Incomplete Contracts and the Boundaries of the Multinational Firm

(Preliminary and Incomplete)

Nathan Nunn[∗‡]
Harvard University and NBER

Daniel Trefler[∗§]
University of Toronto, CIFAR and NBER

June 2008

ABSTRACT: [To be written]

Key words: Intra-firm trade; Incomplete contracts; Vertical integration; Outsourcing; Fixed-costs

JEL classification: F14, F23, L14, L33

[∗]Both authors thank Pol Antràs , Jan De Leocker, Rob Feenstra, Elhanan Helpman, Gene Grossman, Wilhelm Kohler, Marc Melitz, and seminar participants at the NBER. We are especially grateful to Joel Blit for facilitating construction of the data base.

[‡]Department of Economics, Harvard University, 1805 Cambridge Street, Cambridge, MA 02138, U.S.A. (e-mail: `nnunn@fas.harvard.edu`; website: `http://www.economics.harvard.edu/faculty/nunn`).

[§]Rotman School of Management and Department of Economics, University of Toronto, Toronto, Ontario, M5S 3E6, Canada (e-mail: `dtrefler@rotman.utoronto.ca`; website: `http://www.rotman.utoronto.ca/~dtrefler`).

1. Introduction

Recently, a rich research agenda examining the determinants of intra-firm trade has developed. Unlike the older literature on international trade in the presence of imperfect competition e.g., Helpman and Krugman (1985), the current literature provides a much more intellectually satisfying notion of what constitutes a firm. It thus provides us with deeper insights into which elements of international trade are done internally to the firm (multinational or intra-firm trade) and which are done outside the boundaries of the firm. Seminal contributions include McLaren (2000), Antràs (2003, 2005), Grossman and Helpman (2002, 2003, 2004, 2005), and Antràs and Helpman (2004).

In this paper we are particularly interested in a strand of the literature that examines the relationship between a multinational firm and its supplier. Each contributes a customized input that is non-contractible. As a result, there is a classic hold-up problem and the multinational must decide whether to vertically integrate its supplier or outsource to its supplier. One narrow strand of the literature – the one we will be dealing with – treats the difference between these two organizational forms as the difference between the outside options of the multinational in the event that the hold-up problem cannot be resolved through bargaining. This treatment of the difference between vertical integration and outsourcing originates with Antràs (2003) and appears again in Antràs and Helpman (2004, 2008).

These papers yield three important insights about the determinants of the share of total U.S. imports that are intra-firm i.e., the share that is imported by U.S. multinationals from their foreign affiliates. First, Antràs (2003) argues that when the U.S. headquarters firm provides the bulk of the non-contractible inputs, underinvestment in inputs is reduced by highly incentivizing the headquarters firm. Vertical integration provides such incentives because it allows the headquarters firm to control at least some of the supplier's inputs even if bilateral bargaining breaks down. In contrast, when the foreign supplier provides the bulk of the non-contractible inputs, the foreign supplier must be highly incentivized. This is done by outsourcing: outsourcing strips the headquarters firm of any control over the supplier's inputs and thus strengthens the bilateral bargaining position of the supplier. In short, the share of U.S. imports that are intra-firm is increasing in the share of (non-contractible) inputs provided by the U.S. headquarters firm. This logic is a specific instance of the larger property rights approach to the firm e.g., Grossman and Hart (1986).

A second prediction about intra-firm trade appears in (Antràs and Helpman, 2008). While

Antràs (2003) and Antràs and Helpman (2004) assume that inputs are completely non-contractible, Antràs and Helpman (2008) allow inputs to be partially contractible. According to their model what now matters is not only headquarter intensity, but headquarter intensity and the contractibility of headquarter investments. They show that it is the interaction of headquarter intensity with the share of investments that are non-contractible that matters once partial contractibility of inputs is allowed for. The logic of this is a straightforward extension of Grossman-Hart-Moore. Previously, a higher headquarter intensity increased the need to incentivize the headquarter and the result was vertical integration an more intra-firm trade. When some investments are contractible it is now only necessary to only incentivize the non-contractible investments. Therefore, now what matters is the share of headquarter investments that are non-contractible relative to supplier investments that are non-contractible. That is, what matters is headquarter intensity multiplied by the share of headquarter inputs that are not contractible. An increase in this measure, increases the need to incentivize the headquarter and increases the chance of vertical integration and of intra-firm trade.

The final prediction about the share of total U.S. imports that is intra-firm is developed in Antràs and Helpman (2004). They start with the well-known fact that firms display heterogeneous productivities e.g., Bernard and Jensen (1997). Antràs and Helpman (2004) also argue that the fixed costs of producing abroad are lower when outsourcing to a foreign supplier than when using foreign direct investment (vertical integration). Since only the most productive firms capture the market share needed to offset the high costs of vertical integration, not all firms identified by Antràs (2003) as candidates for vertical integration will in fact integrate. Only the most productive will. Thus, the share of U.S. imports that are intra-firm will be large when two conditions are simultaneously satisfied: (*i*) the share of inputs provided by the headquarters firm is large (as in Antràs 2003) and (*ii*) firm productivity is high.

Using data on U.S. intra-firm and arm's-length imports for 5,705 products imported from 220 countries, we examine these determinants of the share of U.S. imports that are intra-firm. Our conclusions mirror the three predictions listed above. (*i*) In terms of the Antràs (2003) mechanism, we find support for the role of the share of headquarter inputs. This support is stronger than that found in the only two extant empirical studies of the issue, namely, Antràs (2003) and Yeaple (2006).[1] (*ii*) We find strong support for the Antràs and Helpman (2004) prediction that intra-firm

[1] Feenstra and Hanson (2005) also provide support for the property-rights approach in an international trade context.

trade is largest where headquarter inputs are important *and* productivity is high. (*iii*) We also find support for the Antràs and Helpman (2008) prediction about increased internalization due to an improved contracting environment for the supplier's inputs.

The paper is organized as follows. Section 2 examines the predictions of Antràs (2003) and Antràs and Helpman (2008), and section 3 examines the predictions from Antràs and Helpman (2004). Section 4 concludes.

2. The Boundary of the Firm and the Role of η (Antràs, 2003; Antràs and Helpman, 2008)

We begin by reviewing the salient features of the Antràs (2003) and Antràs and Helpman (2004, 2008) models from the perspective of the empirical work to follow. Since the models have been presented and summarized elsewhere, we only review its most important features here.[2]

A U.S. firm produces a brand of a differentiated variety j of a product in industry i. Demand is generated by CES preferences. To produce the good, the firm *must* use two inputs, those produced by the U.S. firm ($h_i(j)$ for *h*eadquarters) and those produced by a foreign supplier ($m_i(j)$ for inter*m*ediates). Output of the final good is given by a Cobb-Douglas production function with two key parameters: a Hicks-neutral productivity parameter $\theta(j)$ that is variety (i.e., match) specific and the cost share of the input provided by the firm η_i, which is an industry-specific parameter. Specifically, production is given by:

$$q_i(j) = \theta(j) \left(\frac{h_i(j)}{\eta_i} \right)^{\eta_i} \left(\frac{m_i(j)}{1-\eta_i} \right)^{1-\eta_i} . \tag{1}$$

In Antràs (2003) and Antràs and Helpman (2004) it is assumed that the two inputs are entirely customized and not contractible. This assumption is relaxed in Antràs and Helpman (2008). Customization raises quality to a threshold which allows the final good to be sold to consumers. Unfortunately, for the U.S. firm and its foreign supplier, the investments in customization are non-contractible and they have no value outside of the relationship. Thus, there is a standard hold-up problem. After the investments in customization have been made there is renegotiation over how the *ex post* quasi-rents from the relationship will be shared.

[2] See Helpman (2006) and Nunn and Trefler (2008).

The timing of the game played by the U.S. firm and its foreign supplier is as follows. The two match and the U.S. firm chooses the organizational form. Then investments in customized inputs are made. Finally, the initial contract is renegotiated and, if there is agreement, the product is sold.

Let β be the generalized Nash share of the *ex post* quasi-rents from the relationship that go to the U.S. firm. The U.S. firm receives this share plus its outside option. The organizational form – vertical integration versus outsourcing – chosen by the headquarter alters the outside option of the U.S. firm. Regardless of the organization form, the supplier earns nothing and its outside option is 0. This is also the case for the U.S. firm under outsourcing. However, for a headquarter that has vertically integrated with its supplier, if bargaining breaks down, then the headquarter can still produce some output by 'forcing' its now-disgruntled supplier to do at least some work. Vertical integration is therefore a way for the firm to improve its outside option.

