain are the opposite of those in Proposition 1, and inconsistent with our empirical evidence. It is worth reiterating that although the proof above relies on the condition $\beta_O > 1 - \alpha$, this is only because if $\beta_O < 1 - \alpha$, integration would be the optimal organizational form for all stages $m \in [0, 1]$.

A-2 Data Appendix

A-2.1 Descriptive Statistics

Table A-1: Firm Characteristics

	10th	Median	90th	Mean	Std Dev
A: Firm variables					
All (320,254 obs.)					
Number of Establishments (incl. self)	1	1	1	1.22	3.44
Number of countries (incl. self)	1	1	1	1.05	0.62
Number of integrated SIC codes	1	2	3	1.95	2.21
Year started	1948	1984	1999	1976.84	24.68
Log (Total employment)	3.045	3.807	5.557	4.088	1.080
Log (Sales in USD) (288,627 obs.)	12.522	15.202	17.059	14.803	2.573
Log (Sales/Employment) (288,627 obs.)	8.479	11.429	12.545	10.731	2.635
MNCs only (6,370 obs.)					
Number of Establishments (incl. self)	2	3	17	8.48	22.74
Number of countries (incl. self)	2	2	6	3.47	3.64
Number of integrated SIC codes	2	5	17	8.10	11.88
Year started	1917	1968	1995	1960.29	33.88
Log (Total employment)	3.912	5.737	8.522	6.031	1.788
Log (Sales in USD) (5,891 obs.)	15.895	17.997	20.934	18.208	1.978
Log (Sales/Employment) (5,891 obs.)	11.229	12.145	13.040	12.110	0.921
Integrated inputs (All firms)					
Total Requirements, tr_{ij}	0.000086	0.003861	0.053442	0.019241	0.052952
Total Requirements, tr_{ij} (excl. self-SIC)	0.000035	0.001232	0.009568	0.006774	0.036741
B: From Input-Output Tables					
Total Requirements, tr_{ij}	0.000006	0.000163	0.002322	0.001311	0.008026
Baseline Upstreamness measure (mean)	1.838	3.094	4.285	3.097	0.955
C: Ratio-Upstreamness measures					
Baseline (mean)	0.494	0.561	0.691	0.590	0.141
Baseline (random pick)	0.495	0.561	0.692	0.590	0.141
Ever-integrated inputs only	0.583	0.656	0.803	0.692	0.179
Manufacturing inputs only	0.548	0.633	0.798	0.657	0.174
Exclude parent sic, manufacturing only	0.590	1.100	2.128	1.269	0.625

Notes: Panels A and C are tabulated for the sample of 320,254 firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. In Panel A, the total requirements summary statistics for “Integrated inputs” are computed over the set of integrated input by parent primary industry pairs pooled across firms in our D&B WorldBase sample; there are 666,656 such pairs, with the count equal to 336,168 if the self-SIC is removed from consideration. In Panel B, the summary statistics are computed over the tr_{ij} coefficients in the 1992 U.S. Input-Output Tables, over all input (i) and output (j) SIC industry pairs for which j is in manufacturing and $tr_{ij} > 0$ (416,349 observations); i includes both manufacturing and non-manufacturing inputs. In Panel C, the Ratio-Upstreamness measures under “mean” and “random pick” refer to the treatment adopted for non-manufacturing inputs when mapping from the original IO1992 to SIC codes.

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

	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Quintile 1	0.455	0.057	0.057	0.060	0.041
Quintile 2	0.057	0.011	0.007	0.006	0.004
Quintile 3	0.057	0.007	0.010	0.007	0.005
Quintile 4	0.060	0.006	0.007	0.013	0.008
Quintile 5	0.041	0.004	0.005	0.008	0.007


Notes: Probability matrix constructed using the subset of 34,651 firms that have integrated at least two manufacturing inputs other than the parent industry self-SIC. For the a -th row and b -th column, we compute the probability that any two randomly drawn integrated manufacturing input SICs of the firm in question come from the a -th and b -th quintiles of $upst_{ij}$ values, where j is the SIC output industry of the firm and the quintiles are taken over all SIC manufacturing inputs i . A simple average of the probabilities across all 34,651 firms is reported.

Table A-3: Industry Characteristics

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

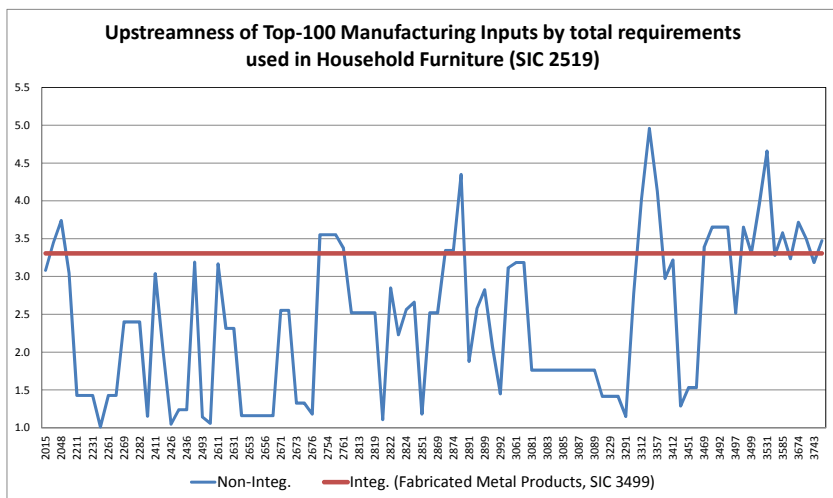
Notes: Summary statistics taken over the 459 SIC manufacturing industries, except for: (i) the “BEC cons. & cap.” elasticity, which is available for only 305 industries; and (ii) the “BEC cons. only” elasticity, which is available for 219 industries. The “contractibility” and “upstream contractibility”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

Figure A-1: Integration Decisions in the Complements Case: An Example



Notes: The figure plots the measure $upst_{ij}$ for the top-100 manufacturing inputs of SIC 3732 (Boat Building and Repairing). This sector exhibits an above-median ρ_j and $\rho_j - \alpha_j$ value regardless of the variant of the demand elasticity proxy considered. The labels on the horizontal axes reflect the SIC codes of these top-100 inputs for SIC 3732. The example is based on a firm in our sample that has only one integrated manufacturing input other than the self-SIC; the upstreamness of this input (SIC 3519) is illustrated by the bold horizontal line.

