

Chapter 5

Contracts and Sourcing: Evidence

In Chapter 3, I overviewed several empirical studies exploring the significance of weak contract enforcement for the export decisions of firms and for the structure of international trade flows. It has become customary to appeal to this empirical literature when motivating the role of contractual frictions in the global sourcing decisions of firms. As demonstrated in Chapter 4, however, imperfect contracting affects the international organization of production in ways distinct from those in which it shapes exporting decisions. This was illustrated by the different variants of the global sourcing model developed in Chapter 4, which highlighted the importance of various factors for predicting the differential effect of weak contracting on trade flows of different types of intermediate inputs.

The goal of this chapter is to develop empirical tests of this global sourcing model using detailed data on U.S. imports by product and source country. I will first use the import data aggregated across source countries to explore the determinants of the cross-industry variation in the extent to which U.S. firms rely on domestically produced inputs versus foreign inputs in their production processes. This specification is motivated by equation (4.12) in Chapter 4, which solved for the share of spending on imported manufacturing inputs over total manufacturing input purchases in a particular industry. The equation related this share to several parameters of the model including trade costs, productivity dispersion, demand elasticities, and the level of contractual enforcement as captured by Γ , which in turn was shown to depend on institutional variables as well as on other primitive parameters of the model. Below, I will review some of the key predictions of the models in Chapter 4 before assessing their empirical validity.

I will next exploit the cross-country dimension of U.S. import data to provide richer tests of the model. Before doing so, however, it will prove necessary to develop a multi-country version of the model that illustrates how cross-country variation in institutional quality shapes the relative propensity of U.S. firms to source particular types of inputs from different countries. This model will build on the multi-country model of sourcing developed toward the end of Chapter 2 and will deliver an explicit formula relating the volume of U.S. imports of a particular input v from a particular country j to trade costs between the U.S. and j , the wage rate in j , an aggregate measure of labor productivity in j , and an index Γ_j of contractual efficiency in j which is analogous to the parameter Γ derived in the two-country sourcing models in Chapter 4. Crucially, the (re)derivation of Γ_j will highlight the differential effect of weak contracting on U.S. imports of different products depending on particular characteristics of the product being traded and of the industry purchasing those inputs. This will motivate empirical tests along the lines of Nunn (2007) and Levchenko (2007), which will relate U.S. imports of a particular input v from a particular country j to the interaction of industry and country characteristics, while controlling for product and country-year fixed effects.

In order to build intuition on this difference-in-differences approach, consider the following motivating example. Chile and Argentina are two countries that are fairly equidistant to the U.S., and had very similar levels of physical capital per worker and of educational attainment in the period 2000-2005.¹ Nevertheless, Chile is recorded as having a significantly higher level of contract enforcement than Argentina does, with the difference in their “rule of law” being 1.90 standard deviations in the underlying measure. In fact, over 2000-05, Chile was ranked 22nd out of 134 countries in terms of this measure of institutional quality, while Argentina was ranked 95th.

Perhaps for this reason, and despite the fact that both population and GDP in Argentina over that period were more than twice as large as in Chile, the latter country actually featured larger manufacturing exports to the U.S. than the former (US \$2.58 billion vs. \$2.38 billion), a difference that persists after netting out manufacturing exports related to some key primary products in these two countries, such as copper, petroleum and aluminium. The higher market share of Chile versus Argentina in U.S. imports is, however, very weak evidence of the importance of contract enforcement for trade flows, as there

¹Out of 134 countries with data on these variables, the differences in these variables across these two countries were 0.09 standard deviations for distance, 0.02 for physical capital, and 0.10 for average years of schooling.

might be a myriad of alternative country characteristics that distinguish these two countries and that might be relevant for their differential exports to the United States.

To better identify the causal role of institutions on trade patterns, one can exploit the cross-industry variation in the data to see whether the depressing effect of bad institutions on trade is disproportionately large precisely in the type of industries in which the theory suggests the effect should be disproportionately large. We shall term these industries “contract intensive” and we will use the models developed in Chapter 4 to suggest different proxies for contract intensity.

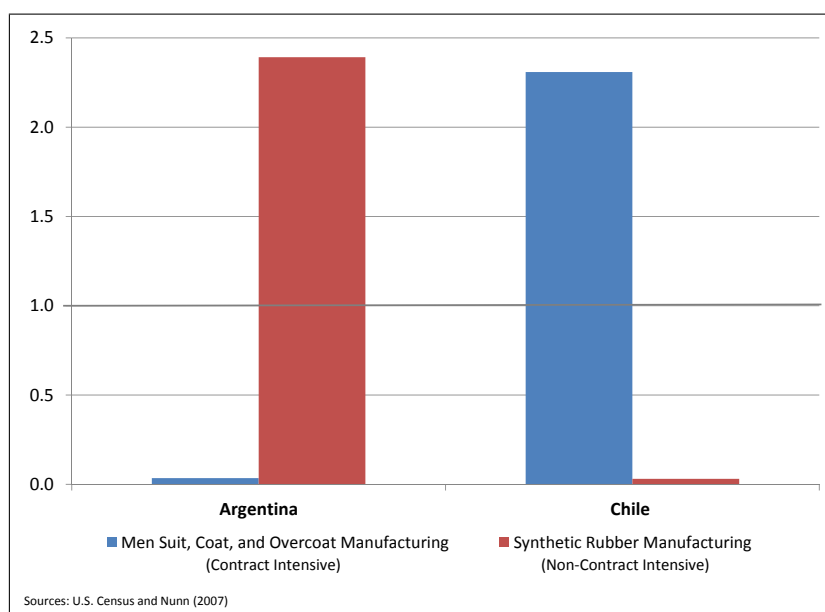


Figure 5.1: Industry Market Share in U.S. Imports relative to Average Market Share

As an example, consider two manufacturing industries with very different levels of Nunn’s (2007) input relationship-specificity measure discussed in Chapter 3. On the one hand, in the six-digit North American Industry Classification System (NAICS) industry 315222 (‘Men’s and Boys’ Cut and Sew Suit, Coat, and Overcoat Manufacturing’), Nunn (2007) estimates that 75% of that industry’s intermediate inputs are relationship-specific. On the other hand, in the NAICS six-digit industry 325212 (‘Synthetic Rubber Manufacturing’), this same percentage is only 19%. This suggests that the former industry is much more contract intensive than the latter, and thus U.S. buyers might be particularly inclined to purchase this industry’s manufactured

goods from Chile relative to Argentina. The logic of comparative advantage suggests in turn that Argentina should be a more attractive source than Chile of low contract-intensive goods, such as synthetic rubber manufacturing.

Figure 5.1 confirms this logic by plotting Argentina's and Chile's market share in each of these two industries, while normalizing these shares by each country's aggregate market share in U.S. manufacturing imports. As is clear from the figure, Argentina exports virtually no men's suits, coats and overcoats to the U.S., while its market share in synthetic rubber is 2.4 times its aggregate market share. Conversely, Chile features virtually no exports of synthetic rubber to the U.S., while its market share in men's suits, coats and overcoats is 2.3 times its aggregate market share.

Of course, skeptical readers might argue that this is just a conveniently picked example, so I will develop empirical tests below that exploit this identification strategy in a more systematic manner using *all* available U.S. import data. Before specifying these tests, I will briefly discuss some of the pros and cons of using U.S. import data to test our global sourcing model.

Using U.S. Import Data: Pros and Cons

The variants of the global sourcing model developed in Chapter 4 focus on the decisions of firms regarding the location from which intermediate inputs are sourced. Hence, firm-level data would appear to be the ideal laboratory for testing these models. Nevertheless, firm-level data on the sourcing decisions of firms are not readily available, and most of the datasets that have been used for this purpose in the literature do not provide a sufficiently rich picture of the variation in the sourcing decisions of firms across inputs and locations.² I will instead conduct tests at the product level which exploit the extent to which different types of manufactured goods are sourced from particular foreign countries or from domestic producers in the United States. It is important to emphasize that these tests are well grounded in theory. As demonstrated in Chapters 2 and 4, and further illustrated below, by solving for the sectoral equilibrium in which a continuum of differentiated final-good producers make sourcing decisions, one can aggregate these producers' decisions and obtain predictions for the relative market share of all countries (including the U.S.) in the purchases of the different intermediate inputs sourced by U.S.-based firms. In sum, I will use *sectoral*-level data to test *sectoral*-level predictions.

²Some datasets, such as the Spanish Encuesta sobre Estrategias Empresariales (ESEE) employed in other parts of the book, only record the firm-level of imported inputs aggregated across inputs and foreign sources.

Although the empirical analysis in this chapter could have been conducted with product-level trade data from multiple countries, I will restrict the analysis to import data from a single country, the United States. I do so to facilitate a comparison with the intrafirm trade empirical analysis in Chapter 8, but also because data availability would constrain me from performing some of the sensitivity tests described below for other countries. I will employ U.S. import data for the period 2000-2011 collected by the U.S. Bureau of Customs and Border Protection and publicly available from the U.S. census website. Although the data are available at the extremely detailed ten-digit Harmonized Tariff Schedule classification system (featuring nearly 17,000 good categories), I will work with more aggregated data to be able to match the trade data with a host of other industry-level variables which are only available at more aggregated levels (more on this below).³ The regressions below will exploit variation on import volumes associated with up to 390 manufacturing sectors, 232 countries and 12 years, which results in up to 1,085,760 observations. Many specifications will, however, feature fewer observations due to data limitations on some key explanatory variables, as explained in more detail below.

Mapping the rich theoretical predictions from the models in Chapter 4 to U.S. import data poses at least three additional difficulties. First, the theory demonstrates that characteristics of *both* the final-good producing firm (such as the elasticity of demand it faces) and of the inputs being purchased (such as their relationship-specificity) are relevant for the choice of location from which inputs are sourced. Unfortunately, publicly available U.S. trade statistics are reported based *only* on the sector or industry category of the good being transacted and do not contain information on the sector that is purchasing the good. To give a specific example, while one observes U.S. imports of synthetic rubber (NAICS 325212) from Argentina, a breakdown of these import volumes into those purchased by plants manufacturing footwear, plastic bottles, or tires is not available to researchers. Nevertheless, as described in Antràs and Chor (2013), one can use information from U.S. Input-Output Tables to provide an educated guess of such a breakdown.

A second limitation of U.S. product-level import data is that they do not identify precisely the end use of the good being imported, and thus it is not straightforward to distinguish between import flows corresponding to intermediate inputs and those corresponding to finished products. For instance, although part of U.S. imports of men's suits and coats (NAICS 315222) can

³As described in the Data Appendix, I do, however, use the ten-digit data to isolate the intermediate input component of U.S. imports.

certainly be treated as inputs bought by U.S. manufacturing and service firms in a variety of sectors (this is an informed guess based on inspection of the U.S. Input-Output Use Tables), one would expect that a significant share of U.S. imports in this sector constitute finished products sold to consumers perhaps via the retail sector. The mapping between the latter type of imports and the models in Chapter 4 is certainly a bit of a stretch and generates additional problems in the measurement of the characteristics of these buying sectors. For this reason, in the tests below I will implement the methodology developed by Wright (2014) to attempt to isolate the intermediate input component of U.S. imports.

