

GLOBAL PRODUCTION:
A CONTRACTING PERSPECTIVE

Pol Antràs

April 2014

Contents

Preface	v
1 Made in The World	1
2 Workhorse Models	25
3 Contracts and Export Behavior	57
4 Contracts and Global Sourcing	95
5 Contracts and Sourcing: Evidence	131
6 The Internalization Decision	173
7 Internalization: Empirical Evidence	175
8 Some Implications for Policy	177
A Theoretical Appendix	179
B Data Appendix	193

Preface

This book has grown out of the *CREI Lectures in Macroeconomics* that I delivered in Barcelona in June of 2012. I am grateful to the Series Editor, Hans-Joachim Voth, for inviting me to deliver these lectures and for encouraging me to accept the invitation. Part of my initial reticence, I can now admit, was related to the fact that I viewed the invitation to give these prestigious lectures as a flagrant case of ‘home bias’. At the time, the CREI Lecture series committee was composed of a coauthor of mine (Hans-Joachim Voth), one of my Ph.D. advisors (Jaume Ventura) and two of my favorite teachers during my undergraduate studies at Universitat Pompeu Fabra (Antonio Ciccone and Jordi Galí). Regardless of their motivations, it was an honor to have been selected as the 2012 CREI Lectures speaker. I am grateful for the comments and feedback I received during my lectures and also for CREI’s hospitality during my many visits there.

The book is largely aimed at graduate students and researchers interested in learning about recent developments in the field of International Trade. I have attempted, however, to make the style of the book a bit less terse than is standard in professional journals and graduate-level textbooks. This may alienate some technically-oriented readers, but will hopefully encourage some advanced undergraduate students and trade practitioners to venture into the material in this book. Chapter 1, in particular, provides an overview of the topics covered in later chapters at a highly accessible level. The book contains an extensive Theoretical Appendix, which will hopefully earn me the forgiveness of some mathematically-inclined readers. Although it would be hard to sell this book as being a set of Lectures in *Macroeconomics*, I hope that some of the material in this lecture will appeal to researchers in that field, as well as readers interested in Organizational Economics and Applied Contract Theory.

Although the bulk of the contents of this book has appeared in some

form or other in academic journals, many chapters of this book include new and original work. For instance, the multi-country global sourcing model introduced in Chapter 2 and further developed in Chapters 5, and 6 is a novel contribution of this book (although, it is being explored in parallel work by Antràs, Fort and Tintelnot, 2014). Similarly, I am not aware of the existence of multi-country models of limited commitment of the type developed in Chapter 3. Many of the empirical parts of the book are original as well, although they build heavily on previous work in terms of both methodology and data sources.

I have taught most of the material in these lectures at Harvard but also at Study Center Gerzensee, the London School of Economics, Penn State University, the University of Zurich and Northwestern University. I have found that between four and five 90-minute lectures are generally sufficient to cover the contents of this book. I am grateful to all these institutions for their hospitality and to the lecture participants for many useful comments.

Although I have attempted to provide a broad overview of the topics in this lecture, the spirit of the CREI Lectures dictated me to have my own work feature prominently in this book. For this reason, my greatest debt is to my coauthors on the papers overviewed in this book, including Daron Acemoglu, Davin Chor, Fritz Foley, Esteban Rossi-Hansberg, Bob Staiger, Steve Yeaple, and especially Elhanan Helpman. I am also particularly grateful to my colleagues Elhanan Helpman and Marc Melitz for many stimulating discussions that have shaped my thinking on the material of this book. My interest in the contracting aspects of global production dates back to my Ph.D. years at MIT and I am indebted to Daron Acemoglu, Gene Grossman, Bengt Holmström and Jaume Ventura for their encouragement during those initial phases of this intellectual adventure. I am also very grateful to the Bank of Spain for generously funding my research at that crucial early stage.

Turning my lecture slides into a book manuscript has proved to be much harder and time-consuming than I first anticipated. Mireia Artigot, Teresa Fort, Elhanan Helpman and Felix Tintelnot read different parts of the first draft of this book and provided very useful feedback and corrections. I am also grateful to Eric Unverzagt for his careful editorial assistance.

Several colleagues have kindly shared their data for some of the empirical material in this book. These include Andrew Bernard, Davin Chor, Robert Johnson, Nathan Nunn, Mike Waugh, and Greg Wright. I have also benefited from the outstanding research assistance of Ruiqing Cao, Yang Du, Alonso de Gortari, and Boo-Kang Seol during various periods over which this book

was written. They are of course not responsible for any mistakes left in the manuscript. Finally, I am forever grateful to my wife Lucia and my daughters Daniela and Martina for their patience during the many hours I have spent mulling over and writing this book.

Chapter 1

Made in The World

It does not feel like a long time ago when I began my undergraduate studies in Economics at Universitat Pompeu Fabra (UPF), the same institution that hosts the Centre de Recerca en Economia Internacional (CREI), where these Lectures were delivered. It was 1994 and I felt I lived in a truly global economy. The music I listened to and the movies I watched were mostly British or American. Most of the clothes I wore were manufactured abroad, some of them in rather exotic places such as Morocco or Taiwan. My favorite beer was Dutch. At UPF, about half of my teachers were foreign, a third of the classes were taught in English and most of the textbooks were the same ones used in universities around the globe.

In hindsight, it seems pretty clear, however, that the world had not yet witnessed the full advent of globalization. What has changed since 1994? *First* and foremost, the last two decades have brought a genuine information and communication technology (ICT) revolution that has led to a profound socioeconomic transformation of the world in which we live. The processing power and memory capacity of computers have doubled approximately every two years (as implied by Moore's law), while the cost of transmitting a bit of information over an optical network has decreased by half roughly every nine months (a phenomenon often referred to as Butter's law). The number of internet users has increased by a factor of 100, growing from around 20 million users in 1994 to more than 2,000 million users in 2010 (see World Development Indicators). As a result of these technological developments, the cost of processing and transmitting information at long distances has dramatically fallen in recent times. Consider the following example: the 3.3GB file containing my favorite movie of 1994, *Pulp Fiction*, can be downloaded

today from Amazon.com in about 11 minutes and 16 seconds using a standard broadband connection with a download speed of 5 megabits per second. In 1994, downloading that same file using a dial-up connection and the state-of-the-art modem, which allowed for a maximum speed of 28.8 kilobits per second, would have kept your phone line busy for at least 33 hours and 23 minutes!¹

Second, during the same period, governments have continued (and arguably intensified) their efforts to gradually dismantle all man-made trade barriers. This process dates back to the initial signing of the General Agreement on Tariffs and Trade (GATT) in 1947, but it has experienced a revival in the 1990s and 2000s with the gradual expansion of the European Union, the formation of the North American, Mercosur, and ASEAN free trade agreements, the signing of a multitude of smaller preferential trade agreements under the umbrella of GATT's Article XXIV, and China's accession to the World Trade Organization (WTO), just to name a few. As a consequence, the world's weighted average tariff applied on traded manufactured goods fell from 5.14% in 1996 to 3.03% in 2010 (see World Development Indicators).²

Third, political developments in the world have brought about a remarkable increase in the share of world population actively participating in the process of globalization. These changes largely stemmed from the fall of communism in Eastern Europe and the former Soviet Union, but also from an ensuing ideological shift to the right in large parts of the globe. Thus, not only did former communist countries embrace mainstream capitalist policies, but these policies themselves became more friendly towards globalization, as exemplified by the deepening of trade liberalization mentioned in the last paragraph, but also by a notable relaxation of currency convertibility and balance of payments restrictions in several low and middle-income countries.³

¹Paraphrasing a memorable quote from Samuel L. Jackson's character in *Pulp Fiction*, download speeds today and in 1994 "ain't the same [*freaking*] ballpark. They ain't the same league. They ain't even the same [*freaking*] sport."

²Technological developments since 1994 have also reduced the quality- (or time-) adjusted costs of transporting goods across countries (see Hummels, 2007), while investments in infrastructure in less developed economies have also contributed to spreading the effects of globalization across regions in those countries.

³The late 1990s also saw the emergence of a left-leaning anti-globalization movement, which drew particular attention during the 1999 WTO meetings in Seattle. There is little evidence, however, of this movement having led to any significant slow down in the process of globalization (see, for instance, Harrison and Scorse, 2010).

The Slicing of the Value Chain

One of the manifestations of these three developments in the world economy has been a gradual disintegration of production processes across borders. More and more firms now organize production on a global scale and choose to offshore parts, components or services to producers in foreign and often distant countries. The typical “Made in” labels in manufactured goods have become archaic symbols of an old era. These days, most goods are “Made in the World.”

A variety of terms have been used to refer to this phenomenon: the “slicing of the value chain”, “fragmentation of the production process”, “disintegration of production”, “delocalization”, “vertical specialization”, “global production sharing”, “unbundling”, “offshoring” and many more (see Feenstra, 1998). I shall use these terms interchangeably throughout the book.⁴

The case of Apple’s iPad 3 tablet nicely illustrates the magnitude of this new form of globalization. The slim and sleek exterior of the tablet hides a complex manufacturing process combining designs and components provided by multiple suppliers with operations in various countries. Although Apple does not disclose detailed information on its input providers, a clear picture of the global nature of the iPad 3 production process emerges when combining information from tear-down reports (such as those published by isuppli.com and ifixit.com) with various press releases.⁵ For instance, it is well-known that the tablet itself is assembled in China (and since 2012 also in Brazil) by two Taiwan-based companies, Foxconn and Pegatron. The revolutionary retina display is believed to be manufactured by Samsung of South Korea in its production plant in Wujiang City, China. The distinctive touch panel is produced (at least, in part) by Wintek, a Taiwan-based company

⁴At times, I will also use the buzzword “outsourcing”, but I will do so only when referring to *arm’s-length* sourcing relationships, that is instances of fragmentation in which the firms exchanging parts are not related (i.e., integrated). Outsourcing is often observed not only in foreign but also in domestic vertical relationships.

⁵Facing strong criticism over the working conditions in its suppliers’ factories, Apple released a full list of its 156 global suppliers early in 2012 (see http://images.apple.com/supplierresponsibility/pdf/Apple_Supplier_List_2011.pdf). Teardown reports further facilitate a mapping between the iPad parts and their respective producers. Press releases sometimes also identify particular suppliers with specific iPad 3 components (see, for instance, Forbes’ “Batteries Required?” available at <http://www.forbes.com/global/2010/0607/best-under-billion-10-raymond-sung-simple-technology-batteries-requried.html>).

that also owns plants in China, India and Vietnam, while the case is provided by another Taiwanese company, Catcher Technologies, with operations in Taiwan and China. A third important component, the battery pack, also originates in Taiwan and is sold by Simplo Technologies and Dynapack International. Apart from these easily identifiable parts, the iPad 3 incorporates a variety of chips and other small technical components provided by various firms with headquarters and R&D centers in developed economies and manufacturing plants scattered around the world. A non-exhaustive list includes (again) Korea's Samsung, which is believed to manufacture the main processor (designed by Apple), U.S.-based Qualcomm supplying 4G modules, and Italo-French STMicroelectronics contributing key sensors.⁶

Apple's sourcing strategies are far from being an isolated example of a global approach to the organization of production. In fact, the increasing international disintegration of production processes has been large enough to be salient in aggregate statistics. During the 1990s and early 2000s, when this phenomenon was still in its infancy, researchers devised several approaches to measuring the quantitative importance of global production sharing.⁷ Feenstra and Hanson (1996*b*), for instance, used U.S. Input-Output tables to infer the share of imported inputs in the overall intermediate input purchases of U.S. firms; they found that this share had already increased from 5.3% in 1972 to 11.6% in 1990. Campa and Goldberg (1997) found similar evidence for Canada and U.K., but surprisingly not for Japan, where the reliance on foreign inputs appeared to have declined between 1974 and 1993. Hummels, Ishii and Yi (2001) instead constructed a measure of vertical specialization capturing the value of imported intermediate inputs (goods and services) embodied in a country's exported goods and found that it already accounted for up to 30% of world exports in 1995, having grown by as much as 40% since 1970.