This difference in the headquarter's outside options under the two organizational forms leads to a trade-off. Although vertical integration allows the firm to grab a larger share of the pie, it also potentially leads to a smaller pie because of increased underinvestment by the supplier. This is modeled mathematically as follows. Let $k = V, O$ subscripts denote the organizational form with V for vertical integration and O for outsourcing. Recall that β is the share of the *ex post* quasi-rents that goes to the firm. Let $R_k(j)$ be the revenue generated when there is an agreement. If there is no agreement the firm can only sell a portion δ of the final output. With CES preferences and constant markup $1/\alpha$, this generates a revenue of $\delta^\alpha R_V(j)$. Therefore, the firm receives its outside option $\delta^\alpha R_V(j)$ plus a share β of the quasi-rents $(R_V(j) - \delta^\alpha R_V(j))$. That is, the firm receives $[\delta^\alpha + \beta(1 - \delta^\alpha)]R_V(j)$. Let $\beta_V = \delta^\alpha + \beta(1 - \delta^\alpha)$ be the firm's share of revenues under vertical integration. Under outsourcing, the outside option is 0 and the quasi rents are $R_O(j)$ so that the firm receives $0 + \beta(R_O(j) - 0) = \beta R_O(j)$. Let $\beta_O = \beta$ be the firm's share of total revenues under outsourcing. The upshot of all this is the central result that the organizational form alters the U.S. firm's share of revenue. In particular, $\beta_V > \beta_O$.

Both the U.S. firm and the foreign supplier invest and hence each must worry about the other's underinvestment. Where η_i is large, the surplus generated by the relationship is particularly sensitive to the amount of investment undertaken by the U.S. firm. To reduce the degree of underinvestment by the U.S. firm, the firm must be given a large share of the revenue. This share is largest under vertical integration because $\beta_V > \beta_O$. This is a specific instance of the Grossman and Hart (1986) property-rights theory of the firm where residual control rights are allocated to the

U.S. firm. In contrast, when η_i is small, the surplus generated by the relationship is particularly sensitive to the amount of investment undertaken by the supplier. To reduce supplier underinvestment, the supplier must be given a large share of the revenue. Outsourcing accomplishes this because $1 - \beta_O > 1 - \beta_V$.

In Antràs (2003) it is assumed that $\theta(j) = 1$ for all headquarter-supplier pairs. The results when this assumption is relaxed is analyzed in Antràs and Helpman (2004). We describe this case in section 3.

With the assumption that $\theta(j) = 1$, Antràs (2003) shows that there is a unique value of η_i – call it η_c – such that the U.S. firm prefers vertical integration for $\eta_i > \eta_c$ and prefers outsourcing otherwise.

Hypothesis 1 *There exists a unique cut-off η_c with the following property. If $\eta_i > \eta_c$ then the firm will vertically integrate with the supplier. If $\eta_i < \eta_c$ then the firm will outsource from the supplier.*

We begin by testing hypothesis 1. The following section describes the data used to test hypothesis 1 and subsequent hypotheses which are derived from Antràs and Helpman (2004, 2008).

Data Sources

To investigate hypothesis 1 we use data on intra-firm and total trade from the U.S. Census Bureau. Importers bringing goods into the United States are required by law to report whether or not the transaction is with a related party. This information allows us to identify whether imports are intra-firm (related party) or at arm's-length (non-related party). See the appendix for details. The trade data are at the 6-digit Harmonized System (HS6) level for the year 2005. We are grateful to Andy Bernard for drawing our attention to these data. See Bernard, Jensen, and Schott (2005) for an example of how the data have been used.

Our key dependent variable is intra-firm imports as a share of total U.S. imports. Let i index industries and let M_i^V be the value of intra-firm U.S. imports in industry i. The V superscript is for vertical integration. Let M_i^O be the value of arm's-length U.S. imports in industry i. The O subscript is for outsourcing. $M_i^V + M_i^O$ is total U.S. imports and $M_i^V/(M_i^V + M_i^O)$ is intra-firm imports as a share of total U.S. imports in industry i.

Antràs (2003) takes a stand on how we measure the headquarter intensity of an industry η_i, arguing that capital investments can be provided by the headquarter, but not labor investments. In subsequent papers, the headquarter intensity has also been interpreted to also potentially include inputs such as R&D, advertising, and managerial skill (see Antràs and Helpman, 2004; Antràs and Helpman, 2008; Yeaple, 2006). At the outset, we do not take a stance on which characteristics best measure inputs provided by the headquarter. Instead, we construct measures of each type of input and let the data speak, indicating which measure produces results that are consistent with the model. We construct three measures that potentially capture the headquarter intensity of an industry. The first is the original measure proposed in Antràs (2003), capital intensity. We also construct measures of skill-intensity and R&D intensity to capture managerial inputs and R&D inputs, both of which are potentially inputs provided by the headquarter. As a test of the model we also construct a factor intensity measure for which it is likely that the input is *not* provided by the headquarter. The measure is material intensity.

Capital, skill, and material intensities are constructed using data on the factor intensity of production in each industry, which are from the U.S. Census Bureau's 2005 *Annual Survey of Manufactures*. We use U.S. factor intensities, assuming that they are correlated with the factor intensity of production on other countries. For each 6-digit NAICS industry we collect information on annual capital expenditures, wages of production workers and non-production workers, and total expenditures on materials. Using this information we construct measures of capital intensity denoted K_i/L_i, skill-intensity S_i/L_i and material intensity M_i/L_i. Capital intensity K_i/L_i is measured as the natural log of capital expenditures divided by all worker wages. Similarly, material intensity M_i/L_i is measured as log material expenditures divided by workers wages. Skill intensity S_i/L_i is the log ratio of non-production worker wages to total worker wages.

Because the *Annual Survey of Manufactures* does not include information on R&D expenditures, the data used to construct R&D intensity are from a different data source. This information is obtained from the *Orbis* database, which a database that has information on over 30 million companies worldwide. Data are collected from over 40 different information providers. Both private and public companies are listed in the database, with over 99% of the companies being private. The database is constructed by Bureau van Dijk Electronic Publishing. We measure R&D intensity, which we denote RD_i/Q_i, by the natural log of global R&D expenditures divided by firm sales in each industry.[3] The data are from the most recent for which firm level data are available,

[3]If we do not take the natural log of K_i/L_i, S_i/L_i, and RD_i/Q_i, these measures are left skewed. Taking the natural log results in a distributions that are more normally distributed.

typically either 2006 or 2007.

Examining Hypothesis 1 (Antràs, 2003)

Antràs (2003) examined hypothesis 1 using BEA data on intra-firm U.S. imports as a share of total U.S. imports. As we noted above, he related this share to capital intensity, his proxy for η_i. Antràs (2003) worked at the 2-digit SIC level with 28 industries. We start by examining his relationship with roughly 300 NAICS 6-digit industries. In particular, we consider the following cross-industry regression:

$$\frac{M_i^V}{M_i^V + M_i^O} = +\gamma_H \frac{H_i}{L_i} + \gamma_R \frac{RD_i}{Q_i} + \gamma_M \frac{M_i}{L_i} + \gamma_K \frac{K_i}{L_i} + \varepsilon_i \quad (2)$$

where $\frac{M_i^V}{M_i^V + M_i^O}$ is the share of U.S. imports in industry i that are intra-firm; H_i/L_i is log of Non-production worker wages to total worker wages; RD_i/Q_i is the log of R&D expenditures divided by sales; M_i/L_i is the log of expenditures on materials divided by total worker wages; K_i/L_i is the log of capital expenditures divided by total worker wages

Estimates of equation (2) appear in table 1. In all specifications the coefficient for capital intensity is positive and statistically significant. The capital intensity result confirms the findings of Antràs (2003), Yeaple (2006) and Bernard, Jensen, Redding, and Schott (2008). As well, the estimated coefficients for skill-intensity and R&D intensity are also positive and statistically significant. The positive coefficient for R&D intensity confirms the finding of Antràs (2003) and Yeaple (2006). The positive coefficient for skill-intensity differs from Antràs (2003) and Yeaple (2006) who finds no statistically significant relationship. However, the finding is consistent with the positive coefficient for skill-intensity found by Bernard *et al.* (2008). The difference in the results may be that Antràs (2003) and Yeaple (2006) use data from a much smaller sample of countries and with much more aggregate data than either us or Bernard *et al.* (2008).

Because we report standardized 'beta' coefficients, one can easily assess and compare the magnitudes of the coefficients for the capital, skill and R&D intensity measures. According to the estimates of column (1), a one standard deviation increase in capital results in a .167 standard deviation increase in the share of intra-firm imports. This is an economically large effect. The estimated coefficients for skill and R&D intensity, which are even larger, are .172 and .318, respectively.

Recall that in Antràs (2003) (as well as in subsequent models) the input being produced $q_i(j)$ is a customized variety j of a good in industry i. Because the input is specific to the variety,

Table 1. Headquarter intensity and intra-firm trade: testing Antràs (2003), looking across industries.