A third key concern with U.S. import data relates to the fact that even when one is confident that an import flow into the U.S. reflects the exchange of an intermediate input, this does not ensure that such a flow is associated with an importing decision of the headquarters of a firm based in the U.S., as the models in Chapter 4 focus on. In particular, it is natural to imagine that some of these transactions are related to the headquarters or parent of a foreign multinational company shipping intermediate inputs to one of its affiliates in the United States. Nunn and Treffer (2013*b*) have suggested a correction for this phenomenon that uses data from Bureau van Dijk's Orbis Database to identify the set of countries for which this concern might be particularly salient. I will explain this correction in more detail below and will implement it in some robustness tests.

An additional limitation of using U.S. import data is that they only capture those sourcing decisions that entail goods being shipped back to the United States. In practice, some large U.S. firms have global value chains in which parts and components are shipped across foreign locations and then only shipped back to the U.S. after being assembled abroad, as is the case of the iPad 3 discussed in Chapter 1. For this reason, U.S. imports generally underrepresent the involvement of U.S. firms in global sourcing strategies. I will not attempt to correct for these third-market effects in the empirical exercises to be performed below, but at the same time it is not clear to me in which direction this phenomenon biases the results to be shown.

I have thus far focused on describing some limitations of U.S. import data when serving as a proxy for the relative propensity of U.S. firms to source particular types of inputs from particular countries. Empirically testing the models in Chapter 4 will also require constructing variables related to some of the key parameters driving these decisions in those models. This will naturally raise additional challenges, but it is best to postpone their discussion until we have revisited the main theoretical predictions to be tested.

Cross-Industry Tests: Complete-Contracting Model

I will begin by implementing empirical tests of the variants of the two-country sourcing model in Chapter 4. Although the ultimate goal of this exercise is to explore the contractual determinants of the global sourcing decisions of firms, it will prove useful to first devote some time to an empirical analysis of the benchmark version of the model with complete contracts.

Remember from Chapter 2 that, assuming a Pareto distribution of productivity across producers, we solved for an industry equilibrium in which the share of spending on *imported* manufacturing inputs over total manufacturing input purchases in a particular industry is given by

$$\Upsilon_O = \frac{\left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}{\left(\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D}\right)^{\kappa-(\sigma-1)} - 1 + \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}.$$

where

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)}.$$

As discussed in Chapter 2, the share of imported inputs Υ_O is predicted to increase in wage differences (w_N/w_S), productivity dispersion ($1/\kappa$) and the elasticity of substitution σ , and to decrease in (relative) fragmentation barriers (f_O/f_D , τ) and headquarter intensity (η). We can write this succinctly as

$$\Upsilon_O = \Upsilon_O \left(\underset{+}{w_N/w_S}, \underset{-}{\tau}, \underset{-}{f_O/f_D}, \underset{-}{\kappa}, \underset{+}{\sigma}, \underset{-}{\eta} \right). \quad (5.1)$$

Consider now a particular type of input v that is purchased by firms in different sectors of the Northern economy, which in our empirical application we will associate with the United States. The extent to which U.S. firms procure this input domestically or from foreign sources is then shaped by the key parameters of the model, as summarized in (5.1). Other things equal, the relative importance of imports in the total sales of that input should be higher whenever the input can be produced relatively more cheaply abroad and whenever it can be imported with relatively low trade barriers. Furthermore, equation (5.1) suggests that the relative prevalence of imports of input v should be higher whenever sectors purchasing that input feature high degrees of productivity dispersion, high price elasticities of demand or high levels of headquarter intensity.

A simple way to proxy for the share Υ_O for input v is by computing the ratio of U.S. imports to total U.S. absorption in that particular industrial product v , where U.S. absorption is defined as the sum of shipments by U.S. producers of that good v plus U.S. imports minus U.S. exports of that good. This measure is closely related to what the literature refers to as an import penetration ratio, but I will work below to attempt to refine the measure to better capture intermediate input shipments rather than total shipments.

Table 5.1. The Ten Industries with the Least and Most Offshoring Intensity

10 Least offshoring intensive: lowest Υ_O		10 Most offshoring intensive: highest Υ_O	
.000	Ready-Mix Concrete Manufacturing	.899	Luggage Manufacturing
.001	Fluid Milk Manufacturing	.905	Men's & Boys' Cut and Sew Shirt
.002	Manifold Business Forms Printing	.919	Men's & Boys' Cut and Sew Shirt
.002	Rolled Steel Shape Manufacturing	.924	Plastics, Foil, & Coated Paper Bag
.002	Manufactured Mobile Home Manuf	.926	Infants' Cut and Sew Apparel Ma
.003	Sheet Metal Work Manufacturing	.936	Fur and Leather Apparel Manuf
.003	Guided Missile & Space Vehicle Ma	.952	All Other General Purpose Mach
.004	Poultry Processing	.959	Jewelers' Material and Lapidary
.005	Ice Cream and Frozen Dessert Ma	.966	Women's Footwear (exc. Athletic)
.005	Soybean Processing	.996	Other Footwear Manufacturing

Sources: U.S. Census, NBER-CES Manuf. database and Annual Survey of Manufactures

The left panel of Table 5.1 reports the ten industries with the lowest average offshoring share Υ_O over the period 2000-11. The right-panel of this same table reproduces the ten industries with the highest offshoring shares over the same period 2000-11. These shares are computed by combining import and export data from the U.S. Census website with total shipments data from the NBER-CES Manufacturing database (for 2000-09) and from the Annual Survey of Manufacturing (for 2010-11). Although all the data are available at the six-digit NAICS level, small adjustments were necessary to deal with minor changes in industrial classifications over time (see the Data Appendix for details). For a few industries and years, the share Υ_O turns out to be either negative or higher than one. This is due to the fact that, for 3.3% of the observations (156 out of 4,680), the recorded value of total shipments bizarrely falls short of the value of U.S. exports.⁴ I drop these few observations when computing the averages in Table 5.1.

⁴This in turn might be explained by how the Annual Survey of Manufactures allocates shipments across industry categories for multi-product firms, or by the fact that some

As is clear from Table 5.1, most sectors in the left panel of the table produce goods that are relatively difficult or expensive to ship across borders. Conversely, most sectors in the right panel belong to the apparel sector, which are associated with low trade costs and much lower production costs abroad than in the U.S. It is also clear from Table 5.1 that many of the sectors with high offshoring shares appear to produce almost exclusively final goods, which helps motivate our attempt below to restrict the analysis to imports of intermediate inputs.

Having computed these offshoring shares for the period 2000-11, Table 5.2 presents a simple set of benchmark regressions that attempt to explain these shares using cross-industry variation in (i) freight costs and U.S. tariffs to capture trade frictions; (ii) various proxies for headquarter intensity, (iii) a measure of within-industry productivity dispersion; and (iv) a proxy for the elasticity of demand σ . To better interpret the quantitative importance of the results, all the coefficients in the regressions tables correspond to ‘beta’ coefficients. Furthermore, because the industry controls do not vary across countries or years, I cluster the standard errors at the industry level. Before discussing the results, let me briefly outline the data sources while relegating most details to the Data Appendix.

Sectoral measures of freight costs were downloaded from Peter Schott’s website (see Schott, 2010, for further documentation), while tariff data correspond to applied tariffs from the World Integrated Trade Solution (WITS) database maintained by the World Bank. Both of these trade cost variables are averaged across exporting countries and over all years in 2000-11 in which they were available. With regards to headquarter intensity, I follow the bulk of the literature and proxy for it with measures of capital, skill and R&D intensity of U.S. manufacturing firms. More specifically, capital intensity and skill intensity were computed from the NBER Manufacturing Database, while R&D intensity corresponds to the logarithm of the average R&D expenditures to sales ratio as computed by Nunn and Trefler (2013*b*) using the Orbis database. The measure of productivity dispersion is obtained from Nunn and Trefler (2008), who constructed it based on the standard deviation of log trade flows within the Harmonized System ten-digit sub-industries associated with a sector. Finally, the elasticity of demand was computed based on the widely used U.S. import demand elasticities for Harmonized System (HS) ten-digit products computed by Broda and Weinstein (2006).⁵

exports of manufactured goods are conducted by non-manufacturing firms which add value to them before shipping.

⁵All of the variables used in regressions in Table 5.2 were either downloaded at or converted into the six-digit NAICS classification at which the international trade was

Having discussed the variables included and their sources, we can now turn to describing the results in Table 5.2. I begin in column (1) with a simple regression of offshoring shares on trade costs and proxies for headquarter intensity, using all available data, which corresponds to 390 industries and 12 years, for a total of 4,680 observations. Column (1) confirms that industries with large freight and insurance costs are associated with lower offshoring shares, as predicted by the theory. The effect is significant both in statistic and economic terms: an increase in one standard deviation of transport costs reduces the offshoring share by 0.22 standard deviations. Conversely, the evidence for a negative effect of man-made trade barriers is much weaker, as the coefficient on tariffs is actually positive, though statistically and economically insignificant. This puzzling result might be explained by a reverse-causality bias, as political-economy theories of tariff formation emphasize a positive effect of import penetration ratios on the desired level of protection of a sector. Ideally, one would attempt to correct for this simultaneity bias along the lines of Treffer (1993*b*), but I will leave this for future research. This is in part because, despite the existence of an endogeneity bias, some of the refined specifications below will record a negative and significant effect of tariffs. The remaining three coefficients of column (1) of Table 5.2 confirm a negative effect of the three measures used to proxy headquarter intensity, but the statistical significance of each of these coefficients is very low.

In column (2), I repeat the same regression but dropping the 156 observations for which the offshoring share Υ_O falls outside the interval $[0, 1]$. This improves the fit of the regression, but affects the estimates only slightly, with the exception of capital intensity, which now appears to be close to significant at standard confidence levels. From now on, I work with the sample of observations with $\Upsilon_O \in [0, 1]$. In column (3), I break up the effect of capital intensity into the independent effect of capital equipment and capital structures. Interestingly, equipment intensity appears to have a very significant negative effect on offshoring shares, while structures affect these shares positively. Both effects are highly significant in both economic and statistical terms. This result is intuitive if one interprets headquarter intensity as representing the relative importance of the type of capital investments for which the “North” appears to have the largest comparative advantage. Indeed, Mutreja (2013) documents that cross-country dispersion in capital equipment is much larger than the dispersion in capital structures and that

available in its original form, as explained in more detail in the Data Appendix. The entire dataset and Stata program codes used in the empirical analysis in this chapter and in Chapter 8 are available for download at <http://scholar.harvard.edu/antras/books>.

the ratio of equipment to structures is much higher in rich than in poor countries.⁶ The effects of productivity dispersion and the elasticity of demand on offshoring shares are analyzed in Column (4) of Table 5.2. The coefficients on these variables are small in magnitude and are imprecisely estimated.

