

Kohler, Wilhelm; Smolka, Marcel

Working Paper

Productivity and Firm Boundaries

CESifo Working Paper, No. 7214

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Kohler, Wilhelm; Smolka, Marcel (2018) : Productivity and Firm Boundaries, CESifo Working Paper, No. 7214, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/185412>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Productivity and Firm Boundaries

Wilhelm Kohler, Marcel Smolka

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Productivity and Firm Boundaries

Abstract

This paper develops and applies a test of the property rights theory of the firm in the context of global input sourcing. We use the model by Pol Antràs and Elhanan Helpman, “Global Sourcing,” *Journal of Political Economy*, 112:3 (2004), 552-80, to derive a new prediction regarding how the productivity of a firm affects its choice between vertical integration and outsourcing and how this effect depends on the relative input intensity of the production process. The prediction we derive hinges on less restrictive assumptions than industry-level predictions available in existing literature and survives in more realistic versions of the model featuring multiple suppliers and partial vertical integration. We present robust firm-level evidence from Spain showing that, in line with our prediction, the effect of productivity works more strongly in favor of vertical integration, and against outsourcing, in more headquarter-intensive industries.

JEL-Codes: F120, F190, F230, L220, L230.

Keywords: global sourcing, incomplete contracts, property rights theory, firm productivity.

Wilhelm Kohler
University of Tübingen
Mohlstrasse 36
Germany – 72074 Tübingen
wilhelm.kohler@uni-tuebingen.de

Marcel Smolka
Aarhus University
Fuglesangs Allé 4, Building 2632
Denmark - 8210 Aarhus
msmolka@econ.eu.dk

August 2018

This is a substantially revised version of “Global Sourcing of Heterogeneous Firms: Theory and Evidence” published as CESifo Working Paper No. 5184 (January 2015). We would like to thank Pol Antràs, Hartmut Egger, Peter Egger, Anders Laugesen, Hong Ma, and Jens Südekum for helpful comments and discussions. Preliminary versions of this paper have been presented at various universities and conferences. Financial support from the Volkswagen Foundation under the project “Europe’s Global Linkages and the Impact of the Financial Crisis: Policies for Sustainable Trade, Capital Flows, and Migration” is gratefully acknowledged. Thanks are also due to the Tuborg Research Centre at Aarhus University for facilitating mutual research visits to work on this paper.

1 Introduction

Multinational firms (MNFs) play a vital role in the global economy.¹ What explains the existence and the boundaries of MNFs? To answer this question, trade economists have focused attention on the property rights theory of the firm (Grossman & Hart, 1986; Hart & Moore, 1990), and they have developed models able to explain the fundamental choice between intra-firm trade and trade at arm's length. In this paper we develop and apply a test of the most influential model in this literature: the global sourcing model by Antràs & Helpman (2004) (henceforth AH). This model embeds property rights theory into a standard model of trade with heterogeneous firms. The contribution of our paper is to confront the theory with data by deriving a novel prediction from the model and by testing this prediction on Spanish firm-level data.

At the heart of the AH model is a hold-up problem between a headquarter firm and its input supplier due to incomplete contracts and relationship-specific inputs. The firm seeks to minimize the deadweight loss associated with this problem by trading off investment incentives through an optimal allocation of ownership and control rights over the supplier's assets. This boils down to the classical choice between vertical integration and outsourcing (involving either intra-firm trade or trade at arm's length, respectively). A key insight of the model with two inputs—a headquarter input and a supplier input—is that this choice is governed by the productivity of the firm, not in isolation, but in interaction with the relative intensity of these inputs in the production process.

On the theoretical side, our contribution lies in deriving a new firm-level prediction from the model that characterizes the relative attractiveness of vertical integration as a function of the firm's productivity. The prediction we derive hinges on less restrictive assumptions than industry-level predictions available in existing literature and is robust to extending the model in several important directions. The starting point of our analysis is the fact that in supplier-intensive industries, where the supplier input carries sufficient weight, outsourcing generates higher variable profits than vertical integration. The opposite holds true in headquarter-intensive industries, where the headquarter input looms large enough in the production process. This result is not surprising, as it reflects a key insight of property rights theory: ownership rights are optimally assigned to the party undertaking the more important investment. However, the corresponding implications for the effect of firm productivity on the choice between vertical integration and outsourcing, in our view, have not received sufficient attention in the literature on global sourcing. In particular, since the firm's productivity magnifies *any* variable profit advantage, its effect can clearly go either way: favoring outsourcing in supplier-intensive industries, and vertical integration in headquarter-intensive industries.

We then go one step further and undertake a comprehensive comparative static analysis of the productivity effect. In particular, we show how the effect responds to (small) changes in the relative input intensity of the production process. Importantly, we consider the full range of possible input intensities in this analysis. This is a crucial step towards an empirical test of the model, as it allows formulating expectations about how the (firm-level) productivity effect varies across industries with different supplier intensities. For a plausible parameter subspace of the model, we find a monotonic relationship implying a productivity effect that is generally more favorable to outsourcing (and less so to vertical integration) in more supplier-intensive industries.

The novelty of our theoretical analysis lies in a crucial departure from the exposition in AH and the related literature. AH derive predictions at the industry-level, describing the distribution of

¹Recent estimates suggest that as much as 80% of total world trade takes place in global value chains orchestrated by MNFs; see UNCTAD World Investment Report 2013 at <http://unctad.org/en/pages/PressRelease.aspx?OriginalVersionID=113>.

ownership decisions across firms in relation to various industry-specific parameters (in particular the supplier intensity, the distribution of productivity, and the fixed cost of sourcing). What matters in their set-up is whether a firm enjoys higher profits under vertical integration or outsourcing. In contrast, we develop a different and novel perspective on the model shifting the focus away from the *industry* towards the *firm*. To do so, we extend the analysis in AH to investigate, not just whether vertical integration or outsourcing promises higher profits, but also how large the profit difference is going to be. This is crucial for the transition from theory to empirics, where a larger profit difference predicted by the model translates into a higher *probability* for a firm to choose vertical integration over outsourcing. Our approach has three main advantages. First, the profit difference as a measure of the attractiveness of vertical integration can be investigated empirically using latent variable models that are consistent with the principle of profit maximization and familiar from the discrete choice literature. Secondly, because changes in the profit difference induced by changes in productivity are independent of the fixed cost of sourcing, our results regarding the effect of productivity on the choice between vertical integration and outsourcing, unlike the industry-level predictions in existing literature, do not hinge on any particular fixed cost ranking. And finally, since our unit of analysis is the firm rather than the industry, we do not need any assumption on the distribution of productivity across firms. This is different from the industry-level predictions in Antràs & Helpman (2004) which require that productivity be distributed Pareto across firms.

To investigate the robustness of our theoretical results, we consider three extensions of the AH model that allow for many input suppliers per firm. These extensions bridge the gap between the stylized single-supplier case of the AH model and empirical data, where firms frequently use multiple input suppliers at the same time and decide to outsource some of them while integrating others (Kohler & Smolka, 2011). First, we embed a continuum of perfectly contractible inputs into the AH model, while maintaining the assumption of a single non-contractible input. Secondly, and more importantly, we consider a continuum of perfectly non-contractible inputs. We do so employing two different set-ups: the first assumes that these inputs are produced simultaneously, as in Schwarz & Südekum (2014), the second imposes a sequencing of production for the non-contractible inputs, as in Antràs & Chor (2013). Overall, these extensions reveal a remarkable robustness of our basic prediction regarding the productivity effect and demonstrate that the specific details of the assumed production technology and contracting environment are largely inconsequential for our analysis. That being said, we find that a clear conceptual distinction is needed between the extensive margin of integration (whether or not a firm integrates at least one supplier) and the intensive margin of integration (the share of suppliers integrated within the firm boundary). While both margins are potentially influenced by a firm's productivity, we show that it is only the extensive margin where ownership allocations are governed by firm productivity in the same way as in the AH model. Thus, an important lesson from our analysis is that a carefully designed empirical model must use separate processes to model the two margins: one process determines *participation* in vertical integration, and a second process determines the consequent *extent* of vertical integration.²

We test the prediction regarding the productivity effect on ownership allocations using data on Spanish manufacturing firms taken from the Encuesta Sobre Estrategias Empresariales (ESEE). This is the only data set we know of that provides firm-level panel information on the choice between vertical integration and outsourcing for both, domestic and foreign input sourcing. Since our data

²This observation is consistent in spirit with an analysis of U.S. import data by Bernard et al. (2010). Although the observational unit in their empirical analysis is not the firm, but a pair of product and source country for U.S. imports, they carefully distinguish between the extensive margin (whether there are positive intra-firm imports for a given product-country pair) and the intensive margin (the share of intra-firm imports). Corcos et al. (2013) similarly distinguish between extensive and intensive margins of integration in their analysis of French firm-level data.

distinguish between 20 different industries based on the NACE-2009 classification, the firms in our sample operate in industries that differ markedly in their production technology, covering labor-intensive activities such as textiles production as well as capital-intensive activities such as metal and chemical manufacturing. Using this variation in our empirical analysis allows us to investigate heterogeneity in the effect of firm productivity across industries, and to see whether the empirical pattern of heterogeneity is consistent with the property rights theory of the firm.

Our empirical analysis provides strong support in favor of the property rights theory of the firm. We find that the effect of firm productivity on the likelihood of vertical integration does indeed exhibit marked differences across industries. The pattern of heterogeneity that we find is as predicted by the sourcing models in Antràs & Helpman (2004), Antràs & Chor (2013), and Schwarz & Südekum (2014). Our results thus suggest that contractual imperfections distort the sourcing of inputs in the global economy, and that firm boundaries emerge as a response to mitigate this distortion in the way predicted by property rights theory. We find that a higher productivity leads to a strong increase in vertical integration in headquarter-intensive industries, but has no or even a negative effect in supplier-intensive industries. In fact, as we move along the distribution of supplier intensities, the productivity effect becomes gradually less favorable to vertical integration. The strongest effect, found towards the bottom of the distribution of supplier intensities, implies that a doubling of productivity increases the probability of vertical integration by 20 to 30 percentage points. Strikingly, we find the exact same pattern of productivity effects across the two sourcing locations: the domestic and the foreign economy.

Our paper complements and extends a growing literature on firm boundaries in global sourcing. The seminal work by Antràs (2003) documents that capital-intensive goods and goods imported from capital-abundant countries are often traded within the boundaries of multinational firms, whereas labor-intensive goods and goods imported from labor-intensive countries are traded at arm's length. These patterns are consistent with a Helpman-Krugman model of international trade that takes a property rights view on firm boundaries. Consequently, the model in Antràs (2003) and its extensions in Antràs & Helpman (2004, 2008) have inspired various empirical investigations of the determinants of intra-firm trade at the product-, industry-, and country-level, such as for example different degrees of productivity dispersion across firms within industries (Yeaple, 2006; Nunn & Treffer, 2008, 2013), or different degrees of input contractibility (Nunn & Treffer, 2008; Bernard et al., 2010).³ While it is commonly understood that these studies provide broad empirical support for the property rights theory of the firm and the AH model, it should be noted that this interpretation crucially depends on (i) the fixed cost being larger for intra-firm sourcing than for outsourcing, and (ii) productivity being distributed Pareto within each industry; see Antràs & Helpman (2004). Our paper complements these studies by zooming in on firm-level variation in sourcing and productivity, and by providing an empirical test of the model that dispenses with this potentially restrictive twin assumption.

A further strand of the literature exploits firm-level variation in the data, in order to investigate the relationship between firm productivity and ownership allocations in input sourcing, as we do in this paper. This literature consistently finds firm productivity to play a significant role in shaping firm boundaries in the global economy, as suggested by the AH model. The prime example is a careful analysis of a cross-section of French firms by Corcos et al. (2013).⁴ They document that

³Fernandes & Tang (2012) test a variant of the AH model using Chinese export processing data. In their set-up vertical integration means a foreign firm acquiring ownership in a Chinese assembly plant. The first to study Chinese export processing trade from a property rights view are Feenstra & Hanson (2005).

⁴While the Spanish data we use in this paper are given at the firm-level, the French data are not only disaggregated by firm, but also by product type and source country. This allows the researchers to investigate firm-, country-, and product-level determinants of the choice between vertical integration and outsourcing.

more productive as well as more capital- and skill-intensive firms have a higher propensity to source inputs intra-firm rather than at arm's length.⁵ This finding is interpreted as strong support in favor of property rights theory as modelled in Antràs (2003) and Antràs & Helpman (2004). Defever & Toubal (2013) analyze the same French firm-level data source as Corcos et al. (2013).⁶ Interestingly, they find that it is the less productive rather than the more productive firms that are more strongly inclined towards vertical integration. They argue that this finding matches the predictions of the AH model under the assumption that outsourcing suffers from a fixed cost disadvantage relative to vertical integration.

While we do share the overall conclusions reached in this strand of the literature, our paper opens up a new angle for interpreting the empirical results obtained. More specifically, our analysis clarifies that a productivity effect favoring one or the other ownership allocation does not, as such, represent sufficient evidence in favor or against the AH model. Either outcome is perfectly consistent, not only with the AH model, but also with whatever fixed cost ranking prevails. If one wants to test the key mechanism of the AH model at the firm-level, what matters, instead, is how the productivity effect varies across industries with different relative input intensities. To the best of our knowledge, the empirical analysis we present in this paper is the first that allows answering this question.⁷

The structure of our paper is as follows. In the next section we adopt a firm-level view on the AH model and derive firm-level propositions amenable to empirical testing, first for the baseline case with a single supplier and then for the general case with a continuum of suppliers. In Section 3 we present the data set we use in the empirical analysis and describe salient features regarding firms' global sourcing decisions. Section 4 discusses our estimation strategy and presents the results. Section 5 concludes.

2 The property rights theory of global sourcing

2.1 Single non-contractible input supplier

Model assumptions. The economy produces a given number of differentiated goods and a standardized, numéraire good. Within any industry, firms (or headquarters) produce differentiated varieties

⁵Similar results regarding the effects of both productivity and capital intensity are found in firm-level data from Japan (Tomiura, 2007) and Italy (Federico, 2012). In Kohler & Smolka (2011, 2012), using the Spanish ESEE data, we show that *on average across industries* vertical integration firms are more productive than outsourcing firms. In Kohler & Smolka (2014), we demonstrate that this is due to firms self-selecting into sourcing strategies based on their productivity. A recent paper by Alfaro et al. (forthcoming) uses firm-level data for many countries to examine ownership allocations with a special focus on sequential production stages. They find broad empirical support for a generalized version of the model presented in Antràs & Chor (2013). See Del Prete & Rungi (2017) for empirical evidence suggesting an important role for the position of the headquarter firm along the value chain.

⁶There are at least three other interesting studies on intra-firm trade that use the French firm-level data. Carluccio & Fally (2012) study ownership allocations in relation to financial market imperfections, which in their model interact with contractual imperfections as familiar from property rights theory. Carluccio & Bas (2015) establish a link between intra-firm trade and labor market institutions in source countries. And Naghavi et al. (2015) focus on the issue of product complexity and intellectual property rights protection in source countries.

⁷Defever & Toubal (2013) propose an empirical model whose implications can be judged against our prediction regarding the productivity effect. In particular, they interact firm productivity with the (firm-specific) share of supplier inputs that are *relationship-specific* (and thus not easily contractible), and they find that productivity works more strongly in favor of outsourcing, and against vertical integration, when firms use these inputs more intensively in production. This finding exactly matches our prediction regarding the productivity effect if we consider variation in the share of supplier inputs that are *non-contractible*, as in Antràs & Helpman (2008), rather than variation in the overall supplier intensity, as we do in this paper.

of a final good using two types of inputs: inputs that are provided by headquarters themselves, and inputs that the headquarters cannot provide and must therefore be sourced from input suppliers. For want of better terms, the first type of input is called headquarter input and the second type is called the manufacturing component. Both inputs are essential in production of the final good, according to the following production function:

$$Q = \theta \left(\frac{H}{1-\zeta} \right)^{1-\zeta} \left(\frac{M}{\zeta} \right)^{\zeta}, \quad (1)$$

where Q denotes the quantity of final output, while H and M denote the headquarter input and the manufacturing component, respectively. We refer to $\zeta \in [0, 1]$ as the supplier intensity and to $1 - \zeta$ as the headquarter intensity of the production process, treated as industry-specific variables, while total factor productivity θ is a firm-specific variable. We use ℓ to denote the inverse unit cost of the manufacturing component. Without loss of generality, we normalize the unit cost for the headquarter input to unity. Due to product differentiation, each firm has price setting power which is modelled through a constant perceived price elasticity of demand, assumed the same for all firms and denoted by $\varepsilon > 1$.

The two agents cannot write enforceable contracts specifying the quantity and exact quality of H and M . This is justified by the complexity of the inputs in question, meaning that some of the relevant characteristics of the inputs cannot be fully described with finite contracting cost, and by a lack of third party verifiability. A further assumption is that due to customization the two inputs are relationship-specific. Once the headquarter and the supplier have incurred the cost of producing certain levels of H and M , using them according to (1) and selling the output Q generates revenue in excess of what the levels H and M are able to generate, in their entirety, through their respective outside options. Thus, there is an ex post quasi-rent from using the two inputs inside this relationship that the two parties may share, provided that they agree upon a sharing rule in *ex post* bargaining. Anticipation of this sharing rule dilutes incentives to provide inputs. In other words, investment in inputs H and M is subject to a *hold-up* problem.

This is where the property rights view of vertical integration comes in. Following Grossman & Hart (1986) and Hart & Moore (1990), the AH model defines vertical integration ($j = v$) as the headquarter acquiring residual property rights in the manufacturing input M , ensuring control of this input should ex post bargaining break down. A production relationship where the headquarter has no such property rights is referred to as outsourcing ($j = o$). In this case either party's outside option is zero. Residual ownership of M , in contrast, affords the headquarter a positive outside option and, thus, a larger revenue share in the ex post bargaining game. This improves the incentive to invest in H , but lowers the incentive for M . Writing m_j for the revenue share accruing to the headquarter under ownership allocation j , we have $m_v > m_o$. There is a large mass of potential input suppliers, each with a zero ex ante outside option.