	Dependent variable: $M^V_i/(M^V_i+M^O_i)$				
	(1)	(2)	(3)	(4)	(5)
ln R and D / Sales	0.318***	0.310***	0.212***	0.273***	0.290***
ln Skilled Labor / Worker	0.172***	0.208***	0.236***	0.302***	0.240***
ln Materials / Worker	0.098	0.091	0.173**	0.155**	0.023
ln Capital / Worker	0.167***				
ln Buildings / Worker		-0.111	0.005	-0.029	
ln Machinery / Worker		0.269***			
ln Computers / Worker			0.115		
ln Autos / Worker			-0.245***		
ln Other Machinery / Worker			0.197**	0.207**	0.206**
ln Non-Specific Machinery / Worker				-0.072	
ln Non-Specific Capital / Worker					-0.061
Number of Observations	298	294	200	200	209
R-Squared	0.20	0.23	0.31	0.27	0.24

Notes: The dependent variable $M^V_i/(M^V_i+M^O_i)$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an 6-digit NAICS industry. Standardized 'beta' coefficients are reported. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Non-Specific Machinery' is Computers+Autos. 'Non-Specific Capital' is Buildings+Computers+Autos. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos.

and therefore to the headquarter-supplier match, it is assumed that investments in h and m are also relationship-specific. That is, they have no value outside of the match. This is an important assumption about the investments in the production of the input which has been taken for granted in previous empirical work such as Yeaple (2006) and Nunn and Trefler (2008). In this section we examine this assumption by disaggregating capital into capital with a low value outside of the relationship (relationship-specific) and capital with higher value outside of the relationship (less relationship-specific). To do this we again rely on data from the 2005 *Annual Survey of Manufactures*. The survey reports total capital divided into the following categories: (*i*) buildings, (*ii*) computers, (*iii*) automobiles, and (*iv*) other machinery.

Of these categories of capital, buildings, computers, and automobiles have a higher outside value than other forms of capital i.e., other machinery. Buildings can be resold and used for the production of other goods. This is in contrast to specialized machinery, which has significantly less use outside of its intended production process. Similarly, computers and automobiles, can be resold and have a use outside of the relationship. Computers include standard desktop computers, and automobiles include only standard vehicles that are driven on roads. That is, the category does not include specialized vehicles. Both computers and automobiles have much greater outside uses relative to machinery used directly in the production.

We construct multiple measures of capital intensity using the capital expenditure data form the 2005 *ASM* and include these in equation (2). The results are reported in table 1. Each capital intensity measure is constructed as the natural log of the relevant capital expenditures divided by total worker wages.

We begin by divided capital between buildings and machinery, which is total capital minus buildings. If capital is an input provided by the headquarter as asserted by Antràs (2003), and if machinery is relationship-specific, then a greater machinery intensity to increase the need to incentivize the headquarter, and we should see more vertical integration and intra-firm imports in the industry. As reported in column 2, this is indeed what we observe in the data. The coefficient for machinery intensity is positive and statistically significant. Note however, that if buildings are not relationship-specific as we argue, then an increase in the building intensity of an industry does not increase the need to incentivize the headquarter and therefore it should not increase the share of U.S. imports that are intra-firm. This is what we observe in the data. The coefficient for building intensity is not positive and significant like machinery intensity. Instead, it has not statistically significant effect on the share of imports that are intra-firm.

We examine this result further in column (3), where we include more fine measures of capital intensity. We include separately variables measuring the factor intensity of buildings, computers, and automobiles and other machinery (total capital minus buildings, computers and autos) in the estimating equation. Because the value of buildings, computers and automobiles outside of the relationship is likely higher than other forms of capital, we do not expect a positive relationship between these factor intensities and the share of trade that is intra-firm. The estimation results are consistent with this. Only the estimated coefficient for other machinery intensity is positive and statistically significant.

In columns (4) and (5) we report alternative specifications. In column (4) we aggregate computers and automobiles and include this measure, which we label 'Non-specific machinery', along with buildings and other machinery in the estimating equation. In column (5) we include an aggregated measure, called 'Non-specific capital', which is the sum of buildings, computers and automobiles, along with other machinery in the estimating equation. In both specifications, coefficient for residual capital intensity measure 'other machinery' is the only coefficient that is positive and statistically significant.

Overall, the results of columns (2)–(5) of table 1 show that only for the relationship-specific capital intensity measure do we find the predicted positive relationship with the share of imports that are intra-firm.

A problem with our cross-industry specification is that it assumes that we can aggregate across exporting countries. Yet as Schott (2004) notes, this may be seriously misleading because an HS6 good produced in a poor country may be very different from an HS6 good produced in a rich country. To address this, we turn to a different approach from Antràs (2003) and Yeaple (2006). Let M_{Vic} be the value of U.S. intra-firm imports of good i that are imported from country c. Let M_{ic}^{O} be the corresponding value of arm's-length U.S. imports. Then $M_{ic}^{V} + M_{ic}^{O}$ is total U.S. imports of good i from country c and $M_{ic}^{V}/(M_{ic}^{V} + M_{ic}^{O})$ is the intra-firm import share of good i imported from country c. We estimate a regression that pools across industries and countries:

$$\frac{M_{ic}^{V}}{M_{ic}^{V} + M_{ic}^{O}} = \gamma_c + \gamma_H \frac{H_i}{L_i} + \gamma_R \frac{RD_i}{Q_i} + \gamma_M \frac{M_i}{L_i} + \gamma_K \frac{K_i}{L_i} + \varepsilon_{ic} \tag{3}$$

where $\frac{M_{ic}^{V}}{M_{ic}^{V}+M_{ic}^{O}}$ is the share of U.S. imports that are intra-firm; γ_c denotes country fixed effects; and as before H_i/L_i is log of Non-production worker wages to total worker wages; RD_i/Q_i is the log of R&D expenditures divided by sales; M_i/L_i is the log of expenditures on materials divided by total worker wages; K_i/L_i is the log of capital expenditures divided by total worker wages

In equation (2) we control for exporter heterogeneity by allowing for country fixed effects γ_c. Further, i now subscripts HS6 products rather than 6-digit NAICS industries.

Estimates appear in table 2. Because our variables of interest K_i/L_i, RD_i/Q_i and S_i/L_i only vary at the 6-digit NAICS industry level, while the unit of observation is a country and HS6 good, we report standard errors clustered at the 6-digit NAICS level. The estimates, consistent with the cross-industry estimates of table 1, show that the capital, R&D, and skill intensity of an industry are positively correlated with the share of intra-firm trade. The estimates in columns (2)–(5) further show that this positive relationship is driven solely by the part of capital which is more relationship-specific. Overall, these results combined with the cross-industry results of table 1 provide support for hypothesis 1.

Allowing for partially contractible investments (Antràs & Helpman, 2008)

In Antràs and Helpman (2008) the assumptions of complete relationship-specificity and non-contractibility of investments are relaxed. Instead, it is assumed that a fraction μ_h and μ_m of the inputs provided by the headquarter and supplier are either not relationship-specific or are contractible. (Note that only one of these two conditions need to be met to alleviate underinvestment.) In other words, for the fraction $1-\mu_h$ and $1-\mu_m$ of h and m inputs (respectively), the investments are relationship-specific and ex post enforceable contracts cannot be written.

In equation (10) of Antràs and Helpman (2004), the headquarter's profit maximizing bargaining share as a function of η is given. Slight manipulation of this expression yields:

$$\beta_h^* = \frac{\eta_h(1-\alpha\eta_m) - \sqrt{\eta_h\eta_m(1-\alpha\eta_h)(1-\alpha\eta_m)}}{\eta_h - \eta_m} \tag{4}$$

This can be compared to the optimal bargaining share from Antràs and Helpman (2008), where investments are partially contractible. In Antràs and Helpman (2008), the optimal bargaining share is given by:

$$\beta_h^* = \frac{\omega_h(1-\alpha\omega_m) - \sqrt{\omega_h\omega_m(1-\alpha\omega_h)(1-\alpha\omega_m)}}{\omega_h - \omega_m} \tag{5}$$

where $\omega_h \equiv (1-\mu_h)\eta$ and $\omega_m \equiv (1-\mu_m)(1-\eta)$.

Comparing equations (4) and (5), it is clear that the expression for the headquarter's optimal bargaining share has the same functional form in the two papers except that when the fraction μ_h

Table 2. Headquarter intensity and intra-firm trade: testing Antràs (2003), looking across industries and countries.

	Dependent variable: $M^V_{ic}/(M^V_{ic}+M^O_{ic})$				
	(1)	(2)	(3)	(4)	(5)
ln R and D / Sales	0.043***	0.062***	0.051***	0.064***	0.068***
ln Skilled Labor / Worker	0.074***	0.090***	0.121***	0.130***	0.119***
ln Materials / Worker	0.029*	0.023	0.019	0.018	0.023
ln Capital / Worker	0.043***				
ln Buildings / Worker		-0.046***	-0.034*	-0.039*	
ln Machinery / Worker		0.078***			
ln Computers / Worker			0.000		
ln Autos / Worker			-0.052***		
ln Other Machinery / Worker			0.095***	0.087***	0.077***
ln Non-Specific Machinery / Worker				-0.031	
ln Non-Specific Capital / Worker					-0.047**
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Observations	119,621	118,973	92,611	92,611	94,749
Number of Clusters	298	294	200	200	209
R-Squared	0.12	0.12	0.13	0.13	0.13

Notes: The dependent variable $M^V_{ic}/(M^V_{ic}+M^O_{ic})$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an HS6-country pair. All equations include country fixed effects. Standardized 'beta' coefficients are reported. Standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Non-Specific Machinery' is Computers+Autos. 'Non-Specific Capital' is Buildings+Computers+Autos. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos.

of headquarter investments are contractible then η is replaced with $(1-\mu_h)\eta$. Similarly, $1-\eta$ is also replaced with $(1-\mu_m)(1-\eta)$.

The intuition for this is straightforward. In Antràs (2003) η increases vertical integration because because it increases the importance of the headquarter's underinvestment and the need to incentivize the headquarter. When part of the headquarter's investments are contractible, then η no longer measures the underinvestment of headquarter. This is given by the fraction of investments that are non contractible $(1-\mu_h)$ multiplied by η.