The work of Johnson and Noguera (2012*a*, 2012*b*) constitutes the state of the art in the use of Input-Output tables to quantify the importance of global production sharing and its evolution in recent years. The main innovation of their methodology is in the attempt to compute a *global* Input-Output table from which one can back out the value-added and intermediate input

⁶A more extensive list can be found at: <http://www.chipworks.com/en/technical-competitive-analysis/resources/blog/the-new-ipad-a-closer-look-inside/>.

⁷The task is complicated by the fact that data on trade flows of goods is collected on a gross output basis, without regard to the particular sources of the value added embodied in these goods.

contents of gross international trade flows. In particular, their VAX ratio (the value-added to gross-value ratio of exports) is an appealing inverse measure of the importance of vertical specialization in the world production: the lower is this measure, the larger is the value of imported inputs embodied in exports.⁸ As is clear from Figure 1.1, their VAX ratio has declined rather significantly since 1970 with about two-thirds of the decline occurring after 1990. Johnson and Noguera (2012*b*) show that this decline is explained solely by increased offshoring within manufacturing. Furthermore, they also find that global production sharing has grown disproportionately in emerging economies and also appears to increase following the signing of regional trade agreements.




Figure 1.1: Ratio of Value Added to Gross Exports (VAX), 1970-2009

Two limitations of the fragmentation measures discussed so far are that they rely on fairly aggregated Input-Output data and that they impose strong proportionality assumptions to back out the intermediate input component of trade. A different approach to measuring the degree to which production processes are fragmented across countries was first suggested by Yeats

⁸In a very recent paper, Koopman, Wang and Wei (2014) devise a methodology that nicely nests Johnson and Noguera's (2012) VAX measure with the vertical specialization measures developed by Hummels, Ishii and Yi (2001).

(2001), and consists in computing the share of trade flows accounted for by SITC Rev.2 industry categories that can be safely assumed to contain *only* intermediate inputs (as reflected by the use of the word “Parts of” at the beginning of the category description). It turns out that all these industries are in the “Machinery and Transport Equipment” industrial group (or SITC 7). Yeats (2001) found that intermediate input categories accounted for about 30% of OECD merchandise exports of machinery and transport equipment in 1995, and that this share had steadily increased from its 26.1% value in 1978. A limitation of Yeats’ measure is that, by focusing on industries composed *exclusively* of inputs, it naturally understates the importance of input trade. This might explain why when updating this methodology to present times, one finds little evidence of a further increase in this share.⁹

An alternative to categorizing trade flows as *either* final goods or intermediate inputs is to attempt to calculate a more continuous measure of the “upstreamness” of the goods being traded. This is the approach in Antràs, Chor, Fally and Hillberry (2012), who use Input-Output data to construct a weighted index of the average position in the value chain at which an industry’s output is used (i.e., as final consumption, as direct input to other industries, as direct input to industries serving as direct inputs to other industries, and so on), with the weights being given by the ratio of the use of that industry’s output in that position relative to the total output of that industry. Intuitively, the higher is this measure, the more removed from final good use (and thus the more upstream) is that industry’s output. The Data Appendix contains a lengthier discussion of the construction of this index.¹⁰ Antràs, Chor, Fally and Hillberry (2012) use the measure to characterize the average upstreamness of exports of different countries in 2002, but it can also be employed to illustrate how the upstreamness of world exports has evolved in recent years. As shown in Figure 1.2, world exports became significantly

⁹Other authors attempting to compute the share of intermediate inputs in world trade using alternative methodologies have also found little evidence of a trend in the series (see, for instance, Chen, Kondratowicz and Yi, 2005, or Miroudot, Lanz and Ragoussis, 2009). I have obtained similar results when computing the relative growth of overall trade and input trade using the classification of goods developed by Wright (2014). As argued by Johnson and Noguera (2012*b*), even when taking this finding at face value, it is not necessarily inconsistent with the observed rise in indices of vertical specialization, which better capture the use of imported inputs in producing goods *that are exported*.

¹⁰This upstreamness index was independently developed by Antràs and Chor (2013) and Fally (2012), and its properties were further studied in Antràs, Chor, Fally and Hillberry (2012).

more upstream in recent years, particularly in the period 2002-08. The patterns are in line with those illustrated in Figure 1.1, and also suggest an increase predominance of input trade in world trade. Although a significant share of the observed increase in upstreamness is related to an increase the relative weight of petroleum-related industries, even when netting those out, one observes a significant upward trend in upstreamness (see Figure 1.2). Interestingly, both Figures 1.1 and 1.2 identify a disproportionate fall in global production sharing relative to the overall fall in world trade during the early years of the recent ‘great recession’.




Figure 1.2: Average Upstreamness of World Exports, 1998-2009

Old and New Theories

The noticeable expansion in input trade has also captured the attention of international trade theorists eager to bridge the apparent gap between the new characteristics of international trade in the data and the standard representation of these trade flows in terms of final goods in traditional and new trade theory.

One branch of this new literature has focused on incorporating the notion of fragmentation in otherwise neoclassical models with homogeneous goods,

perfectly competitive markets and frictionless contracting. Key contributions include Feenstra and Hanson (1996a), Jones (2000), Deardorff (2001), and Grossman and Rossi-Hansberg (2008). The main idea in these contributions is that the production process (as represented by an abstract mapping between factors of production and final output) can be decomposed into smaller parts or stages that are themselves (partly) tradable. Different authors assign different labels to these parts: some refer to them as intermediate inputs, others call them vertical production stages, while others view them as tasks. Regardless of the interpretation of the process under study, a common lesson from this body of work is that the possibility of fragmentation generates nontrivial effects on productivity, and that these endogenous changes in productivity in turn deliver novel predictions for the effects of reductions in trade costs on patterns of specialization and factor prices. Antràs and Rossi-Hansberg (2009) elaborate on this broad interpretation of this branch of the literature and also offer more details on the specific results of each of these contributions.¹¹

As insightful as this body of work has proven to be, it seems clear that modeling global production sharing as simply an increase in the tradability of homogeneous inputs across countries misses important characteristics of intermediate input trade. Prominent among these features is the fact that parts and components are frequently customized to the needs of their intended buyers (remember our example above with the iPad 3). In other words, the disintegration of the production process is more suitably associated with the growth of trade in *differentiated* (rather than homogeneous) intermediate inputs.¹²

¹¹Another common feature of the theoretical frameworks developed in these papers is that the number of primitive factors of production is assumed to be small, and normally equal to two. Another branch of the literature has developed perfectly-competitive, frictionless models in which offshoring results from the assignment of a population of a large number of heterogeneous agents into international hierarchical teams (see Kremer and Maskin, 2006, or Antràs, Garicano and Rossi-Hansberg, 2006).

¹²Admittedly, there does not exist much evidence to substantiate this claim. Antràs and Staiger (2012) offer a back-of-the envelope quantification applying the methodology suggested by Schott (2004) to identify international trade in intermediate goods and using the “liberal” classification of Rauch (1999) to distinguish between differentiated and homogeneous goods. They find that the share of differentiated inputs in world trade more than doubled between 1962 and 2000, increasing from 10.56% to 24.85% of world trade. Behar and Freund (2011) show that during the late 1990s and 2000s, intermediate inputs traded within the EU became more sophisticated and involved more relationship-specific

Another important characteristic of global production networks is that they necessarily entail intensive *contracting* between parties located in different countries and thus subject to distinct legal systems. In a world with perfect (or complete) contracting across borders, this of course would be of little relevance. Unfortunately, this is not the world we (or at least, I) live in. Real-world commercial contracts are incomplete in the sense that they cannot possibly specify a course of action for *any* contingency that could arise during the course of a business relationship. Of course, the same can be said about domestic commercial transactions, but the cross-border exchange of goods cannot generally be governed by the same contractual safeguards that typically accompany similar exchanges occurring within borders.

Given the subject of this book, it is worth pausing to describe in more detail some of the factors that make international contract enforcement particularly problematic.

Contracts in International Trade

A first natural difficulty in contractual disputes involving international transactions is determining which country's laws are applicable to the contract being signed. In principle, the parties can include a *choice-of-law* clause specifying that any dispute arising under the contract is to be determined in accordance with the law of a particular jurisdiction, regardless of where that dispute is litigated. Nevertheless, many international contracts do not include that clause and, in any case, it is up to the court of law adjudicating a dispute to decide whether it will uphold the expressed desire of the parties. If the court is not familiar with the law specified in the contract, as may often occur in international transactions, the court might decide to rule on the basis of its own law, or they may inadvertently apply the desired foreign law incorrectly.

A second difficulty relates to the fact that even when local courts are competent (in a legal sense), judges may be reluctant to rule with regard to a contract dispute involving residents of foreign countries, especially if such a ruling would entail an unfavorable outcome for local residents. The evidence on the home bias of local courts is mixed, but even those authors advocating that a formal analysis of case law does not support the hypothesis of biases against foreigners readily admit a widespread belief of the existence of such

investments (in the sense of Nunn, 2007).

xenophobic biases (see Clermont and Eisenberg, 2007).

A third complication with international contracts relates to the enforcement of remedies stipulated in the court's verdict. For instance, the court might rule in favor of a local importer that was unsatisfied with the quality of certain components obtained from an exporter, and the verdict might require the exporter to compensate the importer for any amount already paid for the components, as well as for any court or even attorney fees incurred. An issue arises, however, if the exporter does not have any assets (say bank accounts or fixed assets) in the importer's country. In such case, it is not clear that the exporter will feel compelled to accept the verdict and pay the importer.

In recent years, there have been several coordinated attempts to reduce the contractual uncertainties and ambiguities associated with international transactions. A particularly noteworthy example is the United Nations Convention on Contracts for the International Sale of Goods (or CISG), or Vienna Convention, which attempts to provide a set of uniform rules to govern contracts for the international sale of goods. The idea is that even when an international contract does not include a choice-of-law clause, parties whose places of business are in different signing countries can rely on the Convention to protect their interests in courts. As ambitious as the CISG initiative is, it has arguably fallen short of its objectives. For instance, several countries or regions (most notably, Brazil, Hong Kong, India, South Africa, Taiwan, and the United Kingdom) have yet to sign the agreement. Furthermore, a few of the signing countries have expressed reservations and choose not to apply certain parts of the agreement. Finally, it is not uncommon for private parties to explicitly opt out of the application of the Convention, as allowed by its Article 6. The reluctance to unreservedly embrace the Convention has been associated with the somewhat vague language of the text, which might foster the natural inclination of judges to interpret the Convention through the lens of the laws of their own State.¹³

Another attempt to ameliorate the perceived contractibility of international transactions consists in resorting to international arbitration. More specifically, an international trade contract can include a (so-called) *forum-of-*

¹³The Institute of International Commercial Law at Pace Law School maintains a website (<http://www.cisg.law.pace.edu/>) with comprehensive information on the CISG, including a database of thousands of legal cases in which the Convention was invoked. The details of these cases offer a vivid account of the nature of contractual disagreements in international trade.

law clause establishing that a particular arbitrator, such as the International Chamber of Commerce (ICC) in Paris, will resolve any contractual dispute that may arise between the parties. International arbitration is appealing because it avoids the aforementioned uncertainties associated with litigation in national courts. It is also relatively quick and parties benefit from the fact that arbitrators tend to have more commercial expertise than a typical judge. Furthermore, arbitration rulings are confidential and are generally perceived to be more enforceable than those of national courts because they are protected by the Convention on the Recognition and Enforcement of Foreign Arbitral Awards, also known as New York Convention. Despite its attractive features, international arbitration is rarely used in practice because its cost is too high for most firms to bear.¹⁴

One might argue that even when explicit contracts are incomplete and perceived to be unenforceable, parties in international transactions can still resort to implicit contracting to sustain ‘cooperation’. Nevertheless, it is particularly hard to render international commercial relationships self-enforcing. On the one hand, international parties are less likely to meet face-to-face and to transact on a repeated basis than domestic parties, in part due to distance and trade costs, but also due to shocks (such as exchange rate movements) that can quickly turn efficient relationships into inefficient ones. On the other hand, the possibility of collective or community enforcement is hampered again by distance but also by the fact that parties might have different cultural and societal values. In sum, and in the words of Rodrik (2000), “ultimately, [international] contracts are often neither explicit nor implicit; they simply remain incomplete.”