As usual in this set-up, either ownership allocation may come with a specific fixed cost F_j . The fixed cost are an important element in the AH model, as they pin down the sorting of firms in industry equilibrium (along with other model parameters, in particular the supplier intensity of the industry as well as the dispersion of firm productivity). However, since we adopt a firm-level view in our analysis, it is not the industry equilibrium as such that we are interested in, and the fixed cost ranking will therefore be immaterial for our analysis. This will become evident below.

Set-up for decision making. Headquarters decide upon $j = v, o$, anticipating individually optimal input supplies depending on the revenue shares received by the headquarter and the input supplier, respectively. Under the assumptions made, the revenue generated by a production relationship is a concave function of output, $R(Q)$. Given ownership allocation j , input levels are

$H_j = \operatorname{argmax}_H \{m_j R(Q) - H\}$ and $M_j = \operatorname{argmax}_M \{(1 - m_j)R(Q) - M/\ell\}$, where Q is determined as in (1). The participation constraint for the input supplier involves a lump-sum transfer from the supplier to the headquarter equal to $(1 - m_j)R(Q_j) - M_j/\ell$, where Q_j is output corresponding to inputs H_j and M_j . Therefore, the headquarter's (maximum) profit associated with ownership allocation j is equal to $R(Q_j) - H_j - M_j/\ell - F_j$. In the following, we shall focus on the interaction between the industry's supplier intensity ζ and the firm's productivity level θ . We therefore write the headquarter's profit as $\Pi_j(\zeta, \Theta)$, where $\Theta := \theta^{\varepsilon-1}$. It is known from Antràs & Helpman (2004) that

$$\Pi_j(\zeta, \Theta) = Z_j(\zeta)\Theta - F_j, \quad (2)$$

where

$$Z_j(\zeta) := A \left(1 - \frac{\varepsilon - 1}{\varepsilon} (m_j(1 - \zeta) + (1 - m_j)\zeta) \right) \left(m_j^{1-\zeta} (\ell(1 - m_j))^\zeta \right)^{\varepsilon-1}. \quad (3)$$

In this definition, A captures the general equilibrium interrelationship between different sectors, and it stands for the equilibrium size of the industry in question; see Appendix A.

Existing literature based on the AH model focuses on industry-level predictions that follow from threshold, or cut-off levels of Θ implicitly determined by $\Pi_v(\zeta, \Theta) = \Pi_o(\zeta, \Theta)$.⁸ We extend this literature by exploring the attractiveness of vertical integration as given by

$$\Delta_m \Pi(\zeta, \Theta) := \Pi_v(\zeta, \Theta) - \Pi_o(\zeta, \Theta) = \Delta_m Z(\zeta)\Theta - \Delta_m F, \quad (4)$$

where $\Delta_m Z(\zeta) := Z_v(\zeta) - Z_o(\zeta)$ and $\Delta_m F := F_v - F_o$. The term $\Delta_m Z(\zeta)$, if positive, is a measure of the *strategic advantage of integration*. If it is negative, the strategic advantage lies with outsourcing. The term $\Delta_m F$, which can be positive or negative, measures the fixed cost disadvantage of integration.

The fact that the sign of $\Delta_m Z(\zeta)$ is ambiguous reflects a non-monotonic relationship between the function Z and the headquarter's revenue share m_j . Vertically integrating the supplier (and thus acquiring control rights in the input produced by the supplier) gives the headquarter a larger ex-post share of the production revenue, $m_v > m_o$. However, anticipating a lower ex-post revenue share for itself, the supplier will bring a lower quantity of its input to the production relationship, thereby reducing the overall revenue. Hence, efficiency considerations command vertical integration only if the supplier's input is not too important for the production relationship as a whole (i.e. if the industry's supplier intensity ζ is not too high). This is the central trade-off generated by the hold-up problem in both Antràs (2003) and Antràs & Helpman (2004).

The effect of firm productivity. The effect of firm productivity on the firm's incentive to integrate rather than outsource its input supplier is found by examining the responsiveness of $\Delta_m \Pi(\zeta, \Theta)$ with respect to changes in Θ : $\partial \Delta_m \Pi(\zeta, \Theta) / \partial \Theta = \Delta_m Z(\zeta)$.

Proposition 1 (productivity effect with a single non-contractible input supplier).

- (a) *The effect of firm productivity on $\Delta_m \Pi(\zeta, \Theta)$ is heterogeneous across industries and of ambiguous sign: In headquarter-intensive industries, $\zeta < \zeta^*$, a higher productivity increases $\Delta_m \Pi(\zeta, \Theta)$, thus favoring vertical integration, and conversely in supplier-intensive industries, $\zeta > \zeta^*$, where a higher productivity favors outsourcing. For the knife-edge case of $\zeta = \zeta^*$ the effect is zero.*
- (b) *The effect of firm productivity on $\Delta_m \Pi(\zeta, \Theta)$ is monotonically decreasing in $\zeta \in [\underline{\zeta}, \bar{\zeta}]$, being more favorable to outsourcing in more supplier-intensive industries, with $\underline{\zeta} < \zeta^* < \bar{\zeta}$. For $\zeta \in [0, 1]$, the effect is potentially non-monotonic in ζ .*

⁸This requires profits to be positive. If $\Pi_v(\zeta, \Theta) = \Pi_o(\zeta, \Theta) < 0$ then the corresponding productivity level is of course irrelevant.

- (c) If A is independent of ζ , symmetry in Nash bargaining as well as input costs implies a productivity effect on $\Delta_m \Pi(\zeta, \Theta)$ which is monotonically decreasing in $\zeta \in [0, 1]$. A strong enough asymmetry in input costs that favors production of the supplier input, $\ell > \ell^*$, implies a piecewise monotonic productivity effect on $\Delta_m \Pi(\zeta, \Theta)$ which is increasing in $\zeta \in [0, \underline{\zeta}]$, but decreasing in $\zeta \in [\underline{\zeta}, 1]$.

Proof. The productivity effect is found as $\partial \Delta_m \Pi(\zeta, \Theta) / \partial \Theta = \Delta_m Z(\zeta)$. Proposition 1 and Lemma 3 in Antràs (2003) imply that the ratio $Z_v(\zeta) / Z_o(\zeta)$ is monotonically decreasing in ζ , with a unique threshold ζ^* implicitly defined through $Z_v(\zeta) / Z_o(\zeta) = 1$. Hence, the difference $\Delta_m Z(\zeta)$ is strictly positive for $\zeta < \zeta^*$, strictly negative for $\zeta > \zeta^*$, and equal to zero for $\zeta = \zeta^*$. This proves part (a) of the proposition. Part (b) means that

$$\frac{\partial}{\partial \zeta} \Delta_m Z(\zeta) = \int_{m_o}^{m_v} \frac{\partial^2 Z(\zeta)}{\partial \zeta \partial m} dm$$

is always negative in an interval $[\underline{\zeta}, \bar{\zeta}]$ around ζ^* , but may be positive outside this interval. We prove this in Appendix B. Part (c) invokes symmetry in both Nash bargaining and input costs, hence $m_v > m_o = 1/2$ and $\ell = 1$. In Appendix B we show that in this case $\partial Z_o(\zeta) / \partial \zeta = 0$. Since $Z_v(\zeta) / Z_o(\zeta)$ decreases monotonically in ζ , $Z_v(\zeta)$ must do so too. This proves the monotonicity result. In Appendix B we also prove the piecewise monotonicity result for asymmetric input costs by evaluating the above integral in the relevant parameter subspace. \square

Figure 1. The effect of firm productivity on $\Delta_m \Pi(\zeta, \Theta)$




Figure 1 illustrates Proposition 1. For concreteness, we assume a fixed cost disadvantage of integration: $\Delta_m F > 0$. The solid lines depict the difference in profits between vertical integration and outsourcing, $\Delta_m \Pi(\zeta, \Theta)$, as a function of Θ . They show that the productivity effect can go either way, favoring outsourcing in supplier-intensive industries, and vertical integration in headquarter-intensive industries. The downward-sloping line is for a supplier-intensive industry, $\zeta_1 > \zeta^*$, where $\Delta_m Z(\zeta_1) < 0$

and thus the strategic advantage lies with outsourcing. An increase in productivity magnifies this advantage, as $\Theta'' > \Theta'$ implies $\Delta_m Z(\zeta_1)\Theta'' < \Delta_m Z(\zeta_1)\Theta'$. In headquarter-intensive industries, $\zeta_2 < \zeta^*$, the strategic advantage lies with integration, $\Delta_m Z(\zeta_2) > 0$, as depicted by the upward-sloping line. Again, this advantage is magnified by the firm's productivity: $\Delta_m Z(\zeta_2)\Theta'' > \Delta_m Z(\zeta_2)\Theta'$. As we have stressed above, this logic is entirely independent of the fixed cost ranking, because the fixed cost disadvantage of integration shows up as a level effect that favors outsourcing equally across firms regardless of productivity.

Proposition 1 is related to, but different from, the results on the productivity-dependent sorting of firms established in Antràs & Helpman (2004, 564-570). What matters for the sorting in industry equilibrium is the mere *sign* of $\Delta_m \Pi(\zeta, \Theta)$. As we can see in Figure 1, this means that in a supplier-intensive industry all firms choose outsourcing, whereas in a headquarter-intensive industry some firms (those with productivity below Θ') choose outsourcing, while others (those with productivity above Θ') choose integration. Proposition 1, in contrast, looks at the *magnitude* of $\Delta_m \Pi(\zeta, \Theta)$, and it is independent of the fixed cost, as these enter the expression for profits in an additively-separable way. This is crucial for our empirical analysis, where we allow for idiosyncratic stochastic disturbances at the firm-level, so that it is the magnitude of $\Delta_m \Pi(\zeta, \Theta)$ that matters; see Whinston (2003, 7-8).

The monotonicity result in part (b) of Proposition 1 is an intuitive implication of the property rights theory of the firm. It tells us that within the interval $\underline{\zeta} < \zeta^* < \bar{\zeta}$ the profit difference lines in Figure 1 rotate clockwise when we increase the supplier intensity of the industry, mitigating the productivity effect in headquarter-intensive industries with $\underline{\zeta} < \zeta < \zeta^*$, but reinforcing it in supplier-intensive industries with $\bar{\zeta} > \zeta > \zeta^*$. In other words, within the interval $(\underline{\zeta}, \bar{\zeta})$ the productivity effect is more favorable to outsourcing (and less so to vertical integration) in more supplier-intensive industries. What is perhaps surprising is the potential non-monotonicity mentioned in the second part of (b). As we discuss in detail in Appendix B, non-monotonicity arises in institutional setups where the values of m_v and m_o are both very low, so that in highly headquarter-intensive industries the headquarter incentives are suboptimally low under *either* ownership allocation (misalignment effect); similarly for high values of both m_v and m_o in highly supplier-intensive industries.⁹

If we rule out general equilibrium effects of ζ through A , then we can establish stronger results by imposing symmetry on Nash bargaining as well as input costs; see part (c) of Proposition 1. Under these conditions, we find a productivity effect that is monotonically decreasing in ζ over the entire interval $\zeta \in [0, 1]$. This means that the profit difference lines in Figure 1 will never rotate counter-clockwise when moving towards a more supplier-intensive industry. This is intuitive, because symmetry of the Nash bargaining game implies a sufficiently large $m_v > m_o = 1/2$, so that the misalignment effect responsible for non-monotonicity in part (b) does not arise. However, it may reappear if the supplier input becomes sufficiently cheap relative to the headquarter input (final statement in part (c) of the proposition). Details on this result are found in the proof of part (c) in Appendix B.

Adding multiple fully contractible inputs. Based on Antràs & Helpman (2008) it is relatively straightforward to see that Proposition 1 goes through when we maintain the basic two-input setup of Antràs & Helpman (2004), but allow the two inputs to be partially contractible rather than fully non-contractible. We now show that the proposition is also robust to including a set of multiple fully

⁹One might expect the potential non-monotonicity outside of $[\underline{\zeta}, \bar{\zeta}]$ to be but an artefact of the scaling of inputs by the Cobb-Douglas exponents in the production function; see Equation (1). This is however not the case, as we have verified. This analysis is available from the authors upon request.

contractible inputs. To see this, consider the following production technology:

$$Q = \theta \left[\left(\frac{H}{1-\zeta} \right)^{1-\zeta} \left(\frac{M}{\zeta} \right)^{\zeta} \right]^{1-\tau} \left[\left(\int_0^1 G(n)^{\delta} dn \right)^{1/\delta} \right]^{\tau}. \quad (5)$$

The expression in the first bracketed term is familiar from the AH model, and the expression in the second bracketed term is a standard CES production function with a continuum of *fully contractible* inputs indexed by $n \in (0, 1)$. The parameter $\delta \in (0, 1)$ governs the degree of substitutability between these inputs, and the upper tier production function is Cobb-Douglas, with output elasticities $\tau \in (0, 1)$ and $1 - \tau$, respectively. We assume that the contractible inputs must be sourced and cannot be produced by the headquarter, and we assume they are outsourced for reasons exogenous to the model. These assumptions accommodate the fact, observed in our data (as we shall see in the next section), that firms typically engage in vertical integration *in addition* to outsourcing.

There is now an additional stage in which the headquarter chooses quantities of the contractible inputs $G(n)$. This choice is made after the headquarter has determined the ownership allocation j for the production of the non-contractible input M , and before the headquarter offers a (non-enforceable) contract to potential suppliers of non-contractible inputs. The remaining stages of the game are familiar from the AH model, in particular the non-cooperative choice of quantities H and M by the headquarter and the supplier, respectively, and the Nash bargaining over the surplus generated from this production relationship in the final stage.

Solving this multi-stage game by backward induction, it is straightforward that the presence of contractual frictions distorts input provision not only for the non-contractible inputs, but also for the contractible ones. Interestingly, in contrast to Antràs & Helpman (2008), we observe *underinvestment* in the contractible inputs, not *overinvestment*. The reason is that suboptimal investment in the non-contractible inputs (a standard result of the hold-up problem) reduces the marginal benefits of investment in the contractible inputs (something the headquarter anticipates when investing into the contractible activities).¹⁰ In spite of this underinvestment, it is evident that the profit-maximizing mix of contractible inputs $G(n)$ is not affected by the lack of contractibility regarding inputs H and M . Moreover, as regards investment in the non-contractible inputs as well as the resulting ownership allocation, this extended model inherits all mechanisms relating to the productivity effect across different industries that we have identified above. In the interest of space, we relegate further details to Appendix C.

2.2 Multiple suppliers of non-contractible inputs

It is far less straightforward to see what happens if we allow for not just one, but many suppliers of *non-contractible* inputs. We now deal with two extensions of the AH model, both featuring multiple non-contractible inputs, but one assuming that these inputs are produced simultaneously, the other assuming sequentiality in input production. We demonstrate that the thrust of Proposition 1 is, by and large, robust to these model extensions.

Adding non-contractible inputs under simultaneous input production. Following Schwarz & Südekum (2014), we consider a continuum of input suppliers for non-contractible inputs appearing

¹⁰Differently from Antràs & Helpman (2008), our production function features separability of contractible activities from those that are non-contractible. Hence the result of underinvestment in both types of activities.

in a Cobb-Douglas production function as follows:

$$Q = \theta \left(\frac{H}{1-\zeta} \right)^{1-\zeta} \left(\int_0^1 \left(\frac{M(n)}{\zeta} \right)^\delta dn \right)^{\zeta/\delta}. \quad (6)$$

Relative to the AH model, the input M is replaced by a continuum of supplier inputs that are aggregated according to a standard CES production function. Thus, $M(n)$ is the input quantity produced by supplier n . Due to zero contractibility, there will be multilateral bargaining between the headquarter and all input suppliers, instead of bilateral bargaining with a single supplier. In analytical terms, this is a major change since Nash bargaining as a means to share the surplus no longer seems plausible. Therefore, Schwarz & Südekum (2014) resort to Shapley values as an alternative solution to the problem of revenue sharing. A novel insight generated by such a set-up is that a high degree of substitutability between different inputs (high value of $\delta \in (0, 1)$) ameliorates the hold-up distortion in that it increases the ex-ante incentives for both, the headquarter and the input suppliers, whatever the ownership allocation; see Antràs (2016, Chapter 4). The reason is that with a high elasticity of substitution between different non-contractible inputs, each supplier's bargaining threat is low and the headquarter's ex post revenue share is high. Furthermore, with a high elasticity of substitution the suppliers' chosen input levels are highly sensitive to their respective revenue shares.

The headquarter's choice variable is now the *extent* of vertical integration, i.e., the share of suppliers the firm wants to integrate into the boundaries of control. We denote this variable by $\xi \in [0, 1]$. To establish properties of the productivity effect that are similar in spirit to Proposition 1, we compare maximum profits achievable under *some* vertical integration, $\xi > 0$, with those achievable under *full* outsourcing, $\xi = 0$. Before we can do this, we must recognize an important change relative to the AH model deriving from the fixed cost of sourcing. These are now plausibly assumed to accrue *per supplier*. As a result, the total fixed cost borne by the firm depend on ξ . Because we have normalized the mass of suppliers to one, assuming symmetry we can write the total fixed cost as $F(\xi) = \xi F_v + (1 - \xi) F_o$.

Let $\bar{\Pi}(\xi, \zeta, \Theta) = \bar{Z}(\xi, \zeta) \Theta - F(\xi)$ denote profits as a function of the extent of vertical integration ξ , given the industry's supplier intensity ζ and the firm's productivity level Θ .¹¹ We then define $\xi_v \in (0, 1]$ as maximizing the profit difference between *some* vertical integration and *full* outsourcing, i.e., ξ_v maximizes $\bar{\Pi}(\xi, \zeta, \Theta) - \bar{\Pi}(0, \zeta, \Theta) = [\bar{Z}(\xi, \zeta) - \bar{Z}(0, \zeta)] \Theta - \xi(F_v - F_o)$. For an interior solution, $\xi_v \in (0, 1)$, we write $\xi_v = \xi_v(\zeta, \Theta)$, with $\xi_v(\zeta, \Theta)$ determined by

$$\frac{\partial \bar{Z}(\xi_v, \zeta)}{\partial \xi} \Theta = F_v - F_o. \quad (7)$$

Since we are comparing *some* integration with full outsourcing, we describe a case where profit maximization actually requires zero integration as $\xi_v \rightarrow 0$ with $[\bar{\Pi}(\xi_v, \zeta, \Theta) - \bar{\Pi}(0, \zeta, \Theta)] \rightarrow 0$. For a corner solution with $\xi_v = 1$, Equation (7) is replaced by the usual complementary slackness condition.