Assume for the moment that, as in Antràs (2003), the fixed costs of vertical integration and the fixed costs of outsourcing are equal. The only effect that the choice of organizational form has is on β, the share of the ex post surplus that accrues to the headquarter. From Antràs and Helpman (2008) we then have the following hypothesis.

Hypothesis 2 *There exists a unique cut-off η_{hc} with the following property. If $\eta_i > \eta_{hc}$ then the firm will vertically integrate with the supplier. If $\eta_i < \eta_{hc}$ then the firm will outsource from the supplier. Further the cut-off η_{hc} is higher the higher is μ_h and the lower is μ_m.*

From hypothesis 2 it follows that as before the share of trade that is intra-firm should be increasing in η. As well, the share of intra-firm trade should be decreasing in the share of headquarter-investments that are not contractible μ_h. We test this prediction of Antràs and Helpman (2008) by focusing on the contractibility of headquarter investments, μ_h. Using our disaggregated capital expenditure data from the 2005 *ASM*, we construct two industry-specific measures of μ_h. The first measure of μ_i, which we denote $x_i^{\mu 1}$ is that fraction of capital expenditures on buildings:

$$x_i^{\mu 1} = \frac{\text{Buildings}_i}{\text{Total Capital}_i}.$$

The second measure is the fraction of expenditures on buildings, automobiles, and computers:

$$x_i^{\mu 2} = \frac{\text{Buildings}_i + \text{Autos}_i + \text{Computers}_i}{\text{Total Capital}_i}.$$

The measure of contractibility of capital investments is different from other measures of contractibility that exist in the literature. The measure of contractibility in Nunn (2007) is the share of intermediate inputs used in production that can be bought and sold in thick markets. Bernard *et al.* (2008) use the share of U.S. imports controlled by wholesale and retail firms, measured at the HS10 level. The measures constructed here are not measures that attempt to quantify how 'contractible' goods in an industry are. Therefore it is conceptually very different from the measure from Bernard

Table 3. Partially contractible capital and intra-firm trade: testing Antràs and Helpman (2008).

Dep var: $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$	Industry-Country Regressions		Industry Regressions	
	(1)	(2)	(3)	(4)
ln Capital/Worker, K_i/L_i	0.035***	0.042**	0.157**	0.134**
Buildings/Total Capital, $x_i^{\mu 1}$	-0.150***		-0.114**	
(Buildings+Autos+Computers)/Total Capital, $x_i^{\mu 2}$		-0.001		-0.150***
All Control Variables	Yes	Yes	Yes	Yes
Number of Observations	119,621	119,621	298	298
Country Fixed Effects	Yes	Yes	n/a	n/a
Number of Clusters	298	298	n/a	n/a
R-Squared	0.12	0.12	0.22	0.20

Notes: The dependent variable $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an HS6-country pair. Standardized 'beta' coefficients are reported. In columns 1 and 2, standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. The `Industry-Country Regressions' of columns 1 and 2 include country fixed effects. `All Control Variables' includes H_i/L_i, M_i/L_i, and RD_i/Q_i.

et al. (2008). Instead the measures here quantify the share of capital investments typically made in an industry that have value outside of the relationship i.e., that are not relationship-specific. It is more similar to the measure from Nunn (2007), which quantifies the relationship-specificity of the intermediate inputs used in the production process. Similar to this measure, $x_i^{\mu 1}$ and $x_i^{\mu 2}$ quantify the relationship-specificity of capital used in the production process.

Estimates of our baseline equations (2) and (3), with the x_i^{μ} measures included, are reported in table 3. Columns (1) and (2) report estimates of the country-industry level regression equation (3) and columns (3) and (4) report the estimates of industry level equation (2).

Overall, the estimation results are consistent with the predictions of Antràs and Helpman (2008). In all specifications, the estimated coefficients for capital intensity are positive and statistically significant, and the coefficients for the x_i^{μ} variables are negative and significant in three of the four specifications.

Who's the Parent and Who's the Subsidiary?

The results of Antràs (2003) and Antràs and Helpman (2004, 2008) are based on a specific environment. A key assumption is that the intermediate input is produced in the foreign subsidiary's country and is then shipped to the headquarter country. We have been assuming that the U.S. is the headquarter and the foreign country is the subsidiary, and that when we observe intra-firm imports into the U.S. these are imports being shipped from the foreign subsidiary to the U.S. headquarter. In reality, these imports could be imports being shipped from a foreign parent to a U.S. subsidiary. If this is the case, then this environment is very different from that modeled in Antràs (2003) and Antràs and Helpman (2004, 2008).

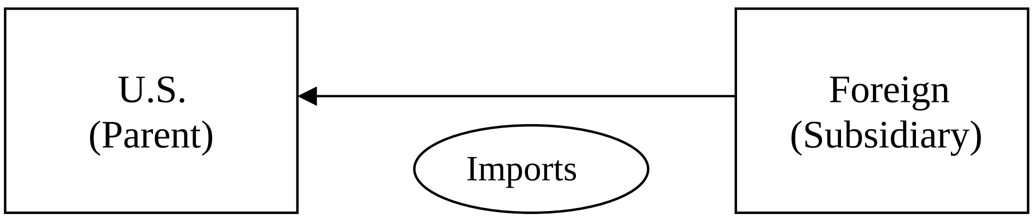

Figure 1. U.S. Imports: The ownership structure consistent with Antràs (2003) and Antràs and Helpman (2004, 2008).

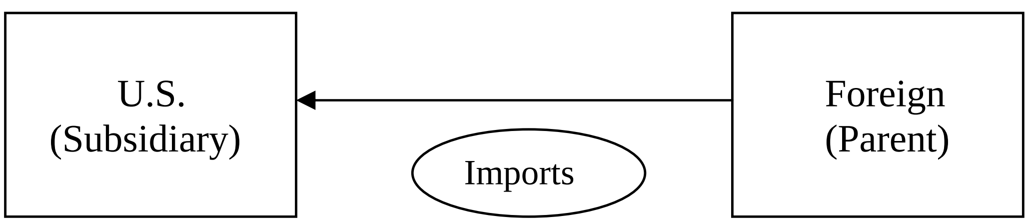

Figure 2. U.S. Imports: The ownership structure not consistent with Antràs (2003) and Antràs and Helpman (2004, 2008).

If, for example, the good is produced by a foreign headquarter and shipped to a U.S. subsidiary, then the crucial assumption that the headquarter can only provide headquarter inputs when producing the input is no longer realistic. In Antràs (2003) it is assumed that the headquarter cannot provide labor inputs because it is not familiar with foreign labor markets and because it is difficult to have a managerial presence in the foreign country (see Antras, 2003, p. 1379). But if the input is being produced in the headquarter's country then these arguments no longer apply. There is no reason that the headquarter cannot provide *all* of the inputs required to produce the intermediate input. It is also no longer clear that capital is the input provided by the headquarter, and therefore there is no longer a reason to expect capital intensity to be correlated with vertical integration.

Figures 1 and 2 report the two possible scenarios, when examining U.S. imports. Figure 1 shows the case where goods are being imported from a foreign subsidiary to a U.S. headquarter. This is the environment being modeled in Antràs (2003) and Antràs and Helpman (2004, 2008). Figure 2 shows the other case where a good is being shipped from a foreign parent to a U.S. subsidiary. This case, we have argued does not fit the models being examined here.

To distinguish empirically between the two cases shown in figures 1 and 2, we use data from *Orbis*. The database reports plants and subsidiaries that are part of global multinational firms. The database provides information about the name and country of the headquarter and subsidiaries that are listed, as well as the industry that they are classified under. Using this information we identify all subsidiary headquarter pairs in which either the subsidiary or the headquarter are from the U.S. Pairs in which both the subsidiary and the headquarter are from the U.S. are excluded. Therefore, within the subsidiary-headquarter pairs examined, one of the two is from a non-U.S. country. We then divide the pairs by foreign country and calculate for each foreign country (*i*) the number of pairs for which the headquarter is from the U.S., (*ii*) the number of pairs for which the subsidiary is from the U.S., and (*iii*) the fraction of pairs for which the U.S. is the headquarter. This information is reported in table 4 for the countries with the share of pairs for which the U.S. is the parent below 75%. As shown, only 18 countries fall into this category. This indicates that for the vast majority of countries, the U.S. is generally the headquarter. The countries for which this is not the case tend to be developed countries, such as Italy, Sweden, Switzerland and Japan.

We use the information from table 4 to identify countries for which intra-firm imports are more likely to be from a foreign subsidiary to a U.S. parent (the case from figure 2. We construct a restricted sample of countries for which the share of U.S. headquarters is above 50%. This amounts to removing Iceland, Italy, Finland, Liechtenstein, and Switzerland from the sample.

Table 5 reproduces the estimates of table 8 using this more restricted sample of countries. As shown, the results are qualitatively identical.[4]

An alternative strategy is to calculate the fraction of partnerships for which the U.S. is the

[4]The results are also similar if a more restrictive samples are examined. For example, if one omits all 18 countries the results are very similar.

Table 4. Nationality of the Parent by Country.