Although contractual risks are also of relevance for the exchange of final goods (see Chapter 3), the detrimental effects of imperfect international contract enforcement are likely to be particularly acute for transactions involving intermediate inputs. This is so for at least two reasons. First, input transac-

¹⁴It may be instructive to illustrate this claim with some figures. Using the arbitration cost calculator available from the ICC website, the estimated cost of arbitration (involving a single arbitrator) would be \$5,401 for a \$10,000 dispute (or a 54% cost-to-dispute-amount ratio), \$15,425 for a \$100,000 dispute, \$61,094 for a \$1 million dispute, and \$170,799 for a \$10 million dispute (or a mere 1.7% cost-to-dispute-amount ratio). It is thus little surprise that there were only 796 ICC arbitrations requests in 2011 and that the amount in dispute was under one million U.S. dollars in only 22.7% of these cases (see <http://www.iccwbo.org/products-and-services/arbitration-and-adr/arbitration/cost-and-payment/cost-calculator/>).

tions are often associated with relatively long time lags between the time an order is placed (and the contract is signed) and the time the goods or services are delivered (and the contract is executed). Second, parts and components often entail significant relationship-specific investments and other sources of lock-in on the part of both buyers and suppliers, which make contractual breaches particularly costly. As argued above, suppliers often customize their output to the needs of particular buyers and would find it hard to sell those goods to alternative buyers, should the intended buyer decide not to abide by the terms of the contract. Similarly, buyers often undertake significant investments whose return can be severely diminished by incompatibilities, production line delays or quality debasements associated with suppliers not going through with their contractual obligations.¹⁵

Firm Responses to Contractual Insecurity

When designing their global sourcing strategies, firms face two key decisions. The first one concerns the *location* of the different stages in the value chain and involves deciding in which country or region will firms conduct R&D and product development, where parts and components should be produced, what is the best place to assemble the finished good, and so on. The second key decision relates to the extent of *control* that firms exert over these different production stages. For instance, firms may decide to keep these production stages within firm boundaries, thus engaging in foreign direct investment (FDI) when the integrated entity is in a foreign country. Other firms may be less inclined to keep a tight control over certain stages and thus choose to contract with suppliers or assemblers at arm's-length.

Neoclassical models of fragmentation focus exclusively on the first of these decisions and emphasize that fragmentation will emerge as part of a competitive equilibrium whenever firms find it cost-minimizing to break up production process across countries. The source of the cost-advantage associated with fragmentation varies by model, sometimes it stems from differences in relative factor endowments across countries (which, for instance, naturally

¹⁵A third more specific reason for which input trade might be perceived to be less contractually secure relates to the fact that Article 3 of CISG explicitly excludes from the applicability of the Convention situations in which “the party who orders the goods undertakes to supply a substantial part of the materials necessary for such manufacture or production,” thus making the Convention less relevant for sustaining cooperation in global production sharing networks.

confer comparative advantage in labor-intensive stages to relatively labor-abundant countries), while other times they are motivated by technological differences across countries.

Neoclassical models are silent on the issue of control. This is not because these models assume perfect competition, constant returns to scale, or homogeneous goods. Instead, the key assumption that renders those models (and just about *any* model in the field of International Trade) vacuous when tackling the notion of control is the assumption of perfect or complete contracting. Indeed, if firms could foresee all possible future contingencies, and if they could costlessly write contracts that specify in an enforceable manner the course of action to be taken in all of these possible contingencies, then firms would no longer need to worry about “controlling” the workers, the internal divisions or the supplying firms with whom they interact in production. The complete contract would in fact confer *full* control to the firm regardless of the ownership structure that governs the transactions between all these producers. In other words, and as Coase (1937) anticipated more than seventy-five years ago, firm boundaries are indeterminate in a world of complete contracts.¹⁶

In the real world, however, contracts are very much incomplete and especially so in international transactions, where as argued above, the enforceability of contracts is particularly questionable. In response to this perceived contractual insecurity, firms spend a substantial amount of time and resources figuring out the best possible way to *organize* production in the global economy. In some cases, foreseeing that producers located in a particular country might not feel compelled to follow through with their contractual obligations, firms contemplating doing business in that country might decide to do so within their firm boundaries, either by setting up a new, wholly- or partially-owned affiliate or by acquiring a controlling stake in an existing firm in that country. In some circumstances, however, the lack of contract enforceability might precisely turn firms to independent suppliers for the procurement of parts because such an arrangement might elicit the best performance from foreign producers. In other words, it is important to keep in mind that internalization is a double-edged sword: it may partly protect the integrating

¹⁶It is worth stressing that even in the presence of product differentiation and market power, firm boundaries remain indeterminate when contracts are complete. For example, the often-cited double-marginalization rationale for vertical integration rests on the assumption that firms and suppliers cannot sign simple two-part tariff contracts, and as such, it also constitutes an incomplete-contracting theory of firm boundaries.

party from the vagaries of international contracting, but it might dilute the incentives to produce efficiently of the integrated party, which is now more tightly controlled and has less power in the relationship (cf., Grossman and Hart, 1986).

The boundaries of firms in the world economy are thus the result of the (constrained) optimal decisions of firms attempting to organize production in the most profitable way possible. A recurring theme of this book is that much can be learned from a theoretical and empirical study of the *fundamental* forces that appear to shape whether international transactions are internalized or not, independently of the firm or sector one is studying.

Practitioners (and perhaps some academics too) might react skeptically to the idea that low-dimensional models can possibly capture the reasoning behind the complex and idiosyncratic decisions of firms in the world economy. Business school cases often highlight the peculiarities of particular organizational decisions, making it hard to envision that much can be gained from extrapolating from those particular cases. The fact that comprehensive datasets on the integration decisions of firms are not readily available might have only compounded this belief, as most empirical studies of integration decisions rely on data from specific industries or firms.¹⁷

A Comparative Advantage of Trade Statistics

An advantage of studying the global integration decisions of firms is that data on international transactions are particularly accessible due to the widespread existence of official records of goods and services crossing borders. For instance, it is well-known that researchers can easily access data on U.S. imports from any country of the world at the remarkably detailed ten-digit Harmonized Tariff Schedule classification system, which consists of nearly 17,000 categories.¹⁸ A less well-known fact is that, in some countries, these same detailed country- and product-level data contain information on the extent to which trade flows involve related parties or non-related parties. Most notably, the “U.S. Related Party Trade” data collected by the U.S. Bureau

¹⁷See Baker and Hubbard (2003) for a particularly careful study using data from the trucking industry and Lafontaine and Slade (2007) for a broad survey of the empirical literature on vertical integration.

¹⁸Downloading this data from the NBER website, one can readily verify that in 2001 France exported \$15,747 worth of frozen potatoes to the United States (HTS code 2004.10), yet none of those were French fries (HTS code 2004.10.8020).

of Customs and Border Protection and managed by the U.S. Census Bureau provides data on related and non-related party U.S. imports and exports at the six-digit Harmonized System (HS) classification (which consists of over 5,000 categories) and at the origin/destination country level. This amounts to hundreds of thousands of observations *per year* on the relative prevalence of integration versus nonintegration across products and countries.¹⁹

What do these data tell us about the global sourcing strategies of firms? The first thing that one notices when using U.S. related party trade data is how predominant intrafirm transactions are in U.S. trade. In 2011, intrafirm imports of goods totaled \$1056.2 billion and constituted a remarkable 48.3 percent of total U.S. imports of goods (\$2,186.9 billion). In fact, the share of intrafirm trade has been higher than 46.5 percent in every year since 2000. On the export side, related-party exports are also pervasive, with their share in total U.S. exports ranging from 28 to 31 percent in recent years. These figures illustrate the importance of multinational firms for U.S. trade.²⁰

A second evident feature of the data is that the share of U.S. intrafirm imports varies widely across countries. On the one hand, in 2011 intrafirm imports equalled 0 for 10 countries and territories (including Cuba), all exporting very low volumes to the U.S. On the other hand, in that same year the share of intrafirm trade reached a record 89.6 percent for U.S. imports from Western Sahara. Leaving aside communist dictatorships and disputed territories, and focusing on the 50 largest exporters to the U.S., Figure 1.3 illustrates that the share of intrafirm trade still varies significantly across countries, ranging from a mere 2.4 percent for Bangladesh to an astonishing 88.5 percent for Ireland.

Similarly, the share of intrafirm trade varies widely depending on the type of product being imported. Again, the raw data contain infrequently

¹⁹The U.S. Related Party Trade data are publicly available at: <http://sasweb.ssd.census.gov/relatedparty/>. This website permits downloading the data at the six-digit NAICS level. The finer six-digit Harmonized System (HS) data are available from the U.S. Census for a fee.

²⁰In contrast, Atalay, Hortacsu and Syverson (2013) study intrafirm shipments across U.S. multiplant firms and find that these constitute a very small share of total shipments, a finding that they interpret as indicating that firm boundaries are shaped by issues related to the the transfer of intangible inputs, rather than of physical goods. However, as argued above, contractual insecurity in the exchange of physical inputs is much more significant in international transactions than in domestic ones (especially, in the U.S. case), and thus firm boundaries might well be shaped by different factors in cross-border relationships than in domestic ones.




Figure 1.3: Share of U.S. Intrafirm Imports for Largest 50 U.S. Exporters in 2011

traded goods with shares equal to 0 and 100, but even when focusing on the top 20 six-digit HS manufacturing industries by importing volume, in Figure 1.4 one observes significant variation in the share of intrafirm trade, which ranges from a share of 11.4 percent for U.S. imports of sweaters, pullovers and sweatshirts made of cotton (HS 611020) to 98.8 percent for imports of automobiles with engines of more than 3000 cc (HS 870324). This variation persists even when focusing on much narrower sectors. As shown in Figure 1.5, when analyzing imports across subcategories of the four-digit Harmonized System sector 8708 (‘Parts and accessories of motor vehicles’), the share of intrafirm trade still ranges from 19.8 percent for drive axles (HS 870850) to 71.2 percent for steering wheels (HS 870894). It is thus clear that U.S.-based producers appear to source different auto parts under quite different ownership structures.

As a final illustration of the richness and variation in the data, consider the six-digit HS industry with the largest share of intrafirm imports in Figure 1.5, namely HS 870894 (Steering Wheels, columns and boxes for motor vehicles). Figure 1.6 reports the share of intrafirm trade for all 56 countries with positive exports to the U.S. in that sector. As is clear from the graph, even




Figure 1.4: Share of U.S. Intrafirm Imports in Top 20 Industries by U.S. Import Volume in 2011




Figure 1.5: Variation in the Share of Intrafirm Trade within HS Sector 8708 (Auto Parts) in 2011




Figure 1.6: Variation in the Share of U.S. Intrafirm Imports within HS 870884 (Steering Wheels) in 2011

when focusing on a narrowly-defined component, a similar pattern to that in Figure 1.3 emerges, with U.S.-based producers appearing to source particular inputs quite differently depending on the location from which these products are bought. Imports from 17 of the 56 countries are exclusively transacted at arm's-length, while one country (Liechtenstein) sells steering wheels to the U.S. almost exclusively within multinational firm boundaries. The remaining 38 countries feature shares of intrafirm trade fairly uniformly distributed between 0 and 100 percent.

The large variation in the relative importance of intrafirm transactions across types of goods and countries might seem to validate the skeptics' view that the decision to integrate or outsource foreign production processes is largely driven by idiosyncratic factors that cannot possibly be captured by parsimonious models of the organization decisions of firms. If that were the case, however, not only would we observe large variation in the share of intrafirm trade, but we would also expect this variation to be uncorrelated with simple industry or country-level variables. As first demonstrated by Antràs (2003), the evidence suggests otherwise. Chapter 7 will describe in detail several stylized facts regarding the intrafirm component of trade. As a sneak

preview, Figures 1.7 and 1.8 illustrate that the share of intrafirm imports in total U.S. imports is significantly higher, the higher the U.S. capital intensity in production of the good being imported and is also significantly higher, the higher the capital-labor ratio of the exporting country. These scatter plots suggest that, as argued above, there may indeed be some common fundamental factors that shape the integration decisions of firms across sectors and countries. The theories of internalization exposted in Chapters ?? and 6 will attempt to shed some light on these factors and will provide a valuable lens through which to study the intrafirm trade data in a more formal and structured manner.