Finally, we denote our key object of interest, the maximum value function of the profit difference as

$$\Delta_\xi \bar{\Pi}(\zeta, \Theta) := \bar{\Pi}_v(\zeta, \Theta) - \bar{\Pi}_o(\zeta, \Theta) = \Delta_\xi \bar{Z}(\zeta) \Theta - \Delta_\xi F, \quad (8)$$

¹¹We use new functional symbols $\bar{\Pi}$ and \bar{Z} , respectively, to indicate that Schwarz & Südekum (2014) employ Shapley values instead of Nash bargaining. Consequently, the term $\bar{Z}(\xi, \zeta)$ is not directly comparable to $Z(\zeta)$ as defined in (3). Moreover, in their key results, Schwarz & Südekum (2014) do not focus on firm heterogeneity. However, the crucial parallel is that in their model profits appear as a product of $\bar{Z}(\xi, \zeta)$ and a market size variable. In our model, the market size variable A and firm productivity Θ are isomorphic in the expression for profits; see (2) and (3). This allows us to write $\bar{\Pi}(\xi, \zeta, \Theta)$ as above. Note, however, that $\bar{Z}(\xi, \zeta)$ still includes a market size variable which we continue referring to using the symbol A .

where $\bar{\Pi}_v(\zeta, \Theta) := \bar{\Pi}(\xi_v, \zeta, \Theta)$ and $\bar{\Pi}_o(\zeta, \Theta) := \bar{\Pi}(0, \zeta, \Theta)$. Moreover, in this expression, we have $\Delta_\xi \bar{Z}(\zeta) := \bar{Z}_v(\zeta) - \bar{Z}_o(\zeta)$ with $\bar{Z}_v(\zeta) := \bar{Z}(\xi_v, \zeta)$ and $\bar{Z}_o(\zeta) := \bar{Z}(0, \zeta)$, and $\Delta_\xi F := \xi_v(F_v - F_o)$.

The term $\Delta_\xi \bar{Z}(\zeta)$ must be thought of as the analogue to the strategic advantage of integration (or outsourcing) in the AH model. Importantly, however, while $\Delta_m Z(\zeta)$ in the AH model is independent of the firm's productivity, this is no longer the case for $\Delta_\xi \bar{Z}(\zeta)$, which depends on Θ through ξ_v . Using these definitions, we can compute the productivity effect as

$$\begin{aligned} \frac{\partial \Delta_\xi \bar{\Pi}(\zeta, \Theta)}{\partial \Theta} &= \frac{\partial \bar{Z}_v}{\partial \xi} \frac{\partial \xi_v}{\partial \Theta} \Theta + \Delta_\xi \bar{Z}(\zeta) - \frac{\partial \xi_v}{\partial \Theta} (F_v - F_o) \\ &= \Delta_\xi \bar{Z}(\zeta), \end{aligned} \tag{9}$$

where the second equality follows from the fact that by definition ξ_v maximizes $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$; see Equation (7).

Proposition 2 (productivity effect with multiple non-contractible input suppliers).

- (a) *The effect of firm productivity on $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ is heterogeneous across industries and of ambiguous sign: If the fixed cost of sourcing are sufficiently large, then in headquarter-intensive industries, $\zeta < \zeta^*$, a higher productivity (weakly) increases $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$, thus favoring some vertical integration over full outsourcing, and conversely in supplier-intensive industries, $\zeta \geq \zeta^*$, where a higher productivity (weakly) decreases $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$, thus favoring full outsourcing.*
- (b) *If A is independent of ζ and the fixed cost of sourcing are sufficiently large, then the effect of firm productivity on $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ is monotonically decreasing in ζ in the neighborhood of ζ^* , being more favorable to full outsourcing in more supplier-intensive industries. For $\zeta \in [0, 1]$, the effect is potentially non-monotonic in ζ .*
- (c) *The effect of firm productivity on the optimal extent of vertical integration $\xi_v(\zeta, \Theta)$ depends on the fixed cost ranking. For $F_v > F_o$, we have $\partial \xi_v / \partial \Theta \geq 0$, whereas for $F_v < F_o$ we have $\partial \xi_v / \partial \Theta \leq 0$. For the knife-edge case of $F_v = F_o$ the effect is zero.*

Proof. See Appendix D. □

The proposition rests on proposition 1.1 in Schwarz & Südekum (2014), which assumes no productivity differences across firms and no fixed cost of sourcing. In the present notation, Schwarz & Südekum (2014) show that there exists a threshold value $\zeta^* \in (0, 1)$ such that for any $\zeta < \zeta^*$ there exists a unique value $\xi > 0$ which maximizes $\bar{Z}(\xi, \zeta)$. We indicate this value by ξ_v^* . Moreover, for any $\zeta \geq \zeta^*$, we have $\xi_v^* = 0$. Therefore, strategic considerations dictate *some* vertical integration in headquarter-intensive industries ($\zeta < \zeta^*$), and *full* outsourcing in supplier-intensive industries ($\zeta \geq \zeta^*$). Proposition 2 generalizes these insights to a case with differences in both productivity and fixed cost, in order to establish how the productivity effect varies across industries with different supplier intensities. One important insight is that parts (a) and (b) of our proposition hold true regardless of the fixed cost ranking. This is crucial for our empirical analysis where the fixed cost of sourcing are unobservable.

To illustrate, we first consider a fixed cost disadvantage of vertical integration ($F_v > F_o$), as depicted in Figure 2. In this case, the productivity effect works in favor of some vertical integration in headquarter-intensive industries, but equals zero in supplier-intensive industries. This is illustrated by the upward-sloping and flat solid lines for $\zeta_2 < \zeta^*$ and $\zeta_1 \geq \zeta^*$, respectively. In headquarter-intensive industries strategic considerations would command a share of integrated suppliers equal to $\xi_v^* > 0$. However, the fixed cost make firms choose $\xi_v < \xi_v^*$, with more productive firms attaining

values of ξ_v closer to ξ_v^* , since the fixed cost weigh less heavily on them.¹² As a result, more productive firms enjoy a bigger profit advantage of vertical integration $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$, as illustrated by the upward-sloping solid line for ζ_2 . This is crucial for our result regarding the productivity effect in part (a) of the proposition. In the limit, as Θ approaches infinity, the optimal extent of vertical integration approaches ξ_v^* , depicted by the dashed line. In the opposite case of a supplier-intensive industry, $\zeta \geq \zeta^*$, the fixed cost disadvantage of vertical integration is reinforced by a strategic disadvantage, and we have $\xi_v \rightarrow 0$ and thus $\Delta_\xi \bar{\Pi}(\zeta, \Theta) \rightarrow 0$. Hence the solid horizontal line for ζ_1 throughout.

Figure 2. The effect of firm productivity on $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ with $F_v > F_o$





Figure 3 illustrates the alternative case where vertical integration features lower fixed cost than outsourcing ($F_v < F_o$). As before, in headquarter-intensive industries ($\zeta < \zeta^*$) strategic considerations would command a strictly positive optimal share of integrated suppliers equal to $\xi_v^* > 0$. However, now vertical integration is also more attractive in terms of the fixed cost. Hence, firms will choose $\xi_v > \xi_v^*$, but ξ_v will be lower, and thus closer to ξ_v^* , for firms with a higher productivity. Note that deviating from ξ_v^* is costly; from a strategic perspective it means “overdoing” integration. Firms with a sufficiently low productivity may be prompted to exploit the fixed cost advantage of vertical integration to an extent where $\Delta_\xi \bar{Z}(\zeta)$ turns negative: on strategic grounds they would even be better off under full outsourcing, but this is overcompensated by the large weight that their low productivity puts on the fixed cost advantage. In Figure 3 this case shows up as a downward-sloping solid line for ζ_2 and close to zero levels of Θ . In Proposition 2 this case is ruled out by requiring sufficiently high levels of fixed cost that prevent such low-productivity firms from entering production in the first place. For firms entering into production, the profit advantage of vertical integration $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ is thus increasing in productivity, as illustrated by the upward-sloping solid line. By logic familiar from

¹²For a high enough fixed cost disadvantage of integration, firms with low enough productivity are driven into the corner solution with $\xi_v \rightarrow 0$. For these firms, we have $\Delta_\xi \bar{\Pi}(\zeta, \Theta) \rightarrow 0$, as illustrated by the flat segment of the profit difference line for ζ_2 in the figure.

Figure 2, the optimal degree of vertical integration again approaches ξ_v^* if we let Θ go to infinity. In supplier-intensive industries ($\zeta \geq \zeta^*$) the strategic advantage lies with outsourcing, whence in this case the $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ -line is sloping downwards and converging to $\xi_v^* = 0$ for $\Theta \rightarrow \infty$.

In terms of Figures 2 and 3, part (b) of Proposition 2 describes how the slopes of the $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ -lines respond to marginal changes in the supplier intensity ζ . By analogy to the single-supplier case in Figure 1, an increase in the supplier intensity around ζ^* makes the convex parts of these lines rotate in a clockwise fashion. The bottom line is therefore that the productivity effect works more strongly in favor of full outsourcing in more supplier-intensive industries.

Figure 3. The effect of firm productivity on $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ with $F_v < F_o$



Adding non-contractible inputs under sequential input production. If technology imposes a sequencing of input production, then revenue sharing may again reasonably be modelled through (sequential) Nash bargaining. Antràs & Chor (2013) provide an in-depth analysis of this case. Following their analysis, assume the following production technology:

$$Q = \theta \left(\frac{H}{1 - \zeta} \right)^{1 - \zeta} \left(\int_0^1 \left(\frac{M(n)}{\zeta} \right)^\delta I(n) dn \right)^{\zeta/\delta}, \quad (10)$$

where production stages are indexed by n , with smaller index numbers referring to production stages further upstream in the production process. The function $I(n)$ is an indicator function that takes on the value one if input n is produced after all inputs from earlier production stages $n' < n$ have been produced, and zero otherwise. The headquarter input is produced after input suppliers have been hired, but before any input production takes place. Importantly, and differently from Schwarz & Südekum (2014), the headquarter firm engages, sequentially, in bilateral Nash bargaining over each supplier's incremental contribution to total revenue.

Antràs & Chor (2013) show that in such a model the ownership allocation along the value chain crucially depends on the price elasticity of demand $\varepsilon = 1/(1 - \alpha)$ and the technology parameters δ and ζ . In particular, maintaining the assumption of ex-ante transfers between the headquarter and suppliers, and abstracting from any fixed cost differences between vertical integration and outsourcing, it turns out that when $\zeta\alpha > \delta$ the headquarter will outsource *all* stages of production, whereas when

$\zeta\alpha < \delta$ it will always integrate *some* stages of production. Specifically, with $\zeta\alpha < \delta$ the entire sequence of stages falls into two distinct subranges, with stages in the upstream range being integrated and stages in the downstream range being outsourced. The intuition is that, other things equal, with a relatively low elasticity of demand (low α), coupled with a high degree of substitutability across stages (high δ), exploitation of market power requires that production of upstream stages be “curbed” through vertical integration.

This result and its underlying set-up in Antràs & Chor (2013) now allow us to reproduce our analysis of the case of simultaneous input production. Using notation familiar from above, we define $\Delta_\xi \widehat{\Pi}(\zeta, \Theta) := \widehat{\Pi}_v(\zeta, \Theta) - \widehat{\Pi}_o(\zeta, \Theta) = \Delta_\xi \widehat{Z}(\zeta)\Theta - \Delta_\xi F$, where a caret indicates the case of sequential as opposed to simultaneous input production indicated by a bar above. The variable ξ_v^* now indicates a cutoff-value of n separating upstream production stages, $n < \xi_v^*$, where strategic considerations dictate integration from stages $n > \xi_v^*$ governed by outsourcing, absent any fixed cost consideration. The variable ξ_v has a corresponding interpretation, taking fixed cost considerations into account. With this reinterpretation it is straightforward that the exact same results of Proposition 2 also apply to multiple input suppliers with sequential production. Part (a) of the proposition can now be written as: $\Delta_\xi \widehat{Z}(\zeta)\Theta \geq 0$ for $\zeta < \zeta^*$ and $\Delta_\xi \widehat{Z}(\zeta)\Theta \leq 0$ for $\zeta \geq \zeta^*$ with $\zeta^* = \delta/\alpha$, and accordingly for parts (b) and (c).

Before we proceed to the empirical part, we briefly pause to summarize our results for multiple non-contractible input suppliers as follows: (1) More productive firms integrate a larger share of their suppliers, provided that integration features a fixed cost disadvantage relative to outsourcing, and conversely if the fixed cost disadvantage lies with outsourcing. (2) Surprisingly, however, if we examine the effect of productivity on the difference in maximum profits between *some* integration and *full* outsourcing, then we observe a striking similarity to the result of the simple AH model: The productivity effect is heterogeneous across industries and works in favor of *some* integration in headquarter-intensive industries, and in favor of *full* outsourcing in supplier-intensive industries. Importantly, the full generality of this result only requires a certain minimum *absolute* level of fixed cost; it does not require any assumption on the *relative ranking* of fixed cost between outsourcing and integration. (3) Around the cut-off value ζ^* , the productivity effect works more strongly in favor of full outsourcing, and against integration, in more supplier-intensive industries.

3 Data

3.1 Firm-level survey

To test the theoretical predictions empirically, we use data from the “Encuesta Sobre Estrategias Empresariales” (ESEE). This is an annual firm-level survey conducted since 1990 by the “Sociedad Estatal de Participaciones Industriales” (SEPI), a public foundation based in Madrid. The ESEE data have a panel structure and include about 1,900 firms each year (all of which are active in the manufacturing sector in Spain).

Unlike most other firm-level data sets we are aware of, the ESEE data contain unique information on firms’ sourcing behavior in both the domestic and the foreign economy. This information was first included in 2006, which is why we focus our empirical analysis on the years 2006-2015. The two key questions in the survey are:¹³

¹³The survey questionnaire is distributed in Spanish and available for download at <http://www.fundacionsepi.es/esee/sp/svariables/indice.asp>.

1. *Of the total amount of purchases of goods and services that you incorporate (transform) in the production process, indicate—according to the type of supplier—the percentage that these represent in the total amount of purchases of your firm in [year].*
 - (a) *Spanish suppliers that belong to your group of companies or that participate in your firm’s joint capital.* [yes/no] / [if yes, then percentage rate]
 - (b) *Other suppliers located in Spain.* [yes/no]/[if yes, then percentage rate]
2. *For the year [year], indicate whether you imported goods and services that you incorporate (transform) in the production process, and the percentage that these imports—according to the type of supplier—represent in the total value of your imports.* [yes/no]
 - (a) *From suppliers that belong to your group of companies and/or from foreign firms that participate in your firm’s joint capital.* [yes/no]/[if yes, then percentage rate]
 - (b) *From other foreign firms.* [yes/no]/[if yes, then percentage rate]

This information allows us to construct a variable measuring the firm’s share of inputs sourced through vertically integrated suppliers as $\xi = (\text{Intra-firm sourcing})/(\text{Intra-firm sourcing} + \text{Outsourcing})$.¹⁴ We do this separately for domestic sourcing and for foreign sourcing using answers to questions 1 and 2, respectively.

The sampling design of the ESEE data implies an oversampling of large firms relative to small firms. Specifically, the initial selection of firms in 1990, the first year the survey was taken, followed a two-tier sampling scheme. Survey questionnaires were sent out to all firms with more than 200 workers (large firms), and to just a subset of firms with 10 to 200 workers (small firms). Small firms were selected through a stratified, proportional, and systematic sampling with a random seed, and the stratification was based on industry affiliation and firm size group.¹⁵

In our empirical analysis we focus on firms that have been in the sample since 2006, as this is the first year the survey includes information on sourcing, and we apply sampling weights in order to account for the specific sampling design of the ESEE data, in particular the stratification and oversampling of large firms. More specifically, we weight each firm by its inverse probability of being in the ESEE sample in 2006 using information on the total number of manufacturing firms in Spain (by industry-size-group cluster) from the Spanish Instituto Nacional de Estadística (INE).¹⁶

3.2 Firm-specific productivity

A pivotal variable in our empirical analysis is a firm’s productivity level Θ . Following Delgado et al. (2002), who analyze the same Spanish firm-level data, we construct an index measure of total factor

¹⁴This is a slight abuse of notation, as this variable is based on input *values* and is thus not identical to the variable ξ in Section 2.2, which measures the *number* of integrated suppliers as a share of the total number of suppliers. We ignore this difference here, as it will be inconsequential for our empirical analysis which focuses exclusively on the extensive margin of vertical integration.

¹⁵The survey distinguishes 20 different industries and six different size groups in terms of the number of workers employed (10-20; 21-50; 51-100; 101-200; 201-500; >500). Industries are defined according to sets of products at the NACE-2009 level; see Table E.1 in Appendix E. Prior to 2009, the ESEE data used the older NACE-1993 classification. We use concordance information provided by the SEPI foundation in order to account for this; see https://www.fundacionsepi.es/esee/en/evARIABLES/i_Cambio_clasificacion_sectorial_ESEE.asp. More information on the survey and its sampling properties are available in English from SEPI’s website at <http://www.fundacionsepi.es/esee/en/epresentacion.asp>.

¹⁶This information is available at <http://www.ine.es/dynt3/inebase/index.htm?padre=51>.

productivity (TFP).¹⁷ The index is constructed as the log of the firm’s output minus a cost-share weighted sum of the log of the firm’s inputs. In order to guarantee that comparisons between any two firm-year observations are transitive, each firm’s inputs and outputs must be expressed as deviations from some reference point. In contrast to the multilateral index used in Caves et al. (1982), which uses as a single reference point a hypothetical firm with input cost shares that equal the arithmetic mean cost shares over all observations and output and input levels that equal the geometric mean of output and the inputs over all observations, we use a separate hypothetical-firm reference point for each firm size group and then chain-link the reference points together. This is a useful extension in our context as it takes into account the sampling properties of the data set, in particular the two-tier sampling scheme that distinguishes between large and small firms and the stratification of the data based on industry affiliation and firm size group. The main advantage of the index approach is that we can directly compute TFP without estimating parameters of an underlying production function.