Figure A-2: Integration Decisions in the Substitutes Case: An Example



Notes: The figure plots the measure $upst_{ij}$ for the top-100 manufacturing inputs of SIC 2519 (Household Furniture). This sector exhibits a below-median ρ_j and $\rho_j - \alpha_j$ value regardless of the variant of the demand elasticity proxy considered. The labels on the horizontal axes reflect the SIC codes of these top-100 inputs for SIC 2519. The example is based on a firm in our sample that has only one integrated manufacturing input other than the self-SIC; the upstreamness of this input (SIC 3499) is illustrated by the bold horizontal line.

A-2.2 Construction of the Industry Controls

Import demand elasticities: Based on the U.S. HS10 product import demand elasticities estimated by Broda and Weinstein (2006). These are mapped into SIC categories using concordance weights based on total U.S. imports between 1989-2006 from Feenstra *et al.* (2002). For each HS10 code missing an elasticity value, we assigned a value equal to the trade-weighted average elasticity of the available HS10 codes with which it shares the same first nine digits. This was done successively up to codes that share the same first two digits, to assign as many HS10 codes with elasticities as possible. The elasticity for each 4-digit SIC code is then calculated as the trade-weighted average over its constituent HS10 elasticities. After these steps, 61 out of the 459 4-digit SIC manufacturing codes remain without elasticities, as these codes are not used in the U.S. import records. This arises because customs is unable to distinguish the source industry of certain goods on the basis of their physical specimen; for example, it cannot distinguish SIC2011 (Meat Packing Plants) from SIC2013 (Sausages and Other Prepared Meats). In such instances, U.S. customs assigns all the goods value to one of the possible SIC codes, and excludes the others. Table 3 in Feenstra *et al.* (2002) provides a list of such excluded codes and their corresponding destination codes, allowing us to compute a trade-weighted elasticity value of the respective destination codes to obtain an elasticity for each excluded code. There were 51 4-digit SIC codes that were successfully assigned in this way. For the remaining 10 4-digit SIC codes, 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.

Contractibility: Following the methodology in Nunn (2007), which in turn relies on the Rauch (1999) classification of goods as either homogeneous, reference-priced, or differentiated. Rauch's original classification is in SITC Rev 2. Based on Feenstra *et al.* (2002), we obtained a master-list of HS by SITC Rev 2 by SIC triplets. The Rauch codings for each SITC Rev 2 category are then associated to all the HS10 products that fall under it. For each SIC 4-digit code, we calculated the specificity of the SIC industry as the fraction of HS10 constituent codes classified as neither reference-priced nor traded on an organized exchange. The procedure described above for import demand elasticities is used to assign the specificity values for missing 4-digit SIC manufacturing codes. The Nunn (2007) measure of contract-intensity of each 4-digit SIC code is then calculated as a direct requirements weighted-average over the specificities of the inputs purchased, using direct requirements coefficients from the 1992 U.S. Input-Output Tables. We take one minus the contract-intensity to get a measure of contractibility.

Factor intensities: From the NBER-CES Manufacturing Industry Database (Becker and Gray, 2009). Skill intensity is the log of the number of non-production workers divided by total employment. Equipment capital intensity and plant capital intensity are respectively the log of the equipment and plant capital stock per worker. Materials intensity is the log of materials purchases per worker. These are computed as averages over 2001-2005, using the annual data for 4-digit SIC industries. For a small number of industries without 2001-2005 data, we used an average over an earlier in-sample window: for SIC 3292 (Asbestos), a 1986-1990 average was used, while for SIC 2411, 2711, 2721 2731, 2741, 2771, and 3732, a 1991-1995 average was used. One further variable – value-added over total shipments – was constructed in the same manner.

R&D intensity: From Nunn and Treffer (2013), who calculated R&D expenditures to total sales on an annual basis for HS6 products, using U.S. firms in the Orbis dataset. For HS6 products missing an R&D intensity value, a procedure analogous to that described above for the import demand elasticities was used, to first assign a value using the trade-weighted average over HS codes that share the same first five digits, and successively until the same first two digits. These are then converted to 4-digit SIC codes using a trade-weighted average R&D intensity of constituent HS6 codes; all concordance weights are based on total U.S. imports between 1989-2006, from Feenstra *et al.* (2002). The procedure described above for import demand elasticities is used to assign the R&D intensity values for missing 4-digit SIC manufacturing codes.

A-3 Robustness Checks

A-3.1 Cross-Firm Regressions

In this Appendix, we document the robustness tests we conduct on our cross-firm regressions.