Figure 1.7: The Share of Intrafirm U. S. Imports and Capital Intensity

While several features of the U.S. related-party trade database make it particularly attractive to empirical researchers, it has some important limitations. Some of the shortcomings of the data relate to the extent to which the characteristics of the data permit a formal test of the theories of internalization developed later in the book, so it is convenient to postpone that discussion until after we have covered those theories in Chapters ?? and 6. Other potential limitations are more fundamental, so it is important to tackle them upfront.

The U.S. database defines ‘related-party imports’ as import transactions

that the data are not subject to sampling error, since an indicator of whether the transaction involves or not related parties is required for *all* import or export transactions recorded by the U.S. Bureau of Customs and Border Protection. Despite this requirement, importers and exporters do not always report that information in their shipment documents. Luckily, these transactions are categorized on the data tables as “nonreported,” so it is easily verified that these account for a very low share of trade volumes (for instance, just 1.4 percent of total imports in 2011). One might also worry about non-sampling errors related to the imputation of trade values for undocumented shipments and for low-valued transactions (which are sometimes estimated). Nevertheless, quality assurance procedures are performed at every stage of collection, processing, and tabulation, thus there is no reason to believe that these data are any less reliable than U.S. customs data on trade flows.²²

One way to gain reassurance regarding the usefulness of the data is to see whether it delivers patterns that are consistent with what one would expect based on independent and reliable sources of data. For example, from a quick search of press releases from recent years, one learns that in 2005, Boston-based Gillette Company completed the construction of a 120 million-euro plant in Łódź (Poland), which manufactures disposable razors and other shaving products.²³ Although production was mostly directed to the European market, it seems reasonable to assume that some of the products produced in the plant were shipped back to the U.S., a transaction that would naturally occur within firm boundaries. As shown in Figure 1.9, it is reassuring to observe that the share of intrafirm imports in total U.S. imports from Poland of NAICS code 332211, which is dominated by non-electric razors and razor blades, went up dramatically around the time of the plant opening, jumping from essentially 0 percent in 2004 to close to 100 percent from 2005 onwards.

Back to the Location Decision

We have emphasized above that the internalization decisions of firms in the global economy cannot be understood without appealing to contractual frictions and we have also illustrated the importance of these frictions in the real world. It seems natural, however, to posit that imperfect contracting not only shapes the ownership structure decisions of firms but might also

²²Ruhl (2013) provides a useful overview of alternative U.S. intrafirm trade data sources.

²³See http://www.paiz.gov.pl/nawosci/?id_news=502.




Figure 1.9: Share of Intrafirm Imports of NAICS 332211 (Razors) from Poland

impact their geographical location decisions. As emphasized by neoclassical models of offshoring, profit-maximizing firms will organize production in a cost-minimizing manner, but the effective costs of doing international business are not solely explained by the factors highlighted by neoclassical theory. Certainly, wages will, other things equal, tend to be relatively lower in relatively labor-abundant countries. And, other things equal, costs of production will also tend to be relatively low in countries or regions where the technologies used in production are particularly advanced. Yet, firms might be reluctant to offshore production lines to low-wage countries where suppliers are unreliable and tend not to honor their contracts, and where local courts are unlikely to effectively enforce contracts. Similarly, firms might be unwilling to operate in countries in which their advanced technologies could be effectively deployed (given the existence of local complementary factors), but in which the contractual environment might not provide enough security to firms, both in terms of quality contracting but also in terms of the risk of intellectual property rights expropriation.

A key factor that makes contractual aspects important for sourcing decisions is the existence of huge variation among countries in judicial quality

and contract enforcement. Empirical researchers often resort to easily accessible measures of the quality of the rule of law which are themselves based on weighted averages of various indices of the perceived effectiveness and predictability of courts in different countries. An advantage to these widely used measures, such as the ‘Rule of Law’ variable produced by the Worldwide Governance Indicators, is that they capture broad features of the contracting environment. A disadvantage is that they are partly based on subjective assessments rather than objective measures of institutional quality. Furthermore, they may provide a useful ordinal measure of legal quality but they are less well equipped to help quantify the existence of cross-country heterogeneity in judicial quality and contract enforcement.

Djankov, La Porta, Lopez-De-Silanes and Shleifer (2003) have proposed an ingenious alternative measure of judicial quality which is narrower in nature but more powerful in illustrating the relevance of differences in the legal system across countries. In particular, Djankov al. (2003) estimate for 109 countries the time it takes a plaintiff using an official court to evict a nonpaying tenant and to collect a bounced check. Figure 1.10 depicts the second of these two variables, which is more likely to be of relevance for firms considering doing business in a particular country. Their estimated total duration of a legal procedure aimed at collecting a bounced check ranges from 7 days in Tunisia to 1003 in Slovenia. Even when focusing on the 43 of the top 50 largest exporters to the U.S. for which they provide data, the estimated duration ranges from 39 days for the Netherlands to 645 days for Italy.

The extent of contractual insecurity not only varies across countries (or jurisdictions) but it naturally also varies depending on the characteristics of the goods being transacted. For instance, basic goods with low levels of differentiation and which are traded in relatively thick markets can be relatively safely procured even from countries with weak contracting institutions. Conversely, transactions involving highly complex or differentiated goods will tend to be much more ‘contract dependent’ and one would expect firms to be significantly more sensitive to the institutional environment when choosing the country from which to procure those goods.

A Brief Road Map

This book will study the various ways in which the contracting environment shapes the location and internalization decisions of firms in the global econ-




Figure 1.10: Duration of a Legal Procedure Aimed at Collecting a Bounced Check

omy. I will focus first on an analysis of the location decision and how it is affected by contracting factors, and only in Part III of the book will I allow firms to optimally decide the extent of control they want to exert over production processes. This does not follow the chronological order in which these topics were developed in the literature, but I will adopt this sequencing for pedagogical reasons.

Before diving into the world of incomplete contracts, it is necessary however to provide an overview of the ‘complete-contracting’ frameworks that will serve as the basis or skeleton for the models to be developed in future chapters. A succinct overview of these models is offered in Chapter 2, to which I turn next. Readers familiar with Melitz’ (2003) classic paper and its various extensions might want to jump straight to Part II of the book, starting on Chapter 3.

Chapter 2

Workhorse Models

The field of International Trade has witnessed a true revolution in recent years. Firms rather than countries or industries are now the central unit of analysis. The workhorse trade models used by most researchers both in theoretical work as well as in guiding empirical studies were published in the 2000s. The purpose of this Chapter is to provide a succinct account of the rich intellectual history of the field and to offer an overview of these modern workhorse models. While these benchmark frameworks ignore contractual aspects, they constitute the backbone of the models developed later in the book, so it is important to gain an understanding of their key features.

Two Centuries of Trade Theory

The recent revolution in the International Trade field would perhaps not be apparent when browsing the leading undergraduate textbooks covering the basics of international trade and investment. Neoclassical trade theory still constitutes the core of what we teach to college students. This should not be surprising: the concept of comparative advantage is as relevant today as it was almost two hundred years ago when David Ricardo initiated the formal modeling of foreign trade in his *Principles of Political Economy and Taxation* (1817). The first hundred and seventy years of the International Trade field were largely devoted to refining Ricardo's rudimentary description of the gains from specialization. The benchmark two-good, two-country Ricardian model found in most introductory textbooks is the culmination of an intellectual endeavor to which John Stuart Mill, Frank Graham, and Lionel McKenzie contributed key advances.

Starting with the work of Eli Heckscher and his disciple Bertil Ohlin, another branch of the neoclassical theory studied models in which comparative advantage is endogenously shaped by the interaction of differences in relative factor abundance across countries and differences in relative factor intensities across sectors. The formalization of the so-called Heckscher-Ohlin model is often associated with the great Paul Samuelson, but Abba Lerner, Ronald Jones and Alan Deardorff should also be credited for particularly significant contributions.¹

The core theorems of neoclassical trade theory – the Heckscher-Ohlin, the Stolper-Samuelson and the Rybczynski theorems – are the product of these intellectual efforts. These beautiful and incredibly sharp results still shape to date the way that most economists think about the determinants and consequences of international trade flows. Why is China the single largest exporter to the U.S.? How does trade with China affect the relative pay of skilled and unskilled workers in the U.S.? How does immigration affect sectoral employment in the U.S.? You would be hard-pressed to answer these questions without appealing to the insights of neoclassical theory.

Neoclassical trade models deliver sharp results but also make strong assumptions. The benchmark models assume a very low number of goods and factors, often only two of each. In higher-dimensional environments, the classical theorems become much less beautiful and much less sharp.² More importantly, in neoclassical models, technology is typically assumed to feature constant returns to scale and market structure is characterized by perfect competition, thus making these frameworks of limited use for firm-level studies of international trade. Indeed, in neoclassical trade theory it is not firms but rather countries that trade with each other.

Trade theory witnessed a first revolution in the late 1970s and early 1980s when a group of young trade economists, led by Paul Krugman and Elhanan Helpman, developed new models attempting to account for some empirical patterns that were hard to reconcile with neoclassical theory. Most notably,

¹A lucid exposition of neoclassical trade theory with extensive references can be found in Jones and Neary (1984).

²It is important to emphasize, however, that the implications of the theory for the net factor content of trade – the so-called Vanek (1968) equations – have been shown to be robust to variation in the number of goods and factors. It is no surprise then that beginning with the seminal work of Leamer (1984), empirical testing of the Heckscher-Ohlin model has largely focused on these factor content predictions (see Treffer, 1993*a*, 1995, Davis and Weinstein, 2001 and Treffer and Zhu, 2010).

traditional theory rationalized the existence of mutually beneficial intersectoral trade flows stemming from cross-country differences in technology or endowments. In the real world, however, the bulk of trade flows occurs between countries with similar levels of technological development and similar relative factor endowments, and a significant share of world trade is accounted for by two-way flows within fairly narrowly-defined sectors (i.e., ‘intraindustry’ instead of intersectoral trade).

This new wave of research, which was dubbed ‘new trade theory’, emphasized the importance of increasing returns to scale, imperfect competition, and product differentiation in accounting for these salient features of the data. Intuitively, even two completely identical countries will find it mutually beneficial to trade with each other as long as specializing in particular *differentiated* varieties of a sector’s goods allows producers to expand their sales and operate at lower average costs, as would naturally be the case whenever technology features *economies of scale*. The relevance of *imperfect competition* for these theories stems from the simple fact that (internal) economies of scale are inconsistent with perfect competition.

A key hurdle facing the pioneers of new trade theory was the absence of a generally accepted modeling of product differentiation and imperfect competition. While there is only one way in which goods can be perfectly homogeneous, there are many ways in which products can be differentiated. Differentiation can arise because individual consumers enjoy spreading their income across different varieties of particular goods (as in the case of cultural goods), or because different consumers prefer to consume different varieties or qualities of the same good (as with tablets or cars). Even when focusing on one of these modeling approaches, there remains the issue of how to mathematically characterize product differentiation in preferences. Similarly, there is only one way in which markets can be perfectly competitive, while there are various possible approaches to modeling imperfect competition.

There are two main reasons why ‘new trade theory’ was able to overcome these difficulties and become mainstream in a relatively short period of time. First, researchers quickly converged in the use of *a particular* modeling of product differentiation and market structure associated with Krugman (1979, 1980), who in turn borrowed from Dixit and Stiglitz (1977). This served the important role of providing a common language for researchers in the field to communicate among themselves. Still, the heavy use of specific functional forms in representing preferences and technology was viewed with some reser-

vations by the old guard in the field.³

The second key factor in the success of new trade theory was the publication of a landmark treatise by Helpman and Krugman (1985). This concise book established the generality of most of the insights from Krugman's work and also illustrated how the new features of new trade theory could be embedded into neoclassical trade theory. As a result, these new hybrid models could explain the features of the data that motivated the new models, while at the same time preserving the validity of some of the salient results from neoclassical theory, such as the Vanek (1968) equations characterizing the factor content of trade. With the publication of this manuscript, the walls of resistance came tumbling down, new trade theory became the new paradigm, and Krugman's modeling choices gained a prominent spot in the toolbox of trade theorists (and of applied theorists in other fields).