Table 5.2. Determinants of U.S. Offshoring Shares

Dep. Var.: $\frac{Imp}{Imp+Shipments-Exp}$	(1)	(2)	(3)	(4)	(5)	(6)
Freight Costs	-0.217** (0.041)	-0.271** (0.044)	-0.280** (0.045)	-0.275** (0.044)	-0.025** (0.004)	-0.052** (0.009)
Tariffs	0.038 (0.073)	0.046 (0.089)	0.015 (0.073)	0.012 (0.075)	0.001 (0.006)	0.003 (0.010)
Log(R&D/Sales)	-0.027 (0.052)	-0.004 (0.051)	0.025 (0.051)	0.023 (0.051)	-0.001 (0.005)	-0.008 (0.010)
Log(Skilled/Unskilled)	-0.000 (0.000)	-0.023 (0.048)	-0.043 (0.049)	-0.045 (0.049)	-0.002 (0.005)	-0.006 (0.010)
Log(Capital/Labor)	-0.049 (0.042)	-0.082 (0.053)				
Log(Capital Equip/Labor)			-0.484** (0.121)	-0.466** (0.120)	-0.037** (0.010)	-0.074** (0.015)
Log(Capital Struct/Labor)			0.411** (0.115)	0.393** (0.114)	0.032** (0.010)	0.067** (0.016)
Productivity Dispersion				-0.002 (0.086)	0.003 (0.006)	0.007 (0.020)
Elasticity of Demand				0.050 (0.063)	0.002 (0.005)	0.005 (0.006)
Sample Restrictions	None	$\Upsilon_{O} \in [0,1]$	$\Upsilon_{O} \in [0,1]$	$\Upsilon_{O} \in [0,1]$	$\Upsilon_{O} \in [0,1]$	$\Upsilon_{O} \in (0,1]$
Fixed Effects	Year	Year	Year	Year	Ctr/Year	Ctr/Year
Observations	4,680	4,524	4,524	4,524	1,085,537	312,929
R-squared	0.063	0.092	0.140	0.142	0.203	0.196

Standard errors clustered at the industry level in all columns. +, *, ** denote 10, 5, 1% significance.

⁶Using data from the Annual Survey of Manufactures (2002-2010), I have further experimented with breaking up capital equipment into expenditures on (i) automobiles and trucks for highway use, (ii) computers and peripheral data processing equipment, and (iii) all other machinery and equipment computers. The effect of autos and 'other equipment' is negative and significant, while that of computers is positive and very significant. The fact that there appears to exist a higher propensity to offshore in sectors that make more intensive use of computers is interesting and intuitive.

In columns (5) and (6) of Table 5.2, I exploit the full cross-sectoral *and* cross-country variation of the import data. I first compute sectoral offshoring shares at the exporter-country level, by replacing total sectoral imports in the numerator of Υ_O with sectoral imports from a particular country j . I then study the cross-product variation in Υ_O within particular exporting countries by introducing source-country-year fixed effects into the regressions. Later in the chapter, I will motivate in more detail the benefits of exploiting the cross-country dimension of the data for identification purposes. For the time being, it suffices to point out that one might be concerned that the patterns we observed in columns (1)-(4) reflect the attractiveness of particular countries as sources of imports, and the fact that these locations are particularly good at producing goods that happen to be cheap to transport or feature a high equipment capital intensity. By introducing source-country-year fixed effects, one is then better able to isolate the effect of sectoral-level characteristics on the relative propensity to offshore or source from domestic producers.

In column (5), I use the full sample of 1,085,760 offshoring shares, except for the mere 223 observations (0.02%) for which Υ_O remains either negative or higher than one. A simple comparison of columns (4) and (5) indicates that the qualitative nature of the results is largely unaffected by the use of country-specific offshoring shares. Note, however, that the (beta) coefficients are around 10 times smaller than in the purely cross-sectoral regressions, and broadly suggest a rather small economic significance of the industry covariates. Part of the reason for this is that 71.1% of the country/industry/year observations feature zero imports into the United States (thus implying $\Upsilon_O = 0$), and hence the simple OLS (linear probability) model I have specified does not ensure the best possible fit of the data. In column (6), I restrict the sample to those observations for which imports flows are positive, and the coefficients roughly double in size relative to those in column (5). I realize, of course, that this is not the proper econometric way to handle the zeroes in the data. In my (weak) defense, however, this approach is fairly standard in the literature.

Overall, the results in Table 5.2 provide mixed evidence in favor of our benchmark global sourcing model with complete contracting. On the one hand, we are able to confirm the negative effect of freight costs and certain proxies of headquarter intensity (most notably, capital equipment) on offshoring shares. On the other hand, the negative effect of other plausible proxies for headquarter intensity appears to be much less precisely estimated. In addition, productivity dispersion and the elasticity of demand affect offshoring shares positively as predicted by the theory, but both the economic as well as the statistical significance of these results is very small.

Cross-Industry Tests: Sample Restrictions

At some level, it should not be too surprising that the empirical evidence is mixed. After all, earlier in this chapter I highlighted three serious caveats associated with the use of U.S. imports in the construction of offshoring shares. Let me next briefly outline how one can refine the above tests to partially address these limitations.

Consider first the concern that the import data identify only the sector to which the good being transacted belongs. This led me above to correlate the degree to which the purchases of an industrial sector's goods come from abroad versus the U.S. with characteristics of that same sector. In light of our global sourcing model, this seems the natural thing to do when studying the effect of freight costs and tariffs, but it is less justifiable for headquarter intensity, and simply erroneous with regards to the elasticity of demand. More specifically, the headquarter intensity parameter η captures the relative importance of the inputs provided by U.S. headquarters and their suppliers, and thus it would seem more appropriate to construct measures of headquarter intensity of the industry *buying* those inputs. Even more clearly, the parameter σ shaping Υ_O in (5.1) is related to the elasticity of demand in the industry buying and not selling those inputs. Unfortunately, the U.S. Census Related-Party data, and publicly available trade statistics more generally, do not contain information on the industry classification of the importing firm. Still, following the approach in Antràs and Chor (2013), one can use interindustry flow data from the U.S. input-output tables to compute industry variables (e.g., proxies for η and σ), related to the *average* industry buying inputs belonging to a particular industry category. The interested reader is referred to the Data Appendix for more details on the construction of these 'buyer' variables.

Although one could construct a buyer version of the Nunn and Trefler (2008) productivity dispersion measure, taking a weighted average of a series of dispersion measures is less likely to provide an accurate measure of the dispersion of the *average* buying industry. Furthermore, as pointed out by Nunn and Trefler (2013*b*), the comparative static relating offshoring shares to the parameter κ will apply regardless of whether size dispersion stems from productivity dispersion across buyers or across sellers. While the latter type of heterogeneity is missing in the model, it could easily be introduced. For these reasons, I follow Nunn and Trefler (2008) in restricting attention to productivity dispersion measures associated with the sector to which the good being imported belongs.

The need to filter the data through an Input-Output table forces me

to abandon the use of the NAICS six-digit industry classification (at which the trade data is reported) and switch to 2002 Input-Output industry codes (IO2002), which is a slightly coarser classification. As a result, I am left with data on 253 IO2002 manufacturing industries instead of the 390 NAICS industries in Table 5.2.⁷ Column (1) of Table 5.3 reports cross-industry results analogous to those in column (4) of Table 5.2 but with the IO2002 classification instead of NAICS classification. Comparing these two columns, we see that the change in industry classification has a relatively modest effect on the estimates. Freight costs continue to have a negative and significant effect on offshoring shares, while the evidence for the other parameters of the model is mixed. The main differences in the two columns are that the negative effect of equipment capital intensity is now statistically significant only at the 7% level, and that the effect of R&D intensity now appears positive and significant at the 10% level.⁸

In column (2) of Table 5.3, I introduce buyer versions of the elasticity of demand and of the proxies for headquarter intensity. The above discussion might have given the impression that the mismeasurement of some key determinants of offshoring shares was responsible for the mixed performance of the model in Table 5.2. Nevertheless, introducing buyer industry variables into the regression only has a minor effect on the estimates. In fact, the coefficient on capital equipment is slightly reduced and loses even more of its significance, and overall, freight costs is the only variable that appears to statistically affect offshoring shares.⁹

Let us now tackle the second limitation of U.S. product-level import data related to the fact that it does not explicitly distinguish between final goods and intermediate inputs. So far, I have included all U.S. imports of manufacturing goods in the construction of offshoring shares, but it seems sensible to attempt to restrict the sample to intermediate input purchases. In part, this is because the model we have laid out is one in which a U.S. final-good manufacturer is deciding on the optimal location of production of the inputs

⁷As explained in the Data Appendix, late in the production of this book, Davin Chor alerted me to the fact that one IO2002 industry appears to have an unrealistically high R&D intensity. Fortunately, the results were found to be virtually unaffected when excluding this industry from the analysis. Such an outlier does not appear in the NAICS-level dataset.

⁸The total number of observations (2986) corresponds to $253 \times 12 = 3,036$ minus 50 observations (1.6%) for which Υ_O falls outside the interval $[0, 1]$.

⁹The correlation between the buyer and seller versions of these key industry variables is high and ranges from 0.74 for capital buildings to 0.89 for the elasticity of demand. This, in turn, partly reflects the disproportionate importance of within-industry commodity flows in Input-Output tables.

it combines in production. In response to this, one might argue that it would suffice to relabel some of the objects in the theory to make the model apply to a U.S. retailer deciding on whether to offshore the production of the manufacturing goods it markets. However, this would make it impossible (with available data) to construct average buyer versions of some of the key determinants of offshoring shares, since I only have access to data on U.S. manufacturing firms. For these reasons, it is worth putting some effort in purging out final-good purchases from the data.

Table 5.3. Refined Determinants of U.S. Offshoring Shares

Dep. Var.: $\frac{Imp}{Imp+Shipments-Exp}$	(1)	(2)	(3)	(4)	(5)	(6)
Seller Industry Freight Costs	-0.315** (0.058)	-0.295** (0.056)	-0.269** (0.058)	-0.273** (0.053)	-0.025** (0.005)	-0.054** (0.012)
Seller Industry Tariffs	-0.025 (0.068)	-0.013 (0.068)	-0.083* (0.031)	-0.088** (0.029)	-0.007** (0.002)	-0.012 (0.008)
Log(R&D/Sales)	0.088+ (0.053)	0.095 (0.072)	0.033 (0.085)	0.034 (0.082)	0.006 (0.008)	0.007 (0.016)
Log(Skilled/Unskilled)	-0.021 (0.062)	-0.036 (0.073)	0.054 (0.073)	0.039 (0.072)	0.004 (0.007)	-0.009 (0.016)
Log(Capital Equip/Labor)	-0.293+ (0.161)	-0.221 (0.163)	-0.113 (0.152)	-0.141 (0.155)	-0.005 (0.014)	-0.048 (0.031)
Log(Capital Struct/Labor)	0.261+ (0.151)	0.108 (0.150)	0.085 (0.148)	0.113 (0.151)	0.003 (0.013)	0.037 (0.029)
Productivity Dispersion	0.016 (0.071)	0.048 (0.064)	0.093 (0.069)	0.118+ (0.069)	0.016* (0.007)	0.031* (0.015)
Elasticity of Demand	-0.023 (0.072)	-0.042 (0.082)	-0.032 (0.080)	-0.045 (0.084)	-0.005 (0.005)	-0.004 (0.019)
Sample Restrictions	$\Upsilon_{O \in [0,1]}$	$\Upsilon_{O \in [0,1]}$	W	W+NT	W+NT	W+NT+
Fixed Effects	Year	Year	Year	Year	Ctr/Year	Ctr/Year
Buyer vs Seller Industry Controls	Seller	Buyer	Buyer	Buyer	Buyer	Buyer
Observations	2,986	2,986	2,510	2,513	582,811	148,879
R-squared	0.149	0.148	0.147	0.156	0.200	0.198

W and NT stand for the Wright (2014) and Nunn and Trefer (2013*b*) sample corrections. Standard errors clustered at the industry level. +, *, ** denote 10, 5, 1% significance.