Let each firm i at time t produce output Q_{it} using the set of inputs G_{nit} (with input cost shares given by W_{nit}) where $n = 1, 2, \dots, N$. Details about the measurement of these variables are relegated to Appendix F. The TFP index for firm i in industry s and size group g in year t is defined as:

$$\begin{aligned} \ln \theta_{it} = & (\ln Q_{it} - \overline{\ln Q}_{sg}) - \sum_n \frac{1}{2} (W_{nit} + \overline{W}_{nsg}) (\ln G_{nit} - \overline{\ln G}_{nsg}) \\ & + (\overline{\ln Q}_{sg} - \overline{\ln Q}_s) - \sum_n \frac{1}{2} (\overline{W}_{nsg} + \overline{W}_{ns}) (\overline{\ln G}_{nsg} - \overline{\ln G}_{ns}), \end{aligned} \quad (11)$$

where bars indicate mean values of the respective variables: $\overline{\ln Q}_{sg}$, for example, is the geometric mean of output over all firms in industry s and size group g (across all years), and $\overline{\ln Q}_s$ is the geometric mean of output over all firms in industry s (across all size groups and years). Notice that the hypothetical-firm reference point is industry-specific. Hence, when we pool observations of different industries, productivity differences across industries cancel out.

Equation (11) measures the proportional TFP difference for firm i at time t relative to a given hypothetical-firm reference point. To gain some intuition, consider the terms on the right-hand side. The first line makes a within-size-group comparison, measuring productivity differences between firm i at time t and a hypothetical firm that represents the “average” firm in size group g . The second line makes a comparison across size groups, comparing the productivity of the hypothetical firm in size group g with the productivity of the “average” firm in the industry. Therefore, comparisons between any two observations are always transitive, no matter whether these observations are drawn from the same size group or not.

3.3 Industry-specific supplier intensity

The key variable at the industry-level is the supplier intensity of production ζ . This variable is not directly observed. On a very fundamental level, ζ reflects the extent to which the input suppliers are bound to bear the cost of production, and $1 - \zeta$ reflects the cost share borne by the headquarter firms. What determines the extent of cost sharing between input suppliers and headquarter firms? Antràs (2003) argues that the cost of physical capital are easier to share than the cost of labor inputs, and that headquarter firms primarily provide (or pre-finance) machinery and specialized tools and equipment, or assist their suppliers in the acquisition of capital equipment and raw materials (as reported in


¹⁷See Good et al. (1997) for an excellent overview of index number approaches to productivity measurement.

Dunning (1993, 455-456)).¹⁸ Therefore, in the logic of property rights theory, the headquarter input carries more weight in the production of capital-intensive goods than in the production of labor-intensive goods. It has thus become common practice in empirical work to proxy the headquarter intensity of production, $1 - \zeta$, by an industry-specific measure of capital intensity; see Antràs (2003), Yeaple (2006), and Nunn & Treffer (2008). We proceed similarly in our empirical analysis, and use the (reversed) scale of industry-specific capital intensities to represent ζ . If z_s is the capital intensity of industry s in the raw data, we have $\zeta_s = (z_{\max} - z_s)/(z_{\max} - z_{\min})$, where min and max indicate the minimum and maximum values found in the data. For ease of interpretation, ζ_s is thus normalized to the unit interval $[0, 1]$. The industry-specific capital intensity is given by the average capital intensity we observe in the industry across the years 2006-2015.¹⁹ The industries Leather & Footwear and Textile & Wearing Apparel plausibly emerge as the most supplier-intensive industries. The industries Beverages and Ferrous Metals & Non-Ferrous Metals, in contrast, are the most headquarter-intensive industries.

3.4 Basic facts

Figure 4 displays the evolution of firms engaged in different sourcing strategies. In this figure, we define sourcing strategies in a mutually inclusive way, so that a firm counts for more than one sourcing strategy if it reports multiple ways of sourcing. We distinguish between domestic integration, domestic outsourcing, foreign integration, and foreign outsourcing (corresponding to questions 1(a), 1(b), 2(a), and (2b) in Section 3.1 above, respectively). The figure shows pronounced differences in the share of firms choosing a particular sourcing strategy. It also shows that these fractions remain roughly constant over time. Domestic outsourcing is almost universally used (more than 90% of firms), followed by foreign outsourcing (30%), domestic integration (7-8%), and foreign integration (3-4%). Thus, as far as the relative importance of sourcing strategies is concerned, we find a pattern similar to the ones observed for other industrialized countries such as Japan (Tomiura, 2007) and Italy (Federico, 2010, 2012).

Figure 4. Sourcing strategies of Spanish manufacturing firms (2006-2015)



¹⁸Other references consistent with this idea and discussed in Antràs (2003) are Milgrom & Roberts (1993), Aoki (1990, 25), and Young et al. (1985).

¹⁹The capital intensity is defined as the value sum of real estate, construction and equipment (measured in prices of 2006) over the average number of workers during the year.

Figure 5 shows that the use of sourcing strategies strongly depends on firm size. The figure displays the fractions of firms engaged in different sourcing strategies in 2015, the most recent year of our sample. It does so separately for the six different firm size groups, with relative frequencies indicated by bars. As firms employing 10 to 50 workers represent more than 80% of all firms in Spanish manufacturing, their use of sourcing strategies closely resembles the picture displayed in Figure 4. However, larger firms show a markedly stronger engagement in both foreign sourcing and vertical integration. In particular, strategies of foreign integration are used by almost 25% of the very large firms (those with more than 500 employees), while for foreign outsourcing and domestic integration the numbers are even higher, at more than 70% and 50%, respectively.

Figure 5. Sourcing strategies by firm size in 2015


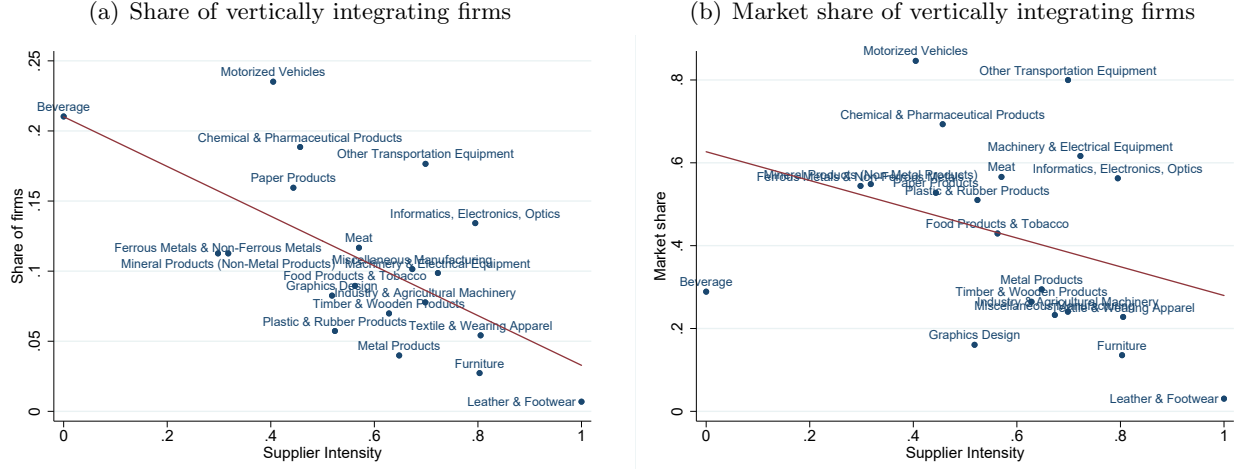


Figure 6(a) illustrates that, consistent with Antràs & Helpman (2004), firms using vertically integrated production relationships are strongly concentrated in headquarter-intensive industries (those producing beverages, certain metal products, mineral products, motorized vehicles, etc.). Outsourcing relationships, in contrast, are considerably more important, relative to vertical integration, in supplier-intensive industries (those producing leather & footwear, textiles, and furniture).²⁰ The between-industry differences we observe in the data are significant. In the Leather & Footwear industry, the most supplier-intensive industry, the share of firms sourcing inputs intra-firm is less than 3%. For Beverages, the most headquarter-intensive industry, this share is more than 21%. In Figure 6(b), we conduct a similar exercise as in 6(a), but we now look at the market share of those firms relying on vertically integrated production relationships. As predicted by Antràs & Helpman (2004), there is a negative relationship between this market share and the supplier intensity of the industry. Overall, our findings for Spain resemble those for the U.S. For instance, figure 1.7 in Antràs (2016) documents that the share of related-party imports in total U.S. imports is the higher, the higher the capital intensity of the industry.

²⁰We construct the share of vertically integrating firms for each year from 2006 to 2015 and then average across all years.

Figure 6. Prevalence of vertical integration by industry (2006-2015)



4 Empirical analysis

4.1 Main analysis

Empirical model. Our key objective is a firm-level model that is able to capture the basic trade-off between vertical integration and outsourcing as postulated by the property rights theory of the firm. We now propose such a model and apply it to our data. In the following we index firms by i , industries by s , and time by t . The choice we wish to explain is whether a firm engages in *some* vertical integration ($\xi_{it} > 0$) as opposed to *full* outsourcing ($\xi_{it} = 0$). One of the crucial features of our data set is that this choice is observed, not only for the foreign economy, but also for the domestic economy. This allows us to test the predictions of the property rights theory of the firm, as summarized in Propositions 1 and 2, for sourcing in both the foreign and the domestic economy. The theory posits that the share ξ_{it} is a function of the firm's productivity, Θ_{it} , as well as the supplier intensity of production, ζ_s . In the simple AH model with a single input supplier, ξ_{it} takes on but two values: zero (for outsourcing) and one (for vertical integration). However, as we have seen above, values of ξ_{it} lying strictly in the interior of $(0, 1)$ are common in our data (especially for domestic sourcing).

This requires a clear conceptual distinction between the extensive and the intensive margin of vertical integration. The extensive margin refers to the decision whether *some* inputs should be sourced through vertical integration or not ($\xi_{it} > 0$ vs. $\xi_{it} = 0$). The intensive margin, in contrast, relates to *how much* should be sourced through vertical integration, conditional on $\xi_{it} > 0$. Importantly, as we have shown in our theoretical analysis, the two margins are not governed by the exact same mechanism and thus involve different predictions that rest on different assumptions. One needs a flexible econometric framework using separate processes to model the two margins: one process determines *participation* in vertical integration and a second process explains consequent *extent* of vertical integration. For this reason, we focus our empirical analysis on the extensive margin of vertical integration (the participation decision), where we do not need to impose any assumption on the unobservable fixed cost ranking between vertical integration and outsourcing (as demonstrated by Propositions 1 and 2).

Define a binary indicator variable $\mathbb{1}(\xi_{it} > 0)$ equal to one for observations with strictly positive values of vertical integration ($\xi_{it} > 0$), and equal to zero for observations where all inputs are sourced

from independent suppliers ($\xi_{it} = 0$).²¹ Denote the set of firms belonging to industry s by \mathcal{I}_s . We then adopt the following latent variable formulation:

$$\mathbb{1}(\xi_{it} > 0) = \begin{cases} 1 & \text{if } \Delta_m \Pi(\zeta_s, \Theta_{it}) + \Delta_m \pi_{it} > 0, \\ 0 & \text{otherwise.} \end{cases}, \quad i \in \mathcal{I}_s, \quad (12)$$

where $\Delta_m \Pi(\zeta_s, \Theta_{it})$ is the *systematic* component of the profit difference (as described in detail in Section 2.1), and where we add a *random* component $\Delta_m \pi_{it} := \pi_{v,it} - \pi_{o,it}$. This random component summarizes all effects unrelated to the economic mechanism highlighted by the property rights theory of the firm. According to (12), crossing the zero threshold for the sum of both the systematic and the random component leads to participation in vertical integration.

For concreteness, we use the profit difference of the single-supplier case, $\Delta_m \Pi(\zeta_s, \Theta_{it})$, in our latent variable formulation in (12). A key insight from our theoretical analysis is that the mechanism and the variables governing the participation decision are the same no matter whether we allow for multiple suppliers or not. This means that, without qualification, we can replace $\Delta_m \Pi(\zeta_s, \Theta_{it})$ in (12) by $\Delta_\xi \bar{\Pi}(\zeta_s, \Theta_{it})$ or by $\Delta_\xi \hat{\Pi}(\zeta_s, \Theta_{it})$ as introduced in Section 2.2 above. What matters here is that a higher productivity (weakly) increases the systematic component of the profit difference if $\zeta < \zeta^*$, and (weakly) decreases it if $\zeta \geq \zeta^*$, and this holds true in both the single-supplier case and the multiple-supplier case. Note the great generality of our approach: we do not have to take any stance on whether production is sequential or simultaneous in nature, nor do we have to impose any assumption on the sign or magnitude of the fixed cost difference between integration and outsourcing.

In order to allow for the productivity effect to vary along the distribution of supplier intensities, we interact firm productivity with the industry's supplier intensity in our regression framework. The model we estimate then reads as follows:

$$\begin{aligned} \Pr(\mathbb{1}(\xi_{it} > 0) | \cdot) &= \Pr(\Delta_m \Pi(\zeta_s, \Theta_{it}) + \Delta_m \pi_{it} > 0) \\ &= \Pr(\Delta_m Z(\zeta_s) \Theta_{it} - \Delta_m F_{st} + \Delta_m \pi_{it} > 0) \\ &= \Pr(\lambda_0 \cdot \Theta_{it} + \lambda_1 \cdot \Theta_{it} \times \zeta_s + \gamma_s + \gamma_t > -\Delta_m \pi_{it}), \quad i \in \mathcal{I}_s, \end{aligned} \quad (13)$$

where λ_0 and λ_1 are the parameters of interest, and γ_s and γ_t are industry and year fixed effects, respectively. These fixed effects absorb the main (i.e., non-interaction) effect of ζ_s , as well as the level effect from the fixed cost difference $\Delta_m F_{st}$, which enter γ_s and γ_t in the final line of (13). Since $\Delta_m F_{st}$ is a constant term that enters the model in an additive way, we place no restriction on how it varies across industries. In fact, both the sign and the magnitude may be industry-specific. Moreover, $\Delta_m F_{st}$ is allowed to change over time as long as the change occurs uniformly across industries.

As far as the fixed cost are concerned, the multiple-supplier case requires slightly more restrictive assumptions. Here, the fixed cost difference appears as $\Delta_\xi F = \xi_{it}(F_v - F_o)$ instead of just $\Delta_m F_{st} = (F_v - F_o)_{st}$, where $\Delta_\xi F$ is expected to depend on Θ_{it} and ζ_s through ξ_{it} . Hence, the difference $\Delta_\xi F$ is now part of the function $\Delta_\xi \Pi(\zeta_s, \Theta_{it})$ whose shape we evaluate by including Θ_{it} and its interaction with ζ_s in the model. While there is still no need to impose any restriction on the sign or magnitude of the difference $F_v - F_o$, we must assume that it does not vary across industries or through time and therefore write $\Delta_\xi F$ without subscripts s and t .

The probability of vertical integration depends on the distribution assumed for the random component $\Delta_m \pi_{it}$. We proceed with two alternative assumptions. First, we assume that this term is uniformly distributed between two values $-L$ and L , with $L > 0$. This gives rise to the *linear probability model* (LPM) with $\Pr(\mathbb{1}(\xi_{it} > 0) | \cdot) = \frac{\Delta_m \Pi(\zeta_s, \Theta_{it}) + L}{2L}$ for $-L \leq \Delta_m \Pi(\zeta_s, \Theta_{it}) \leq L$. For

²¹We report summary statistics of all variables used in our empirical analysis in Table E.2 in Appendix E.

$\Delta_m \Pi(\zeta_s, \Theta_{it}) > L$ the choice probability of vertical integration is one, and for $\Delta_m \Pi(\zeta_s, \Theta_{it}) < -L$ it drops to zero. Secondly, and alternatively, we assume that $\Delta_m \pi_{it}$ is normally distributed with mean zero and variance σ^2 . This gives rise to the *Probit model* with $\Pr(\mathbb{1}(\xi_{it} > 0) | \cdot) = \Phi(\Delta_m \Pi(\zeta_s, \Theta_{it})/\sigma)$, where $\Phi(\cdot)$ denotes the standardized cumulative normal distribution. Under the given set of assumptions, we can pool the data across all years to obtain consistent parameter estimates.

In terms of the property rights theory of the firm, the single most important parameter in the model in (13) is λ_1 . This parameter indicates how the productivity effect responds to changes in the industry's supplier intensity (all else equal). We are also interested in the sign of λ_0 and the magnitude of λ_0/λ_1 . Against the backdrop of Propositions 1 and 2, we expect a positive estimate of λ_0 and a negative estimate of λ_1 . This is clear from Figures 1 to 3. Estimates of λ_0/λ_1 indicate how strongly the profit difference lines in the figures rotate clockwise when we gradually increase the supplier intensity. For $\lambda_0 > 0$ and $\lambda_1 < 0$, a ratio of λ_0/λ_1 strictly lower than one (in absolute value) would imply that in sufficiently supplier-intensive industries a higher productivity actually discourages vertical integration and increases the probability of full outsourcing. This requires a downward-sloping profit difference line, as depicted in Figures 1 and 3. A ratio equal to or larger than one (in absolute value) would mean that the profit difference line is never downward-sloping in our data.²²

In the following we do not focus on the magnitudes of the estimated coefficients of Θ_{it} and $\Theta_{it} \times \zeta_s$, as these identify λ_0 and λ_1 only up to a scale parameter (chosen at will) in the underlying distribution function for $\Delta_m \pi_{it}$. We will therefore abstain from commenting on the absolute values of these coefficients altogether. The scale parameter cancels out when we look at λ_0/λ_1 instead, so that this ratio is identified.

A final remark on the model in (13) concerns linearity and monotonicity of the productivity effect. Since we model variation in the productivity effect across industries through a simple interaction term between Θ_{it} and ζ_s , we are forcing the data into a rigid parametric relationship. By construction, the productivity effect depends linearly, and thus monotonically, on ζ . We see this model as a first-order approximation of the true relationship, and provide estimates of a more flexible model that allows for both non-linearity and non-monotonicity later on. In fact, our theoretical analysis suggests a non-linear and potentially non-monotonic response of the productivity effect to changes in the supplier intensity. Our estimates will reveal that the key message of our empirical analysis based on the simple model in (13) remains entirely unaffected.