In Table A-4, we show that the results are robust to examining different subsamples of firms. In Column (1), we restrict the regression sample to single-establishment firms, while in column (2), we focus on domestic firms (these being either single-establishment firms or multi-establishment firms with plants in only one country). In both these cases, we continue to find significant effects on the quintile elasticity dummies, as well as similar patterns on the interaction terms with $UpstCont_j$, i.e., a negative and significant coefficient for the first-quintile interaction, but the opposite sign for the fifth-quintile interaction. Column (3) focuses on parents that have establishments in more than one country, i.e., multinational firms. The empirical findings remain largely intact, despite the fact that the number of observations decreases substantially with this cut of the dataset.³⁴

In Table A-5, we consider several variables that have appeared elsewhere in the literature on firm-level vertical integration. Column (1) adds the share of direct input use in the production of j that could be obtained from within firm boundaries; for each parent, this is the sum of the direct coefficients of the inputs in $I(p)$ (see Acemoglu *et al.* 2009, and Alfaro *et al.* 2016). Column (2) controls for the share of total requirements value that each parent could in principle source from an overseas affiliate, together with a set of country fixed effects that indicate whether the parent has an establishment located in the country in question. Column (3) tests for whether the results might be driven by double marginalization motives, wherein parent firms would have an incentive to integrate inputs that exhibit a low demand elasticity, for which the markups charged by arm’s length suppliers would be higher. We control here for the (log) tr_{ij} -weighted average of the demand elasticity of inputs used by industry j . In addition, we include a tr_{ij} -weighted covariance of the input demand elasticity and $upst_{ij}$, to see if the correlation between these elasticities and production line position might matter. (Here, the demand elasticity associated with each input is computed using only those constituent HS10 products classified as intermediates by the UN BEC.) Our results remain robust to the inclusion of these variables, even when they are jointly entered into the regression (column (4)). Interestingly, the weighted covariance between the input elasticity and upstreamness has a coefficient with the expected sign (negative and significant), consistent with the interpretation that the presence of demand-*inelastic* inputs upstream in the production process would be associated with more upstream integration.³⁵

A key issue that we give due consideration to is how to designate the primary output industry of multi-product firms. In Table A-6, we present several alternative treatments of parent firms that could be active as output producers in multiple manufacturing industries. We first verify whether the patterns are similar when limiting the sample to parents that have only one manufacturing SIC code, i.e., that do not report any secondary SIC manufacturing activities (columns (1) and (2)). Alternatively, we can designate the output industry j to be the SIC code of the parent (among the up to six codes reported) that is the most proximate to final demand, on the basis of the upstreamness measure of Fally (2012) and Antràs *et al.* (2012) (columns (3) and (4)). Last but not least, we have constructed R_{jpc} taking in turn each secondary manufacturing SIC code as the parent’s output industry j . The regression in (22) is then run, pooling across the multiple R_{jpc}

³⁴The results are also unaffected if we expand the sample by lowering the employment threshold to a minimum of 10 employees, or if we restrict the sample to parents labeled as “global ultimates” (results available upon request).

³⁵We have also explored the robustness of our results to the inclusion of several controls related to various dimensions of input contractibility, such as: (i) the contractibility of the output industry j itself; (ii) a tr_{ij} -weighted average of the contractibility of the inputs used by j ; and (iii) a set of interactions between each quintile dummy and a tr_{ij} -weighted variance of the contractibility of the inputs used by j . The results are available upon request.

values per parent (columns (5) and (6)); two-way clustered standard errors by SIC output industry and by parent firm are reported (Cameron *et al.* 2011). Overall, our regression findings remain stable under each of these approaches to account for multi-product firms.

In Table A-7, we report several checks based on alternative constructions of the ratio-upstreamness dependent variable. The version of R_{jpc} in column (1) is based on $upst_{ij}$ values obtained from a random pick when the mapping from I-O to SIC codes yielded multiple matches for a non-manufacturing input i . In column (2), we limit the set $S(j)$ in the construction of R_{jpc} to those inputs for which we observe at least one parent firm in j in our sample integrating the input in question (“ever-integrated” inputs). To be more precise, an input i is considered to be “ever-integrated” in the production of j if there exists a firm with either primary or secondary SIC code in j that has integrated the input i ; this definition is more conservative in terms of which inputs it designates to be “never-integrated”, i.e., always sourced from outside firm boundaries. This is a particularly useful check in light of the sparse nature of integration highlighted in Section 2.2.C, as this variant of the ratio-upstreamness measure should in principle exclude from consideration those inputs for which integration is not feasible in industry j , for example because of high technological or regulatory costs. In column (3), we alternatively restrict $S(j)$ to the set of manufacturing inputs used by industry j , $S^m(j)$. Column (4) further drops the parent SIC from $S^m(j)$, to explore the sensitivity of the results to the default treatment thus far where the parent SIC is always viewed as an integrated input. (There is a decrease in the number of available observations in column (4), since this variant of the ratio-upstreamness measure can only be computed for those parent firms that have integrated at least one other manufacturing input apart from the parent’s primary SIC code.) Our findings are broadly retained, with the main exception being the final column of Table A-7. There, $UpstCont_j$ does reduce the propensity to integrate upstream in the first quintile (the substitutes case), but the point estimates for the fifth-quintile interactions (the complements case) are not significantly different from zero. Note, however, that the overall effect of being in quintile-5 (when evaluated at the median in-sample value of $UpstCont_j$) remains negative and significant, with the p-value from this coefficient test being 0.0043; in other words, the results in column (4) are still very much consistent with the earlier prediction P.1 (Cross).

Table A-8 explores a further robustness check where we restrict the construction of R_{jpc} to larger and hence more relevant inputs. The findings are largely intact when using those inputs with $tr_{ij} \geq 0.001$ (columns (1)-(2)), even though this threshold already exceeds the median total requirements coefficient in the 1992 U.S. I-O Tables (see Table A-1). We lose some precision in our estimates when applying a higher minimum threshold of either $tr_{ij} \geq 0.01$ (columns (3)-(4)) or $tr_{ij} \geq 0.05$ (columns (5)-(6)), but that should not come as a surprise as more than 95% of the inputs are discarded in these later exercises.

We make two further observations to round off this appendix section related to the cross-firm regressions. First, following up on column (2) of Table A-7, we have verified that the value-chain position of the “never-integrated” inputs is not systematically correlated with the quintiles of the various demand elasticity proxies adopted in this paper. In particular, this means that the greater propensity to outsource upstream stages in the complements case is not arising simply because “never-integrated” stages tend to be clustered upstream in high demand elasticity industries. This check is implemented in a series of cross-industry regressions in Table A-9, where the dependent variable is a (log) tr_{ij} -weighted average upstreamness of inputs that are “never-integrated” by firms in industry j relative to the tr_{ij} -weighted average upstreamness of the “ever-integrated” inputs in that industry.