In recent years, international trade theory has witnessed a second revolution which in many respects parallels the one witnessed thirty years ago. As in the case of new trade theory, and consistently with Kuhn's (1996) description of the structure of scientific revolutions, the need for a new paradigm was fueled by the discovery of a series of new empirical facts that were inconsistent with new trade theory models. To understand these inconsistencies, it is important to note that in Krugman-style models, all firms within a sector are treated symmetrically. Although firms produce differentiated products, they do so under a common cost function, and all varieties enter symmetrically into demand with an elasticity of substitution between any pair of varieties that is constant and common for any pair. As a result, firm behavior within an industry is 'homogeneous'. Furthermore, under the common assumption of iceberg (or ad valorem) trade costs, new trade theory models deliver the stark implication that all firms within a differentiated-good sector will export their output to every single country in the world.

In the 1990s, a wave of empirical papers using newly-available longitudinal plant and firm-level data from various countries demonstrated the existence of significant levels of heterogeneity in revenue, productivity, factor inputs, and trade behavior across firms within sectors. In fact, in some cases, heterogeneity in performance was shown to be almost as large within sectors

³As an illustration of this resistance, Krugman's 1979 seminal article was rejected by the *Quarterly Journal of Economics* in 1978 and was subsequently salvaged by Jagdish Bhagwati at the *Journal of International Economics* despite two negative referee reports (see Gans and Shepherd, 1994; note however that Ethier, 2001, offers a slightly less glorifying account of Bhagwati's role in rescuing the paper at the JIE).

than across sectors (see, for instance, Bernard et al., 2003). With regards to export behavior, studies found that only a small fraction of firms engage in exporting, and that most exporting firms sell only to a few markets. This so-called extensive margin of trade has been shown to be important in order to understand variation in aggregate exports across destination markets. Several studies have also documented that exporters appear to be systematically different from non-exporters: they are larger, more productive, and operate at higher capital and skill intensities. In addition, firm heterogeneity has been shown to be of relevance for assessing the effects of trade liberalization, as those episodes appear to lead to market share reallocations towards more productive firms, thereby fostering aggregate productivity via new channels.

Access to micro-level data has also served to confirm the importance of multinational firms in world trade. For instance, according to 2009 data from the Bureau of Economic Analysis, 75 percent of the sales by U.S. firms in foreign markets is carried out by foreign affiliates of U.S. multinational enterprises (MNEs), and only 25 percent by exports from the U.S. (Antràs and Yeaple, 2013). Furthermore, not only do intrafirm trade flows constitute a very significant share of world trade flows (as mentioned in Chapter 1), but an important share of the volume of arm's-length international trade is accounted for by transactions involving multinational firms as buyers or sellers. For instance, data from the U.S. Census Bureau indicates that roughly 90 percent of U.S. exports and imports flow through multinational firms (Bernard et al., 2009). New trade theory did not ignore the importance of multinational firms or intrafirm trade in the world economy (see Helpman, 1984, or Helpman and Krugman, 1985, Chapter 12 and 13), but by focusing on complete-contracting, homogenous-firm models, it was unable to account for central aspects of multinational activity, such as the rationale for internalizing foreign transactions and the existence of heterogeneous participation of firms in FDI (or affiliate) sales and in global sourcing.⁴

Motivated by these new empirical findings, recent trade theory has been developed in frameworks that incorporate intraindustry firm heterogeneity. The seminal paper in the literature is that of Melitz (2003), which follows closely the structure of Krugman (1980). Although Melitz's framework fea-

⁴The fact that firms engaged in FDI sales and in importing appear to be distinct from other firms has been documented, among others, by Helpman, Melitz and Yeaple (2004) and Bernard, Jensen and Schott (2009). In addition, Ramondo, Rappoport and Ruhl (2013) have recently documented that U.S. intrafirm trade appears to be highly concentrated among a small number of large foreign affiliates.

tures no multinational activity, no global sourcing and no contractual frictions, it is natural to begin our incursion into theoretical territory with a variant of his model.

A Multi-Sector Melitz Model

Consider a world consisting of J countries that produce goods in $S+1$ sectors using a unique (composite) factor of production, labor, which is inelastically supplied and freely mobile across sectors. One sector produces a homogenous good z , while the remaining S sectors produce a continuum of differentiated products. Preferences are identical everywhere in the world and given by:

$$U = \beta_z \log z + \sum_{s=1}^S \beta_s \log Q_s, \quad (2.1)$$

with $\beta_z + \sum_{s=1}^S \beta_s = 1$ and

$$Q_s = \left(\int_{\omega \in \Omega_s} q_s(\omega)^{(\sigma_s-1)/\sigma_s} d\omega \right)^{(\sigma_s/\sigma_s-1)}, \quad \sigma_s > 1. \quad (2.2)$$

It is worth pausing to discuss the specific assumptions we have already built into the model. The preferences in (2.1) feature a unit elasticity of substitution across sectors, so industry spending shares are constant. Within differentiated-good sectors, the preferences in (2.2) are of the Dixit-Stiglitz type: there is a continuum of varieties available to consumers and these enter preferences symmetrically and with a constant, higher-than-one-elasticity of substitution between any pair of varieties. These assumptions are special, but they are standard in the International Trade field. In particular, the preferences in (2.1) and (2.2) are a strict generalization of those in Krugman (1980) and Melitz (2003), which correspond to the case $\beta_z = 0$ and $S = 1$.⁵ I incorporate multiple differentiated-good sectors because this will facilitate the derivation of cross-sectional predictions, while the presence of a homogeneous-good sector will simplify the general equilibrium aspects of the

⁵To be precise, the last section of Krugman (1980) develops a two-industry model featuring cross-country differences in demand patterns.

model. I will however consider the Krugman-Melitz, one-sector version of the model at times in the book. It would be valuable to follow the approach of Helpman and Krugman (1985) and work out the robustness of the results below to more general preference structures, but I will not attempt to do so in this book.⁶

Given (2.1), consumers in country j will optimally allocate a share β_z of their spending E_j to good z and a fraction β_s to differentiated-good sector s . I will use the subscripts i and j to refer to countries, with i denoting producing/exporting countries and j denoting consuming/importing countries. In order to keep the notation as neat as possible, I will drop the subscript s associated with differentiated-good sectors and their sector-specific parameters. Similarly, and although the model is dynamic (time runs indefinitely), I will omit time subscripts throughout since I will focus on describing stationary equilibria.

Within a representative differentiated-good sector then consumers allocate spending across varieties to maximize Q in (2.2), which gives rise to following demand for variety ω in country j :

$$q_j(\omega) = \beta E_j P_j^{\sigma-1} p_j(\omega)^{-\sigma}, \quad (2.3)$$

where $p_j(\omega)$ is the price of variety ω , P_j is the ideal price index associated with (2.2),

$$P_j = \left[\int_{\omega \in \Omega_j} p_j(\omega)^{1-\sigma} d\omega \right]^{1/(1-\sigma)}, \quad (2.4)$$

and Ω_j is the set of varieties available to consumers in j .

Consider next the supply side of the model. The homogenous good is produced with labor under conditions of perfect competition, and according to a constant-returns-to-scale technology which is allowed to vary across countries. In particular, output is equal to

$$z_i = L_{zi}/a_{zi}, \quad (2.5)$$

⁶As it will become apparent, however, the Cobb-Douglas assumption in (2.1) is of little relevance for the main results derived in future chapters of the book. Also, the literature has developed versions of the Melitz (2003) model with alternative, specific functional forms for the aggregate industry index Q_m (see, for instance, Melitz and Ottaviano, 2008, or Novy, 2013). Relaxing the assumption of a continuum of varieties would severely complicate the analysis by introducing strategic pricing interactions across firms within an industry.

where L_{zi} is the amount of labor in country i allocated to the production of good z , and a_{zi} is country i 's unit labor requirement in that sector. The homogeneous good z is freely tradable across countries and will serve as the numéraire in the model.

The differentiated-good industries are instead monopolistically competitive. Each variety is produced by a single firm under a technology featuring increasing returns to scale, and there is free entry into each industry. The existence of internal economies of scale stems from the presence of three types of fixed costs. First, the process of entry and differentiation of a variety entails a fixed cost of f_{ei} units of labor in country i . Second, production of final-good varieties in country i entails an overhead cost equal to f_{ii} units of country i 's labor. Finally, firms in country i need to incur an additional fixed 'market access' cost equal to f_{ij} units of labor in order to export in country $j \neq i$. These fixed export costs capture costs associated with marketing and distributing goods in foreign markets that need to be incurred regardless of the volume exported. I will specify these costs in terms of the exporting country's labor, but not much would change if they were specified in terms of the importing country's labor. Notice, that we do not assume that $f_{ij} > f_{ii}$, but the latter type of fixed costs of production need to be incurred before the firm can sell in any market.

The fixed cost parameters f_{ei} , f_{ii} and f_{ij} are common for all firms within an industry. Intraindustry heterogeneity stems from differences in the marginal cost of production faced by firms. In particular, after incurring the fixed cost of entry f_{ei} , firms learn their productivity level φ , which determines their marginal cost of production, $1/\varphi$, in terms of labor. These productivity levels are drawn independently from a cumulative distribution function $G_i(\varphi)$ which is assumed Pareto with shape parameter $\kappa > \sigma - 1$, so

$$G_i(\varphi) = 1 - \left(\frac{\varphi_i}{\varphi}\right)^\kappa, \quad \text{for } \varphi \geq \underline{\varphi}_i > 0. \quad (2.6)$$

The marginal cost of servicing foreign markets is further magnified by 'iceberg' trade costs such that $\tau_{ij} > 1$ units of output need to be shipped from country i for 1 unit to make it to country j . The firm productivity parameter φ is time invariant, but firms face a common, exogenous probability $\delta \in (0, 1)$ of being subject to a (really) bad shock that would force them to exit, which keeps the value of the firm bounded for any φ .

When selling to local consumers, firms need not incur variable trade costs ($\tau_{ii} = 1$) nor market access costs in excess of the production fixed cost f_{ii} .

Under the mild assumption that any firm with positive production sells some amount of output in their domestic market, we can then succinctly express the cost for a firm with productivity φ of producing q units of output in country i and selling them in country j as

$$C_{ij}(q) = \left(f_{ij} + \frac{\tau_{ij}}{\varphi} q \right) w_i, \quad (2.7)$$

where note that the formula applies both for foreign ($i \neq j$) as well as for domestic sales ($i = j$).

This completes the description of the model. Before discussing some features of the equilibrium, it is worth briefly relating the model above to other ones in the literature. The structure of the model is most closely related to that of the multi-sector Melitz models in Arkolakis, Demidova, Klenow and Rodríguez-Clare (2008) and Helpman, Melitz and Rubinstein (2008).⁷ The original model in Melitz (2003) corresponds to the particular case in which $\beta_z = 0$ and $S = 1$, and parameters are fully symmetric across countries, so $f_{ei} = f_e$, $f_{ii} = f$, $f_{ij} = f_X$, $\tau_{ij} = \tau$ and $L_i = L$, where L_i is the stock of labor in country i .⁸ As hinted above, the seminal paper of Krugman (1980) – except for its last section – is also a special case of the framework above, in which on top of the assumptions in Melitz (2003), there are no fixed marketing costs $f_X = 0$ and the distribution of productivity $G_i(\varphi)$ is degenerate, so firms are homogenous.⁹

Selection into Exporting

I next illustrate how this simple model is able to explain some of the firm-level exporting facts discussed above. Given the isoelastic demand in (2.3), firms will charge a price in each market in which they sell equal to a constant markup $\sigma/(\sigma - 1)$ over the marginal cost of servicing that market. As a

⁷Melitz and Redding (2013a) have recently used a model with a very similar structure to navigate the literature on heterogeneous firms and trade. Chaney (2008) also develops a multi-sector Melitz framework but does not allow for free entry.

⁸The above model is less general than Melitz (2003) in that I impose that $G(\varphi)$ is Pareto, while he considers a general cumulative probability distribution.