Results. Table 1 presents model estimates for two different panels, Panel A looking at sourcing in the domestic economy and Panel B looking at sourcing in the foreign economy. For each panel we estimate an LPM as well as a Probit model. For either model, we pool the data across all years and firms. All specifications include year fixed effects as well as industry fixed effects capturing the influence of time-invariant industry-specific variables such as skill intensity or R&D intensity. Statistical inference is based on robust standard errors clustered by firm.

Consistently for both panels, we find that the coefficients of Θ_{it} and $\Theta_{it} \times \zeta_s$ are estimated with a positive and a negative sign, respectively, and they are always significantly different from zero, $\hat{\lambda}_0 > 0$ and $\hat{\lambda}_1 < 0$. The estimates of λ_0/λ_1 lie between -0.77 and -1.13 and are quite stable, not only across estimators, but also across sourcing locations. Importantly, a formal pairwise comparison across any of the two estimators and sourcing locations reveals that in no case we can reject the null hypothesis that the corresponding ratios are identical. This suggests that the choice of estimator (LPM vs.

²²This could be the case because all industries in our data are headquarter-intensive ($\zeta < \zeta^*$) or because the fixed cost ranking favors outsourcing over integration ($F_v > F_o$), as assumed in Figure 2. Without additional assumptions, we are not able to discriminate between these two possibilities.

Table 1. Main estimation results—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	0.361*** (0.082)	1.405*** (0.385)
TFP \times Supplier intensity	-0.441*** (0.124)	-1.239* (0.673)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.819*** (0.062)	-1.134 (0.339)
N	12205	12205
R^2 or Pseudo R^2	0.075	0.109
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	0.433*** (0.115)	2.463*** (0.682)
TFP \times Supplier intensity	-0.553*** (0.176)	-3.214*** (1.160)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.782*** (0.079)	-0.766** (0.093)
N	5896	5754
R^2 or Pseudo R^2	0.079	0.107

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

Probit) is not a crucial one in our application. Overall, the results thus confirm the twin predictions that the productivity effect works in favor of vertical integration for a low supplier intensity, and that this effect becomes watered down as the supplier intensity increases. The considerable variation that we find across industries supports the predictions of property rights theory, as summarized in Propositions 1 and 2. Consistently with our propositions, estimates of $|\lambda_0/\lambda_1|$ significantly below one (as found in Panel B for the foreign economy) indicate that the productivity effect might actually work in favor of full outsourcing, and against integration, for a sufficiently high supplier intensity. The estimates for the domestic economy are not fully conclusive in this regard, as the LPM suggests a ratio of $|\lambda_0/\lambda_1|$ significantly below one, while the Probit estimates do not allow us to reject that the ratio is equal to one.

4.2 Robustness analysis

We now investigate the robustness of our baseline estimates with respect to important modifications of our empirical model and changes in the estimation sample. We describe each step of this analysis in turn and briefly summarize the main results obtained. We always report the results from both the LPM and the Probit model. For convenience, we report the estimates of the ratio λ_0/λ_1 ; in all cases reported below the signs of the estimated parameters λ_0 and λ_1 are the same as in the baseline models above.

Industry-year fixed effects. In the baseline models above, we include industry fixed effects as well as year fixed effects. Since the passage of time could work differently across industries regarding firms' propensity to integrate, we augment the model to include industry-and-year fixed effects. Table 2 tells us that our results are extremely robust to this modification; see row (A) in either panel. All estimates of λ_0/λ_1 are very close to the corresponding values in our baseline estimates in Table 1.

Firm-level control variables. In our baseline models, all firm-level heterogeneity other than in productivity Θ_{it} is absorbed by the error term. Consistency of our estimates requires that any unobserved firm-level heterogeneity relevant for the choice between integration and outsourcing be uncorrelated with productivity. Arguably, there are dimensions of heterogeneity among firms that may play a role for firms' integration propensity, yet are correlated with productivity. Although we lack clear theoretical guidance on what dimensions of firm heterogeneity are important in this context, the candidates put forward in the literature (e.g. Corcos et al. (2013)) and adopted in this robustness check are the firm's capital intensity (capital-to-labor ratio), the skill intensity (proxied by the average wage), the R&D intensity (R&D-expenditure over sales) and, finally, the export intensity (exports over sales). All of these firm-level controls enter our regression equation in log-values, and most of the time they are significant.²³ Interestingly, Table 2 reveals that, while the message of Table 1 is clearly upheld, the estimates of λ_0/λ_1 are significantly reduced (in absolute value) relative to the baseline estimates; see row (B) in either panel. This reinforces our basic argument that the effect of productivity can be ambiguous: it works in favor of outsourcing in supplier-intensive industries, and in favor of vertical integration in headquarter-intensive industries.

Sample restricted to headquarters. In our baseline estimations, we include all firms in the sample irrespective of whether they are majority-owned by some other firm or not (i.e., our sample includes both headquarter firms and affiliated firms). This can be a problem in our analysis, because affiliated firms might be bound to receive inputs from their headquarters, which would count towards intra-firm

²³While all variables enter positively and (mostly) significantly in all models, the R&D intensity enters positively and significantly for sourcing in the domestic economy, but negatively and significantly for sourcing in the foreign economy. These results hold true regardless of the estimator used (LPM or Probit).

Table 2. Robustness analysis—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
(A) Industry-year fixed effects	-0.832*** (0.064)	-1.255 (0.450)
(B) Firm-level control variables	-0.645*** (0.043)	-0.698** (0.119)
(C) Sample restricted to headquarters	-0.717*** (0.055)	-0.736*** (0.093)
(D) Outliers with respect to Θ	-0.819*** (0.062)	-1.134 (0.339)
(E) Outliers with respect to ζ	-0.816** (0.058)	-1.001 (0.230)
(F) Combined (A)-(E)	-0.644*** (0.040)	-0.612*** (0.058)
(G) Combined (A)-(E) + firm fixed effects	-0.598** (0.166)	
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
(A) Industry-year fixed effects	-0.784*** (0.077)	-0.772** (0.092)
(B) Firm-level control variables	-0.488*** (0.061)	-0.510*** (0.052)
(C) Sample restricted to headquarters	-0.713*** (0.046)	-0.694*** (0.065)
(D) Outliers with respect to Θ	-0.782*** (0.079)	-0.766** (0.093)
(E) Outliers with respect to ζ	-0.747*** (0.059)	-0.774** (0.102)
(F) Combined (A)-(E)	-0.674*** (0.050)	-0.650*** (0.040)
(G) Combined (A)-(E) + firm fixed effects	-0.599*** (0.141)	

Notes: The table reports estimates of λ_0/λ_1 obtained by estimating the model in (13). In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. (A): Including industry-year fixed effects. (B): Including the capital-to-labor ratio, the average wage, the R&D intensity, and the export intensity (all in logs) as firm-level controls. (C): Excluding firms majority-owned by some other firm. (D): Excluding the bottom and the top one percentile of firms in the productivity distribution of each industry-year combination. (E): Excluding the Beverages industry and the Leather & Footwear industry (bottom and top outliers in supplier intensity distribution). (F): Combining (A)-(E). (G): Adding firm fixed effects to (F) by within-transforming the data; see the model in (14). Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from minus one at the 10%, 5%, 1% levels, respectively.

sourcing in our survey data.²⁴ In order to address this problem, we split the sample into firms not majority-owned by other firms (headquarter firms), and firms under majority ownership (affiliated firms), drawing the line at 50 percent ownership. The property rights model then predicts that we should see robustness of our results for headquarter firms, whereas for affiliated firms our model predictions will hold only by chance. Table 2 shows the results obtained for the sample restricted to headquarter firms, and we see reassuring robustness; see row (C) in either panel. The complementary sample of affiliated firms yields results that are strikingly different. First, we do not observe a clear sign pattern as regards λ_0 and λ_1 , and secondly, we do not find a significant interaction between the productivity of the firm and the supplier intensity of the industry in any of the models we estimate (i.e., the estimated coefficient of the interaction term $\Theta_{it} \times \zeta_s$ is never statistically different from zero); see Table G.1 in Appendix G.

Outliers with respect to Θ . To see if our estimation results are driven by productivity outliers in the sample, we re-estimate our baseline models excluding the most productive as well as the least productive firms (top and bottom 1%) in each industry-year pair from the sample. We do not find that they are.

Outliers with respect to ζ . The Beverages industry is by far the least supplier-intensive industry in our sample ($\zeta_s = 0$), while the Leather & Footwear industry is, by a large margin, the most supplier-intensive one ($\zeta_s = 1$). Excluding these industries from our sample does not change any of our conclusions from above.

Combined modifications. To further test the robustness of our results, we also estimate the model with all of the above modifications included at once, i.e., we include industry-year fixed effects as well as firm-level control variables, restrict the sample to headquarter firms, and exclude outliers with respect to both productivity and supplier intensity. The results are robust and remarkably consistent across the domestic and the foreign economy, as well as across the two estimators used; see row (F) in either panel.

Firm fixed effects. Finally, we want to go one step further in tackling unobserved heterogeneity at the firm-level. We do this by augmenting the last model specification with firm fixed effects that absorb any time-invariant firm characteristics in the estimation. We restrict this fixed effects (FE) estimation to the linear model, because the Probit model suffers from the incidental parameters problem when the number of parameters to be estimated increases with the number of firms N . Formally, the FE model we estimate reads as

$$\Pr(\mathbb{1}(\xi_{it} > 0)|\cdot) = E(\mathbb{1}(\xi_{it} > 0)|\cdot) = \frac{\lambda_0 \cdot \Theta_{it} + \lambda_1 \cdot \Theta_{it} \times \zeta_s + \gamma_i + \gamma_{st} + \beta \mathbf{X}_{it} + L}{2L}, \quad i \in \mathcal{I}_s, \quad (14)$$

where γ_i is a firm fixed effect, γ_{st} is an industry-year fixed effect, and \mathbf{X}_{it} is a vector of time-varying firm-level controls (with a corresponding vector of parameters β to be estimated). A word of caution regarding the FE estimator is in order. This estimator identifies the productivity effect from within-firm variation alone. This leaves little identifying variation in our application, due to relatively few firms changing their sourcing strategy over time.²⁵ While this is reflected in larger standard errors in the estimation, the point estimates of λ_0/λ_1 decrease only slightly (in absolute value) and are not statistically different from the model without firm fixed effects; see row (G) in either panel of the table. This clearly suggests that unobserved firm heterogeneity is not driving our results regarding the productivity effect.

²⁴It might also be that affiliated firms are unable to exercise full discretion over their ownership decisions in relation to other suppliers.

²⁵In our baseline sample, we have 286 firms changing from $\xi = 0$ to $\xi > 0$ (or vice versa) for sourcing in the domestic economy, and 173 such firms for sourcing in the foreign economy.

4.3 Instrumental variables (IV) estimation

We now address concerns of reverse causality that could imply that our estimates are biased due to endogeneity. In particular, it could be that the specific ownership structure the firm adopts affects the firm’s measured productivity.²⁶ Since we lack exogenous variation in productivity (e.g. variation induced by policy shocks), we adopt an internal instrumentation strategy to establish causality running from productivity to sourcing behavior, and not the other way around. In particular, we exploit the panel dimension of our data set and use past productivity to explain current productivity in the first stage. Since we have two potentially endogenous variables in our model, we also include past productivity interacted with the supplier intensity of production as a second excluded instrument. A big advantage of applying this strategy on our data is that we observe most of our variables (including TFP) also in pre-sample years (i.e., before 2006). The loss of observations is thus minimal even if we go back in time by several periods. The key assumptions underlying our instrumentation strategy are (i) that past productivity is correlated with current productivity, and (ii) that past productivity is uncorrelated with the firm’s current sourcing behavior (conditional on current productivity).

We first adopt the baseline specification from Table 1 and pool the data across all years and firms using five-year lags as instruments in both the LPM and the Probit model. Our first-stage regressions attest to a statistically significant positive (partial) correlation. Its significance is also reflected in relatively high values for the first-stage F statistics. The second-stage estimations in Table 3 strengthen our previous result that the productivity effect strongly depends on the supplier intensity of the industry, whether we look at domestic or foreign sourcing. Reassuringly, the point estimates of λ_0/λ_1 are very similar to our baseline estimates in Table 1, and this holds true for both the LPM and the Probit model.

To examine the robustness of these results, we also adopt an alternative specification that controls for unobserved firm heterogeneity (IV-FE model) and uses the first three lags of productivity (along with the corresponding interactions with the industry’s supplier intensity) as instruments. Table G.2 in Appendix G reveals no evidence that the instruments are weak or invalid, as the first-stage F statistics are reasonably large, and we cannot reject the joint null hypothesis that our instruments are valid instruments based on the Hansen J statistic. The IV-FE estimates strengthen our previous results since the estimated values of λ_0/λ_1 lie in the close neighborhood of what we find in both our baseline estimates and the pooled IV estimates.

4.4 Extensions

In this section we present several important extensions of our baseline empirical models. First, we take further industry variation into account that could influence our results. Secondly, we control for general equilibrium effects. And finally, we allow for non-linearity and non-monotonicity.

Further industry interactions. Our baseline specification chooses the inverse capital intensity as the preferred industry-specific measure of the supplier intensity. However, the fundamental reasoning behind this choice may to some extent also be applied to other industry characteristics, like the skill intensity or the R&D intensity; see Antràs (2003), Yeaple (2006), and Nunn & Treffer (2008). In order to see whether including these alternative measures changes our results, we run additional estimations reported in Tables 4 and 5. Table 4 augments the model in (13) by including an interaction term with the R&D intensity, while Table 5 includes an interaction term with the skill intensity. In either

²⁶See Hortaçsu & Syversen (2007) and Forbes & Lederman (2010) for evidence on the effect of vertical integration on productivity in the cement and the airline industry in the United States, respectively.

Table 3. Instrumental variables estimation—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	0.940*** (0.182)	3.256*** (0.721)
TFP \times Supplier intensity	-1.166*** (0.274)	-2.879** (1.256)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.806*** (0.046)	-1.131 (0.262)
Kleibergen-Paap Wald rk F statistic	58.04	
N	9231	9231
R^2	0.039	
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	1.058*** (0.204)	4.906*** (0.955)
TFP \times Supplier intensity	-1.320*** (0.305)	-6.117*** (1.639)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.801*** (0.056)	-0.802** (0.079)
Kleibergen-Paap Wald rk F statistic	77.59	
N	4728	4635
R^2	0.061	

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. TFP and TFP \times Supplier intensity are instrumented with their five-year lags. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

Table 4. Model augmented with R&D intensity—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	0.368*** (0.094)	1.289*** (0.437)
TFP \times Supplier intensity	-0.444*** (0.129)	-1.200* (0.674)
TFP \times R&D intensity	0.005 (0.021)	-0.084 (0.130)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.829** (0.068)	-1.074 (0.316)
N	12205	12205
R^2 or Pseudo R^2	0.076	0.110
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	0.420*** (0.128)	2.186*** (0.738)
TFP \times Supplier intensity	-0.539*** (0.183)	-2.912** (1.190)
TFP \times R&D intensity	-0.005 (0.028)	-0.160 (0.170)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.780** (0.086)	-0.751** (0.098)
N	5896	5754
R^2 or Pseudo R^2	0.079	0.108

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

Table 5. Model augmented with skill intensity—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	0.260** (0.131)	1.052 (0.782)
TFP \times Supplier intensity	-0.395*** (0.128)	-1.089 (0.740)
TFP \times Skill intensity	0.0458 (0.053)	0.157 (0.317)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.657 (0.209)	-0.965 (0.465)
N	12205	12205
R^2 or Pseudo R^2	0.077	0.112
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	0.257 (0.201)	2.054* (1.068)
TFP \times Supplier intensity	-0.487*** (0.186)	-3.077*** (1.151)
TFP \times Skill intensity	0.0762 (0.078)	0.178 (0.389)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.527 (0.289)	-0.668 (0.240)
N	5896	5754
R^2 or Pseudo R^2	0.079	0.107

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

case the new interaction terms are insignificant, and our key result on how the productivity effect varies across industries with different supplier intensities remains intact. Hence, we may conclude that these measures do not appear to have major explanatory power in relation to the productivity effect at the extensive margin of vertical integration, over and above what is captured by our measure of the supplier intensity of the production process.

Controlling for industry size A . We assume in Proposition 1 (c) as well as in Proposition 2 (b) that the term A in Equations (2) and (3) is invariant with respect to the supplier intensity of production ζ . The term A , essentially a measure of industry size, captures the general equilibrium interrelationship between industries, and enters maximum profits in a multiplicative way; see Appendix A. Hence, the role of industry size, if any, cannot adequately be absorbed by industry fixed effects. More specifically, to the extent that A is correlated with the supplier intensity, the cross-industry heterogeneity in the productivity effect that we find in our baseline estimates might stem from differences in industry size, in addition to differences in the supplier intensity of production. To allay this concern, we have augmented the baseline models to allow for the productivity effect to vary, not only with supplier intensity, but also with industry size. Estimation of this augmented model reveals that the interaction effects between TFP and A are mostly positive but insignificant; see Table G.3 in Appendix G. Importantly, our results on the interaction between TFP and supplier intensity are not significantly affected.²⁷

Modeling non-linearity as well as non-monotonicity. In the baseline models above, we model variation in the productivity effect across industries through an interaction term between Θ_{it} and ζ_s . This imposes a rigid parametric relationship such that the effect of productivity changes linearly, and thus monotonically, with ζ . We now provide estimates of a more flexible model that allows for non-linearity as well as non-monotonicity. This modeling strategy can be motivated by our theoretical analysis, which demonstrates a non-linear and potentially non-monotonic response of the productivity effect to changes in the supplier intensity.