Second, we have also performed similar robustness tests on the within-sector, cross-firm regressions in Table 5. (These results are available in full upon request.) These findings are entirely robust when adopting the alternative constructions of the ratio-upstreamness measure seen earlier in Table A-7, such as that based

on “ever-integrated” inputs or on manufacturing inputs only. We also obtain similar results when restricting these regressions to just single-establishment, domestic, or multinational firms, as in Table A-4, although we do lose statistical significance slightly. This should not come as a surprise, as the regressions related to firm heterogeneity in log output per worker leverage off the contrast in integration patterns across firms at different productivity levels within an industry, and some of this contrast is lost when focusing on specific subsamples of firms.

Table A-4: Cross-Firm Regressions: Different Subsamples

Dependent variable:	Log Ratio-Upstreamness		
	Single-plant firms (1)	Domestic firms (2)	Multinationals (3)
Ind.(Quintile 2 $Elas_j$)	-0.0461 [0.0445]	-0.0487 [0.0432]	-0.0870*** [0.0288]
Ind.(Quintile 3 $Elas_j$)	-0.0630* [0.0338]	-0.0681** [0.0330]	-0.0787*** [0.0279]
Ind.(Quintile 4 $Elas_j$)	-0.1625*** [0.0284]	-0.1619*** [0.0278]	-0.1103*** [0.0268]
Ind.(Quintile 5 $Elas_j$)	-0.1638*** [0.0299]	-0.1649*** [0.0294]	-0.1206*** [0.0330]
Upstream Contractibility $_j$			
× Ind.(Quintile 1 $Elas_j$)	-1.8620*** [0.4612]	-1.8635*** [0.4498]	-1.5014*** [0.3691]
× Ind.(Quintile 2 $Elas_j$)	-0.7401 [0.8055]	-0.7030 [0.7713]	0.2330 [0.3979]
× Ind.(Quintile 3 $Elas_j$)	-0.4965 [0.3919]	-0.4335 [0.3899]	0.2476 [0.2838]
× Ind.(Quintile 4 $Elas_j$)	0.6749*** [0.2162]	0.6890*** [0.2117]	0.5686** [0.2484]
× Ind.(Quintile 5 $Elas_j$)	1.1025*** [0.2321]	1.1195*** [0.2286]	0.9941*** [0.2949]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.1000]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y
Firm controls	Y	Y	Y
Parent country dummies	Y	Y	Y
Observations	117,956	141,617	2,490
No. of industries	219	219	199
R ²	0.2990	0.3027	0.2467

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

Table A-5: Cross-Firm Regressions: Additional Controls

Dependent variable:	Log Ratio-Upstreamness			
	(1)	(2)	(3)	(4)
Ind.(Quintile 2 $Elas_j$)	-0.0429 [0.0414]	-0.0491 [0.0430]	-0.0492 [0.0403]	-0.0418 [0.0386]
Ind.(Quintile 3 $Elas_j$)	-0.0549* [0.0305]	-0.0683** [0.0328]	-0.0532* [0.0308]	-0.0384 [0.0293]
Ind.(Quintile 4 $Elas_j$)	-0.1601*** [0.0253]	-0.1613*** [0.0277]	-0.1437*** [0.0230]	-0.1444*** [0.0213]
Ind.(Quintile 5 $Elas_j$)	-0.1546*** [0.0269]	-0.1642*** [0.0292]	-0.1666*** [0.0258]	-0.1565*** [0.0233]
Upstream Contractibility $_j$				
× Ind.(Quintile 1 $Elas_j$)	-1.6826*** [0.4083]	-1.8554*** [0.4451]	-1.6147*** [0.3643]	-1.4820*** [0.3275]
× Ind.(Quintile 2 $Elas_j$)	-0.6775 [0.7338]	-0.6876 [0.7626]	-0.5599 [0.7994]	-0.6227 [0.7701]
× Ind.(Quintile 3 $Elas_j$)	-0.5875 [0.3681]	-0.4186 [0.3854]	-0.4597 [0.4041]	-0.6614* [0.3966]
× Ind.(Quintile 4 $Elas_j$)	0.5891*** [0.1714]	0.6850*** [0.2105]	0.6457*** [0.2157]	0.5434*** [0.1890]
× Ind.(Quintile 5 $Elas_j$)	0.9582*** [0.2165]	1.1183*** [0.2272]	1.1302*** [0.2518]	0.9516*** [0.2393]
Vertical Integration Index $_p$	-1.1296*** [0.2065]			-1.1144*** [0.2044]
Foreign integrated tr. share $_p$		-1.0690*** [0.1330]		-0.2034* [0.1214]
Log (Input Elasticity) $_j$			-0.2999*** [0.1099]	-0.2853*** [0.1024]
Wtd. Cov. of Input Elasticity $_j$ and upstreamness $_{ij}$			-0.4963*** [0.1718]	-0.4330*** [0.1555]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.0000]	[0.0000]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Subsidiary country dummies	N	Y	N	Y
Observations	144,107	144,107	144,107	144,107
No. of industries	219	219	219	219
R ²	0.3526	0.3079	0.3204	0.3655