In order to do so, I build on the methodology developed by Wright (2014). I relegate most details to the Data Appendix, but in a nutshell, Wright's (2014) approach employs a U.S. Census industry concordance between ten-digit HS codes and five-digit End-Use codes to categorize highly

disaggregated commodities into final goods and intermediate inputs. With that information, one can then remove from the sample all ten-digit HS codes associated with final good production, and then reaggregate the data to the IO2002 level to have a proxy for intermediate input import and export flows.¹⁰

Implementing this methodology naturally reduces the volume of trade differentially across industries and also leads to the loss of observations associated with industries that are composed *entirely* of final goods, such as IO 2002 industry 335222 ('Household Refrigerator and Home Freezer Manufacturing'). The Data Appendix contains the full list of dropped final-good industries.¹¹ This procedure can be used to compute imports and exports of intermediate inputs, but the offshoring share formula also requires that we adjust U.S. shipments in order to constrain these to reflect U.S. intermediate input sales. I do so by multiplying U.S. shipments in each industry by the average 'Wright' factor applied to trade flows in that industry over 2000-11. Before discussing the effects of implementing this refinement on the estimates, in Table 5.4 I report the ten sectors with the lowest and highest Wright-corrected average offshoring share Υ_O over the period 2000-11. It is reassuring to compare these sectors with those in Table 5.1 and notice that most of the consumer good sectors in that earlier table are no longer in the sample.

The results of applying this filter to our empirical model explaining offshoring shares are shown in column (3) of Table 5.3. A first noteworthy result is that the effect of tariffs is six times larger than in column (2) and is now statistically significant at the 5% level. Furthermore, the effect of productivity dispersion doubles in size and is close to significant at the 10% level. On the negative side, the effect of the buyer proxies of headquarter intensity and the elasticity of demand continue to have an imprecisely estimated impact on offshoring shares, and even the sign of some of these coefficients is the opposite of that implied by the model.

¹⁰I follow Wright (2014) in also removing industries that purely process raw materials so we can more comfortably treat inputs as differentiated.

¹¹A sector is dropped whenever it comprises zero aggregate imports of intermediate inputs in each year with an offshoring share in $[0, 1]$. There are 39 industries for which this is true (see Table B.1 in the Data Appendix). Column (3) of Table 5.3 drops 526 observations, which is more than $39 \times 12 = 468$ because 58 additional observations have offshoring shares outside the interval $[0, 1]$.

Table 5.4. The Ten Industries with the Least and Most *Corrected* Offshoring Intensity

10 Least offshoring intensive: lowest Υ_O		10 Most offshoring intensive: highest Υ_O	
.000	Ready-Mix Concrete Manufacturing	.651	Computer Storage Devices
.004	Support Activities for Printing	.661	Metal Cutting Machine Tools
.006	Asphalt Paving Mix. & Block Manuf	.664	Electr. Capacitors & other Inductors
.007	Textile and Fabric Finishing Mills	.668	Electronic Connectors
.007	Concrete Pipe/ Brick / Block Manuf	.748	Optical Instruments & Lens
.008	Sign Manufacturing	.764	Doll, Toy, Game Manufacturing
.015	Asphalt Shingle & Coating Materials	.773	Leather & Hide Tanning & Finishing
.015	Ornamental & Architectural Metal	.831	Other General Purpose Machinery
.016	Motor Vehicle Body Manufacturing	.838	Pulp Mills
.017	Paperboard Mills	.882	Audio & Video Equipment Manuf

Sources: Same as in Table 5.1 plus a sample adjustment based on Wright (2014)

In column (4) of Table 5.3, I experiment with one additional refinement of the empirical test. Our discussion above regarding the effects of ‘buyer’ headquarter intensity, productivity dispersion and demand elasticities relied on an interpretation of U.S. intermediate input imports as being associated with U.S. headquarters importing goods from foreign suppliers. A nontrivial share of these imports, however, consists of shipments from foreign headquarters to their U.S. affiliates or to U.S. unaffiliated parties. Arguably, the rationale for these transactions might not be best interpreted through the lens of the models developed in Chapters 2 and 4. For this reason, in column (4), I follow Nunn and Treffer (2013*b*) in checking the robustness of our results to a restricted sample that better fits the spirit of our global sourcing model. More specifically, Nunn and Treffer (2013*b*) use data from Bureau van Dijk’s Orbis Database to identify all subsidiary headquarter pairs in which either the subsidiary or the headquarter are from the United States. They find that there are only 18 countries (see their Table 4) such that the share of pairs for which the U.S. firm is the parent is below 75 percent. Furthermore, only for 5 of these 18 countries is this share below 50 percent. It is thus advisable to present results in which these five countries (Iceland, Italy, Finland, Liechtenstein, and Switzerland) are removed from the sample. I have also experimented with dropping all 18 countries with a share below 75% and the results are not materially changed.¹²

¹²The full list of 18 countries includes Iceland, Italy, Finland, Liechtenstein, Switzerland, Sweden, Taiwan, Belgium, Bermuda, Norway, Denmark, Korea, Japan, Spain, Israel, Austria, France, and Germany.

As is clear from column (4) of Table 5.3, this Nunn-Trefler correction has a qualitatively similar effect as the previous refinements. The negative effect of tariffs and the positive effect of productivity dispersion are larger and more precisely estimated than in previous columns (and the latter effect is now significant at the 10% level), but supporting evidence for a negative effect of buyer headquarter intensity and positive effect of demand elasticities on offshoring remains elusive.

In the last two columns of Table 5.3, I return to the use of both cross-industry and cross-country variation in offshoring shares, while applying the Wright and Nunn-Trefler corrections to trade flows and U.S. shipments in the construction of the shares. Again, the specifications include country-year fixed effects, so the purpose here is to compare offshoring shares across industries, while controlling for time-varying unobserved country characteristics. Arguably, even when one is interested in purely cross-industry variation, it is advisable to fit the data through this straighter jacket. Similarly to Table 5.2, in column (5) I include all (Wright and Nunn-Trefler adjusted) offshoring shares in the interval $[0, 1]$, while in column (6) I drop all observations in column (5) with zero U.S. imports and thus zero offshoring shares.¹³ As in Table 5.2, the results are qualitatively similar to the regressions exploiting only the cross-industry dimension of the data, although the economic size of these effects is greatly reduced, while their statistical significance is generally enhanced. Note, in particular, that the effect of buyer productivity dispersion is now significant at the 5% level in both columns (5) and (6).¹⁴ The results on the elasticity of demand and headquarter intensity remain mixed, although the effect of capital equipment in column (6) is negative and very close to significant at the 10% level.

Put together, the estimates in Table 5.3 provide supporting evidence for some of the key predictions of the benchmark complete-contracting benchmark model. In particular, offshoring shares appear to be significantly higher for goods that are relatively cheap to transport (due to low trade costs or low tariffs) and for goods purchased by sectors featuring high productivity dispersion. We have also found some evidence for a negative effect of some proxies of ‘buyer’ headquarter intensity on offshoring shares, though these effects are not particularly robust. Disappointingly, we have found little ev-

¹³To make sense of the number of observations in column (5), note that the sample now excludes 5 countries and 39 industries, and 125 of the remaining offshoring shares are lower than 0 or higher than 1. We thus have $(232 - 5) \times (253 - 39) \times 12 - 125 = 582,811$ observations.

¹⁴It is worth stressing that the standard errors in these regressions are clustered conservatively at the industry level.

idence suggesting a positive effect of buyer demand elasticities on offshoring shares. We next turn to the incomplete-contracting version of the global sourcing model to see whether one can make sense of the mixed effects of the tests performed above, and to formally test the distinctive predictions that arise from the modeling of contractual frictions.

Cross-Industry Tests: Incomplete-Contracting Model

Recall that towards the beginning of Chapter 4, I derived a formula – see equations (4.12) and (4.13) – for the share of spending on imported manufacturing inputs over total manufacturing input purchases in a particular industry in the presence of contractual frictions. I did so, however, under the strong assumptions of complete contracting in the North, ‘totally’ incomplete contracting in the South, a single input, symmetric bargaining, and no financial constraints. As I show in the Theoretical Appendix, this formula can be readily extended to *all* the variants of the model developed in Chapter 4. In those variants of the model, I expressed firm profits under domestic sourcing and under offshoring as

$$\pi_D(\varphi) = (w_N)^{1-\sigma} B\Gamma_D\varphi^{\sigma-1} - f_D w_N$$

and

$$\pi_O(\varphi) = ((w_N)^\eta (\tau w_S)^{1-\eta})^{1-\sigma} B\Gamma_O\varphi^{\sigma-1} - f_O w_N,$$

respectively, where Γ_D and Γ_O denote the levels of contractual efficiency associated with domestic sourcing and offshoring, respectively. The general formula for the share of offshored intermediate inputs is then given by

$$\Upsilon_O = \frac{\frac{\Gamma_O}{\Gamma_D} \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}{\left(\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D}\right)^{\kappa-(\sigma-1)} - 1 + \frac{\Gamma_O}{\Gamma_D} \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}, \quad (5.2)$$

where

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\frac{\Gamma_O}{\Gamma_D} \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)}. \quad (5.3)$$

In a manner analogous to the complete-contracting case, we can summarize the dependence of the share Υ_O on the parameters of the model by

$$\Upsilon_O = \Upsilon_O \left(\underset{+}{w_N/w_S}, \underset{-}{\tau}, \underset{-}{f_O/f_D}, \underset{-}{\kappa}, \underset{+}{\sigma}, \underset{-}{\eta}, \underset{+}{\Gamma_O/\Gamma_D} \right), \quad (5.4)$$

where the only novel feature relative to (5.1) is the positive dependence of Υ_O with respect to Γ_O/Γ_D .

Although the same formulas (5.2) and (5.3) apply to all the variants of the global sourcing model, it is important to emphasize that the particular values of Γ_D and Γ_O (and how they are shaped by parameters) differ across the various extensions of the model. For reasons that will become clear later in this chapter, the discussion in Chapter 4 was centered around the effects of the deep parameters of the model on the *level* of Γ_D and Γ_O , rather than on the ratio Γ_O/Γ_D . Nevertheless, the Theoretical Appendix contains detailed derivations and proofs of how various parameters affect the ratio Γ_O/Γ_D under the plausible assumption that offshore transactions are associated with a lower degree of contractibility than domestic transactions.