We explore this issue by partitioning the supplier intensity into quintiles of its sample distribution. That is, we recode ζ so that the value 1 represents the 20% of the sample with the lowest values on ζ , the value 2 represents the 20% with the next-lowest values on ζ , and so on, until, finally, the value 5 represents the 20% of the sample with the highest values on ζ . We then dummy up the quintiles and examine the relationship between the productivity effect and the quintile dummies, while controlling for other channels through fixed effects (as before):

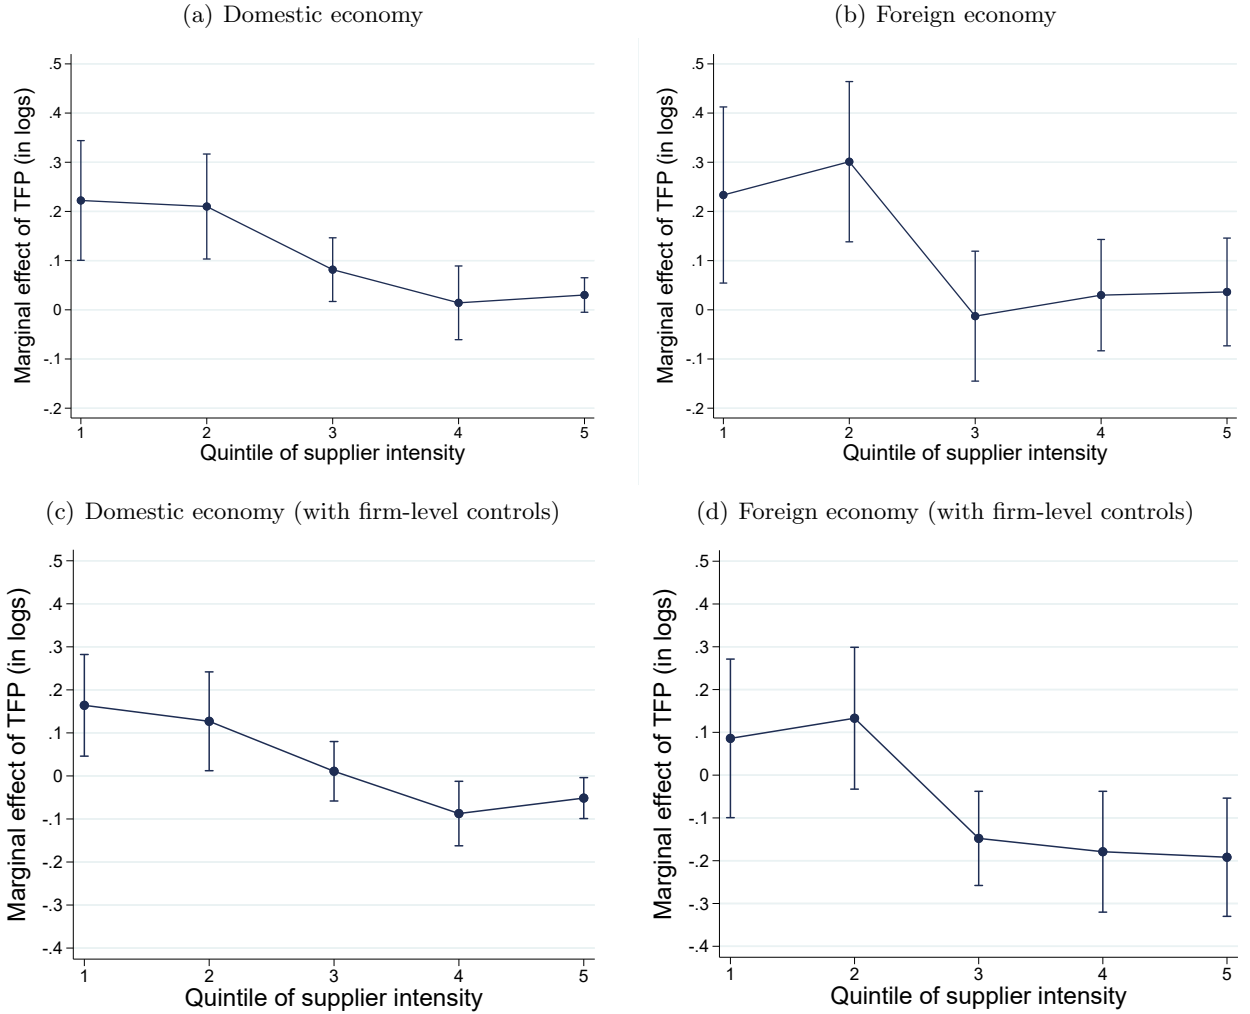
$$\Pr(\mathbb{I}(\xi_{it} > 0) | \cdot) = \Pr \left(\sum_{q=1}^5 \lambda_q \cdot \text{Quintile}_q^s \times \Theta_{it} + \gamma_s + \gamma_t > -\Delta\pi_{it} \right), \quad i \in \mathcal{I}_s, \quad (15)$$

where Quintile_q^s is a dummy variable equal to one if industry s falls into the q th quintile of the distribution of ζ_s , and $\lambda_1, \dots, \lambda_5$ are the parameters of interest. Subfigures (a) and (b) in Figure 7 display the main results from these estimations. These results paint a very consistent picture that provides considerable further support for the property rights theory of the firm. In both the domestic and the foreign economy, we find significantly positive productivity effects in favor of vertical integration at the bottom of the distribution of supplier intensities, and a zero effect at the top of the distribution. The point estimates suggest that, in industries belonging to the two bottom quintiles of the supplier intensity of production, a doubling of firm productivity increases the likelihood of

²⁷We have also experimented with a polynomial regression framework, allowing for a non-monotonic relationship between the productivity effect and industry size. Including interactions with higher-order polynomials of industry size does not change any of our conclusions. The detailed results are available upon request.

vertical integration by 20 to 30 percentage points (in either sourcing location). Interestingly, when we augment the model to include firm-level controls (i.e. capital, skill, R&D, and export intensity), we find a range of high supplier intensities where the productivity effect actually works in favor of full outsourcing (and against vertical integration); see subfigures (c) and (d) in Figure 7. This is similar to our robustness analysis in Table 2.

Figure 7. Productivity effect by quintiles of supplier intensity



5 Conclusions

We believe that our results amount to considerable empirical support of the property rights model of firm boundaries in input provision, as pioneered by Antràs & Helpman (2004). We have derived a novel firm-level prediction from this model, focusing on how productivity gains affect the optimal allocation of ownership rights in input sourcing. The empirical results we present in this paper clearly demonstrate that this prediction finds empirical support in Spanish firm-level data on both, domestic and foreign input sourcing.

The novelty of our prediction draws from a simple change in perspective: Instead of looking only at whether or not vertical integration delivers larger profits than outsourcing, we argue that

a firm-level decision model should examine the magnitude of the profit difference between vertical integration and outsourcing. Doing so, we find that the property rights model predicts differential effects of productivity gains on this profit difference, depending on the differential mix of supplier inputs and headquarter inputs (the supplier intensity of production). Two features stick out, setting our prediction apart from industry-level predictions. First, changes in the profit difference induced by gains in productivity are potentially non-monotonic in the supplier intensity of production. And secondly, such changes are always independent of the fixed cost associated with different ownership structures.

Arguably, testing such a prediction on firm-level data requires that it is formulated under empirically meaningful assumptions. We therefore generalize our prediction to the case of multiple input suppliers. A key finding here is that the mechanism determining whether a firm chooses to integrate at least some of its suppliers (the extensive margin of integration) is fundamentally different from the mechanism determining the share of integrated supplies (the intensive margin of integration). More specifically, in this multi-supplier setup, while the extensive margin response of firms to productivity gains continues to be independent of the fixed cost ranking, the intensive margin response crucially depends on the fixed cost ranking. This is a general conclusion relevant for future empirical research on the property rights model. Since the fixed cost ranking is an unobservable variable, we restrict our empirical analysis to the extensive margin, and find ample evidence that (i) the estimated productivity effects on firms' sourcing decisions are significant, both statistically and economically, and (ii) their pattern across industries' supplier intensity supports the property rights model.

From a policy perspective, a conclusion to be drawn from our analysis is that the hold-up-induced inefficiency highlighted by the property rights theory of the firm is empirically significant and that the choice of a suitable boundary is an important way for firms to deal with this inefficiency. However, according to property rights theory, this still leaves a certain degree of inefficiency to worry about by policy makers. Moreover, a firm's productivity plays an important role for the efficient firm boundary, and crucially, this effect is asymmetric across firms. Thus, an equal percentage increase in productivity across all firms will alter firms' decisions in favor of vertical integration in industries at the bottom end of the distribution of supplier intensities, while this effect works in favor of outsourcing in industries positioned at the upper end of this distribution.

References

- Alfaro, L., Antràs, P., Chor, D., & Conconi, P. (forthcoming). Internalizing global value chains: A firm-level analysis. *Journal of Political Economy*.
- Antràs, P. (2003). Firms, contracts, and trade structure. *Quarterly Journal of Economics*, 118(4), 1375–1418.
- Antràs, P. (2016). *Global Production: Firms, Contracts, and Trade Structure*. Princeton University Press.
- Antràs, P., & Chor, D. (2013). Organizing the global value chain. *Econometrica*, 81(6), 2127–2204.
- Antràs, P., & Helpman, E. (2004). Global sourcing. *Journal of Political Economy*, 112(3), 552–580.
- Antràs, P., & Helpman, E. (2008). Contractual frictions and global sourcing. In E. Helpman, D. Marin, & T. Verdier (Eds.) *The Organization of Firms in a Global Economy*, chap. 1, (pp. 9–54). Harvard University Press: Cambridge, MA.

- Antràs, P., & Staiger, R. W. (2012). Offshoring and the role of trade agreements. *American Economic Review*, 102(7), 3140–3183.
- Aoki, M. (1990). Toward an economic model of the Japanese firm. *Journal of Economic Literature*, 28(1), 1–27.
- Bergstrand, J. H., Egger, P., & Larch, M. (2013). Gravity redux: Estimation of gravity-equation coefficients, elasticities of substitution, and general equilibrium comparative statics under asymmetric bilateral trade costs. *Journal of International Economics*, 89(1), 110–121.
- Bernard, A. B., Jensen, J. B., Redding, S. J., & Schott, P. K. (2010). Intrafirm trade and product contractibility. *American Economic Review: Papers & Proceedings*, 100(2), 444–448.
- Carluccio, J., & Bas, M. (2015). The impact of worker bargaining power on the organization of global firms. *Journal of International Economics*, 96(1), 162–181.
- Carluccio, J., & Fally, T. (2012). Global sourcing under imperfect capital markets. *Review of Economics and Statistics*, 94(3), 740–763.
- Caves, D., Christensen, L., & Diewert, W. (1982). Output, input and productivity using superlative index numbers. *Economic Journal*, 92, 73–96.
- Corcos, G., Irac, D. M., Mion, G., & Verdier, T. (2013). The determinants of intrafirm trade: Evidence from French firms. *Review of Economics and Statistics*, 95(3), 825–838.
- Defever, F., & Toubal, F. (2013). Productivity, relationship-specific inputs and the sourcing modes of multinationals. *Journal of Economic Behavior & Organization*, 94, 345–357.
- Del Prete, D., & Rungi, A. (2017). Organizing the global value chain: A firm-level test. *Journal of International Economics*, 109, 16–30.
- Delgado, M. A., Fariñas, J. C., & Ruano, S. (2002). Firm productivity and export markets: A non-parametric approach. *Journal of International Economics*, 57(2), 397–422.
- Dunning, J. H. (1993). *Multinational Enterprises and the Global Economy*. Addison Wesley Longman, Inc.
- Federico, S. (2010). Outsourcing versus integration at home or abroad. *Empirica*, 37(1), 47–63.
- Federico, S. (2012). Headquarter intensity and the choice between outsourcing versus integration at home or abroad. *Industrial and Corporate Change*, 21(1), 1–22.
- Feenstra, R. C., & Hanson, G. H. (2005). Ownership and control in outsourcing to China: Estimating the property-rights theory of the firm. *Quarterly Journal of Economics*, 120(2), 729–761.
- Fernandes, A. P., & Tang, H. (2012). Determinants of vertical integration in export processing: Theory and evidence from China. *Journal of Development Economics*, 99(2), 396–414.
- Forbes, S. J., & Lederman, M. (2010). Does vertical integration affect firm performance? Evidence from the airline industry. *The RAND Journal of Economics*, 41(4), 765–790.
- Good, D., Nadiri, M. I., & Sickles, R. C. (1997). Index number and factor demand approaches to the estimation of productivity. In M. H. Pesaran, & P. Schmidt (Eds.) *Handbook of Applied Econometrics, Volume II-Microeconometrics*, chap. 1, (pp. 14–80). Oxford: Basil Blackwell.

- Grossman, S. J., & Hart, O. D. (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy*, 94(4), 691–719.
- Hart, O., & Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy*, 98(6), 1119–1158.
- Hortaçsu, A., & Syversen, C. (2007). Cementing relationships: Vertical integration, foreclosure, productivity, and prices. *Journal of Political Economy*, 115(2), 250–301.
- Kohler, W., & Smolka, M. (2011). Sourcing premia with incomplete contracts: Theory and evidence. *B.E. Journal of Economic Analysis and Policy*, 11(1), 1–37.
- Kohler, W., & Smolka, M. (2012). Global sourcing: Evidence from Spanish firm-level data. In R. M. Stern (Ed.) *Quantitative Analysis of Newly Evolving Patterns of International Trade*, chap. 4, (pp. 139–189). World Scientific Studies in International Economics.
- Kohler, W., & Smolka, M. (2014). Global sourcing and firm selection. *Economics Letters*, 124(3), 411–415.
- Milgrom, P., & Roberts, J. (1993). Johnson Controls Inc., Automotive Systems Group: The Georgetown, Kentucky Plant. Stanford Graduate School of Business Case S-BE-9.
- Naghavi, A., Spies, J., & Toubal, F. (2015). Intellectual property rights, product complexity and the organization of multinational firms. *Canadian Journal of Economics*, 48(3), 881–902.
- Nunn, N., & Treffer, D. (2008). The boundaries of the multinational firm: An empirical analysis. In E. Helpman, D. Marin, & T. Verdier (Eds.) *The Organization of Firms in a Global Economy*, chap. 2, (pp. 55–83). Harvard University Press: Cambridge, MA.
- Nunn, N., & Treffer, D. (2013). Incomplete contracts and the boundaries of the multinational firm. *Journal of Economic Behavior & Organization*, 94, 330–344.
- Ornelas, E., & Turner, J. L. (2008). Trade liberalization, outsourcing, and the hold-up problem. *Journal of International Economics*, 74(1), 225–241.
- Ornelas, E., & Turner, J. L. (2012). Protection and international sourcing. *Economic Journal*, 122(559), 26–63.
- Schwarz, C., & Südekum, J. (2014). Global sourcing of complex production processes. *Journal of International Economics*, 93(1), 123–139.
- Tomiura, E. (2007). Foreign outsourcing, exporting, and FDI: A productivity comparison at the firm level. *Journal of International Economics*, 72(1), 113–127.
- Whinston, M. D. (2003). On the transaction cost determinants of vertical integration. *Journal of Law, Economics, and Organization*, 19(1), 1–23.
- Yeaple, S. R. (2006). Offshoring, foreign direct investment, and the structure of U.S. trade. *Journal of the European Economic Association*, 4(2-3), 602–611.
- Young, S., Hood, N., & Hamill, J. (1985). Decision-making in foreign-owned multinational subsidiaries in the United Kingdom. ILO Multinational Enterprises Programme Working Paper No. 35, Geneva: International Labor Office.

Appendices

A General equilibrium interrelationship

In the definition of maximum profits in the text, the term A captures the general equilibrium interrelationship between industries. In this appendix we provide more details on the exact meaning of this term in a model of the type considered in Antràs & Helpman (2004).

For easier notation, we abstain from indexing sectors. Suppose that there are quasi-linear preferences over sectoral aggregates X , composed in a CES way of differentiated varieties of final goods, such as the ones produced in the production relationship considered above. Inverse demand of the consumption aggregate is given by $P = X^{\kappa-1}$, where κ is a utility parameter satisfying $0 < \kappa < 1$. Defining $\alpha := (\varepsilon - 1)/\varepsilon$, the term A then emerges as

$$A := X^{*(\kappa-\alpha)/(1-\alpha)} \alpha^{\alpha/(1-\alpha)}, \quad (\text{A.1})$$

where X^* is the *general equilibrium quantity* of aggregate consumption of the differentiated final goods. This is easily verified from Equations (6) and (7) in Antràs & Helpman (2004). Equilibrium industry-specific expenditure may be written as $R^* := X^* P^* = X^{*\kappa}$, whence A may equivalently be written as $R^{*\varepsilon+(1-\varepsilon)/\kappa} \alpha^{\alpha/(1-\alpha)}$.

In Antràs & Helpman (2004), X^* is determined by a zero profit condition, assuming free and costless entry. Note that this type of general equilibrium closure of the model assumes $\alpha > \kappa$ (higher degree of substitution within than between sectors). This implies that maximum profits are falling in X^* (or R^*). The intuition is that a higher X^* is generated through firm entry, so that output (and thus revenue as well as profits) per firm is reduced. This guarantees that entry caused by positive profits eventually leads to a zero profit equilibrium.

B Proof of Proposition 1, parts (b) and (c)

To prove part (b) of Proposition 1, it proves convenient to introduce the notions of supermodularity and submodularity of functions.

Definition 1. (a) The function $H(g, q)$ is called *supermodular with respect to g and q* , if for any two values $g_1 > g_0$ and $q_1 > q_0$ the following is true: $\Delta_g H(q_1) > \Delta_g H(q_0)$, where $\Delta_g H(q) := H(g_1, q) - H(g_0, q)$. (b) The function $H(g, q)$ is called *submodular*, if $-H(g, q)$ is supermodular with respect to g and q . (c) If $H(g, q)$ is twice differentiable, then it is called *supermodular*, if $\partial^2 H / (\partial g \partial q) > 0$, and vice versa for submodularity.