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

Table A-6: Parent Firms with Multiple SIC Output Activities

Dependent variable:	Log Ratio-Upstreamness					
	Single		Most downstream		Firm by mfg. output	
	mfg. output SIC	mfg. output SIC	mfg. output SIC	mfg. output SIC	SIC (two-way cluster)	SIC (two-way cluster)
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Quintile 2 $Elas_j$)	-0.0779 [0.0527]	-0.0419 [0.0464]	-0.0586 [0.0446]	-0.0387 [0.0414]	-0.0744* [0.0426]	-0.0476 [0.0428]
Ind.(Quintile 3 $Elas_j$)	-0.1147** [0.0465]	-0.1021*** [0.0292]	-0.0588 [0.0446]	-0.0218 [0.0458]	-0.0793* [0.0412]	-0.0362 [0.0398]
Ind.(Quintile 4 $Elas_j$)	-0.1671*** [0.0503]	-0.1521*** [0.0305]	-0.1422*** [0.0444]	-0.1455*** [0.0293]	-0.1645*** [0.0411]	-0.1642*** [0.0256]
Ind.(Quintile 5 $Elas_j$)	-0.1789*** [0.0493]	-0.1521*** [0.0306]	-0.1559*** [0.0463]	-0.1481*** [0.0316]	-0.1834*** [0.0431]	-0.1680*** [0.0286]
Upstream Contractibility _j						
× Ind.(Quintile 1 $Elas_j$)		-1.9121*** [0.4691]		-1.5439*** [0.4575]		-1.7766*** [0.4150]
× Ind.(Quintile 2 $Elas_j$)		-0.7892 [0.7723]		-0.4447 [0.6291]		-0.5588 [0.7887]
× Ind.(Quintile 3 $Elas_j$)		0.1059 [0.2068]		-0.8775 [0.6081]		-0.8416 [0.5438]
× Ind.(Quintile 4 $Elas_j$)		0.6619*** [0.2346]		0.6950*** [0.2115]		0.6808*** [0.2039]
× Ind.(Quintile 5 $Elas_j$)		1.1166*** [0.2104]		1.2290*** [0.2640]		1.1637*** [0.2544]
p-value: Q5 at median $UpstCont_j$		[0.0000]		[0.0000]		[0.0000]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y	Y	Y
Observations	97,174	97,174	146,829	146,829	211,232	211,232
No. of industries	219	219	219	219	—	—
R ²	0.2469	0.3308	0.1951	0.2649	0.2204	0.2881

Notes: ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. Columns (1) and (2) restrict the sample to those firms with at least 20 employees that report only one SIC manufacturing output activity, this being their primary SIC industry; robust standard errors clustered by output industry are reported. For columns (3) and (4), we designate as the output industry the SIC manufacturing activity of the firm that has the smallest upstreamness value with respect to final demand, this being the measure developed by Antràs *et al.* (2012); robust standard errors clustered by this output industry are reported. For columns (5) and (6), each observation is a parent firm by SIC output activity pair, and the ratio-upstreamness variable is constructed treating in turn each SIC manufacturing activity as the output industry for the firm in question; two-way clustered standard errors – by parent firm and by SIC output activity – are reported. “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used in the SIC output industry in question. Quintile dummies are used to distinguish firms/plants with primary SIC output that are in high vs low demand elasticity industries; the measure used here is based only on HS10 codes classified as consumption goods in the UN BEC. All columns include the full list of SIC output industry controls, firm-level variables, and parent country dummies that were used in the earlier specifications in Table 2, columns (3)-(6).

Table A-7: Alternative Constructions of Ratio-Upstreamness

Dependent variable:	Log Ratio-Upstreamness			
	Random pick (1)	“Ever-Integrated” inputs (2)	Mfg. inputs only (3)	Mfg. inputs only, drop parent SIC (4)
Ind.(Quintile 2 $Elas_j$)	-0.0481 [0.0428]	-0.0240 [0.0413]	-0.0385 [0.0497]	-0.0262 [0.0926]
Ind.(Quintile 3 $Elas_j$)	-0.0687** [0.0329]	-0.0402 [0.0341]	-0.0786** [0.0394]	-0.0642 [0.0514]
Ind.(Quintile 4 $Elas_j$)	-0.1574*** [0.0277]	-0.1293*** [0.0307]	-0.1825*** [0.0320]	-0.1388** [0.0661]
Ind.(Quintile 5 $Elas_j$)	-0.1652*** [0.0303]	-0.1313*** [0.0261]	-0.1762*** [0.0396]	-0.2958*** [0.0934]
Upstream Contractibility _{j}				
× Ind.(Quintile 1 $Elas_j$)	-1.8583*** [0.4454]	-0.8338*** [0.3137]	-2.1696*** [0.4819]	-1.1117* [0.5749]
× Ind.(Quintile 2 $Elas_j$)	-0.6960 [0.7602]	-0.8880 [0.7960]	-0.9343 [0.9046]	0.0021 [0.8379]
× Ind.(Quintile 3 $Elas_j$)	-0.4193 [0.3873]	0.0377 [0.4977]	-0.2726 [0.4890]	-1.8093* [0.9849]
× Ind.(Quintile 4 $Elas_j$)	0.6473*** [0.2126]	0.9039*** [0.3313]	0.8981*** [0.2504]	-2.5374*** [0.7379]
× Ind.(Quintile 5 $Elas_j$)	1.1816*** [0.2803]	1.3664*** [0.2992]	1.1370*** [0.3822]	-0.0754 [1.1158]
p-value: Q5 at median $UpstCont_j$	[0.0000]	[0.0000]	[0.0000]	[0.0043]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y
Observations	144,107	144,107	143,846	46,992
No. of industries	219	219	219	218
R ²	0.3059	0.1950	0.3311	0.1216