Table 5.5. Effect of Parameters on Γ_O , Γ_D , and Γ_O/Γ_D

	σ	η	ϕ	μ_S	ϵ	ρ
Γ_D	–	Ambiguous	0	0	–	+
Γ_O	–	Ambiguous	+	+	–	+
Γ_O/Γ_D	–	Ambiguous	+	+	–	+

Table 5.5 summarizes some of the key comparative statics associated with the levels of contractual efficiency as well as their ratio. Notice first that a higher elasticity of demand σ of the buying industry is associated with lower contractual efficiency, with the effect being disproportionately large for offshoring relationships relative to domestic ones.¹⁵ Hence, the distortions generated by incomplete contracting appear to be aggravated by a high degree of competition in final-good markets. Next, note that for reasons discussed in Chapter 4, the effect of headquarter intensity η on these indices of contractibility and their ratio is ambiguous.¹⁶ Naturally, when assessing the effects of the elasticity of demand σ and headquarter intensity η on offshoring shares Υ_O , one also needs to take into account the *direct* effects of these parameters in equations (5.2) and (5.3), as summarized in (5.4). The overall effect of σ and η on offshoring shares turns out to be ambiguous because it is the balance of two effects of opposite (or potentially opposite) sign.

¹⁵To be precise, in Chapter 4 we showed that in the extension with limitations on ex-ante transfers, this negative effect required making a mild parametric assumption.

¹⁶For instance, in the version of the model with generalized Nash bargaining, the primitive bargaining power parameter β was key for determining the sign of the dependence of contractual efficiency with respect to headquarter intensity η .

These theoretical ambiguities might explain why, in the regressions in Table 5.3, the empirical proxies for these two parameters did not appear to affect offshoring shares in a robust way as predicted by the complete-contracting benchmark model. Conversely, buyer productivity dispersion $1/\kappa$ and trade costs τ have no indirect effect on offshoring shares working through the size of contractual frictions, and thus the model continues to predict in an ambiguous way the positive effect of dispersion and the negative effect of trade frictions that we confirmed empirically in the last columns of Table 5.3.

Apart from offering a potential rationale for the mixed results in Table 5.3, the incomplete-contracting version of our global sourcing model demonstrates that offshoring shares should also be shaped by certain novel variables that affect profitability only through their effect on the size of contractual inefficiencies. As indicated in Table 5.5, the model predicts that offshoring shares Υ_O should be increasing in the ability of the final-good producer to extract rents from suppliers (ϕ), in the degree of contractibility associated with offshoring (μ_S), and in the degree to which the imported input is substitutable with other inputs in production (ρ). Conversely, offshoring shares should be lower whenever the input being purchased features a relatively high level of customization (ϵ).

I will next attempt to incorporate the effect of these variables into the cross-industry empirical specifications that I developed in Tables 5.2 and 5.3. Before doing so, I first need to briefly discuss how to empirically proxy for these contractual determinants of offshoring shares (see the Data Appendix for more details).

In Chapter 4 we have related the parameter ϕ to the extent to which suppliers face financial constraints that inhibit their ability to make upfront payments to final-good producers. The most widely used industry-level proxies for the importance of financial constraints are the external financial dependence measure of Rajan and Zingales (1998) and the asset tangibility measure of Braun (2002). The idea behind these variables is that financial constraints will be tighter in sectors in which firms' internal cash flows are not a significant source of funding or in which firms' assets are largely intangible and cannot be used as collateral.

The recent empirical literature on trade and contracting institutions has proposed various sector-level proxies for the extent to which contractual frictions reduce production efficiency. Anecdotally, three of these measures originate from the Ph.D. theses of three of the brightest young researchers in International Trade. In Chapter 3, I described in some detail the input specificity measure developed by Nunn (2007), which I will use extensively

below. In a contemporaneous paper, Levchenko (2007) suggested an alternative measure of contractual dependence based on the degree to which firms in a sector use a large number of intermediate inputs in production. More precisely, Levchenko's (2007) measure consists of a sector's Herfindahl index of intermediate input use, computed from U.S. Input-Output Tables for 1992. Also, in the mid-2000s, Costinot (2009) devised a third alternative measure of contractibility related to the complexity of production, as captured by the average training time required to be qualified to work in that sectors (based on PSID survey questions). A fourth and final measure of contractibility I will experiment with below is the one proposed by (the also bright and young!) Bernard, Jensen, Redding and Schott (2010), which builds on the idea that products that are shipped across borders through intermediaries (such as wholesalers) are indirectly revealed to be more contractible. More specifically, their index is computed with U.S. census data as a weighted average of the wholesale employment share of firms importing a particular product. I will for now run with the idea that each of these four measures constitutes an empirical proxy for the degree of contractibility associated with offshoring (μ_S), but below I will highlight some caveats related with that interpretation.¹⁷

In order to explore the role of the degree of customization ϵ in shaping offshoring shares, I will again build on the methodology of Nunn (2007) but will instead build a measure of the average specificity of the good being transacted, rather than of the inputs used in the production of that good. More precisely, and following Antràs and Chor (2013), for each IO2002 sector, I calculate the fraction of ten-digit HS constituent codes classified by Rauch (1999) as neither reference-priced nor traded on an organized exchange (under Rauch's "liberal" classification).

The model also suggests that the degree of input substitutability ρ , which was irrelevant in the complete-contracting framework, should have a positive effect on offshoring shares. Because I will be restricting the sample to imports of intermediate inputs, a simple approach to proxy for ρ is to use the demand elasticity of the good being imported, as estimated by Broda and Weinstein (2006), since this should capture how substitutable that input is vis à vis other inputs. To better capture input substitution rather than differentiation by country of origin, I follow Antràs and Chor (2013) in using demand elasticities estimated at the three-digit industry level rather than at

¹⁷As described in the Data Appendix, each of these four sectoral measures of contract intensity were normalized so that higher levels imply higher contractibility or lower dependence on formal contract enforcement.

the ten-digit product level (see the Data Appendix for more details).

Finally, in Chapter 4, I developed a global sourcing model with sequential production that illustrated the potential role of downstreamness in the offshoring decision. I have not included that comparative static in the Table 5.5 because the sign of that effect depends on the environment in subtle ways, but below I will explore the effect of downstreamness in some specifications. To do so, I will employ the measure of downstreamness developed by Antràs and Chor (2013) and overviewed in the Data Appendix.

In Table 5.6, I report the results of introducing these nine variables (the two proxies for ϕ , the four for μ_S , plus those for ϵ , ρ , and downstreamness) into the specifications in Table 5.3. For simplicity, I focus on incorporating them into the Wright- and Nunn-Trefler-corrected regressions in columns (4), (5) and (6) of Table 5.3, and I do not report the coefficients on the variables already included in that table. These coefficients are, however, virtually unaffected by the inclusion of these contractually motivated variables.¹⁸

In the first three columns of Table 5.6, I report the results of adding these nine new variables *one at a time* to the regressions in columns (4), (5) and (6) of Table 5.3. Thus, although nine coefficients appear in each column, each of them is produced by a different regression. As is clear from the table, these nine variables have a small and imprecisely estimated effect on offshoring shares. In fact, of the 27 coefficients in those three first columns of Table 5.6, only four are statistically significant at the 5% level. Furthermore, the sign of these coefficients is frequently the opposite from the one predicted by theory and different proxies for the same variable often appear with opposite signs. For instance, we expect all four proxies of μ_S to appear with a positive coefficient, but close to half of those estimates are negative. Similarly, asset tangibility has a negative effect on offshoring shares, whereas the model predicts this effect to be positive (since the more tangible the assets, the lower should financial constraints ϕ be).

In columns (4), (5) and (6) of the table I report the results of regressions analogous to those in columns (1), (2) and (3), but in which a single proxy for financial constraints, a single proxy for contractibility, and the proxies for ϵ , ρ , and downstreamness are all included *in the same* specification. I choose to include Rajan and Zingales' financial dependence measure and Nunn's input relationship-specificity measure because they are the most widely used empirical proxies for financial constraints and contractibility. The results in the table speak for themselves and I will not attempt to sugarcoat them. Only

¹⁸The interested reader can consult the whole set of regression coefficients by accessing the data and programs available online at <http://scholar.harvard.edu/antras/books>.

one of those coefficients (for downstreamness in column (6)) appears statistically significant. I have also experimented with simultaneously including all nine contractual variables in the same regression, and the results were equally disappointing.

Table 5.6. Contractual Determinants of U.S. Offshoring Shares

Dep. Var.: $\frac{Imp}{Imp+Shipments-Exp}$	(1)	(2)	(3)	(4)	(5)	(6)
Financial Dependence	-0.010 (0.078)	-0.004 (0.009)	-0.004 (0.018)	-0.037 (0.079)	-0.000 (0.009)	-0.002 (0.019)
Asset Tangibility	-0.179** (0.069)	-0.009 (0.008)	-0.021 (0.017)			
Nunn Contractibility	-0.096 (0.060)	-0.005 (0.007)	-0.011 (0.016)	-0.037 (0.086)	-0.000 (0.008)	0.002 (0.016)
Levchenko Contractibility	-0.118* (0.049)	-0.000 (0.009)	0.004 (0.021)			
Costinot Contractibility	0.130 ⁺ (0.067)	0.008 (0.006)	0.018 (0.013)			
BJRS Contractibility	0.078 (0.071)	0.006 (0.006)	0.022 (0.013)			
Specificity	0.116* (0.055)	0.006 (0.006)	0.013 (0.014)	0.083 (0.080)	0.003 (0.008)	0.003 (0.016)
Input Substitutability	-0.019 (0.064)	-0.003 (0.005)	-0.012 (0.011)	0.008 (0.062)	-0.002 (0.006)	-0.010 (0.013)
Downstreamness	0.096 (0.084)	0.009 (0.007)	0.031* (0.016)	0.055 (0.094)	0.009 (0.008)	0.032 ⁺ (0.017)
Sample Restrictions	W+NT Year	W+NT Ctr/Year	W+NT ⁺ Ctr/Year	W+NT Year	W+NT Ctr/Year	W+NT ⁺ Ctr/Year
Observations	2,513	582,811	148,879	2,513	582,811	148,879
R-squared	≈0.15	≈0.19	≈0.20	0.168	0.200	0.199

Standard errors clustered at the industry level. ⁺, *, ** denote 10, 5, 1% significance. The sign \simeq indicates that this value is roughly the average R-squared across specifications.

Overall, the results in Table 5.6 provide no evidence in support of our contracting models of global sourcing. I next turn to discussing why that might be so, and how one can use alternative approaches to more cleanly evaluate the model.

Limitations and Alternative Approaches

The industry-level tests performed so far have failed to provide much supportive evidence for the importance of contractual factors in determining the global sourcing decisions of U.S. firms. The fact that the earlier tests of the benchmark complete-contracting model did not deliver much more robust results suggest, however, that perhaps part of blame for this rests on the approach we have followed so far. Indeed, the low R^2 in the above regressions indicates that most of the variation in offshoring shares is explained by ‘omitted’ factors. These omitted characteristics might well be correlated with the industry variables included in the regressions above, thus creating biases that have the potential to explain the poor results obtained so far.

Recent empirical literature in international trade has been well aware of these potential biases and has developed alternative strategies to identify the role of factor endowments and institutional factors in shaping comparative advantage and trade flows across countries. A particularly dominant approach builds on the seminal work of Rajan and Zingales (1998) and exploits the idea that industry characteristics should have a differential effect on trade flows (and on input flows in our context) across countries, depending on characteristics of these countries. This difference-in-difference approach was first applied in a trade context by Romalis (2004), who cast the Heckscher-Ohlin model as predicting an effect of capital intensity on export flows that is disproportionately large for physical capital abundant countries. As mentioned in Chapter 3, the recent empirical literature on institutions and trade has emphasized, in a similar vein, that differences in contracting institutions across countries should have a differential effect on trade flows in different sectors, depending on certain characteristics of these sectors.