⁹A hybrid model in the spirit of Helpman and Krugman (1985) could also be derived from our model (whenever $\beta_z > 0$) if we allowed sectors to use two factors of production (say capital and labor) under different factor intensities. Also, our benchmark model could easily be turned into the standard neoclassical Ricardian and Heckscher-Ohlin models by setting $\sigma \rightarrow \infty$ and all fixed costs to 0.

result, the potential operating profits for a firm from i with productivity φ considering servicing a particular market j can be concisely written as

$$\pi_{ij}(\varphi) = (\tau_{ij}w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_{ij} \quad (2.8)$$

where

$$B_j = \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} P_j^{\sigma-1} \beta E_j. \quad (2.9)$$

Notice that $\pi_{ij}(\varphi)$ increases linearly with the transformation of productivity $\varphi^{\sigma-1}$ and that for a sufficiently low φ , $\pi_{ij}(\varphi)$ is necessarily negative. More formally, only the subset of firms from i with productivity $\varphi \geq \tilde{\varphi}_{ij}$, where

$$\tilde{\varphi}_{ij} \equiv \tau_{ij} w_i \left(\frac{w_i f_{ij}}{B_j} \right)^{1/(\sigma-1)}, \quad (2.10)$$

will find it optimal to export to country j . Other things equal, the higher are trade barriers between i and j (τ_{ij} and f_{ij}), the lower will be the share of firms in i choosing to service j . This contrasts with homogeneous firm models, in which all firms from i would sell to all possible markets j .

The model also sheds light on the fact that exporters typically appear to be more productive than non-exporters. In particular, provided that the profit shifter B_j does not vary too much across countries, and given that $f_{ij} > 0$ for all $j \neq i$, firms will find it relatively harder to profitably sell in foreign markets than in their local market. Furthermore, provided that $\tau^{\sigma-1} f_{ij} > f_{ii}$ for all $j \neq i$, it will necessarily be the case that a positive measure of firms sells domestically but does not export. These intuitive results regarding selection into exporting and productivity differences between exporters and non-exporters are depicted in Figure 2.1. The figure exploits the fact that firm profits $\pi_{ij}(\varphi)$ in (2.8) increase linearly with the transformation of productivity $\varphi^{\sigma-1}$. The lower slope of the export profit line reflects the large variable transport costs $\tau_{ij} > 1$ (remember that we are assuming small differences in the profit shifter B_j across countries). In the figure, it is also assumed that $f_{ij} > f_{ii}$, although we have noted above that $\tau^{\sigma-1} f_{ij} > f_{ii}$ is sufficient to obtain selection into exporting.¹⁰

¹⁰Note that no matter how low f_{ij} is, whenever B_j is identical across countries, it will never be the case that a firm from i produces *only* for a particular export market $j \neq i$. This is because in such a case, these *pure* exporters would need to incur the overhead costs f_{ii} on top of the marketing costs f_{ij} , and thus we would have $f_{ij} + f_{ii} > f_{ii}$. Lu




Figure 2.1: Selection into Exporting with Heterogeneous Firms

The Extensive Margin, Gravity and Reallocation Effects

The logic behind the fact that a model with heterogeneous firms and fixed export costs can deliver selection into exporting based on productivity is hardly earth-shattering. The beauty of the Melitz (2003) model resides in the fact that, despite its simple structure, it can account for several additional features documented in empirical studies. These additional results from the model are less central for the set of results emphasized in this book, but nonetheless it is worth discussing them briefly.

Consider first the implications of the model for aggregate exports at the sectoral level. Letting X_{ij} denote aggregate exports from i to j in a representative differentiated-good sector, and denoting by N_i the measure of potential producers from i in that sector (i.e., the set of firms that have paid

(2011) shows that a ‘reverse’ sorting is observed among Chinese manufacturing firms, a fact that she attributes to a particularly low value of B_j in labor-intensive industries in China relative to other countries.

the fixed cost of entry $w_i f_{ei}$), we have

$$X_{ij} = N_i \int_{\tilde{\varphi}_{ij}}^{\infty} \sigma (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} dG_i(\varphi), \quad (2.11)$$

where we have used (2.8) and the fact that export revenues are a multiple σ of $\pi_{ij}(\varphi) + w_i f_{ij}$.

A first point to notice is that variation in exporting across destination markets j is composed of an extensive margin and an intensive margin. In particular, we can write

$$X_{ij} = N_{ij} \cdot \bar{x}_{ij},$$

where $N_{ij} = (1 - G_i(\tilde{\varphi}_{ij})) N_i$ is the actual measure of firms from i selling in j (the extensive margin) and

$$\bar{x}_{ij} = \frac{1}{1 - G_i(\tilde{\varphi}_{ij})} \int_{\tilde{\varphi}_{ij}}^{\infty} \sigma (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} dG_i(\varphi), \quad (2.12)$$

are average firm-level exports (the intensive margin). As first worked out by Arkolakis, Demidova, Klenow and Rodríguez-Clare (2008), when productivity is distributed Pareto as in (2.6), integrating (2.12) and using (2.10) to simplify, delivers

$$\bar{x}_{ij} = \frac{\kappa}{\kappa - \sigma + 1} \sigma w_i f_{ij}, \quad (2.13)$$

and thus the intensive margin is independent of variable trade costs and of market size of the destination country. In other words, the model is consistent with export volumes from i to j being lower for smaller and more distant markets, but the reason for this is very different than in homogeneous firm models à la Krugman. It is not because firms export on average lower volumes to those markets but rather because a smaller set of firms export to those markets. As shown by Chaney (2008), this is not an immaterial distinction, since it critically affects, for instance, how the elasticity of trade flows to trade frictions depends on the elasticity of substitution σ .

Another remarkable feature of the model is that it delivers a modified sectoral version of the gravity equation for trade flows, which has been shown to fit the data rather well. As shown by Melitz and Redding (2013a) (see also Chapter 3 for a related derivation), in the Pareto case, aggregate exports in (2.11) can alternatively be expressed as

$$X_{ij} = \frac{Y_i}{\Theta_i} \left(\frac{\beta E_j}{P_j^{1-\sigma}} \right)^{\frac{\kappa}{(\sigma-1)}} \tau_{ij}^{-\kappa} f_{ij}^{-\frac{\kappa-(\sigma-1)}{\sigma-1}}, \quad (2.14)$$

where Y_i is the aggregate industry output in i (i.e., $Y_i \equiv \sum_j X_{ij}$) and Θ_i is a structural measure of country i 's market potential in that industry.¹¹ Notice that equation (2.14) structurally justifies the use of empirical log-linear specifications for sectoral trade flows with importer-sector and exporter-sector asymmetric fixed effects and measures of bilateral trade frictions. In the one-sector models of Krugman (1980) and Melitz (2003) (i.e., $\beta_z = 0$ and $S = 1$), the model predicts that the gravity equation will hold for aggregate bilateral trade flows across countries, and as shown by Helpman, Melitz and Rubinstein (2008), for estimation purposes, the model serves a very useful role in structurally correcting for the large number of bilateral zero trade flows in the data (we will cover their contribution in more detail in Chapter 3).

In this same one-sector version of the model, Arkolakis, Costinot and Rodríguez-Clare (2012) have derived a neat formula for the welfare effects of trade in terms of two sufficient statistics: the import penetration ratio and the elasticity of imports with respect to variable trade costs. Arkolakis, Costinot and Rodríguez-Clare (2012) have also shown that, remarkably, this formula is identical to the one obtained in the Anderson and van Wincoop (2003), Eaton and Kortum (2002), and Krugman (1980) models.¹²

One final aspect of the model that is worth discussing is its ability to rationalize the reallocation effects following trade liberalization documented by the empirical literature. This is most elegantly derived in the symmetric, one-sector model of Melitz (2003) in which no parametric assumptions on $G(\varphi)$ are imposed. Essentially, what Melitz shows is that reductions in trade costs will not only expand the number and revenues of exporting firms, but will also (via competition effects) reduce the scale of non-exporting firms and will also lead to the exit of a set of producers that were marginally profitable before the reduction in trade costs.¹³ Formally, in terms of the

¹¹In particular, $\Theta_i \equiv \sum_j \left(\frac{\beta E_j}{P_j^{1-\sigma}} \right)^{\frac{\kappa}{\sigma-1}} \tau_{ij}^{-\kappa} f_{ij}^{-\frac{\kappa-(\sigma-1)}{\sigma-1}}$.

¹²Because the import penetration ratio and the ‘trade elasticity’ respond to trade opening in distinct manners in these different frameworks, their results do not necessarily imply, however, that information on the microstructure of these models is irrelevant for assessing the welfare consequences of trade liberalization (see Melitz and Redding, 2013*b*, for more on this).

¹³Even though the size of continuing exporters increases and that of continuing non-exporters shrinks, in the case in which $G(\varphi)$ is Pareto, the extensive margin responses ensure that the average size of exporters and non-exporters will remain unaffected by

notation above, Melitz (2003) shows that reductions in trade costs will not only reduce $\tilde{\varphi}_{ij}$, but will also increase $\tilde{\varphi}_{ii}$ thus forcing firms with productivity marginally above $\tilde{\varphi}_{ii}$ to shut down. As discussed by Baldwin and Forslid (2010) and Arkolakis, Demidova, Klenow and Rodríguez-Clare (2008), this in turn leads to ‘anti-variety’ effects by which the measure of varieties available to consumers will, under plausible conditions, decrease following trade liberalization.

The Melitz (2003) model has been extended in a variety of fruitful ways, ranging from the exploration of alternative demand systems, the introduction of Heckscher-Ohlin features into the model, the modeling of multi-product firms, and many others. Several applications and extensions of the model are reviewed in Melitz and Redding (2013a). I will next focus on an extension of the model that is particularly relevant for the study of the global organization of production, which is the central topic of this book.

Global Sourcing with Heterogeneous Firms

In the Melitz (2003) model the only involvement of firms with foreign markets is via the exportation of final goods produced with local labor. As documented in Chapter 1, the recent process of globalization has led to a disintegration of the production process across borders in which international trade in intermediate inputs has been a dominant feature in the world economy. I next develop a simple variant of the Melitz framework in which firms not only export, but also make global sourcing decisions related to the location and quantity of inputs to buy from different countries.

In order to meaningfully study offshoring, one needs to consider multi-stage production processes, and a natural starting point is a two-stage model. With that in mind, assume that the production of varieties in the differentiated-good sectors now involves two stages, which we will refer throughout the book as *headquarter services* and *manufacturing production*. Headquarter services may include a variety of activities such as R&D expenditures, brand development, accounting, and finance operations, but may also involve high-tech manufacturing or assembly. The important characteristic of this stage in terms of the model is that these activities need to be produced in the same country in which the entry cost f_{ei} was incurred. Manufacturing production can instead be thought of as entailing low-tech manufacturing or assembly

changes in variable trade barriers, as shown in (2.13).

of inputs into a final product. Crucially, we will depart from Melitz (2003) in allowing manufacturing production to be geographically separated from the location of entry and headquarter services provision. This is a highly simplified characterization of the process of offshoring, but we will work to enrich the model later in the book.

Relative to the multi-sector Melitz (2003) framework developed above, the key new decision facing firms is thus whether to maintain plant production in the same country in which entry and headquarter service provision takes place, or whether to offshore that stage. In order to simplify the model and isolate the new insights arising from the modelling of offshoring, we shall assume that there are no costs, fixed or variable, associated with exporting final goods so that the exporting decision is trivial and all firms producing final goods export them worldwide. Conversely, the decision of whether to source locally or engage in offshoring will be nontrivial: offshoring will be associated with a reduction in production costs but will also entail additional fixed and variable transportation costs that might lead some firms to opt out of that strategy.