Country	Number of Relationships with: A U.S. Parent	Number of Relationships with: A Foreign Parent	Share of Relationships with a U.S. Parent
Iceland	4	7	0.36
Italy	485	800	0.38
Finland	89	142	0.39
Liechtenstein	5	6	0.45
Switzerland	503	559	0.47
Sweden	251	230	0.52
Taiwan	141	112	0.56
Belgium	362	256	0.59
Bermuda	315	216	0.59
Norway	124	85	0.59
Denmark	231	161	0.60
Korea	202	131	0.61
Japan	491	309	0.61
Spain	520	319	0.62
Israel	112	63	0.64
Austria	137	61	0.69
France	1,176	521	0.69
Germany	1,571	556	0.74

Notes : The table reports the total number of parent subsidiary-pairs groups into pairs where the U.S. is the headquarter and pairs in which the headquarter is foreign owned. The final column reports the fraction of relationships in which the parent is a U.S. firm. These figures are reported separately for each country. The table only reports the 18 countries with the lowest share of relationship with a U.S. parent. Data are from the *Orbis Database* .

Table 5. Imports into the U.S. among a restricted sample of countries: testing Antràs (2003) and Antràs and Helpman (2008).

Dep var: $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$	Restricted Sample: $HQ^{US}_{c}>HQ^{F}_{c}$				
	(1)	(2)	(3)	(4)	(5)
ln R and D / Sales	0.074***	0.062***	0.051***	0.064***	0.068***
ln Skilled Labor / Worker	0.068***	0.088***	0.117***	0.127***	0.117***
ln Materials / Worker	0.031*	0.025	0.020	0.020	0.024
ln Capital / Worker	0.041**				
ln Buildings / Worker		-0.046***	-0.036*	-0.040*	
ln Machinery / Worker		0.077***			
ln Computers / Worker			0.003		
ln Autos / Worker			-0.052***		
ln Other Machinery / Worker			0.094***	0.085***	0.076***
ln Non-Specific Machinery / Worker				-0.011	
ln Non-Specific Capital / Worker					-0.049***
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Observations	111,665	111,065	86,622	86,622	88,579
Number of Clusters	298	294	200	200	209
R-Squared	0.12	0.12	0.13	0.13	0.13

Notes: The dependent variable $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an HS6-country pair. All equations include country fixed effects. Standardized 'beta' coefficients are reported. Standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Non-Specific Machinery' is Computers+Autos. 'Non-Specific Capital' is Buildings+Computers+Autos. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos.

headquarter for all countries and all industries. Then industry-country pairs can be omitted from the analysis. The results are similar if this strategy is pursued, except we do not have data on the number of U.S. headquarters and foreign headquarters in about 40% of the industry-country pairs. The results are robust to either omitting these missing data observations or including them in the regression.

An alternative scenario that fits the Antràs (2003) and Antràs and Helpman (2004, 2008) models cleanly is the case of exports from a U.S. subsidiary to a foreign parent. Here the interpretation is that an intermediate input is produced by a U.S. subsidiary and then it is *exported* to a foreign parent. Figures 3 shows this case. The only difference is now the foreign country is the parent and the U.S. is the subsidiary.

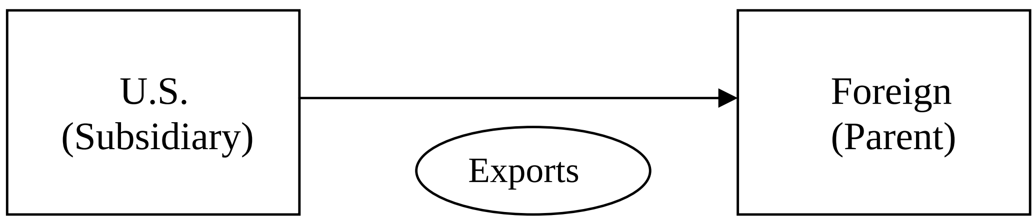

Figure 3. U.S. Exports: The ownership structure consistent with Antràs (2003) and Antràs and Helpman (2004, 2008).

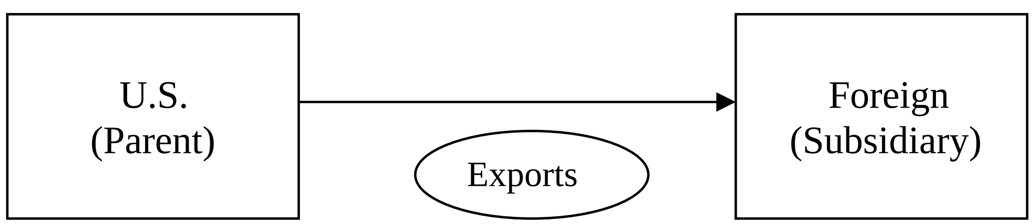

Figure 4. U.S. Exports: The ownership structure inconsistent with Antràs (2003) and Antràs and Helpman (2004, 2008).

The other case is where a U.S. headquarter produces the intermediate input and exports it to the foreign subsidiary. This is shown in 4. The same argument and discussion for figure 2 also applies to figure 4. Therefore, there is no reason to expect the predictions of the models to apply to this case.

We try and distinguish empirically the case shown in figure 3 from the case in figure 4 by examining the determinants of intra-firm exports from the U.S. to countries for which the headquarter is typically the foreign country and the U.S. is the subsidiary. We identify these countries again using the Orbis data. I choose to include in the sample the 15 countries with the lowest share of U.S. headquarters. This is equivalent to choosing a cut-off of 65%.[5]

These result are reported in table 6. Overall, the results are consistent with the results when U.S. imports were examined (reported in table 2 and 5). The share of intra-firm exports from the U.S. are higher in headquarter intensive industries, as measured by capital intensity. This is shown in column (1). Columns (2)–(5) show that this effect is not driven by the non-relationship specific components of capital: buildings, computers and autos.

Again, one can pursue the alternative strategy of constructing country-industry specific measures of the share of relationships for which the parent is the foreign country. Examining exports and restricting observations using this measure also yields results that support hypothesis 1.

3. Productivity Heterogeneity (Antràs and Helpman, 2004)

Antràs and Helpman (2004) relax the assumption of the same productivity across firms. That is, $\theta(j)$ in equation (1) is allowed to vary across firm-supplier relationships. This is accomplished by appending the Antràs (2003) framework onto the Melitz (2003) model.

Let $\overline{\pi}_k$ be variable profits for a firm with $\theta = 1$ that uses organizational form $k = V, O$. Then, as is well known, profits for a firm with productivity θ that adopts organizational form k are linear in $\theta^{\alpha/(1-\alpha)}$:

$$\pi_k(\theta) = \theta^{\alpha/(1-\alpha)}\overline{\pi}_k - F_k \tag{6}$$

where F_k is the fixed costs of offshoring using organizational form k. Antràs (2003) assumes that all firms have the same productivity ($\theta = 1$) and the same fixed costs ($F_V = F_O$). Under these assumptions the Antràs effect states that $\overline{\pi}_V > \overline{\pi}_O$ if and only if $\eta > \eta_c$ i.e., the U.S. firm prefers vertical integration to outsourcing when the firm's share of inputs is large. Heterogeneity of productivity by itself does not alter this conclusion – it simply magnifies the advantages (or disadvantages) of vertical integration.[6]

However, when the fixed costs of outsourcing vary across organizational forms then productivity heterogeneity matters. How heterogeneity matters depends on whether $F_V - F_O$ is positive

[5] The coefficient estimates are similar, but with larger standard errors, if one chooses a cut-off of 50% so that only 5 countries are included in the sample.

[6] That is, if $F_V = F_O$ then $\pi_V(\theta) > \pi_O(\theta) \Leftrightarrow \overline{\pi}_V > \overline{\pi}_O$ and it remain true that the firm prefers vertical integration if and only if $\eta > \eta_c$.

Table 6. Exports into the U.S. among a restricted sample of countries: testing Antràs (2003) and Antràs and Helpman (2008).

Dep var: $X^V_{ic}/(X^V_{ic}+X^O_{ic})$	Restricted Sample: $HQ_c^F/(HQ_c^F+HQ_c^{US})>.65$				
	(1)	(2)	(3)	(4)	(5)
ln R and D / Sales	0.039**	0.024*	0.001	0.012	0.015
ln Skilled Labor / Worker	0.025***	0.063***	0.051**	0.061***	0.064***
ln Materials / Worker	0.010	0.013	0.010	0.007	0.007
ln Capital / Worker	0.061***				
ln Buildings / Worker		0.008	0.009	0.003	
ln Machinery / Worker		0.055***			
ln Computers / Worker			0.037*		
ln Autos / Worker			-0.044***		
ln Other Machinery / Worker			0.068***	0.064***	0.060***
ln Non-Specific Machinery / Worker				0.009	
ln Non-Specific Capital / Worker					0.011
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes
Number of Observations	41,584	41,304	30,678	30,678	31,923
Number of Clusters	294	290	197	197	206
R-Squared	0.07	0.07	0.07	0.07	0.07

Notes: The dependent variable $X^V_{ic}/(X^V_{ic}+X^O_{ic})$ is U.S. intra-firm exports as a share of total exports from the U.S. An observation is an HS6-country pair. All equations include country fixed effects. Standardized 'beta' coefficients are reported. Standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Non-Specific Machinery' is Computers+Autos. 'Non-Specific Capital' is Buildings+Computers+Autos. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos.