In econometric terms, this approach advocates the inclusion of both country and industry fixed effects in regressions predicting trade flows, and tests the validity of models by inspecting the effect of interactions terms composed of industry and country characteristics. From an empirical point of view, adapting this methodology to the global sourcing environment I have been studying is relatively straightforward, particularly given my use of variation both across sectors and countries in some of my specifications above. From a theoretical point of view, this approach is also feasible because the different variants of our global sourcing model deliver comparative statics relating offshoring shares to interactions of parameters, which one might associate with country or industry characteristics. For instance, in Chapter 4, I showed that the positive effect of offshore contractibility μ_S on offshoring shares should be higher whenever the buyer’s demand elasticity σ or the de-

gree of customization ϵ are high, or whenever input substitutability ρ is low. It seems natural to associate offshore contractibility, at least partly, to the quality of contractual institutions in the exporting country, while the parameters σ , ϵ , and ρ can be thought of as being industry characteristics, as in the regressions above. Hence, our global sourcing model implies that the interaction of source country contract enforcement with proxies for σ , ϵ , and ρ , should have predictive power for U.S. imports of particular intermediate inputs from particular source countries.

As natural as this approach may appear, it is, however, not firmly grounded in the two-country model of global sourcing I have developed in Chapter 4. Before returning to the data, I will thus briefly describe a multi-country version of this global sourcing model that provides a semi-structural interpretation of the tests to be performed below.

Multi-Country Framework

Towards the end of Chapter 2, I discussed how to extend the two-country complete-contracting global sourcing model to a multi-country environment. The key for tractability was to follow Eaton and Kortum (2002) in modelling labor productivity as the realization of an extreme-value Fréchet random variable. In order to characterize the intensive and extensive margins of global sourcing, it also proved convenient to consider a richer environment in which production required the completion of a continuum of stages, with each of these stages being potentially produced in a different country. Adapting that framework to an incomplete-contracting environment raises important challenges, so I will instead focus on a version of the model in which each final-good producer procures only one input (as in the two-country model). Furthermore, I will restrict the analysis to a variant of the model in which the firm-level extensive margin of offshoring is not operative.

Let us now discuss the assumptions of the model in more detail. I now consider a framework with J countries in which final-good producers in every country combine locally produced headquarter services with a manufacturing input that can be procured from any of the J countries. To build intuition, let us consider first the complete-contracting version of the model. As in equation (2.24) in Chapter 2, the operating profits associated with a firm based in i using an input manufactured in country j are given by

$$\pi_{ij}(\varphi) = \left((a_{hi}w_i)^\eta (\tau_{ij}a_{mj}w_j)^{1-\eta} \right)^{1-\sigma} B\varphi^{\sigma-1} - f_{ij}w_i, \quad (5.5)$$

where B is now given by

$$B = \frac{1}{\sigma} \left(\frac{\sigma}{(\sigma - 1)P} \right)^{1-\sigma} \beta \sum_{j \in J} w_j L_j$$

and P is the common price index (2.4) for final-good varieties in each country. Remember the parameters a_{hi} and a_{mj} capture the unit labor requirements associated with headquarter service provision and manufacturing production, and they are allowed to vary across countries. Furthermore, while a_{hi} is a technological parameter common across firms based in i , the manufacturing productivity parameters $1/a_{mj}$ for all j are assumed to be firm-specific and drawn from a Fréchet distribution as in Eaton and Kortum (2002), so that

$$Pr(a_{mj} \geq a) = e^{-T_j a^\theta}, \quad \text{with } T_j > 0.$$

These Fréchet draws are assumed to be independent across firms and locations, and also orthogonal to core productivity φ .

As in Chapter 2, a firm obtains a productivity draw from a given country j only after paying the fixed cost f_{ij} of sourcing from country j . I will simplify matters relative to Chapter 2 by assuming that the fixed costs of sourcing f_{ij} are small enough such that all firms from i find it profitable to incur these costs and draw a parameter a_{mj} from each country $j \in J$. This is obviously a strong assumption, but it is worth emphasizing that it does not imply that firms will buy inputs from all countries in the world. In fact, because they only require a single input for production, firms will only buy inputs from a single market. As mentioned before, this will shut down the extensive margin of sourcing at the firm level.

Because firms only learn their particular realizations of $1/a_{mj}$ for each $j \in J$ after they have incurred all sunk costs of offshoring, the choice of location of production of the single manufacturing input will simply maximize the first term of the profit function in (5.5), which is analogous to choosing $j^* = \arg \min_{j \in J} \{\tau_{ij} a_{mj} w_j\}$. Importantly, this is true for all firms in i regardless of their core productivity φ . Appealing to the properties of the Fréchet distribution, we can then conclude that all firms from i will source inputs from country j with probability

$$\chi_{ij} = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\sum_{l \in J} T_l (\tau_{il} w_l)^{-\theta}}. \quad (5.6)$$

Since there are a continuum of firms in i we can then apply the law of large numbers to conclude that χ_{ij} in (5.6) will also constitute the share of

inputs purchased by firms in i that originate in j . Less trivially, but again following in a straightforward manner from the results in Eaton and Kortum (2002), the distribution of the actual price paid for any input turns out to be independent of the actual source j of that input, and thus χ_{ij} in (5.6) also corresponds to country j 's share of all manufacturing input purchases by firms from i .

With this machinery in hand, we can now reintroduce contractual frictions into this multi-country version of the model. Notice that apart from the initial vector of sunk costs of sourcing f_{ij} , all production decisions of and negotiations between final-good producers and suppliers are performed with knowledge of the realization of the vector of cost draws a_{mj} . It is then straightforward to verify that, in all versions of the global sourcing model developed in Chapter 4, the operating profits (net of sunk costs) associated with input manufacturing in a given location j can be written as

$$\pi_{ij}(\varphi) = ((a_{hi}w_i)^\eta (\tau_{ij}a_{mj}w_j)^{1-\eta})^{1-\sigma} \Gamma_{ij} B \varphi^{\sigma-1}, \quad (5.7)$$

where $\Gamma_{ij} < 1$ summarizes the reduction in profitability associated with incomplete contracting in the different versions of the model. For instance, in the version of the model with partial contractibility, we have from equation (4.22) that

$$\Gamma_{ij} = \left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_{ij})} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_{ij})} \left(\frac{1}{2} \right)^\sigma,$$

where μ_{ij} corresponds to the degree of contractibility associated with firms from i sourcing inputs from j .

Even though we have focused on a single-input version of the model, this multi-country version of the model can easily accommodate an extension with a continuum of inputs as long as these inputs are all produced in the same country j under the same labor productivity a_{mj} , in which case the index of contractual efficiency Γ_{ij} becomes (see equation (4.28))

$$\Gamma_{ij} = \left(1 + \frac{1}{\rho} \frac{(\sigma - 1)(1 - \eta)}{\sigma - (\sigma - 1)(1 - \mu_{ij})} \right)^{\sigma - (\sigma - 1)(1 - \mu_{ij})} \left(\frac{\rho\sigma}{\rho\sigma + (\sigma - 1)(1 - \eta)} \right)^\sigma. \quad (5.8)$$

Given that Γ_{ij} is not stochastic, it affects the profit function in equation (5.7) in a manner analogous to the trade cost τ_{ij} and the wage rate w_j , and we can use steps analogous to the ones above for the complete-contracting

case to conclude that the share of intermediate input purchases sourced from country j is given by

$$\chi_{ij} = \frac{T_j \left(\tau_{ij} w_j \Gamma_{ij}^{1/(1-\eta)(1-\sigma)} \right)^{-\theta}}{\sum_{l \in J} T_l \left(\tau_{il} w_l \Gamma_{il}^{1/(1-\eta)(1-\sigma)} \right)^{-\theta}}. \quad (5.9)$$

Hence, conditional on transport costs, technological productivity and wage costs, locations associated with worse perceived contract enforcement from the point of view of firms from i will tend to sell a relatively lower share of the intermediate inputs purchased by firms from country i .¹⁹

Empirical Implementation of the Multi-Country Model

In order to transition back to the empirical exploration of the model, it will prove useful to reintroduce input subscripts v and express the offshoring share χ_{ijv} associated with input v as

$$\chi_{ijv} = \frac{T_{jv} \left(\tau_{ijv} w_{jv} \Gamma_{ijv}^{1/(1-\eta_v)(1-\sigma_v)} \right)^{-\theta}}{\sum_{l \in J} T_{lv} \left(\tau_{ilv} w_{lv} \Gamma_{ilv}^{1/(1-\eta_v)(1-\sigma_v)} \right)^{-\theta}}, \quad (5.10)$$

where I will write Γ_{ijv} as

$$\Gamma_{ijv} = \Gamma \left(\sigma_v, \eta_v, \epsilon_v, \rho_v, \phi_{ij}, \mu_{ij} \right). \quad (5.11)$$

For simplicity, I omit time subscripts throughout, although remember that the estimation uses annual data for 2000-11. Notice from equation (5.10) that I am allowing the absolute advantage parameter T_{jv} and the wage rate w_{jv} to vary not only across countries but also across products. These features would complicate the general equilibrium of the model, but I have restricted attention to industry equilibria, so this cross-sectoral variation can be introduced at little cost. Similarly, I am allowing trade costs (or tariffs) τ_{ijv} to vary both across countries and sectors. Note also that the index of contract efficiency Γ_{ijv} associated with firms from i sourcing from country j inputs of type v is written in (5.11) as a function of product characteristics $(\sigma_v, \eta_v, \epsilon_v, \rho_v)$ and country-pair characteristics (ϕ_{ij}, μ_{ij}) .

¹⁹As in the two-country model, one needs to take a stance of how inputs are priced in an incomplete-contracting framework. To derive equation (5.9), I assume as in Chapter 4 that input expenditures constitute the same multiple of operating profits in all countries.

These choices are not without loss of generality and warrant some discussion. I view it as natural to assume that headquarter intensity η_v and specificity ϵ_v are largely sectoral characteristics, independent of the country that exports an input. Similarly, the parameters σ_v and ρ_v govern substitutability across final goods and inputs, and although these could presumably vary across countries, the data I am using to proxy for them (from Broda and Weinstein, 2006) provides a unique sectoral estimate based on U.S. import data. Conversely, I will for the most part treat the degree of financial constraints ϕ_{ij} and contractibility μ_{ij} as country (or country-pair) characteristics. In some specifications, I will, however, allow the effect of financial constraints and contractibility to be expressed as an interaction of a sector-specific component and a country-pair-specific component, so we can write these as $\phi_{ijv} = \phi_v \times \phi_{ij}$ and $\mu_{ijv} = \mu_v \times \mu_{ij}$.