More formally, the overall costs of producing q units of a final-good variety incurred by a firm with headquarters in country i and manufacturing production in country j (with possibly $j = i$) are given by

$$C_{ij}(q, \varphi) = f_{ij}w_i + \frac{q}{\varphi} (a_{hi}w_i)^\eta (\tau_{ij}a_{mj}w_j)^{1-\eta}. \quad (2.15)$$

As before, φ is a firm-specific productivity parameter. The parameters f_{ij} , τ_{ij} , η , a_{hi} and a_{mj} are instead sector specific but common across firms within a sector s , while the wage rates w_i and w_j vary only across countries. The parameters f_{ij} and τ_{ij} appeared already in the Melitz (2003) model (see equation (2.7)) but their interpretation is somewhat different in the present context. In particular, f_{ij} and τ_{ij} now reflect the fixed and variable trade costs associated with a particular sourcing strategy. Although, we will often associate f_{ij} and τ_{ij} with the costs of transporting intermediate inputs across countries, these parameters can be interpreted more broadly to reflect other technological barriers associated with international fragmentation, such as communication costs, language barriers or search costs. For these reasons, it is now natural to assume not only that $\tau_{ij} > \tau_{ii}$, but also that $f_{ij} > f_{ii}$ whenever $j \neq i$.

Relative to the specification of technology in (2.7), the new parameters are η , a_{hi} and a_{mj} . The first of these captures the headquarter services

intensity (or *headquarter intensity* for short) of the production process, and the associated primal representation of technology (leaving aside fixed costs and trade costs) is a Cobb-Douglas technology in headquarter services h and manufacturing production m :

$$q(\varphi) = \varphi \left(\frac{h(\varphi)}{\eta} \right)^\eta \left(\frac{m(\varphi)}{1-\eta} \right)^{1-\eta}, \quad 0 < \eta < 1. \quad (2.16)$$

Finally, the parameters a_{hi} and a_{mj} capture the unit labor requirements associated with headquarter service provision and manufacturing production and these are allowed to vary across sector and countries reflecting comparative advantage considerations.

Although the benchmark model of offshoring we have developed is quite stylized it is a generalization of a complete-contracting variant of the heterogeneous firm model in Antràs and Helpman (2004). In particular, in Antràs and Helpman (2004) it is further assumed that:

- The world consists of only two countries, North and South.
- The homogenous good z is always produced in both countries but with a higher labor productivity in the North, thus implying that wages rates are fixed at $w^N = 1/a_{zN} > 1/a_{zS} = w^S$.
- The South features either very low productivity in producing headquarter service or very high fixed costs of entry, so that all entry and headquarter service provision occurs in the North, where $a_{hN} = 1$.
- Plant production can be done with the same physical productivity – in particular, $a_{mN} = a_{mS} = 1$ – in both North and South, so offshoring to South offers a production cost advantage.

With these additional assumptions, and simplifying further the notation by denoting $f_{NN} = f_D$, $f_{NS} = f_O$, $\tau_{NS} = \tau$, the total cost of production associated with **D**omestic sourcing and **O**ffshoring can be written, respectively, as

$$C_D(q, \varphi) = \left(f_D + \frac{q}{\varphi} \right) w_N, \quad (2.17)$$

and

$$C_O(q, \varphi) = f_O w_N + \frac{q}{\varphi} (w_N)^\eta (\tau w_S)^{1-\eta}. \quad (2.18)$$

Selection into Offshoring

We can now study the implications of the above framework for the selection of firms into offshoring. For now I will focus on the simplified two-country framework in Antràs and Helpman (2004) since it has featured prominently in the literature, but at the end of the chapter I will discuss how the results can be extended to a multi-country environment.

In light of the cost functions in (2.17) and (2.18), and given that firms charge a price for the final good equal to a constant markup $\sigma/(\sigma - 1)$ over the marginal cost of production, the potential operating profits for a Northern firm with productivity φ associated with Domestic sourcing and Offshoring can be expressed as

$$\pi_D(\varphi) = (w_N)^{1-\sigma} B\varphi^{\sigma-1} - f_D w_N \quad (2.19)$$

and

$$\pi_O(\varphi) = ((w_N)^\eta (\tau w_S)^{1-\eta})^{1-\sigma} B\varphi^{\sigma-1} - f_O w_N, \quad (2.20)$$

respectively, where

$$B = \frac{1}{\sigma} \left(\frac{\sigma}{(\sigma - 1)P} \right)^{1-\sigma} \beta (w_N L_N + w_S L_S)$$

and P is the common price index in (2.4) in each country, given costless international trade in final goods.

As in Melitz (2003), the profit functions $\pi_D(\varphi)$ and $\pi_O(\varphi)$ are linearly increasing in the transformation of productivity $\varphi^{\sigma-1}$ and for a sufficiently low φ , both of these profit levels necessarily take negative values. Hence, upon observing their productivity, the least productive firms in an industry will optimally decide not to produce. Furthermore, the fact that $f_O > f_D$ ensures that for sufficiently low levels of productivity, we have $\pi_D(\varphi) > \pi_O(\varphi)$, and offshoring is not a viable option in situations in which domestic sourcing might be profitable. In fact, whenever

$$f_O > \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)} f_D, \quad (2.21)$$

there always exists a subset of firms in the industry that find it optimal to opt out of offshoring and decide instead to source locally in the North. In order for some firms within the industry to find it optimal to offshore it is




Figure 2.2: Equilibrium Offshoring Sorting with High Wage Differences

necessary to assume that offshoring trade costs τ are low enough to ensure that $w_N > \tau w_S$. This case is depicted in Figure 2.2, which also drawn under the implicit assumption that condition (2.21) holds.

As shown by the Figure, the model features selection into offshoring by which only the most productive firms within an industry find it worthwhile to pay the fixed costs of fragmentation to benefit from the lower production costs associated with manufacturing in the South.¹⁴ In particular, offshoring is the preferred option only for firms with productivity $\varphi \geq \tilde{\varphi}_O$, where

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

The sorting pattern in Figure 2.2 is consistent with the evidence on selection into importing in Bernard, Jensen, Redding and Schott (2007), who show that not only U.S. exporting firms but also U.S. importing firms appear

¹⁴In the much less interesting case in which $w_N < \tau w_S$, no firm in the industry finds it optimal to offshore and if this condition holds for all sectors of the economy, then the South is fully specialized in the production of the homogenous good z .

to be more productive than purely domestic producers. Their results are reproduced in Table 2.1 below, which shows that U.S. manufacturing plants that import employ more workers, sell more, are more productive, pay higher wages and are more capital and skill intensive than plants that do not source abroad.

Table 2.1. Trading Premia in U.S. Manufacturing, 1997

	Exporter Premia	Importer Premia
Log Employment	1.50	1.40
Log Shipments	0.29	0.26
Log Value-Added per Worker	0.23	0.23
Log TFP	0.07	0.12
Log Wage	0.29	0.23
Log Capital per Worker	0.17	0.13
Log Skill per Worker	0.04	0.06

Source: Bernard, Jensen and Schott (2009), Table 8.

More specifically, firms that import appear to be 12% more productive than firms that do not, while the productivity advantage of exporting plants is only of 7%. Furthermore, Bernard, Jensen, Redding and Schott (2007) report that only 14 percent of U.S. manufacturing plants report positive imports (versus 27 percent of plants reporting positive exports), which is again suggestive of the existence of significant fixed costs of importing.

One might be concerned that the patterns observed by Bernard, Jensen, Redding and Schott (2007) do not necessarily support the sorting pattern implied by the theory because it is unclear whether the goods that firms are importing are intermediate inputs, rather than finished goods. In the latter case, one might worry that Table 2.1 is simply picking the role of large intermediaries (wholesalers or retailers) in bringing consumer goods into the United States. The fact that the importer premia reported in Table 2.1 correspond to the operations of U.S. *manufacturing* firms should however dispel that concern. Furthermore, recent work by Fort (2013) using U.S. Census data demonstrates that a similar sorting pattern is observed when focusing on imports of contract manufacturing services, which cover exclusively offshoring of inputs that are customized to specific U.S. firm's production needs. More specifically, Fort (2013) finds that U.S. firms that offshore contract manufacturing services feature on average 13% higher valued-added la-

bor productivity than U.S. firms in the same six-digit NAICS that purchase those services only domestically.

Figure 2.3 provides further confirmation of the superior performance of offshoring firms with 2007 data from the Spanish Encuesta sobre Estrategias Empresariales (ESEE). The dataset distinguishes between firms that purchase inputs only from other Spanish producers and firms that purchase inputs from abroad. As is clear from the picture, the distribution of productivity of firms that engage in foreign sourcing is a shift to the right of that of firms that only source locally.¹⁵




Figure 2.3: Selection into Offshoring in Spain

Determinants of the Prevalence of Offshoring

We can next use this simple model of global sourcing to study the determinants of the relative prevalence of offshoring in an industry. For instance, consider computing the share of spending on *imported* manufacturing inputs over total manufacturing input purchases in a particular industry. Given the Cobb-Douglas technology in (2.16) and the CES preferences in (2.1), manufacturing input purchases will constitute a share $(\sigma - 1)(1 - \eta) / \sigma$ of revenue

¹⁵More details on this dataset and on the measurement of productivity according to the Olley-Pakes method will be provided in Chapter ??.

for all firms, while revenue itself will be a multiple σ of firm operating profits (defined as revenue minus variable costs). Using the profit functions (2.19) and (2.20) and cancelling common terms, we can thus express the share of imported manufacturing input purchases in a given industry as

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

Particularly sharp results can be obtained when assuming that the distribution of firm productivity is Pareto as in equation (2.6) in which case we obtain

$$\Upsilon_O = \frac{\left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}{\left(\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D}\right)^{\kappa-(\sigma-1)} - 1 + \left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)}}. \quad (2.22)$$

where

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\left(\frac{w_N}{\tau w_S}\right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)}. \quad (2.23)$$

As indicated by equation (2.22) and (2.23), the prevalence of offshoring is naturally increasing in the wage gap (w_N/w_S) and decreasing in fragmentation barriers (f_O/f_D , τ). These comparative statics are quite intuitive. Note that the elasticity of substitution σ and the parameter κ governing the thickness of the right tail of the Pareto distribution also have an impact on the prevalence of offshoring in an industry. The intuition for these effects is analogous to that in Helpman, Melitz and Yeaple (2004). In particular, the Pareto parameterization of productivity combined with CES preferences imply that the distribution of sales of all active firms is also Pareto with shape parameter $\kappa/(\sigma-1)$. As a result, a decrease in κ raises the sales and input purchases of firms with productivity $\varphi > \tilde{\varphi}_O$ – i.e., firms that find offshoring optimal – relative to the sales and input purchases of firms with productivity $\varphi \in (\tilde{\varphi}_D, \tilde{\varphi}_O)$ – i.e., firms that source domestically. Because the standard deviation of the logarithm of sales by all active firms in the industry is equal to $(\sigma-1)/\kappa$, this result can be interpreted as indicating that the prevalence of offshoring should be higher in industries with a larger dispersion in firm

size.¹⁶

Note that the elasticity of substitution σ affects positively the share of imported inputs for an additional reason – see the exponent of $w_N/(\tau w_S)$ in (2.22) and (2.23). The intuition for this effect is simpler: the more substitutable final-good varieties are, the more elastic will demand be and the higher will be the incentive of firms to engage in a costly investment (in this case offshoring) to reduce the marginal cost of input provision from w_N down to τw_S .

Back to the Multi-Country Model

Having worked with a simplified two-country model to build intuition, we can now go back to the multi-country environment in which the overall costs of producing q units of a final-good variety faced by a firm with headquarters in country i and manufacturing production in country j are given by equation (2.15). Given CES preferences over final-good varieties, it is then straightforward to show that the operating profits associated with that sourcing strategy are given by

$$\pi_{ij}(\varphi) = ((a_{hi}w_i)^\eta (\tau_{ij}a_{mj}w_j)^{1-\eta})^{1-\sigma} B\varphi^{\sigma-1} - f_{ij}w_i, \quad (2.24)$$

where B is now given by

$$B = \frac{1}{\sigma} \left(\frac{\sigma}{(\sigma-1)P} \right)^{1-\sigma} \beta \sum_j w_j L_j$$

and P is the common price index (2.4) for final-good varieties in each country.

Equation (2.24) illustrates again that the profit levels associated with different sourcing strategies are all linear in $\varphi^{\sigma-1}$ and thus the sourcing decision of firms can be analyzed with graphs analogous to that in Figure 2.2. Of course, with multiple countries the range of possible sorting patterns is much more complex, but we can still derive some general results.