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

Table A-8: Dropping Inputs with Small Total Requirements Coefficients

Dependent variable:	Log Ratio-Upstreamness limited to Inputs with $tr_{ij} \geq \kappa$					
	$\kappa = 0.001$		$\kappa = 0.01$		$\kappa = 0.05$	
	(1)	(2)	(3)	(4)	(5)	(6)
Ind.(Quintile 2 $Elas_j$)	-0.0960** [0.0480]	-0.0772** [0.0318]	-0.0371 [0.0292]	-0.0573* [0.0304]	0.0936 [0.0969]	-0.0691 [0.0604]
Ind.(Quintile 3 $Elas_j$)	-0.1007** [0.0456]	-0.0967*** [0.0306]	-0.0169 [0.0402]	-0.0672* [0.0380]	0.2292* [0.1159]	0.0901 [0.0986]
Ind.(Quintile 4 $Elas_j$)	-0.1584*** [0.0492]	-0.1509*** [0.0308]	-0.1357*** [0.0470]	-0.1416*** [0.0415]	-0.2755* [0.1593]	-0.3050*** [0.1138]
Ind.(Quintile 5 $Elas_j$)	-0.1882*** [0.0493]	-0.1624*** [0.0299]	-0.1109* [0.0597]	-0.1047** [0.0469]	0.0980 [0.1768]	0.0943 [0.1472]
Upstream Contractibility _j						
× Ind.(Quintile 1 $Elas_j$)		-1.8097*** [0.4775]		0.3782 [0.7651]		-2.2377* [1.1530]
× Ind.(Quintile 2 $Elas_j$)		-0.1884 [0.3140]		0.7700 [0.5218]		0.7183 [0.8877]
× Ind.(Quintile 3 $Elas_j$)		0.2387 [0.2381]		1.5469*** [0.4525]		-4.8807** [2.0299]
× Ind.(Quintile 4 $Elas_j$)		0.7049*** [0.2172]		1.2898*** [0.4787]		3.1515*** [0.7650]
× Ind.(Quintile 5 $Elas_j$)		1.4457*** [0.2034]		2.5916*** [0.3724]		4.5668*** [0.7617]
p-value: Q5 at median $UpstCont_j$		[0.0000]		[0.0002]		[0.7171]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Industry controls	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y
Parent country dummies	Y	Y	Y	Y	Y	Y
Observations	139,053	139,053	81,970	81,970	13,677	13,677
No. of industries	219	219	214	214	98	98
R ²	0.3144	0.4308	0.4995	0.6285	0.4950	0.6873

Notes: The sample comprises firms with primary SIC in manufacturing and at least 20 employees in the 2004/2005 vintage of D&B WorldBase. Standard errors are clustered by the parent primary SIC industry; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is the log ratio upstreamness measure constructed when limiting the set of integrated and non-integrated inputs under consideration to those with total requirements coefficient respectively greater than or equal to κ , where $\kappa = 0.001$ in columns (1)-(2), $\kappa = 0.01$ in columns (3)-(4), and $\kappa = 0.05$ in columns (5)-(6). “Upstream Contractibility” is the total requirements weighted covariance between the contractibility and upstreamness of the manufacturing inputs used in the SIC output industry in question. Quintile dummies are used to distinguish firms/plants with primary SIC output that are in high vs low demand elasticity industries; the measure used here is based only on HS10 codes classified as consumption goods in the UN BEC. All columns include the full list of SIC output industry controls, parent country dummies, and the full list of firm-level variables used in the earlier specifications of Table 2, columns (3)-(6). The sample size decreases with higher κ , as firms that do not have any integrated inputs satisfying $tr_{ij} \geq \kappa$ are dropped.

Table A-9: Diagnostic: Upstreamness of Never- vs Ever-Integrated Inputs

Dependent variable:	$\log \left(\frac{\text{Wtd. Avg. } upst_{ij} \text{ of Never-Integrated Inputs}}{\text{Wtd. Avg. } upst_{ij} \text{ of Ever-Integrated Inputs}} \right)$				
	(1)	(2)	(3)	(4)	(5)
Ind.(Quintile 2 $Elas_j$)	0.0263 [0.0201]	0.0213 [0.0191]	-0.0006 [0.0253]	-0.0228 [0.0336]	-0.0142 [0.0326]
Ind.(Quintile 3 $Elas_j$)	0.0284 [0.0214]	0.0139 [0.0203]	-0.0017 [0.0239]	0.0127 [0.0318]	0.0272 [0.0306]
Ind.(Quintile 4 $Elas_j$)	0.0000 [0.0228]	0.0103 [0.0226]	0.0101 [0.0276]	0.0108 [0.0362]	0.0108 [0.0346]
Ind.(Quintile 5 $Elas_j$)	-0.0024 [0.0210]	-0.0033 [0.0224]	-0.0188 [0.0277]	-0.0460 [0.0384]	-0.0371 [0.0385]
Log (Skilled Emp./Workers) $_j$		-0.0346* [0.0198]	-0.0211 [0.0231]	-0.0282 [0.0338]	-0.0264 [0.0345]
Log (Equip. Capital/Workers) $_j$		-0.0896*** [0.0184]	-0.0981*** [0.0223]	-0.1245*** [0.0280]	-0.1237*** [0.0289]
Log (Plant Capital/Workers) $_j$		0.0545*** [0.0183]	0.0742*** [0.0214]	0.0904*** [0.0252]	0.0936*** [0.0257]
Log (Materials/Workers) $_j$		-0.0283 [0.0213]	-0.0243 [0.0245]	-0.0077 [0.0314]	-0.0119 [0.0318]
R&D intensity $_j$		-0.0029 [0.0049]	-0.0107 [0.0067]	-0.0151 [0.0095]	-0.0129 [0.0094]
(Value-added/Shipments) $_j$		-0.3011*** [0.1047]	-0.2459** [0.1158]	-0.1835 [0.1328]	-0.2023 [0.1347]
p-value: F-test, $Elas_j$ quintile coeffs.	[0.3437]	[0.6596]	[0.8364]	[0.2270]	[0.1980]
Elasticity based on:	All goods	All goods	BEC cons. & cap. goods	BEC cons. goods	BEC cons. & α proxy
Observations	459	459	305	219	219
R ²	0.0093	0.2335	0.2408	0.2918	0.2913

Notes: The sample comprises all SIC manufacturing industries. Robust standard errors are reported; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. For each output industry j , the set of never-integrated inputs is the list of inputs i that are never found among the SIC activities of D&B parent firms with either primary or secondary output industry listed as j , while the set of ever-integrated inputs is the list of inputs i that are integrated within firm boundaries by at least one D&B parent firm with primary or secondary industry listed as j . The D&B parent firms considered are those from the 2004/2005 vintage with at least 20 employees. The dependent variable is the log ratio of the weighted average upstreamness of never-integrated to that of ever-integrated inputs, where the weights are proportional to tr_{ij} . Quintile dummies are used to distinguish firms with primary SIC output that are in high vs low demand elasticity industries. Columns (1)-(2) use a measure based on all available HS10 elasticities from Broda and Weinstein (2006); Column (3) restricts this construction to HS codes classified as consumption or capital goods in the UN BEC; Column (4) further restricts this to consumption goods only; Column (5) uses the consumption-goods-only demand elasticity minus a proxy for α to distinguish between the complements and substitutes cases. The p-value reported is that from an F-test with null hypothesis that the coefficients of the Quintile 2 through Quintile 5 $Elas_j$ dummies are jointly equal to zero.