It might be useful to employ equations (5.10) and (5.11) to illustrate some of the limitations of the cross-industry empirical tests developed earlier in this Chapter. Above, we argued that the inclusion of source-country-year fixed effects helped isolate the effect of sectoral-level characteristics on offshoring shares. Indeed, if the inclusion of country-year fixed effects completely partialled out the effect of any country level variable on offshoring shares χ_{ijv} , then all that would be left for identification would be the cross-sectoral variation in the data. It is pretty clear from the formula for χ_{ijv} in (5.10), however, that demeaning offshoring shares within countries (and years) will not eliminate the effects of country-level variables, since these effects interact with industry-level variables. This is in fact true even when we consider a log-transformation of offshoring shares, in which case we can rewrite (5.10) as

$$\ln \chi_{ijv} = \ln T_{jv} - \theta \ln w_{jv} - \theta \ln \tau_{ijv} + \frac{\theta}{(1 - \eta_v)(\sigma_v - 1)} \ln \Gamma_{ijv} + \alpha_{iv}, \quad (5.12)$$

where α_{iv} is an importer/product fixed effect given by

$$\alpha_{iv} = -\ln \left(\sum_{l \in J} T_{lv} \left(\tau_{ilv} w_{lv} \Gamma_{ilv}^{1/(1-\eta)(1-\sigma)} \right)^{-\theta} \right).$$

Only when T_{jv} , w_{jv} , τ_{ijv} , and Γ_{ijv} can all be decomposed into the product of a sector-specific term and a country-specific term, will country-year fixed effects effectively partial out the effect of country-level variables. Nevertheless, it should be clear from the formulas for Γ_{ijv} – such as equation (5.8) – that this decomposability is not a feature satisfied by the index of contractibility Γ_{ijv} . As demonstrated in Chapter 4 and also in the Theoretical Appendix,

the partial derivative of $\ln \Gamma_{ijv}$ with respect to industry characteristics is generally affected by country-level variables, such as μ_{ij} or ϕ_{ij} .

In light of these interaction effects, I will next present the results of empirical specifications that include both product and country-year fixed effects and judge the validity of our global sourcing model based on the predicted effect of the interaction of sector and country characteristics. More specifically, we can express the specification equation as

$$\ln \chi_{ijv} = \alpha_{iv} + \alpha_{ij} + \beta \mathbf{Z}_{ij} \mathbf{z}_v + \gamma \ln \tau_{ijv} + \delta \ln \Gamma_{ijv} + \varepsilon_{ijv}. \quad (5.13)$$

Because in the empirical application we are fixing the importing country to be the U.S., the terms α_{iv} and α_{ij} in this expression are effectively sector and exporting-country-year fixed effects. The effect of technological or factor endowments differences as sources of comparative advantage is captured by the interaction terms $\mathbf{Z}_{ij} \mathbf{z}_v$, that of trade frictions (freight costs and tariffs) is represented by $\ln \tau_{ijv}$, while the vector $\ln \Gamma_{ijv}$ summarizes the effect of interactions of the primitive parameters of the model on the logarithm of the index of contractual efficiency Γ_{ijv} . In light of our results in Chapter 4 and the Theoretical Appendix, we can write $\ln \Gamma_{ijv}$ in terms of the following interaction terms (and predicted effects)

$$\ln \Gamma_{ijv} = \Phi \left(\begin{array}{cccc} \mu_{ij} \times \rho_v, & \mu_{ij} \times \sigma_v, & \mu_{ij} \times \epsilon_v, & \mu_{ij} \times \eta_v, \phi_{ij} \times \eta_v \\ - & + & + & \text{ambiguous} \quad + \end{array} \right). \quad (5.14)$$

Although the log-linear specification in (5.13) is quite standard in the literature, it has the downside of dropping all the observations with zero import flows. I have also experimented with linear specifications in which the dependent variable is the share χ_{ijv} instead of its logarithm. The qualitative results I obtained in those regressions were similar to those reported below, but the R^2 were orders of magnitude smaller than in log-linear specifications. Given that the model I have used to motivate the empirical analysis does not feature an extensive margin of offshoring, and thus captures variation only across positive trade flows, I will focus on discussing the results of the log-linear specification.

A Brief Detour into Previous Empirical Studies

Before formally testing our global sourcing model, I briefly discuss the results of running specifications analogous to the one in (5.13) but with a number of institutional interactions that have been suggested in the literature in recent

years (rather than those suggested by the global sourcing model developed in this book). The goal of this exercise is to document that the results obtained by other authors continue to apply when studying the determinants of U.S. imports, even when restricting the analysis to U.S. imports of intermediate inputs using the sample restrictions described above.

The specifications to be discussed below are most closely related to the work of Nunn (2007) and Levchenko (2007). More specifically, I follow Nunn (2007) in studying log-linear specifications in which trade flows are projected on sectoral and country-year fixed effects, together with Heckscher-Ohlin interactions associated with physical capital and skilled labor and a series of institutional interactions. By focusing on U.S. imports across sectors and countries (rather than worldwide exports of individual countries) I follow the approach in Levchenko (2007). The key novelties of the analysis below are that (i) similarly to Chor (2010), I experiment with the inclusion of a wide set of institutional interactions, and (ii) I attempt to capture the effect of these variables on the global sourcing decisions of U.S. firms, rather than on overall U.S. imports.²⁰

I experiment below with the inclusion of four interactions associated with contracting institutions, two related to financial institutions and one reflecting the role of labor-market institutions. Although I do not theoretically motivate these specifications, the contracting and financial interactions can be thought of as corresponding to the effect of ϕ_{ijv} and μ_{ijv} in the global sourcing model, whenever these are expressed as an interaction of a sector-specific component and a country-pair-specific component ($\phi_{ijv} = \phi_v \times \phi_{ij}$ and $\mu_{ijv} = \mu_v \times \mu_{ij}$).

The four contract enforcement interactions correspond to the product of the exporter's rule of law averaged over 2000-05 (from the Worldwide Governance Indicators) with the contract intensity measures created by Nunn (2007), Levchenko (2007), Costinot (2009), and Bernard, Jensen, Redding and Schott (2010), all of which were described above. These industry variables were normalized so that higher values of these variables are associated with lower dependence on formal contract enforcement (see the Data Appendix).²¹ Following Manova (2012)'s work, the two financial institutions interactions are the product of the exporter's log private credit to GDP ratio averaged over 2000-05 (from the World Development Indicators), with the

²⁰See Nunn and Treffer (2013a) for an overview of the literature on trade and institutions.

²¹In the regressions including Costinot's product complexity measure, I follow Costinot (2009) in always including as a control variable the interaction of this industry measure with the measure of skilled labor abundance used in the Heckscher-Ohlin skilled labor interaction.

external financial dependence measure of Rajan and Zingales (1998) and the asset tangibility measure of Braun (2002) (also discussed above). Finally, the labor-market institutions interaction corresponds to the one developed by Cuñat and Melitz (2012), which is the product of the labor market flexibility measure developed by Botero, Djankov, Porta, Lopez-de Silanes and Shleifer (2004) and a measure of average firm-level sales volatility in an industry (which captures the need for labor reallocations across firms within a sector). In addition, all specifications will include two Heckscher-Ohlin interactions that are constructed with standard measures of physical capital and skilled labor intensity and relative abundance, as described in the Data Appendix. In Table 5.7, I do not report the coefficient on these Heckscher-Ohlin interactions but they generally appear positive and statistically significant in explaining U.S. imports and offshoring shares.

Column (1) of Table 5.7 reports the results of adding the seven institutional interactions *one at a time* to a simple OLS regression of the log of U.S. imports from a given country in a given sector on sectoral and country-year fixed effects and the two Heckscher-Ohlin interactions. All the coefficients in the table are beta coefficients. Due to data availability, some specifications include fewer observations than others, but these differences are small (the number of observations ranges from 180,653 to 196,584). As is clear from the results in column (1), all seven institutional interactions appear with the expected sign, are sizeable in magnitude and are statistically significant at the extremely low significant levels. More precisely, better rule of law increases exports to the U.S. disproportionately less in sectors that are less dependent on formal contract enforcement. Furthermore, higher financial development increases exports to the U.S. disproportionately more in sectors with higher external capital dependence or low asset tangibility, while more flexible labor markets foster exports to the U.S. disproportionately in sectors with high sales volatility.

In column (2) I re-run these seven specifications but applying the Wright and Nunn-Trefler corrections to U.S. imports in an attempt to restrict the analysis to intermediate input purchases by U.S.-based firms. This leads to a loss of about 35% of observations associated with U.S. imports of final goods, and remember that it also modifies U.S. imports differentially across sectors and countries. Despite these modifications, the results in column (2) are quite comparable to those in column (1) and are suggestive of the importance of the seven institutional interactions for the global sourcing decisions of U.S. firms. In the remainder of the Table, the analysis is restricted to these Wright and Nunn-Trefler corrected U.S. imports.

In column (3), I follow Chor (2010) in including the seven institutional

interaction terms in the *same* specification. This leads to a noticeable reduction of the partial effect of each of these interactions, but with the exception of the Rajan-Zingales interaction, all of the other key explanatory variables remain significant.

A natural concern with the results in column (3) is that the included interaction terms simply capture the effect of alternative interaction effects that are omitted from the specification. For instance, one might worry that the interaction of the contract intensity variables and rule of law simply captures the fact that richer countries (which typically have better rule of law) tend to specialize in relatively complex goods (which typically are recorded as being contract intensive) for reasons distinct from contractual considerations. A commonly used way to deal with this concern is to include interactions of all seven industry-level institutional variables with a measure of overall development, such as GDP per capita. In column (4), I report the results associated with that specification. The impact of that robustness test on the estimates is more significant. On the plus side, the Nunn, Levchenko, Braun and Cuñat-Melitz interactions retain their expected sign as well as their economic and statistical significance. The Costinot and Bernard et al. interaction terms also continue to have the expected sign, but are now indistinguishable from zero. More puzzlingly, the addition of these GDP per capita interactions now reverses the sign of the Rajan-Zingales interaction, thus indicating that better financial development is associated with higher exports in sectors with low external dependence.

Table 5.7. Contractual Determinants of U.S. Offshoring Shares

Dep. Var.: $\ln\left(\frac{Imp}{Imp+Ship.-Exp}\right)$	(1)	(2)	(3)	(4)	(5)
Nunn \times Rule	-0.139** (0.012)	-0.175** (0.014)	-0.051** (0.019)	-0.152** (0.033)	-0.134** (0.033)
Levchenko \times Rule	-0.165** (0.009)	-0.166** (0.010)	-0.123** (0.013)	-0.076** (0.026)	-0.087** (0.026)
Costinot \times Rule	-0.242** (0.014)	-0.178** (0.018)	-0.038 ⁺ (0.021)	-0.015 (0.031)	-0.019 (0.032)
BJRS \times Rule	-0.270** (0.016)	-0.178** (0.022)	-0.118** (0.025)	-0.053 (0.045)	-0.048 (0.045)
Rajan-Zingales \times Credit/GDP	0.309** (0.025)	0.272** (0.029)	0.059 (0.037)	-0.200* (0.096)	0.041 (0.044)
Braun \times Credit/GDP	-0.392** (0.030)	-0.400** (0.035)	-0.185** (0.047)	-0.187** (0.054)	-0.169** (0.053)
Firm Volatility \times Labor Flexibility	0.123** (0.025)	0.119** (0.028)	0.076** (0.029)	0.100** (0.029)	0.101** (0.029)
Sample Restrictions	$\Upsilon_{O>0}$	W+NT ⁺	W+NT ⁺	W+NT ⁺	W+NT ⁺
Ctr/Year & Ind Fixed Effects	Yes	Yes	Yes	Yes	Yes
Interactions with GDP pc	No	No	No	Yes	No
Industry Effects \times GDP pc	No	No	No	No	Yes
Observations	$\simeq 190,000$	$\simeq 125,000$	120,034	120,034	120,034
R-squared	$\simeq 0.610$	$\simeq 0.607$	0.622	0.623	0.637

Standard errors clustered at the country/ind. level. ⁺, *, ** denote 10, 5, 1% significance.