We start by proving the first statement in part (b) of Proposition 1. The productivity effect is equal to $(\varepsilon - 1)\Delta_m Z(\ell; \zeta)\theta^{\varepsilon-2}$. This effect being more favorable to outsourcing in more sourcing-intensive industries means that $\Delta_m Z(\ell_h; \zeta)$ is falling in ζ . We know from Lemma 3 in Antràs (2003) that $Z(\ell, m_v; \zeta) / Z(\ell_h, m_o; \zeta)$ is decreasing monotonically in ζ . We write $Z_v(\cdot, \zeta)$ for $Z(\ell, m_v; \zeta)$ and accordingly for $Z_o(\cdot, \zeta)$. Since by definition $\frac{Z_v(\cdot, \zeta^*)}{Z_o(\cdot, \zeta^*)} = 1$, we have $\frac{\partial Z_v(\cdot, \zeta^*)}{\partial \zeta} - \frac{\partial Z_o(\cdot, \zeta^*)}{\partial \zeta} < 0$. Therefore,

$$\frac{\partial}{\partial \zeta} \Delta_m Z(\ell_h; \zeta^*) = \left[\frac{\partial Z_v(\cdot, \zeta^*)}{\partial \zeta} - \frac{\partial Z_o(\cdot, \zeta^*)}{\partial \zeta} \right] < 0. \quad (\text{B.1})$$

Continuity of the function Z implies that in the neighborhood of ζ^* the function $\Delta_m Z(\ell_h; \zeta)$ is unambiguously falling in ζ . Thus, the interval $[\underline{\zeta}, \bar{\zeta}]$ includes ζ^* , which proves the first statement in part (b) of Proposition 1.

We next prove that the effect is potentially non-monotonic in $\zeta \in [0, 1]$ (the second statement in part (b) of the proposition). Given that $\partial \Delta_m Z(\ell; \zeta) / \partial \zeta < 0$ for ζ^* , non-monotonicity arises if there exist values of $m \in [0, 1]$ and $\zeta \in [0, 1]$ for which the function Z is *not* submodular with respect to m and ζ . It proves convenient to use the decomposition

$$\frac{\partial^2 Z}{\partial \zeta \partial m} = Z \frac{\partial^2 \ln Z}{\partial \zeta \partial m} + \frac{\partial \ln Z}{\partial \zeta} \frac{\partial Z}{\partial m}. \quad (\text{B.2})$$

We take each of the three derivatives on the right hand side of (B.2) in turn, starting with

$$\frac{\partial \ln Z}{\partial \zeta} = \frac{\partial \ln z}{\partial \zeta} + \frac{\partial \ln C}{\partial \zeta} + \frac{\partial \ln A}{\partial \zeta}. \quad (\text{B.3})$$

The model does not in any way restrict the value of the final term in this decomposition. Hence, we cannot rule out that $\frac{\partial^2 Z}{\partial \zeta \partial m}$ is positive outside the interval $[\underline{\zeta}, \bar{\zeta}]$, which proves the second statement in part (b) of the proposition.

We next show, however, that $\frac{\partial^2 Z}{\partial \zeta \partial m}$ can be positive even if we assume that $\partial A / \partial \zeta = 0$. We have

$$\frac{\partial \ln C}{\partial \zeta} = (\varepsilon - 1) [\ln \ell + \ln(1 - m) - \ln m], \quad (\text{B.4})$$

$$\text{and } \frac{\partial \ln z}{\partial \zeta} = \frac{\varepsilon - 1}{\varepsilon} \frac{2m - 1}{z}. \quad (\text{B.5})$$

Bringing the terms together, we obtain

$$\frac{\partial \ln Z}{\partial \zeta} = \frac{\varepsilon - 1}{\varepsilon} \frac{2m - 1}{z} + (\varepsilon - 1) [\ln \ell + \ln(1 - m) - \ln m], \quad (\text{B.6})$$

$$\text{and } \frac{\partial^2 \ln Z}{\partial \zeta \partial m} = \frac{\varepsilon - 1}{\varepsilon} \left[\frac{2}{z} - \frac{2m - 1}{z^2} \frac{\partial z}{\partial m} \right] - \left(\frac{\varepsilon - 1}{m - m^2} \right). \quad (\text{B.7})$$

We insert $\frac{\partial z}{\partial m} = \frac{\varepsilon - 1}{\varepsilon} (2\zeta - 1)$ into (B.7) to obtain

$$\frac{\partial^2 \ln Z}{\partial \zeta \partial m} = -\frac{(\varepsilon - 1)}{m - m^2} + \frac{1}{z^2} [2\alpha z - \alpha^2 (2m - 1) (2\zeta - 1)], \quad (\text{B.8})$$

where $\alpha := \frac{\varepsilon - 1}{\varepsilon}$. Inserting for z reduces the bracketed term in (B.8) to $2\alpha - \alpha^2$, so that

$$\frac{\partial^2 \ln Z}{\partial \zeta \partial m} = -\frac{(\varepsilon - 1)}{m - m^2} + \frac{2\alpha - \alpha^2}{z^2} = -\left[\frac{(\varepsilon - 1)}{m - m^2} + \frac{\alpha^2 - 2\alpha}{z^2} \right]. \quad (\text{B.9})$$

This term is strictly negative. To see this, recall that α , m , and ζ all lie strictly between zero and one. Moreover, note that the first term in the square bracket on the right-hand side (call it x) is strictly positive, and the second term (call it y) is strictly negative. For any given value of α , the values of m and ζ that minimize the denominator in y , and thus maximize the (absolute) value of the function y , are $m = 1/2$ and $\zeta = 1/2$. In turn, for any given value of α the value of m that minimizes the value of the function x is $m = 1/2$. Hence, a sufficient condition for the right-hand side of (B.9) to be negative is

$$\frac{4\alpha}{1 - \alpha} + \frac{\alpha^2 - 2\alpha}{(1 - \alpha/2)^2} > 0. \quad (\text{B.10})$$

Straightforward manipulation of this expression yields $2\alpha - \alpha^2 > 0$, which is always true. Hence, the first term in the main decomposition (B.2) is unambiguously negative.

As to the second term in this decomposition, $\frac{\partial \ln Z}{\partial \zeta} \frac{\partial Z}{\partial m}$, we first turn to $\frac{\partial \ln Z}{\partial \zeta}$ as given in (B.6) above. On account of $\ln(1 - m)$, this term converges to minus infinity as m approaches one. Conversely, it converges to infinity as m approaches zero (on account of $-\ln m$). It will also be positive for large enough values of ℓ . The term $\frac{\partial Z}{\partial m}$ may be written as

$$\frac{\partial Z}{\partial m} = A \left[C \frac{\partial z}{\partial m} + z \frac{\partial C}{\partial m} \right]. \quad (\text{B.11})$$

Note that $\partial A / \partial m = 0$, which follows from the fact that derivation with respect to m refers to the difference between m_v and m_o , and not to changes in these ex-post revenue shares..

$$\begin{aligned} \frac{\partial C}{\partial m} &= (\varepsilon - 1) \left[m^{1-\zeta} [\ell(1 - m)]^\zeta \right]^{\varepsilon-2} \\ &\quad \times \left[(1 - \zeta) m^{-\zeta} [\ell(1 - m)]^\zeta - m^{1-\zeta} \zeta [\ell(1 - m)]^{\zeta-1} \ell \right] \\ &= (\varepsilon - 1) C \left(\frac{1 - \zeta}{m - m^2} - \frac{1}{1 - m} \right). \end{aligned} \quad (\text{B.12})$$

Putting things together, we have

$$\begin{aligned} \frac{\partial Z}{\partial m} &= AC \left[\frac{\varepsilon - 1}{\varepsilon} (2\zeta - 1) + z(\varepsilon - 1) \left(\frac{1 - \zeta}{m - m^2} - \frac{1}{1 - m} \right) \right] \\ &= AC(\varepsilon - 1) \left[\frac{2\zeta - 1}{\varepsilon} + \frac{z}{1 - m} \left(\frac{1 - \zeta - m}{m} \right) \right]. \end{aligned} \quad (\text{B.13})$$

This term is again ambiguous. It is positive for sufficiently low values of m and negative for sufficiently high values of m . It now becomes clear that $\frac{\partial \ln Z}{\partial \zeta} \frac{\partial Z}{\partial m}$ is positive for sufficiently low as well as for sufficiently high values of m , whatever the value of $\frac{\partial \ln A}{\partial \zeta}$ (including, as assumed above, $\frac{\partial \ln A}{\partial \zeta} = 0$ of course). Low values render both terms positive, whereas high values render both terms negative. Since the possible values of $\frac{\partial \ln Z}{\partial \zeta} \frac{\partial Z}{\partial m}$ include infinity, it follows that the second term in (B.2) potentially dominates the first, thus leading to a positive value of $\frac{\partial^2 Z}{\partial \zeta \partial m}$. Therefore, Z is not submodular in ζ and m for all values $\zeta \in [0, 1]$ and $m \in [0, 1]$, which proves potential non-monotonicity irrespective of the value of $\frac{\partial A}{\partial \zeta}$.


To obtain an intuition for this potential non-monotonicity, the crucial point to recognize is that the implication of a change in the supplier intensity on the strategic advantage of integration (or outsourcing) is determined by an interaction of m and ζ . Consider an institutional setup that allows for headquarters to appropriate only small ex-post revenue shares (such that both m_o and m_v are small) in industries where efficiency considerations would call for strong headquarter incentives (i.e. in industries with a low ζ). We know that in such a situation increasing the supplier intensity would increase maximum profits under either ownership structure, as this would mitigate the misalignment between m_j and ζ . However, this effect would be felt more strongly under integration, because the effect of mitigating the misalignment is increasing in the *level* of maximum profits, and in the relevant parameter subspace maximum profits are always higher under integration than under outsourcing (on account of $\zeta < \zeta^*$). We may thus have $\partial \Delta_m \Pi(\zeta, \Theta) / \partial \zeta > 0$ for $\zeta < \zeta^*$, so that vertical integration looks relatively more attractive in more sourcing-intensive industries at low levels of ζ . A similar logic applies when m_o , m_v , and ζ are large, but in this case any increase in ζ would *decrease* maximum profits, as this would reinforce the misalignment in a case where $\zeta > \zeta^*$. The logic is illustrated by Figure B.1.

The proof of part (c) in the text demonstrates the monotonicity result obtaining for full symmetry. For the asymmetric case, $\ell > 1$, the potential non-monotonicity stated in part (b) becomes relevant.

We demonstrate this by evaluating the integral $\int_{m_o}^{m_v} \frac{\partial^2 Z}{\partial \zeta \partial m} dm$ for alternative values of ℓ . Figure B.1 depicts isoclines for $\frac{\partial^2 Z}{\partial \zeta \partial m} = 0$ separating the parameter space (m, ζ) into subspaces of submodularity and supermodularity, respectively. These isoclines depend on the value of ε , which we set equal to six (following Bergstrand et al. (2013)). The asymmetric presence of supermodularity subspaces for $\ell > 1$ is what generates *piecewise* monotonicity. As we increase ℓ , the supermodularity subspace for low values of m and ζ becomes ever more dominating. Therefore, for any value of $m_v > 1/2$, the above integral must lead to a positive value, provided that $\ell > \ell^*$ and ζ is sufficiently small. This proves piecewise monotonicity as stated in (c). Given monotonicity for values $m_v > 1/2$ and $\ell = 1$, continuity implies that monotonicity obtains for a sufficiently small deviation of ℓ from symmetry. Hence, ℓ^* is strictly larger than 1. Figure B.2 translates the modularity result for the cross-derivative $\frac{\partial^2 Z}{\partial \zeta \partial m}$ into a corresponding modularity result for the entire integral, plotting $m_v \geq 1/2$ on the horizontal axis.


The intuition for non-monotonicity to reappear also under symmetric Nash bargaining when ℓ becomes large enough is subtle, but it is clear that lower unit costs for the sourced input change the input supply mix towards the supplier and away from the headquarter. This increases the potential for the type of misalignment responsible for non-monotonicity effects to be observed for low values of ζ , where supplier incentives are less important, but decreases it for high values of ζ . It is relatively straightforward to show that this statement holds for both symmetric as well as asymmetric Nash bargaining games, and applies accordingly also to the case of asymmetric input costs that favor the headquarter input rather than the supplier input.

Figure B.1. Proposition 1, part (c): Modularity of Z with respect to m and ζ^\dagger



[†]Note: The figure assumes different values of $\ell \geq 1$ and sets $A = 1$ as well as $\varepsilon = 6$.

Figure B.2. Proposition 1, part (c): Modularity of Z with respect to m and ζ for $m_v > m_o = 1/2^\dagger$



[†]Note: The figure assumes different values of $\ell \geq 1$ and sets $A = 1$ as well as $\varepsilon = 6$.

C Extension of the AH model with multiple input suppliers: Adding fully contractible inputs

In line with Appendix A, we assume an inverse demand function $P = X^{\kappa-\alpha} Q^{\alpha-1}$. Revenue of the final good then emerges as

$$R = X^{\kappa-\alpha} \theta^\alpha \left(\frac{H}{1-\zeta} \right)^{\alpha(1-\zeta)(1-\tau)} \left(\frac{M}{\zeta} \right)^{\alpha\zeta(1-\tau)} \left(\int_{n \in \omega} G(n)^\delta dn \right)^{\alpha\tau/\delta}. \quad (\text{C.1})$$

The profit maximization problem is solved by backward induction. First, it can be shown that profit-maximizing quantities of non-contractible inputs are given by $H^* = \left(\frac{m\alpha(1-\zeta)(1-\tau)}{c_H} \right) R$ and $M^* = \left(\frac{(1-m)\alpha\zeta(1-\tau)}{c_M} \right) R$, where c_H and c_M are the unit costs of producing H and M , respectively.²⁸ The term H^* maximizes headquarter profits $mR - c_H H - \int_{n \in \omega} c(n) G(n) dn$ in the last stage before the bargaining stage, where $c(n)$ denotes the unit costs of producing contractible input $G(n)$. Accordingly, the term M^* maximizes the supplier's profits $(1-m)R - c_M M$.

When choosing the quantities of the contractible inputs $G(n)$, the headquarter maximizes $R -$

²⁸For convenience, this expression omits subscripts indicating the ownership structure of sourcing.

$c_H H^* - c_M M^* - \int_{n \in \omega} c(n) G(n) dn$. The solution satisfies

$$\tilde{G} \equiv \left(\int_{n \in \omega} G(n)^{\delta} dn \right)^{1/\delta} = \left(\frac{\alpha \tau [1 - \alpha(1 - \tau)[m(1 - \zeta) + (1 - m)\zeta]]}{\left(\int_{n \in \omega} c(n)^{\frac{\delta}{\delta-1}} dn \right)^{1-1/\delta}} \right) R. \quad (\text{C.2})$$

The denominator on the right-hand side of this equation gives the minimum cost of producing one unit of \tilde{G} . We will refer to this as c_G in the following. Notice that, if all inputs in the production process were contractible (including H and M), profit-maximization would yield $\tilde{G} = \left(\frac{\alpha \tau}{c_G} \right) R$. Hence, we find that the presence of contractual frictions distorts input provision not only for the non-contractible inputs, but also for the contractible inputs (see the remarks in the body of the paper).

To show that the predictions we derive from the standard AH model survive in our extension, we have to compute maximum operating profits under either ownership structure. Recalling that $\alpha = (\varepsilon - 1)/\varepsilon$, we can compute these as

$$\begin{aligned} \Pi(\ell, m_j; \zeta, \theta) &= A \varphi \left[1 - (1 - \tau) \frac{\varepsilon - 1}{\varepsilon} [m_j(1 - \zeta) + (1 - m_j)\zeta] \right]^{\varepsilon - (\varepsilon - 1)(1 - \tau)} \times \\ &\quad \times \left[m_j^{1 - \zeta} (\ell(1 - m_j))^{\zeta} \right]^{(\varepsilon - 1)(1 - \tau)} \left(\frac{1}{c_G} \right)^{(\varepsilon - 1)\tau} \theta^{\varepsilon - 1}, \end{aligned} \quad (\text{C.3})$$

where $\varphi \equiv (1 - (\varepsilon - 1)\tau/\varepsilon)(1 - \tau)^{(\varepsilon - 1)(1 - \tau)} \tau^{(\varepsilon - 1)\tau}$ collects constant terms, and where we have set $c_H = 1$ and $c_M = 1/\ell$. This expression resembles Equation (2) in the body of the paper.

The key to the results we derive in our proposition is to show that the ratio $\Pi(\ell, m_v; \zeta, \theta)/\Pi(\ell, m_o; \zeta, \theta)$ is decreasing in ζ , and that there is a unique threshold level ζ^* , implicitly given by $\Pi(\ell, m_v; \zeta^*, \theta)/\Pi(\ell, m_o; \zeta^*, \theta) = 1$, for which the headquarter is indifferent between vertical integration and outsourcing. The argument that shows this is involved, but we can follow the steps sketched in Lemma 3 of Antràs (2003) or in Appendix A.4 in Antràs & Helpman (2008). Details of this are available upon request.

D Proof of Proposition 2

Proof of parts (a) and (c): To prove part (a) of Proposition 2 we must show that for $\zeta < \zeta^*$ the difference $\Delta_\xi \bar{Z}(\zeta)$ is non-negative, and that for $\zeta \geq \zeta^*$ it is non-positive, provided that the fixed cost of sourcing are sufficiently large; see Equation (9). To do so, we use proposition 1.1 in Schwarz & Südekum (2014) which demonstrates the existence and uniqueness of a relevant threshold value of ζ equal to ζ^* .

First, we look at industries with $\zeta < \zeta^*$. Schwarz & Südekum (2014) show that in this case there is a *strictly positive* value of $\xi \in [0, 1]$ that maximizes $\bar{Z}(\xi, \zeta)$. We denote this value by ξ_v^* and note that it is unique on account of $\partial^2 \bar{Z}(\xi, \zeta)/\partial \xi^2 < 0$.²⁹ It follows immediately that $\bar{Z}_v^*(\zeta) > \bar{Z}_o(\zeta)$ and $\partial \bar{Z}_v^*/\partial \xi = 0$, where $\bar{Z}_v^*(\zeta) := \bar{Z}(\xi_v^*, \zeta)$.