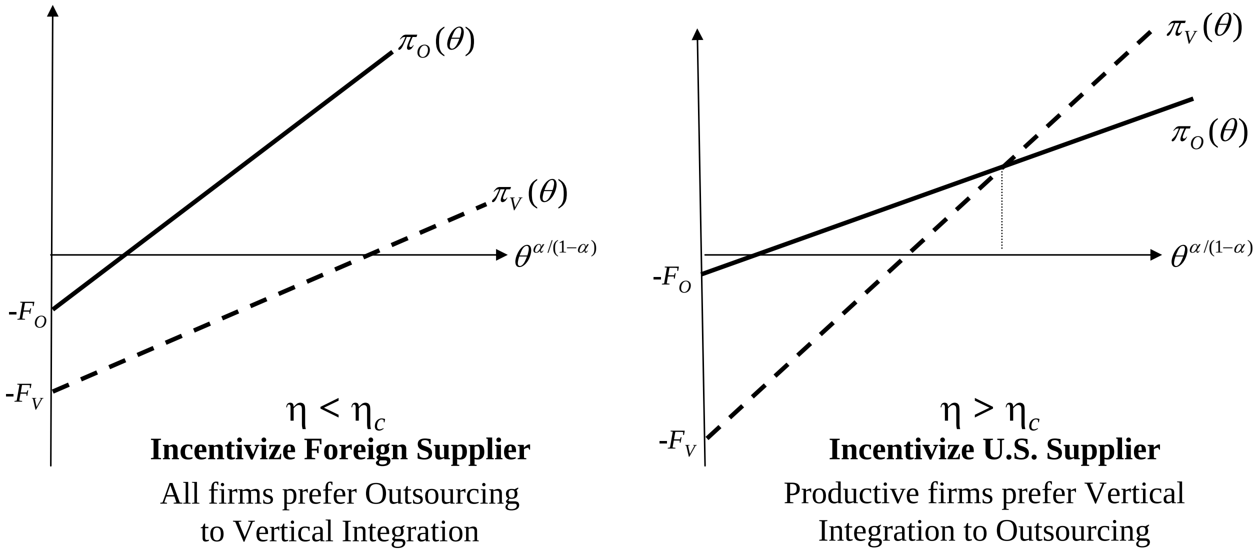

Figure 5. The Outsourcing Decision with Firm-Specific Productivity Differences

or negative. Since the results depend transparently on the sign of $F_V - F_O$, we only present the results under Antràs and Helpman's preferred assumption, namely, $F_V > F_O$. They reason that vertical integration creates a need to supervise the production of intermediate inputs, thus creating managerial overload.

Figure 5 illustrates what happens when heterogeneity is introduced. The figure plots profits under outsourcing $\pi_O(\theta)$ and vertical integration $\pi_V(\theta)$. From Antràs (2003, lemma 3), we know that $\overline{\pi}_V/\overline{\pi}_O$ is increasing in η and equals 1 for $\eta = \eta_c$. This together with equation (6) implies that $\pi_V(\theta)$ is steeper than $\pi_O(\theta)$ for $\eta > \eta_c$ and flatter for $\eta < \eta_c$. From the left-hand panel of figure 5 where $\eta < \eta_c$ and $F_V > F_O$, it must be that outsourcing is always preferred to vertical integration – the Antràs effect and the lower fixed costs of outsourcing both work in favor of outsourcing.

When $\eta > \eta_c$, as in the right-hand panel of figure 5, $\pi_V(\theta)$ is steeper than $\pi_O(\theta)$. It follows that the two curves must cross. Firms with productivity to the right of the crossing point will vertically integrate. Firms to the left will outsource. The tension here is that fixed costs push for outsourcing while the Antràs effect pushes for vertical integration. Since the Antràs effect is greatest for the most productive firms, the Antràs effect dominates for productive firms.

All of this leads to an interesting empirical prediction about the share of U.S. imports that are intra-firm i.e., about $M^V_{ic}/(M^V_{ic} + M^O_{ic})$. The share should depend on an interaction of η_i with $\theta(j)$. In industries for which $\eta_i < \eta_c$, we have that $M^V_{ic} = 0$ so that an increase in the productivity of a match $\theta(j)$ has no effect on its organizational form and $M^V_{ic}/(M^V_{ic} + M^O_{ic}) = 0$. In industries with

$\eta_i > \eta_c$, an increase in $\theta(j)$ increases the likelihood that vertical integration is chosen and therefore will increase $M_{ic}^V/(M_{ic}^V + M_{ic}^O)$.

One implication of this result can be seen by looking across industries. In industries with firms that tend to have large values of $\theta(j)$ we should see a larger share of firms that vertically integrate. Antràs and Helpman (2004) formalize this by assuming that the distribution of productivities θ within an industry is described by the Pareto distribution i.e., by the cumulative distribution function

$$G(\theta) = 1 - \left(\frac{\theta}{b}\right)^{\lambda}, \quad \theta \geq b > 2. \tag{7}$$

Note that for the Pareto distribution the mean and variance of θ are increasing in λ. Consider the case where $\eta > \eta_c$ as in the right-hand panel of figure 5. Antràs and Helpman (2004, page 573) show that an increase in λ (a rise in the mean and variance of θ) leads to a rise in the share of firms that vertically integrate. Empirically, this means a rise in $M_{ic}^V/(M_{ic}^V + M_{ic}^O)$. On the other hand, when $\eta < \eta_c$ as in the left-hand panel of figure 5, an increase in λ has no effect on $M_{ic}^V/(M_{ic}^V + M_{ic}^O)$ which is 0 independent of λ.

Moving to an empirical counterpart of λ, let x_i^θ be some measure of productivity dispersion in industry i. Helpman, Melitz, and Yeaple (2004) and Yeaple (2006) use productivity dispersion x_i^θ as a measure of λ. We can now offer an empirically-oriented hypothesis that comes out of the Antràs-Helpman model.

Hypothesis 3 *Assume $F_V > F_O$. Let x_i^θ be the dispersion of θ across firms within industry i. Then,*

(a) If $\eta < \eta_c$ then dispersion does not affect the intra-firm share of imports: $\frac{\partial M_{ic}^V/(M_{ic}^V + M_{ic}^O)}{\partial x_i^\theta} = 0.$

(b) If $\eta > \eta_c$ then dispersion increases the intra-firm share of imports: $\frac{\partial M_{ic}^V/(M_{ic}^V + M_{ic}^O)}{\partial x_i^\theta} > 0.$

Examining Hypothesis 3 (Antràs & Helpman, 2004)

To construct our measure of productivity dispersion x_i^θ we follow the logic of Helpman *et al.* (2004) and Yeaple (2006). With CES utility and the production function that we have been using, more productive firms have larger total sales. Using firm sales as a measure of firm productivity, Helpman *et al.* (2004) and Yeaple (2006) construct estimates of the dispersion of firm productivity using the standard deviation of firm sales across all firms within an industry. (Their level of industry aggregation allows for only 51 or 52 industries.) Using data on firm sales from the *Orbis* database, we construct measures of the dispersion of firm productivities within 6-digit NAICS industries. Data on firm sales are available annually from 1998 to 2006. Let $Q_{f,i,c}$ denote firm sales of firm f producing in country c and in industry i. Our sample includes all firms with sales data in the Orbis database.

Helpman *et al.* (2004) and Yeaple (2006) use U.S. data to construct an industry-specific measure of productivity dispersion. Note however, that in the theory $\theta(j)$ is pair-specific. Therefore, productivity varies across matches, with part of the match being a U.S. firm and part of the match being a foreign firm. Therefore, a priori U.S. productivity is not the only measure of dispersion that one can use. In our view, because of the existing data limitations we believe that a measure of productivity dispersion in U.S. industries is not the best measure one can use.

As noted, in the model $\theta(j)$ is pair specific. In the model, differences in $\theta(j)$ map directly into firm sales $Q_{f,i,c}$. (Note that this measure varies across both industries and countries.) Although we observe sales at the firm level, we do not observe the trade data at the firm level. This is only observed at the industry-country level. Therefore, we need to think carefully about how we will match our constructed measure of productivity to the trade data.

In the trade data we can identify the location of the headquarter and supplier. We know the country the goods are being shipped from and the country they are being shipped to. We also have a sense of which is most likely the headquarter and which is most likely the subsidiary. (Assume for the moment that we restrict our sample to cases where the headquarter is in the U.S.) Next, we need to think about whether we can identify the industry of the headquarter and the subsidiary. If we can do this, then we can map headquarter and/or subsidiary productivities to the trade data.

The industry of the subsidiary, who produces the input, is easy to identify. The industry of the input and therefore its producer is recorded in the U.S. trade statistics. However, the industry of the headquarter cannot be determined directly from the data unless one assumes that the industry of the headquarter is the same as the industry of the supplier. By using productivity dispersion in the U.S. this is the assumption implicitly being made made in Yeaple (2006). He is assuming part of $\theta(j)$ is determined solely by the U.S. headquarter and that the headquarter's industry is the same as the subsidiaries (which is reflected in the trade data). Given these assumptions, an acceptable measure of productivity dispersion is the productivity dispersion of U.S. firms within an industry.