A-3.2 Robustness: Within-Firm Regressions

In this Appendix, we broaden our discussion of the robustness checks performed on the within-firm regressions.

We report in Table A-10 tests related to limiting the regression to specific subsamples of firms, namely single-establishment, domestic, and multinational firms respectively. This is done first using the median elasticity cutoff specification (in columns (1)-(3)), and then for the quintile cutoff specification (in columns (4)-(6)). Reassuringly, these retain the broad patterns seen in Tables 6 and 7 in the main paper. The one exception where the results are slightly weaker is for the subsample of multinationals when adopting the quintile cutoff regression specification (column (6)); here, the fourth and fifth quintile elasticity dummies interacted respectively with “contractibility up to i ” are positive but statistically insignificant.

In Table A-11, we undertake a number of further checks and illustrate these with the quintile cutoff specifications; the findings here are similarly strong under the median cutoff specification (available on request). In columns (1) and (2), we focus on subsets of firms that feature more interesting variation in their integration patterns. Column (1) drops firms that do not have an integrated manufacturing input (apart from the self-SIC) among the top 100 inputs as ranked by the total requirements value. Alternatively, column (2) retains only those parents that have integrated at least three of their top-100 manufacturing inputs. The findings we obtain from these different subsample cuts of the data turn out to be very similar to those presented already in Table 7. Column (3) adopts a different treatment of the self-SIC code, which is classified mechanically as an integrated input in our regressions. Here, the self-SIC is instead dropped altogether from the estimation. While we lose some statistical significance on the effect of “contractibility up to i ” in the lowest elasticity quintiles, the positive and significant coefficients in the highest quintiles that map to the complements case are preserved. Last but not least, in column (4), we include the full set of quintile elasticity dummies interacted with “contractibility at i ”, where this latter variable is given by: $\frac{tr_{ij}cont_i}{\sum_{k \in S^m(j)} tr_{kj}cont_k}$. In words, this is the component of $ContUpToi_{ij}$ that is accrued at stage i itself. The results indicate that it is indeed the profile of contractibility prior to input i , rather than that at stage i , that matters for explaining integration patterns.

Table A-10: Integration Decisions within Firms (Top 100 Inputs): Different Subsamples

Dependent variable: Quantile:	Indicator variable: Input Integrated?					
	Median-cutoff specification			Quintile-cutoff specification		
	Single-plant (1)	Domestic (2)	Multinationals (3)	Single-plant (4)	Domestic (5)	Multinationals (6)
Upstreamness _{<i>ij</i>}						
× Ind.(Quantile 1 Elas _{<i>j</i>})	0.0003 [0.0020]	0.0001 [0.0021]	0.0024 [0.0025]	0.0008 [0.0014]	0.0008 [0.0015]	0.0061* [0.0032]
× Ind.(Quantile 2 Elas _{<i>j</i>})	0.0060*** [0.0014]	0.0063*** [0.0014]	0.0067** [0.0029]	0.0001 [0.0033]	0.0002 [0.0033]	-0.0004 [0.0031]
× Ind.(Quantile 3 Elas _{<i>j</i>})				0.0027 [0.0021]	0.0017 [0.0025]	0.0046 [0.0046]
× Ind.(Quantile 4 Elas _{<i>j</i>})				0.0063*** [0.0018]	0.0063*** [0.0017]	0.0054 [0.0033]
× Ind.(Quantile 5 Elas _{<i>j</i>})				0.0053*** [0.0018]	0.0058*** [0.0019]	0.0044 [0.0045]
Cont. up to <i>i</i> (in prod. of <i>j</i>)						
× Ind.(Quantile 1 Elas _{<i>j</i>})	0.0102* [0.0059]	0.0101* [0.0061]	0.0192** [0.0080]	0.0090* [0.0046]	0.0096** [0.0047]	0.0294*** [0.0113]
× Ind.(Quantile 2 Elas _{<i>j</i>})	0.0248*** [0.0054]	0.0264*** [0.0055]	0.0337*** [0.0097]	0.0114 [0.0098]	0.0120 [0.0101]	0.0024 [0.0101]
× Ind.(Quantile 3 Elas _{<i>j</i>})				0.0151** [0.0076]	0.0132 [0.0081]	0.0420*** [0.0148]
× Ind.(Quantile 4 Elas _{<i>j</i>})				0.0231*** [0.0088]	0.0233*** [0.0086]	0.0187 [0.0129]
× Ind.(Quantile 5 Elas _{<i>j</i>})				0.0270*** [0.0087]	0.0295*** [0.0091]	0.0283 [0.0173]
Dummy: Self-SIC	0.9333*** [0.0092]	0.9329*** [0.0086]	0.8837*** [0.0101]	0.9333*** [0.0092]	0.9330*** [0.0086]	0.8835*** [0.0100]
Log (Total Requirements _{<i>ij</i>})	0.0051*** [0.0008]	0.0053*** [0.0009]	0.0108*** [0.0014]	0.0051*** [0.0008]	0.0053*** [0.0009]	0.0110*** [0.0014]
p-value: Cont. up to <i>i</i> , Bottom minus Top Quantile	[0.0418]	[0.0271]	[0.2175]	[0.0581]	[0.0464]	[0.9601]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Firm fixed effects	Y	Y	Y	Y	Y	Y
Input industry <i>i</i> fixed effects	Y	Y	Y	Y	Y	Y
Observations	3,608,516	4,525,106	182,607	3,608,516	4,525,106	182,607
No. of parent firms	36,019	45,169	1,823	36,019	45,169	1,823
No. of <i>i-j</i> pairs	21,836	21,836	19,519	21,836	21,836	19,519
R ²	0.5799	0.5715	0.3936	0.5799	0.5715	0.3938