Finally, in column (5) I perform an even more stringent robustness test by incorporating interactions of GDP per capita with sectoral dummies (as advocated by Nunn and Treffer, 2013a). This is motivated by the same concerns as in column (4), but these wide set of interactions allow the level of development to affect U.S. intermediate input imports in each individual sector differentially in an unrestricted way. The results of this test are very similar to those in column (4), except that the Rajan and Zingales interaction term recovers its expected sign (though not in a statistically significant fashion).

Overall, the results in Table 5.7 indicate that even when restricting attention to the cross-country and cross-industry determinants of U.S. imports of intermediate inputs, we find similar results to those obtained in previous contributions studying the institutional determinants of comparative advantage. This fact enhances our confidence in the use of very similar specifications with the goal of testing our global sourcing model, a task to which we turn next.

Back to the Test of the Multi-Country Sourcing Model

We now return to the specification in (5.13) and (5.14) motivated by our global sourcing model. In order to transition to a precise estimating equation, I first let the vector $Z_{ij}z_v$ comprise again an interaction k_vK_j of capital intensity and capital relative abundance of the exporting country and an analogous interaction s_vS_j based on skilled-labor intensity and skilled-labor relative abundance data (see the Data Appendix for more details). Second, I proxy for trade frictions τ_{ijv} with measures of freight costs and tariffs analogous to those used in the cross-industry regressions above, but with those variables being computed at the country and sectoral level with U.S. import data. Finally, I assume that the function Φ in (5.11) is linear in its interactions, so we can succinctly write the empirical specification as the follows:

$$\begin{aligned} \ln \chi_{jv} = & \alpha_v + \alpha_j + \beta_1 k_v K_j + \beta_2 s_v S_j + \gamma_1 freight_{jv} + \gamma_2 tariff_{jv} \\ & + \delta_1 \rho_v \mu_j + \delta_2 \sigma_v \mu_j + \delta_3 \epsilon_v \mu_j + \delta_4 \eta_v \mu_j + \delta_5 \eta_v \phi_j + \varepsilon_{jv}. \end{aligned} \quad (5.15)$$

As mentioned above, in constructing the institutional interactions, I proxy ρ_v , σ_v , and ϵ_v in the same manner as in the cross-industry tests above, while μ_j and ϕ_j correspond to the rule of law and the log private credit over GDP ratio of the exporting country, respectively. Headquarter intensity η_v is measured as the first principal component from a factor analysis of the buyer versions of R&D, equipment capital and skill intensity variables. I have dropped the subscript i referring to the importing country because in the results below, i is always the U.S. Notice also that because the denominator of χ_{ijv} in (5.10) is common for all exporting countries j , one can simply replace the dependent variable $\ln \chi_{jv}$ with the logarithm of U.S. imports from country j in sector v , which is the same dependent variable included in the results in Table 5.7.

In light of the global sourcing model, one would expect the coefficients β_1 , β_2 , δ_2 , δ_3 and δ_5 to be positive, and the coefficients γ_1 , γ_2 , and δ_1 to be negative. The model instead does not provides an unambiguous prediction for the sign of δ_4 (though it suggests that this interaction should affect trade flows, which motivates its inclusion).

The first column of Table 5.8 reports the results of estimating equation (5.15) without the institutional interactions, but with the trade costs variables, which were missing in Table 5.7. Both Heckscher-Ohlin interactions affect U.S. imports in a positive and significant way, with the effect being particularly highly significant for the case of the skilled labor interaction. Freight costs and tariffs in turn have a negative effect on U.S. imports, with the former effect being significant at the 1% level and the latter at 10% level. In column (2) of Table 5.8, I re-run the same specification, but this time

attempting to restrict the sample to U.S. imports of intermediate inputs by U.S.-based firms by applying the Wright and Nunn-Trefler corrections. The effect on the coefficients is rather modest, except for the effect of tariffs, which is now notably larger in absolute terms and statistically significant at the 5% level.

I next experiment with the inclusion of the five interactions motivated by our model to the specification in column (2). I first do so in column (3) by introducing these interactions *one at a time*, so even though all coefficients appear on the same column, it should be understood that these are obtained by running five separate regressions.²² Interestingly, each of the five institutional interactions appears to be highly significant and with a sign consistent with the theory. More specifically, better rule of law appears to foster U.S. imports disproportionately in sectors with lower input substitutability, higher buyer elasticities of demand, higher input specificity, and higher headquarter intensity. Furthermore, higher levels of financial development also have a differentially higher effect on U.S. imports in sectors with higher headquarter intensity.

In column (4), I present results in which the five institutional interactions are included in the *same* regression, together with the Heckscher-Ohlin interactions and the trade cost measures. Analogously to the results obtained when doing the same in Table 5.7, the partial effect of each of the independent variables on U.S. imports of inputs is lower than when included in isolation. Nevertheless, all coefficients retain their theoretically predicted sign and are highly significant, with the exception of the input elasticity times rule of law interaction ($\rho_v \times \mu_j$), which remains negative but is now statistically indistinguishable from zero.

²²This also explains why the Heckscher-Ohlin and trade cost coefficients are not reported in column (3). These omitted coefficients and their standard errors vary slightly across specifications, but they are very close in levels and statistical significance to those reported in column (2), except that the positive and significant effect of the physical capital Heckscher-Ohlin interaction is not robust to the inclusion of the interactions involving headquarter intensity.

Table 5.8. Testing the Global Sourcing Model

Dep. Var.: $\ln\left(\frac{Imp}{Imp+Ship.-Exp}\right)$	(1)	(2)	(3)	(4)	(5)	(6)
K Intensity \times K Abund.	0.120*	0.151*		0.380**	0.357**	0.469
	(0.058)	(0.069)		(0.078)	(0.081)	(0.294)
Skill Inten \times Skill Abund	0.435**	0.467**		0.252**	0.251**	0.118*
	(0.028)	(0.031)		(0.034)	(0.038)	(0.046)
Freight Costs	-0.102**	-0.085**		-0.089**	-0.089**	-0.089**
	(0.018)	(0.010)		(0.010)	(0.010)	(0.010)
Tariffs	-0.015 ⁺	-0.023*		-0.018 ⁺	-0.018 ⁺	-0.015 ⁺
	(0.008)	(0.011)		(0.010)	(0.011)	(0.009)
Input Substit. \times Rule			-0.037**	-0.009	-0.026 ⁺	-0.012
			(0.009)	(0.009)	(0.016)	(0.016)
Demand Elasticity \times Rule			0.026**	0.027**	0.001	-0.002
			(0.008)	(0.008)	(0.012)	(0.016)
Nunn Specificity \times Rule			0.189**	0.164**	0.255	0.224**
			(0.015)	(0.016)	(0.161)	(0.030)
Headq. Inten. \times Rule			0.093**	0.050**	0.050**	0.047**
			(0.010)	(0.012)	(0.013)	(0.012)
Headq. Int. \times Credit/GDP			0.074**	0.045**	0.044**	0.045**
			(0.007)	(0.009)	(0.012)	(0.012)
Sample Restrictions	$\Upsilon_{O>0}$	W+NT ⁺	W+NT ⁺	W+NT ⁺	W+NT ⁺	W+NT ⁺
Ctr/Year & Ind Fixed Eff	Yes	Yes	Yes	Yes	Yes	Yes
Interactions with GDP	No	No	No	No	Yes	No
Industry Effects \times GDP	No	No	No	No	No	Yes
Observations	188,187	128,482	\simeq 127,999	126,068	126,068	126,068
R-squared	0.601	0.619	\simeq 0.621	0.624	0.624	0.641

Standard errors clustered at the country/ind. level. ⁺, *, ** denote 10, 5, 1% significance.

Finally, in columns (5) and (6) I perform the same robustness tests as in columns (4) and (5) of Table 5.7 by first including interactions of the main industry-level institutional variables with GDP per capita, and later by controlling for a whole vector of interactions of sectoral dummies with GDP per capita. The inclusion of these controls reduces the statistical significance of some of the coefficients, most notably that of the interaction of the elasticity of demand with rule of law ($\sigma_v \times \mu_j$), but the results still provide broad support for the empirical validity of some of the key predictions of our global sourcing model.

I have also experimented with specifications that include all the institutional interactions from the previous literature in Table 5.7 with the new ones motivated by our global sourcing model. Consistently with the results above, the freight cost measure as well as the interactions of Nunn specificity with rule of law and headquarter intensity with financial development appear to be very robust and continue to affect U.S. input imports volumes with the expected sign and at high levels of statistical significance. We also generally find a negative and significant effect of tariffs on U.S. imports of intermediate inputs and a positive and significant effect of the interaction of headquarter intensity with rule of law. Conversely, the effects of the interactions of input substitutability and the elasticity of demand with rule of law are much less robust, particularly when adding interactions of industry variables or industry fixed effects with GDP per capita. Also consistently with Table 5.7, the Nunn, Levchenko, Braun and Cuñat-Melitz interactions continue to be robust predictors of U.S. global sourcing decisions, retaining throughout their expected sign as well as their economic and statistical significance. These results can easily be replicated with the dataset and programs downloadable at <http://scholar.harvard.edu/antras/books>.

Concluding Remarks

Overall, it would be hard to argue that the results presented in this chapter provide resounding evidence supporting the empirical validity of the global sourcing model (and its various variants) developed in Chapter 4. Admittedly, some of the key predictions of the framework, such as its cross-industry implications for the determinants of offshoring shares, have been hard to validate with U.S. import data. Nevertheless, I have argued that this is in part due to the fact that these cross-industry specifications might be prone to serious econometric biases. When adopting a cleaner approach exploiting both the cross-country as well as the cross-industry variation in the data, while controlling for industry and country-year fixed effects, we have obtained much more favorable results for the model. For instance, controlling for standard Heckscher-Ohlin effects, U.S. firms appear to rely less on offshoring when trade costs (and particularly freight costs) are high. Variation in contract enforcement across countries also appears to be an important determinant of the observed variation in the propensity of U.S. firms to offshore, with the effect often being differentially higher precisely in those sectors in which the model predicts that the effects should be disproportionately higher.

A critical challenge in the empirical analyses performed in this chapter is that the key industry characteristics (specificity, input substitutability, de-

mand elasticities, headquarter intensity,...) that shape the differential effect of contract enforcement on the profitability of offshoring across sectors are particularly hard to measure in the data. I have devoted significant space and effort to discussing the hurdles that one encounters when mapping these sectoral characteristics with available data. Similar challenges emerged in isolating the intermediate input component of U.S. imports. Although I have attempted to surpass these hurdles, I have obviously done so in imperfect ways, and this might explain some of the less favorable results I obtained when attempting to validate certain predictions of the model. This is obviously a favorable interpretation of the few negative results obtained above, but future research with newer data sources and more ingenious empirical strategies will ultimately verify or falsify that the contractual determinants of offshoring highlighted in this book are indeed a central feature of the global sourcing decisions of firms.