Assuming that the headquarter is in the same industry as the subsidiary may be inaccurate in many cases. Take for example a tire manufacturer that ships tire inputs to a car manufacturer in the United States. The tires being shipped (and the tire manufacturer) are classified in the trade data under the NAICs 6-digit category "Tire manufacturing". For this pair, $\theta(j)$ can arise from either (*i*) the productivity of the Korean tire manufacturer, or (*ii*) the productivity of the U.S. car manufacturer. The productivity dispersion measures that would be relevant would be either (*i*) the productivity dispersion of Korean tire manufacturers, or (*ii*) the productivity dispersion of the U.S. car manufacturer. However, using the productivity dispersion of U.S. tire manufacturers, which is the measure used in Yeaple (2006), would be incorrect in this case.

Note that between measures (*i*) and (*ii*), only (*i*) is easily identifiable. From the trade data we do not know the industry of the importing headquarter, only the industry of the exporting subsidiary. Therefore, for our analysis we use measure (*i*), which is the productivity dispersion in the industry and country of the exporting subsidiary. This implicitly only consider the influence on $\theta(j)$ exerted by the supplier side of the match. This is clearly inaccurate and we are missing at least half of $\theta(j)$ by doing this. However, this assumption is driven by data limitations.[7]

In Yeaple's (2006) empirical analysis the dispersion measures are constructed for 51 industries. As we have argued, an important assumption of the measure is that the industry of the headquarter is the same as the industry of the subsidiary. Looking within more aggregated industries it is much more likely that this assumption is satisfied. Therefore, when working with more aggregate data U.S. productivity dispersion may be a perfectly fine measure to use. However, when moving to the NAICS 6-digit level, it become much less likely that the industry of the supplier is the same as the industry of the headquarter.

Our measure of dispersion is the standard deviation of the log of firm sales $\ln Q_{f,i,c}$ within an industry i and a country c:

$$x^{\theta}_{ic} \equiv \sqrt{\mathbf{V}(\ln Q_{f,i,c})} \tag{8}$$

where $\mathbf{V}$ is the variance operator.[8]

The Orbis database has sales data for 760,000 plants that have sales data in 2005. The average number of plants in a country-industry pair is 90.

To illustrate our measure consider the example of Korea tire manufacturers. One observation of x^{θ}_{ic} will be the log dispersion of firm sales across tire manufacturers in Korea. The model of Antràs and Helpman (2004) predicts that when dispersion (and average productivity) is high among tire manufacturers in Korea, then the productivity of the match $\theta(j)$ is more likely high, and therefore

[7] In future work it is possible to us Input-Output tables to estimate probabilistically the industry of the headquarter.

[8] See Helpman *et al.* (2004, p. 307) for an explanation of how the standard deviation of the log of firm sales recovers the parameter λ.

we are more likely to observe a vertically integrated relationship (if tire manufacturing is a high headquarter intensive η industry).

More formally, we test hypothesis 3 we estimate how the relationship between productivity dispersion x^{θ}_{ic} and intra-firm imports varies with headquarter intensity. We begin by testing very crudely the predictions of hypothesis 3. We estimate equation (3), but now interact our measure of headquarter intensity with our measure of productivity dispersion x^{θ}_{ic}. The estimating equation is:

$$\frac{M^V_{ic}}{M^V_{ic} + M^O_{ic}} = \gamma_c + \gamma_K \frac{K_i}{L_i} + \gamma_H \frac{H_i}{L_i} + \gamma_R \frac{RD_i}{Q_i} + \gamma_M \frac{M_i}{L_i} + \delta_\theta \, x^{\theta}_{ic} + \delta_{\theta\eta} \, x^{\theta}_{ic} \cdot x^{\eta}_i + \varepsilon_{ic} \tag{9}$$

where x^{η}_i denotes a measure of headquarter intensity, either K_i/L_i, H_i/L_i or RD_i/Q_i.

The estimation results are reported in table 7. Each column reports the results using a different measure of headquarter intensity. As shown, when any of our measures of capital are used as the measure of headquarter intensity, the results are consistent with hypothesis 3. The estimated relationship between productivity dispersion and the share of trade that is intra-firm is higher in higher capital intensity. As well, the estimated effect becomes larger as the measure of capital becomes more restricted and excludes capital that is not customized and relationship-specific.

The final two columns of the table show that when the either skill-intensity or R&D intensity are used as the measure of headquarter intensity, the estimates do not provide support for hypothesis 3. Therefore, the model is supported only if we assume that capital is the input provided by the headquarter.

Strictly speaking, from hypothesis 3, we do not expect the second derivative to be linear as imposed in (9). It will be 0 when η is small and positive when η is large. As well, we do not know where the cut-off level η_c will be. Because of this, we pursue the following estimation strategy. We rank our 298 6-digit NAICS industries by headquarters intensity. Headquarters intensity will be measured by either R&D intensity, skill intensity, or capital intensity. Based on this ranking, we divide the 298 industries into 10 deciles of about 300 industries each. Let $p = 1, \ldots, 10$ index deciles, with $p = 1$ being the least headquarters-intensive decile. Finally, let $I^{\eta}_{ip} = 1$ if industry i is in decile p and $I^{\eta}_{ip} = 0$ otherwise.

We consider a regression that allows the relationship between dispersion and intra-firm imports to differ by decile:

$$\frac{M^V_{ic}}{M^V_{ic} + M^O_{ic}} = \gamma_c + \gamma_K \frac{K_i}{L_i} + \gamma_H \frac{H_i}{L_i} + \gamma_R \frac{RD_i}{Q_i} + \gamma_M \frac{M_i}{L_i} + \sum_{p=1}^{10} \gamma_{\eta p} \, I^{\eta}_{ip} + \sum_{p=1}^{10} \gamma_{\theta\eta p} \, (x^{\theta}_{ic} \cdot I^{\eta}_{ip}) + \varepsilon_{ic}$$

Table 7. Productivity dispersion, headquarter intensity and intra-firm trade: testing Antràs and Helpman (2004).

	Dependent variable: $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$				
	(1)	(2)	(3)	(4)	(5)
	Capital	Machinery	Other Machinery	Skill	R & D
Dispersion: x^{θ}_{ic}	0.037***	0.042***	0.044***	-0.018	0.014
Headquarter interaction: $x^{\theta}_{ic} \times x^{\eta}_{i}$	0.014***	0.015***	0.019***	-0.034**	-0.000
Country Fixed Effects	Yes	Yes	Yes	Yes	Yes
All Control Variables	Yes	Yes	Yes	Yes	Yes
Number of Observations	42,101	41,895	33,433	42,101	42,101
Number of Clusters	267	264	189	267	267
R-Squared	0.12	0.12	0.13	0.12	0.13

Notes: The dependent variable $M^{V}_{ic}/(M^{V}_{ic}+M^{O}_{ic})$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an HS6-country pair. All equations include country fixed effects. Standardized 'beta' coefficients are reported. Standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Machinery' is (total) Capital minus Buildings. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos. `All Control Variables' includes M_i/L_i and headquarter intensity measures.

The primary coefficients of interest are the $\gamma_{\theta\eta p}$. Hypothesis 3 states that for low η and hence low p the impact of dispersion should be zero. That is, $\gamma_{\theta\eta p} = 0$ for low p. Hypothesis 3 also states that for high η, and hence high p, the impact of dispersion should be positive. That is, $\gamma_{\theta\eta p} > 0$ for high p. Since we do not know which decile p contains the cut-off η_c we cannot be more precise about what 'low' and 'high' p means. We will let the data answer this.

Table 8 reports the results of equation (3). The table reports three columns each using a different measure of capital as a measure of headquarter intensity.

In column 1 we measure headquarter intensity by capital intensity. The columns 2 and 3 we our two measures of capital that exclude the less relationship-specific components of total capital. Consistent with hypothesis 3 we observe a one time significant jump in the magnitude of the estimated coefficient when moving from the first to tenth deciles. The jump occurs at about the 8th decile. That is, we find that $0 \approx \hat{\gamma}_{\theta\eta 1} \approx \ldots \approx \hat{\gamma}_{\theta\eta 7} < \hat{\gamma}_{\theta\eta 8} \approx \hat{\gamma}_{\theta\eta 9} \approx \hat{\gamma}_{\theta\eta 10}$. F-tests cannot reject the null hypothesis of the equality of any pair of coefficients among $\hat{\gamma}_{\theta\eta 1}$ to $\hat{\gamma}_{\theta\eta 7}$. However, F-tests do reject the null hypothesis of equality between any of $\hat{\gamma}_{\theta\eta 1}$ to $\hat{\gamma}_{\theta\eta 7}$ and either $\hat{\gamma}_{\theta\eta 8}$, $\hat{\gamma}_{\theta\eta 9}$ or $\hat{\gamma}_{\theta\eta 10}$. These results provide dramatic confirmation of hypothesis 3.

One concern with the results of table 8 is the small number of observations within each decile. We have estimated equation (3) using quintiles rather than deciles and find qualitatively identical results. Specifically, we find zero estimated coefficients for low quintiles, and then a significant

Table 8. Deciles of productivity dispersion, headquarter intensity and intra-firm trade: testing Antràs and Helpman (2004).