From the first-order condition in (7) it is clear that ξ_v will only equal ξ_v^* if the fixed cost difference is equal to zero (as assumed by Schwarz & Südekum (2014)), or if productivity is infinitely large, since only then $\partial \bar{Z}_v/\partial \xi = \partial \bar{Z}_v^*/\partial \xi = 0$. In other words, for a finite level of productivity any fixed cost difference drives a wedge between the *strategically optimal* share of integration, ξ_v^* , and the

²⁹This inequality is shown to hold on page 129 in Schwarz & Südekum (2014) under the assumption that $\alpha + \beta < 1$, which is equivalent to $\delta - \epsilon < 0$ in our notation.

profit-maximizing share of integration, ξ_v . As is clear from Equation (7), the *direction* of this wedge depends on the sign of the fixed cost difference. For $F_v - F_o > 0$, we have $\partial \bar{Z}_v / \partial \xi > \partial \bar{Z}_v^* / \partial \xi = 0$, and so firms choose $\xi_v < \xi_v^*$, whereas for $F_v - F_o < 0$ we have $\partial \bar{Z}_v / \partial \xi < \partial \bar{Z}_v^* / \partial \xi = 0$ and $\xi_v > \xi_v^*$. The *magnitude* of this wedge depends on the productivity of the firm, as can be seen by differentiating the first-order condition in (7):

$$\frac{\partial \xi_v(\zeta, \Theta)}{\partial \Theta} = - \frac{\partial \bar{Z}_v / \partial \xi}{\Theta [\partial^2 \bar{Z}(\xi, \zeta) / \partial \xi^2]}. \quad (\text{D.1})$$

Since we know that $\partial^2 \bar{Z}(\xi, \zeta) / \partial \xi^2 < 0$ (see above), the first-order condition implies that this derivative is positive for $F_v - F_o > 0$ and negative for $F_v - F_o < 0$. In words, a higher productivity prompts firms to choose an optimal share of integration, ξ_v , which is closer to ξ_v^* regardless of the sign of the fixed cost difference.

Now assume a fixed cost advantage of outsourcing, leading to $\xi_v < \xi_v^*$. The fact that $\bar{Z}_v^*(\zeta)$ is a maximum value function along with the fact that $\partial^2 \bar{Z}(\xi, \zeta) / \partial \xi^2 < 0$ then implies immediately that $\bar{Z}_v^*(\zeta) > \bar{Z}_v(\zeta) \geq \bar{Z}_o(\zeta)$. Hence, $\Delta_\xi \bar{Z}(\zeta)$ is always non-negative for $F_v - F_o > 0$. Alternatively, assume a fixed cost advantage of integration, leading to $\xi_v > \xi_v^*$. Then for any ξ_v that lies in the neighborhood of ξ_v^* we have the same result as before: $\bar{Z}_v^*(\zeta) > \bar{Z}_v(\zeta) \geq \bar{Z}_o(\zeta)$ so that $\Delta_\xi \bar{Z}(\zeta) \geq 0$. However, for a large enough difference between ξ_v and ξ_v^* we get a reversal in the ranking: $\bar{Z}_v^*(\zeta) > \bar{Z}_o(\zeta) > \bar{Z}_v(\zeta)$ implying a negative difference $\Delta_\xi \bar{Z}(\zeta)$. In light of (D.1), this is ruled out for firms with a sufficiently high productivity so that ξ_v is close enough to ξ_v^* . Since a sufficiently large fixed cost of sourcing prevents low-productivity firms from entering into production, this completes the proof of part (a) of Proposition 2 as far as industries with $\zeta < \zeta^*$ are concerned.³⁰

Next, we look at industries with $\zeta \geq \zeta^*$. In this case we know from Schwarz & Südekum (2014) that the value of $\xi \in [0, 1]$ that maximizes $\bar{Z}(\xi, \zeta)$ is equal to zero: $\xi_v^* = 0$. Hence, $\partial \bar{Z}_v^* / \partial \xi = \partial \bar{Z}_o / \partial \xi < 0$, with equality for $\zeta = \zeta^*$. This implies that $\bar{Z}_v^*(\zeta) = \bar{Z}_o(\zeta) \geq \bar{Z}_v(\zeta)$, which demonstrates that the difference $\Delta_\xi \bar{Z}(\zeta)$ is non-positive as claimed in part (a) of Proposition 2. We now show that this difference is zero for $F_v - F_o > 0$ and is negative for $F_v - F_o < 0$ (assuming a finite level of productivity). For $F_v - F_o > 0$, we know that $\xi_v \rightarrow \xi_v^* = 0$ so that $\bar{Z}_v(\zeta) \rightarrow \bar{Z}_o(\zeta)$ and thus $\Delta_\xi \bar{Z}(\zeta) \rightarrow 0$. For $F_v - F_o < 0$, there are two possibilities. The first possibility is that $\xi_v \rightarrow \xi_v^* = 0$ so that $\Delta_\xi \bar{Z}(\zeta) \rightarrow 0$ (as before), meaning that the strategic advantage of outsourcing outweighs the fixed cost advantage of integration. As is easily verified, this case may arise for high-productivity firms (and is bound to arise for infinitely productive firms), since the fixed cost weigh less heavily on them. The second possibility is that $\xi_v > \xi_v^* = 0$, which immediately implies that $\bar{Z}_v^*(\zeta) = \bar{Z}_o(\zeta) > \bar{Z}_v(\zeta)$ so that $\Delta_\xi \bar{Z}(\zeta) < 0$. This completes the proof of part (a) of Proposition 2.

Part (c) of Proposition 2 is proven by the fact that the sign of the derivative in (D.1) depends on whether $F_v - F_o > 0$ or $F_v - F_o < 0$, as discussed above. The weak inequalities in the proposition derive from the possibility of corner solutions, as $\partial \xi_v / \partial \Theta$ can be zero if $\xi_v \rightarrow 0$ or $\xi_v = 1$.

Proof of part (b): The proposition states that in the neighborhood of ζ^* the following inequality holds:

$$\frac{\partial}{\partial \zeta} \Delta_\xi \bar{Z}(\zeta) = \frac{\partial}{\partial \zeta} [\bar{Z}_v(\zeta) - \bar{Z}_o(\zeta)] \leq 0. \quad (\text{D.2})$$

We write

$$\bar{Z}_v(\zeta) = \bar{Z}_o(\zeta) + \int_0^{\xi_v} \bar{Z}_\xi(\xi, \zeta) d\xi, \quad (\text{D.3})$$

³⁰In the text we focus on situations in which $\xi_v \in (0, 1)$ and $\xi_v^* \in (0, 1]$. It may of course happen that firms are driven into the corner solutions with $\xi_v \rightarrow 0$ or $\xi_v = 1$, but it is easy to verify that this does not change any of our conclusions.

where $\bar{Z}_\xi(\xi, \zeta)$ denotes $\partial \bar{Z}(\xi, \zeta)/\partial \xi$. Equation (D.3) assumes that ζ affects \bar{Z} only through the hold-up mechanisms captured by a change in the extent of integration, and not through the market size variable A which is included in \bar{Z} . Taking derivatives and applying Leibniz' rule, we find that the inequality in (D.2) holds if

$$\frac{\partial}{\partial \zeta} \int_0^{\xi_v} \bar{Z}_\xi(\xi, \zeta) d\xi = \int_0^{\xi_v} \bar{Z}_{\xi\zeta}(\xi, \zeta) d\xi + \bar{Z}_\xi(\xi_v, \zeta) \frac{\partial \xi_v}{\partial \zeta} \leq 0, \quad (\text{D.4})$$

where $\bar{Z}_{\xi\zeta}(\xi, \zeta)$ denotes the cross-derivative of $\bar{Z}(\xi, \zeta)$ with respect to ξ and ζ . We refer to the first term on the right-hand side as the *inframarginal* effect, and to the second term as the *marginal* effect.

To prove that (D.4) holds in the neighborhood of ζ^* , we start from a situation where ζ is slightly below ζ^* . Moreover, we invoke the assumption, as stated in the proposition, that the fixed cost of sourcing are sufficiently large. This implies that only firms with a sufficiently high level of productivity enter into production, so that ξ_v is in the neighborhood of ξ_v^* .

We first note that in the neighborhood of ξ_v^* the *marginal effect* is negligibly small. This follows from $\bar{Z}_\xi(\xi_v^*, \zeta) = 0$ and from continuity of $\bar{Z}_\xi(\xi_v, \zeta)$ in ξ . This is a direct consequence of the Envelope theorem, applied to $\bar{Z}_v(\zeta)$ for an infinitely large productivity level.

A sufficient condition for the *inframarginal effect* to be negative is that $\bar{Z}_{\xi\zeta}(\xi, \zeta) < 0$ for any $\xi \in [0, \xi_v]$. We can determine the sign of this cross-derivative for $\xi = \xi_v$ by differentiating the first-order condition on $\xi_v \in (0, 1)$:

$$\frac{\partial \xi_v}{\partial \zeta} = -\frac{\bar{Z}_{\xi\zeta}(\xi_v, \zeta)}{\bar{Z}_{\xi\xi}(\xi_v, \zeta)}. \quad (\text{D.5})$$

We combine this equation with two results in Schwarz & Südekum (2014). First, $\partial \xi_v^*/\partial \zeta < 0$. By continuity of $\bar{Z}_{\xi\xi}(\xi, \zeta)$ in ξ , this implies that in the vicinity of ξ_v^* we have $\partial \xi_v/\partial \zeta < 0$. Second, $\bar{Z}_{\xi\zeta}(\xi, \zeta) < 0$. In the light of (D.5), the two results together imply that in the neighborhood of ξ_v^* we have $\bar{Z}_{\xi\zeta}(\xi_v, \zeta) < 0$. The same argument implies that $\bar{Z}_{\xi\zeta}(\xi_v^*, \zeta) < 0$. Since ζ is slightly below ζ^* , ξ_v^* is strictly positive but small. The share ξ_v can be larger or smaller than ξ_v^* depending on whether F_v is larger or smaller than F_o . If $\xi_v < \xi_v^*$, then ξ_v is either at its lower bound with $\xi_v \rightarrow 0$, in which case both effects in (D.4) are equal to zero, or it is close to zero, in which case continuity implies that $\bar{Z}_{\xi\zeta}(\xi, \zeta) < 0$ for any $\xi \in [0, \xi_v]$. Consequently, the inframarginal effect and thus the derivative in (D.4) is negative. If $\xi_v > \xi_v^*$, then continuity implies $\bar{Z}_{\xi\zeta}(\xi, \zeta) < 0$ for any $\xi \in [0, \xi_v^*]$ as well as for any $\xi \in [\xi_v^*, \xi_v]$. Hence, the inframarginal effect is negative, and so is the derivative in (D.4).

It is straightforward to apply this reasoning to all values of ζ in the neighborhood of ζ^* . However, if we extend the analysis to values of $\zeta < \zeta^*$ beyond the vicinity of ζ^* , we may observe situations in which $\bar{Z}_{\xi\zeta}(\xi, \zeta)$ switches sign and is positive for low values of ξ . Hence, we cannot rule out that the inframarginal effect is positive. This is why the effect of productivity on $\Delta_\xi \bar{\Pi}(\zeta, \Theta)$ is potentially non-monotonic over the full interval $\zeta \in [0, 1]$, as stated in the proposition.

E Data appendix

Table E.1. Industries in ESEE data

CNAE-2009 Classification	Industry
101	Meat
102-109, 120	Food Products and Tobacco
110	Beverages
131-133, 139, 141-143	Textile
151-152	Leather & Footwear
161-162	Timber & Wooden Products
171-172	Paper Products
181-182	Graphics Design
201-206, 211-212	Chemical & Pharmaceutical Products
221-222	Plastic & Rubber Products
231-237, 239	Mineral Products (Non-Metal Products)
241-245	Ferrous Metals & Non-Ferrous Metals
251-257, 259	Metal Products
281-284, 289	Industry & Agricultural Machinery
261-268	Informatics, Electronics, Optics
271-275, 279	General & Electric Machinery
291-293	Motorized Vehicles
301-304, 309	Other Transportation Equipment
310	Furniture Industry
321-325, 329	Miscellaneous Manufacturing

Note: See http://www.ine.es/daco/daco42/clasificaciones/cnae09/estructura_en.pdf for individual products (or groups of products) belonging to each industry in ESEE data.

Table E.2. Summary statistics

Variable	Description	Obs.	Mean	SD
<u>Firm-specific variables</u>				
$\mathbb{1}(\xi_{it} > 0)$ (Domestic economy)	Dummy variable for a positive domestic integration share	12205	0.17	0.376
$\mathbb{1}(\xi_{it} > 0)$ (Foreign economy)	Dummy variable for a positive foreign integration share	5896	0.241	0.428
TFP	See Section 3.2	12982	-0.022	0.29
Capital intensity	Ratio of tangible fixed assets over the number of effective working hours (in logs)	12982	3.371	1.145
Skill intensity	Proxied by the average wage (in logs)	12982	2.842	0.412
R&D intensity	R&D expenditure over sales (in logs)	12945	0.29	0.581
Export intensity	Exports over sales (in logs)	12968	0.169	0.212
<u>Industry-specific variables</u>				
Supplier intensity	See Section 3.3	20	0.578	0.223
R&D intensity	Average firm-level R&D intensity across the years 2006-2015 (in logs)	20	-0.936	1.148
Skill intensity	Average share of workers with a Master's degree across the years 2006/10/14	20	1.691	0.483
Industry size	Average of total sales across the years 2006-2015 (in logs)	20	22.992	1.111

Note: All values are suitably deflated where necessary. We add one to the firm-level variables R&D intensity and export intensity before taking logs in order to keep zero observations. As for the skill intensity at the industry-level, we note that the share of workers with a Master's degree is not available every year in the data, but every four years.

F Multilateral index of TFP

In order to compute the multilateral index of TFP employed in the paper (see Equation (11)) we use the following variables: *Output*: measured by the value of the firm's annual sales of goods and services plus the value of the change in inventory of output goods (expressed in real terms using a firm-level price index computed from the survey data). *Labor input*: measured by the number of effective work-hours per year reported by the firm. *Material input*: measured by the sum of the firm's goods purchases (including energy and fuel costs as well as intermediate inputs but excluding any services) minus the change in inventory of these goods (expressed in real terms using a firm-level price index computed from the survey data). *Capital input*: computed according to the permanent inventory method, considering land, buildings, and other capital such as equipment and machinery (expressed in real terms using an industry-level price index available from the Spanish Instituto Nacional de Estadística (INE)). *Input cost shares*: measured by the cost share of each input in total input costs (i.e. total cost of labor, materials, and capital). The cost of labor is measured by the sum of gross wages and other compensations, social security contributions paid by the firm, the contributions made to supplementary pension systems, and other labor expenses. The cost of capital are estimated based on the following formula for the user cost of capital (UCC) for asset a :

$$UCC_{it}^a = \frac{P_t}{P_{it}^Q} (\tilde{r}_{it} + \delta^a - \pi_t^e), \quad (\text{F.1})$$

where P_t is the capital price deflator in year t , P_{it}^Q is the firm-level output price deflator in year t , \tilde{r}_{it} is a weighted average of the interest rate the firm pays on its long-term external debt (distinguishing between two sources of finance: financial institutions and other creditors), δ^a is the (economic) depreciation rate of asset category a (distinguishing between land & buildings and other capital goods), and π_t^e is the expected price inflation rate for capital goods (approximated by the actual price inflation rate for capital goods).

G Further empirical results

Table G.1. Estimates on a sample of affiliated firms—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	-0.042 (0.215)	-0.175 (0.608)
TFP \times Supplier intensity	0.302 (0.367)	1.044 (1.108)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.139 (0.554)	-0.168** (0.420)
N	4611	4611
R^2	0.140	0.122
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	0.141 (0.214)	0.760 (0.953)
TFP \times Supplier intensity	-0.142 (0.361)	-0.980 (1.649)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.995 (1.257)	-0.775 (0.451)
N	3162	3102
R^2	0.125	0.113

Notes: The estimation follows Table 1, but the estimation sample is restricted to firms majority-owned by some other firm in Spain or abroad. In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

Table G.2. IV estimation with firm fixed effects—Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy
	LPM
	(1)
TFP	0.360** (0.179)
TFP \times Supplier intensity	-0.541* (0.284)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.666*** (0.116)
Kleibergen-Paap Wald rk F statistic	12.36
Hansen J statistic	3.123
p-value	(0.537)
N	10237
	Panel B: Foreign economy
	LPM
	(1)
TFP	0.475* (0.271)
TFP \times Supplier intensity	-0.627 (0.501)
Estimates of λ_0/λ_1 ($H_0: \lambda_0/\lambda_1 = -1$)	-0.758 (0.235)
Kleibergen-Paap Wald rk F statistic	13.30
Hansen J statistic	1.153
p-value	(0.886)
N	5023

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. TFP and TFP \times Supplier intensity are instrumented with their first three lags. All regressions include firm fixed effects, industry fixed effects, and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.

Table G.3. Controlling for industry size A —Dependent variable: $\mathbb{1}(\xi_{it} > 0)$

	Panel A: Domestic economy	
	LPM	Probit
	(1)	(2)
TFP	0.291* (0.154)	0.735 (0.748)
TFP \times Supplier intensity	-0.405*** (0.155)	-1.040 (0.730)
TFP \times Industry size	0.0741 (0.114)	0.794 (0.683)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.717* (0.148)	-0.707 (0.416)
N	12205	12205
R^2	0.075	
	Panel B: Foreign economy	
	LPM	Probit
	(1)	(2)
TFP	0.410** (0.186)	4.263*** (1.409)
TFP \times Supplier intensity	-0.544*** (0.201)	-4.462*** (1.504)
TFP \times Industry size	0.0254 (0.153)	-1.444 (0.984)
Estimates of λ_0/λ_1 ($H_0 : \lambda_0/\lambda_1 = -1$)	-0.754 (0.155)	-0.955 (0.131)
N	5896	5754
R^2	0.080	

Notes: In Panel A, the dependent variable, $\mathbb{1}(\xi_{it} > 0)$, is an indicator function equal to one if the firm sources at least *some* domestic inputs through vertical integration, and zero if *all* domestic inputs come from independent suppliers (full outsourcing). The same definition applies in Panel B, but for inputs sourced from the foreign rather than the domestic economy. The ratio λ_0/λ_1 is the ratio of coefficients of TFP and TFP \times Supplier intensity. All regressions include industry fixed effects and year fixed effects. Robust standard errors (in parentheses) are clustered at the firm-level. *, **, *** denote estimates significantly different from zero (or significantly different from minus one in the case of λ_0/λ_1) at the 10%, 5%, 1% levels, respectively.