Notes: Each observation is a SIC input by parent firm pair, where the set of parent firms is that from Table A-7, column (4), namely firms with primary SIC industry in manufacturing and employment of at least 20, which have integrated at least one manufacturing input apart from the output self-SIC. Manufacturing inputs ranked in the top 100 by total requirements coefficients of the SIC output industry are included. Standard errors are clustered by input-output industry pair; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is a 0-1 indicator for whether the SIC input is integrated. The “Contractibility up to *i*” measure is the share of the total requirements weighted contractibility of inputs that has been accrued in production upstream of and including input *i* in the production of output *j*. Columns (1)-(3) are based on the median cutoff specification as in Table 6, while columns (4)-(6) are based on the quintile cutoff specification as in Table 7. The median and quintile dummies are based on the elasticity measure constructed using only those HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC. The sample in each column restricts to different subsets of firms from the 2004/2005 vintage of D&B WorldBase, as described in each column heading. All columns include parent firm fixed effects and SIC input industry fixed effects.

Table A-11: Integration Decisions within Firms (Top 100 Inputs): Further Robustness

Dependent variable:	Indicator variable: Input Integrated?			
	# non-self-SIC integ. inputs ≥ 1 (1)	# integ. inputs ≥ 3 (2)	Drop self-SIC (3)	Contractibility at i (4)
Upstreamness $_{ij}$				
× Ind.(Quintile 1 Elas $_j$)	0.0019 [0.0022]	0.0020 [0.0039]	0.0006 [0.0015]	0.0001 [0.0014]
× Ind.(Quintile 2 Elas $_j$)	0.0006 [0.0055]	-0.0038 [0.0098]	0.0001 [0.0035]	0.0001 [0.0033]
× Ind.(Quintile 3 Elas $_j$)	0.0041 [0.0035]	0.0047 [0.0059]	0.0009 [0.0026]	0.0017 [0.0025]
× Ind.(Quintile 4 Elas $_j$)	0.0108*** [0.0026]	0.0115*** [0.0038]	0.0068*** [0.0017]	0.0055*** [0.0018]
× Ind.(Quintile 5 Elas $_j$)	0.0102*** [0.0031]	0.0096** [0.0044]	0.0050** [0.0023]	0.0043** [0.0021]
Cont. up to i (in prod. of j)				
× Ind.(Quintile 1 Elas $_j$)	0.0158** [0.0074]	0.0338*** [0.0128]	0.0071 [0.0049]	0.0094** [0.0046]
× Ind.(Quintile 2 Elas $_j$)	0.0228 [0.0166]	0.0271 [0.0292]	0.0094 [0.0105]	0.0163 [0.0108]
× Ind.(Quintile 3 Elas $_j$)	0.0236** [0.0114]	0.0426** [0.0189]	0.0082 [0.0085]	0.0176** [0.0084]
× Ind.(Quintile 4 Elas $_j$)	0.0393*** [0.0127]	0.0527*** [0.0188]	0.0253*** [0.0086]	0.0221** [0.0093]
× Ind.(Quintile 5 Elas $_j$)	0.0501*** [0.0141]	0.0599*** [0.0200]	0.0256*** [0.0093]	0.0206** [0.0095]
Dummy: Self-SIC	0.9031*** [0.0114]	0.8409*** [0.0165]		0.9312*** [0.0087]
Log (Total Requirements $_{ij}$)	0.0084*** [0.0013]	0.0136*** [0.0021]	0.0056*** [0.0008]	0.0046*** [0.0011]
p-value: Cont. up to i , Quintile 1 minus Quintile 5	[0.0282]	[0.2539]	[0.0798]	[0.2858]
Elasticity based on:	BEC cons.	BEC cons.	BEC cons.	BEC cons.
Firm fixed effects	Y	Y	Y	Y
Input industry i fixed effects	Y	Y	Y	Y
Cont. at i interactions	N	N	N	Y
Observations	3,001,343	700,443	4,662,172	4,707,722
No. of parent firms	29,967	6,995	46,992	46,992
No. of i - j pairs	21,835	20,223	21,633	21,836
R ²	0.4668	0.3750	0.0791	0.5601

Notes: Each observation is a SIC input by parent firm pair, where the set of parent firms is that from Table A-7, column (4), namely firms with primary SIC industry in manufacturing and employment of at least 20, which have integrated at least one manufacturing input apart from the output self-SIC. Manufacturing inputs ranked in the top 100 by total requirements coefficients of the SIC output industry are included. Standard errors are clustered by input-output industry pair; ***, **, and * denote significance at the 1%, 5%, and 10% levels respectively. The dependent variable is a 0-1 indicator for whether the SIC input is integrated. The “Cont. up to i ” measure is the share of the total-requirements weighted contractibility of inputs that has been accrued in production upstream of and including input i in the production of output j . The quintile dummies are based on the elasticity measure constructed using only those HS10 elasticities from Broda and Weinstein (2006) classified as consumption goods in the UN BEC. All columns include parent firm fixed effects and SIC input industry fixed effects. Column (4) further controls for the full set of quintile elasticity dummies interacted with the “Cont. at i ” measure, namely the share of the total-requirements weighted contractibility of inputs accrued at stage i itself (coefficients not reported).