	Dependent variable: $M^V_{ic}/(M^V_{ic}+M^O_{ic})$		
	(1)	(2)	(3)
	Capital	Machinery	Other Machinery
Dispersion x^{θ}_{ic} interacted with:			
$I^{\eta}_{i\,1}$	0.001	0.003	0.001
$I^{\eta}_{i\,2}$	0.009	-0.004	-0.008
$I^{\eta}_{i\,3}$	0.008	0.009	-0.030**
$I^{\eta}_{i\,4}$	-0.012	-0.023	0.005
$I^{\eta}_{i\,5}$	-0.007	0.003	-0.012
$I^{\eta}_{i\,6}$	-0.000	-0.008	-0.003
$I^{\eta}_{i\,7}$	0.008	-0.003	-0.017
$I^{\eta}_{i\,8}$	0.022	0.031**	0.022
$I^{\eta}_{i\,9}$	0.039***	0.038***	0.032**
$I^{\eta}_{i\,10}$	0.027***	0.024**	0.029***
Country Fixed Effects	Yes	Yes	Yes
Decile Fixed Effects, I^{η}_{ip}	Yes	Yes	Yes
All Control Variables	Yes	Yes	Yes
Number of Observations	42,101	41,895	33,433
Number of Clusters	267	264	189
R-Squared	0.13	0.13	0.13

Notes: The dependent variable $M^V_{ic}/(M^V_{ic}+M^O_{ic})$ is U.S. intra-firm imports as a share of total U.S. imports. An observation is an HS6-country pair. All equations include country fixed effects. Standardized 'beta' coefficients are reported. Standard errors are clustered at the 6-digit NAICS industry level. ***, ** and * indicate significance at the 1, 5 and 10 percent levels, respectively. 'Machinery' is (total) Capital minus Buildings. 'Other Machinery' is (total) Capital minus Buildings, Computers, and Autos. I^{η}_{ip} is an indicator variable that equals on if industry i is in the p^{th} decile of headquarter intensity. `All Control Variables' includes H_i/L_i, M_i/L_i, RD_i/Q_i, and the relevant headquarter intensity measures i.e., either capital, machinery or other machinery intensity.

increase in the estimated coefficient: $0 \approx \hat{\gamma}_{\theta\eta 1} \approx \ldots \approx \hat{\gamma}_{\theta\eta 3} < \hat{\gamma}_{\theta\eta 4} \approx \hat{\gamma}_{\theta\eta 5}$.

Overall, the results provides strong support for hypothesis 3. We find that there is indeed a cut-off level of headquarter intensity. For industries with headquarter intensity greater than this cut-off, productivity dispersion increases the share of intra-firm imports. For industries with headquarter intensity below the cut-off, the estimated relationship is much weaker and close to zero.

4. Conclusions

Antràs (2003) proposed that we think of the boundaries of the firm – of the choice between outsourcing on the one hand and vertical integration, foreign direct investment and multinationals on the other – in the property-rights terms of Grossman and Hart (1986). The central assumption of the Antràs approach is that vertical integration allows the U.S. firm to partially control the customized intermediate inputs produced by its foreign supplier. The central implication is that we should see vertical integration in industries that intensively use the headquarter inputs produced by the U.S. firm. We analyzed this implication using Census data on U.S. intra-firm and arm's-length imports of 5,423 products from 210 countries in 2005. As predicted by Antràs we found that , R&D, skill- and capital-intensive industries have a higher ratio of intra-firm imports to total imports. This is true even after controlling for exporter fixed effects. Our results extend those of Antràs (2003) and Yeaple (2006) who used a much smaller number of industries and countries.

Antràs and Helpman (2008) extend the earlier models to allow for partial contractibility of headquarter and subsidiary inputs. We find that, consistent with their model an increase in the non-contractibility of headquarter inputs increases the share of trade that is intra-firm. Put differently, we find that the intensity of relationship-specific capital of an industry increases the share of trade that is intra-firm, while the intensity of capital that is not relationship-specific is not. Consistent with the prediction of their model it is the intensity of relationship-specific capital that increases the share of trade that is intra-firm.

Antràs and Helpman (2004) extends Antràs (2003) and allows for firm-level heterogeneity in productivities and fixed costs that are higher for vertical integration than for outsourcing. The extension implies that the intra-firm share of U.S. imports will be highest for firms with two characteristics: (1) high headquarter intensity η and (2) high productivity θ. We found very strong evidence to support this implication.

Overall, the findings of this paper provide strong support for the central predictions in Antràs (2003) and Antràs and Helpman (2004, 2008) about the share of U.S. imports and exports that are intra-firm.

Appendix A. Data Description

Data on intra-firm and total trade are from the U.S. Census Bureau. The trade data are at the 6-digit Harmonized System (HS6) level and for the year 2005. Each shipment imported into the United States is accompanied by a form which asks about the value of the shipment, the HS10 code and whether or not the transaction is with a related party i.e., whether or not the transaction is intra-firm or at arm's length.

Two parties are related if one owns at least 6% of the other. There is an oft-used alternative data set on intra-firm trade, namely, the BEA's multinationals database. In the BEA data a 10% ownership stake is used to define intra-firm transactions. However, the BEA data provide two pieces of evidence to suggest that a 6% threshold is large enough to ensure a controlling stake. First, only five percent of intra-firm BEA imports involve ownership positions of less than 10%. See table 11 in Mataloni Jr. and Yorgason (2006). Thus, if an ownership position is at least 6% it is likely at least 10%. Second, for a very large proportion of ownership positions in the BEA data, once the position is more than 10%, it is also more than 50%. (Authors' calculations from the BEA data available on the web.) Thus, although the threshold ownership stake is only 6% in our data, it is safe to say that in most cases it is a controlling stake.

The capital intensity K_i/L_i, skill intensity S_i/L_i, and material intensity M_i/L_i measures are constructed using data from the United States Annual Survey of Manufactures. The data are from the same year as the trade data, 2005, with industries classified at the 6-digit NAICS level.

References

Antràs, Pol, "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics*, 118 (2003), 1375–1418.

———, "Incomplete Contracts and the Product Cycle," *American Economic Review*, 95 (2005), 1054–1073.

Antràs, Pol, and Elhanan Helpman, "Global Sourcing," *Journal of Political Economy*, 112 (2004), 552–580.

———, "Contractual Frictions and Global Sourcing," in Elhanan Helpman, Dalia Marin, and Thierry Verdier, eds., *The Organization of Firms in a Global Economy* (Harvard University Press, Cambridge, 2008).

Bernard, Andrew B., and J. Bradford Jensen, "Exporters, Skill Upgrading, and the Wage Gap," *Journal of International Economics*, 42 (1997), 3–31.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "Intra-Firm Trade and Product Contractibility," (2008), mimeo.

Bernard, Andrew B., J. Bradford Jensen, and Peter Schott, "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. that Trade Goods," Working Paper No. 11404, National Bureau of Economic Research, (2005).

Feenstra, Robert C., and Gordon H. Hanson, "Ownership and Control in Outsourcing to China: Estimating the Property-Rights Theory of the Firm," *Quarterly Journal of Economics*, 120 (2005), 729–761.

Grossman, Gene M., and Elhanan Helpman, "Integration versus Outsourcing in Industry Equilibrium," *Quarterly Journal Economics*, 117 (2002), 85–120.

———, "Outsourcing versus FDI in Industry Equilibrium," *Journal of the European Economic Association*, 1 (2003), 317–327.

———, "Managerial Incentives and the International Organization of Production," *Journal of International Economics*, 63 (2004), 237–262.

———, "Outsourcing in a Global Economy," *Review of Economic Studies*, 72 (2005), 135–159.

Grossman, Sanford J., and Oliver D. Hart, "Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 94 (1986), 691–719.

Helpman, Elhanan, "Trade, FDI and the Organization of Firms," *Journal of Economic Literature*, 44 (2006), 589–630.

Helpman, Elhanan, and Paul R. Krugman, *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy* (MIT Press, Cambridge, MA, 1985).

Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, "Export versus FDI with Heterogeneous Firms," *American Economic Review*, 94 (2004), 300–316.

Mataloni Jr., Raymond J., and Daniel R. Yorgason, "Operations of U.S. Multinational Companies: Preliminary Results from the 2004 Benchmark Survey," *Survey of Current Business*, 86 (2006), 37–68.

McLaren, John, "Globalization and Vertical Structure," *American Economic Review*, 90 (2000), 1239–1254.

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (2003), 1695–1725.

Nunn, Nathan, "Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade," *Quarterly Journal of Economics*, 122 (2007), 569–600.

Nunn, Nathan, and Daniel Trefler, "The Boundaries of the Multinational Firm: An Empirical Analysis," in Elhanan Helpman, Dalia Marin, and Thierry Verdier, eds., *The Organization of Firms in a Global Economy* (Harvard University Press, Cambridge, 2008), 47–75.

Schott, Peter, "Across-Product versus Within-Product Specialization in International Trade," *Quarterly Journal of Economics*, 119 (2004), 647–678.

Yeaple, Stephen R., "Offshoring, Foreign Direct Investment, and the Structure of U.S. Trade," *Journal of the European Economic Association*, 4 (2006), 602–611.