For instance, as long as $f_{ij} > f_{ii}$ for all $j \neq i$, so domestic sourcing is the sourcing strategy associated with the lowest fixed costs, the model can only deliver a positive amount of offshoring in an industry whenever $\tau_{ij}a_{mj}w_j < a_{mi}w_i$ for some country $j \neq i$. Importantly, in such a case,

¹⁶Other measures of industry firm size dispersion, such as the Theil index also vary monotonically with $(\sigma-1)/\kappa$.

if firms sourcing domestically and abroad coexist within an industry, then firms that offshore are necessarily larger and more productive than firms that source domestically. In sum, under the plausible condition $f_{ij} > f_{ii}$ for all $j \neq i$, the model continues to predict selection into offshoring based on productivity in a manner consistent with the U.S. import premia in Table 2.1 and with the evidence from Spain depicted in Figure 2.3.

It is also noteworthy that in contrast to the simple two-country model above, this multi-country extension of the model can easily generate two-way intermediate input trade flows across countries. For instance, a given country i can feature a high manufacturing productivity level $1/a_{mj}$ in some industries and a very low one in others. In the former type of industries, this country i may well export inputs to firms with headquarters located in other countries (particularly when country i 's productivity in headquarter provision is low in that industry), while it may well import manufacturing inputs in the latter type of industries.¹⁷

With multiple countries, firms not only decide on whether to offshore manufacturing production or not but also choose the optimal location of production among all possible ones. It is evident that, other things equal, firms based in country i will be drawn to locations j entailing low fixed costs of sourcing f_{ij} and low variable costs of manufacturing, as summarized by $\tau_{ij}a_{mj}w_j$. Some highly productive firms might however be drawn to locations with high sourcing fixed costs as long as those locations offer a particularly favorable marginal cost of input manufacturing.

Figure 2.4 depicts a possible equilibrium in a world of four countries, a 'Home' country i and three 'Foreign' countries j , k and l . Domestic sourcing is the least fixed cost sourcing strategy, and as argued above, this is the preferred option for the least productive among the active firms in the industry. Offshoring to country l entails high fixed costs and also high variable costs (perhaps due to high transportation costs τ_{il} or high productivity-adjusted manufacturing wages $a_{ml}w_l$), and thus no firm finds it optimal to import inputs from l . Country j offers the largest marginal cost savings when offshoring there, but the fixed costs of fragmentation are high there, so only the most productive firms within an industry find it optimal to import inputs from j . Finally, country k is associated with moderate fixed costs of

¹⁷From this discussion, it should be obvious that the two-country model developed above failed to deliver two-way input trade flows because of its assumptions on technology (e.g., ruling out headquarter services provision in the South), and not because it only featured two countries.

offshoring and offers a cost advantage relative to domestic sourcing, so a subset of middle-productivity firms chooses it as their optimal location of manufacturing input production.




Figure 2.4: Selection into Offshoring with Multiple Countries

Naturally, the example illustrated in Figure 2.4 is rather special and, more worryingly, very different sorting patterns could emerge with mild changes in the key productivity and cost parameters. To illustrate this sensitivity, consider the case in which all foreign countries share the same level of offshoring fixed costs or $f_{ij} = f_{iO}$ for all $j \neq i$. It is then clear, that conditional on finding it optimal to offshore, firms headquartered in country i will offshore manufacturing to the location j that minimizes marginal costs, or $j^* = \arg \min_j \{\tau_{ij} a_{mj} w_j\}$. Small changes in any of these parameters could thus lead to discontinuous jumps in the prevalence of imports of inputs from particular countries.

Another limitation of this multi-country model is that it is not well designed to aggregate all firm decisions within an industry in order to guide empirical analyses of the determinants of the relative prevalence of offshoring to particular countries depending on some fundamental parameters of those countries. For similar reasons, the model is not a particularly useful tool for

quantitative analysis, particularly when envisioning a more realistic world with multiple inputs.

Fortunately, below we will be able to make some progress on these limitations by borrowing some neat modeling techniques from a recent paper by Tintelnot (2013), who in turn builds on the seminal work of Eaton and Kortum (2002).¹⁸

Bringing Eaton and Kortum (2002) Inside the Firm

Imagine now that the manufacturing stage of production entails the procurement of a continuum of measure one of inputs indexed by v , rather than just one input as assumed so far. I let these inputs be imperfectly substitutable with each other with a constant and symmetric elasticity of substitution equal to $1/(1-\rho)$. Very little will depend on the particular value of ρ . The cost function associated with producing q units of a final-good variety faced by a firm with headquarters in country i is now given by

$$C_{i\{j(v)\}_{v=0}^1}(q, \varphi) = w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} + \frac{q}{\varphi} (a_{hi} w_i)^\eta \left(\int_0^1 (\tau_{ij(v)} a_{mj(v)}(v) w_{j(v)})^{1-\rho} dv \right)^{(1-\eta)/(1-\rho)}, \quad (2.25)$$

where $j(v)$ corresponds to the country in which input v is produced and $\mathcal{J}_i(\varphi) = \{\hat{j} : j(v) = \hat{j} \text{ for some } v\}$ is the set of locations from which this firm with productivity φ sources inputs.

I will depart from the previous model in allowing the manufacturing productivity parameters $a_{mj(v)}$ to be *firm-specific*, and following Eaton and Kortum (2002), in treating them as the realization of random variables rather than as being deterministic. More formally, by paying the fixed cost f_{ij} of offshoring to country j , a firm headquartered in country i gains the ability of having *any* input v produced in that country j under an input-specific unit labor requirement drawn (independently from other inputs) from the Fréchet distribution

$$Pr(a_{mj}(v) \leq a) = e^{-T_j a^{-\theta}}, \quad \text{with } T_j > 0 \text{ and } \theta > \rho/(1-\rho).$$

¹⁸Tintelnot's (2013) framework does not feature trade in intermediate inputs, but the same tricks he develops can be fruitfully adapted to the current setting. Garetto (2013) also applies the Eaton and Kortum (2002) framework to a global sourcing environment, but does so in a two-country model and with other goals in mind.

As in Eaton and Kortum's (2002) model, T_j governs the industry-level state of technology in country j , while θ determines the variability of productivity draws across inputs, with a lower θ fostering the emergence of comparative advantage *within* input subsectors across countries.

In order to simplify matters, it is assumed that firms only learn their particular realization of $a_{mj}(v)$ after they have incurred all sunk costs of offshoring. Hence, regardless of the different amounts that firms paid to have the ability to source from particular countries, the choice of location of production of any input v will simply solve $j^*(v) = \arg \min_{j(v) \in \mathcal{J}_i(\varphi)} \{\tau_{ij} a_{mj}(v) w_j\}$. Remember that the set $\mathcal{J}_i(\varphi)$ from which $j^*(v)$ is chosen corresponds to the set of countries in which a firm from country i with productivity φ paid the associated fixed costs of offshoring f_{ij} . I will refer to $\mathcal{J}_i(\varphi)$ as the *sourcing strategy* of a firm headquartered in i with productivity φ .

Because this model has many moving parts, it is worth pausing to review the timing of events. Firms in a given sector s (subscripts omitted) initially pay a fixed cost of entry $f_{ei}w_i$ to enter country i and gain the ability to later produce headquarter services there at a unit labor cost equal to $a_{hi}w_i$. After paying this entry cost, firms learn their core productivity φ which affects firm productivity in a Hicks-neutral manner. Firms next select a set of countries $\mathcal{J}_i(\varphi)$ from which to be able to import inputs and pay all fixed offshoring costs $w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij}$. Once those countries have been selected, the firm observes the vector of input-location-specific productivity draws $\{a_{mj}(v)\}_{v \in [0,1]}$ for each $j \in \mathcal{J}_i(\varphi)$. The firm then decides from which country to buy a particular input v , after which headquarter services and manufacturing inputs are produced, and the final good is assembled and sold in world markets.

We have obviously made the model significantly more complicated than it originally was. Some readers might then be wondering: to what effect? To understand the purpose of this added structure, consider first the choice of location of manufacturing inputs, once all offshoring fixed costs have been paid. As argued above, at that point, a firm headquartered in i with productivity φ simply solves $j^*(v) = \arg \min_{j(v) \in \mathcal{J}_i(\varphi)} \{\tau_{ij} a_{mj}(v) w_j\}$. The beauty of the Fréchet distribution (see Eaton and Kortum, 2002) is that the probability that a given location j is chosen for any input v is simply given by

$$\chi_{ij}(\varphi) = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\Theta_i(\varphi)}. \quad (2.26)$$

where

$$\Theta_i(\varphi) \equiv \sum_{k \in \mathcal{J}_i(\varphi)} T_k (\tau_{ik} w_k)^{-\theta} \quad (2.27)$$

summarizes the *sourcing potential* of firm φ from i . With a continuum of inputs, $\chi_{ij}(\varphi)$ corresponds to the fraction of inputs sourced from j conditional on the sourcing strategy $\mathcal{J}_i(\varphi)$. Even more remarkably, the distribution of the actual price paid for any input v turns out to be independent of the actual source j of those inputs (again, see Eaton and Kortum, 2002, for details), which implies that $\chi_{ij}(\varphi)$ in (2.26) also corresponds to country j 's share of all manufacturing input purchases by a firm with sourcing strategy $\mathcal{J}_i(\varphi)$.

Hopefully, the reader is beginning to appreciate that the extra machinery is starting to pay off. According to expression (2.26), and conditional on the set of active locations $\mathcal{J}_i(\varphi)$, sourcing decisions at the level of the firm now vary smoothly with the key parameters of the model. Furthermore, each country's market share in a firm's purchases of intermediates corresponds to this country's contribution to the sourcing potential $\Theta_i(\varphi)$ in (2.27). Countries in the set $\mathcal{J}_i(\varphi)$ with lower wages w_j , more advanced technologies T_j , or lower distance from country i are predicted to have higher market shares in the intermediate input purchases of firms based in country i .

Although it might seem that the core productivity parameter φ no longer plays a relevant role in the model, it is important to stress that the set of 'activated' offshoring locations $\mathcal{J}_i(\varphi)$ is endogenous and will naturally be a function of that core productivity level. To see this, let us then turn to studying the determination of the set $\mathcal{J}_i(\varphi)$.

After choosing the least cost source of supply for each input v , the overall cost function associated with producing q units of a final-good variety can be written, after some nontrivial derivations, as

$$C_i(q, \varphi, \mathcal{J}_i(\varphi)) = w_i \sum_{j \in \mathcal{J}_i(\varphi)} f_{ij} + \frac{q}{\varphi} (a_{hi} w_i)^\eta (\gamma \Theta_i(\varphi))^{-(1-\eta)/\theta}, \quad (2.28)$$

where $\gamma = [\Gamma(\frac{\theta+1-\rho}{\theta})]^{\theta/(1-\rho)}$ and Γ is the gamma function.¹⁹ Note that the addition of a new location to any potential set of active locations necessarily lowers the marginal cost faced by firms. Intuitively, an extra location grants the firm an extra cost draw for all varieties $v \in [0, 1]$. It is thus natural

¹⁹These derivations are analogous to those performed by Eaton and Kortum (2002) to solve for the aggregate price index in their model of final-good trade.

that greater competition among suppliers will reduce the expected minimum sourcing cost $\tau_{ij^*} a_{mj^*}(v) w_{j^*} = \min_{j(v) \in \mathcal{J}_i(\varphi)} \{\tau_{ij} a_{mj}(v) w_j\}$ per intermediate. In fact, the addition of a country to $\mathcal{J}_i(\varphi)$ lowers the expected price paid for *all* varieties v , and not just for those that are ultimately sourced from the country being added to $\mathcal{J}_i(\varphi)$.²⁰

Following analogous steps as in the previous models to solve for the profit function associated with the cost function in (2.25), we can express the profits associated with the optimal sourcing strategy of a firm from country i with productivity φ as the solution to the following problem:

$$\pi_i(\varphi) = \max_{\mathcal{J}_i(\varphi)} \{ (a_{hi} w_i)^{-\eta(\sigma-1)} (\gamma \Theta_i(\varphi))^{(\sigma-1)(1-\eta)/\theta} B \varphi^{\sigma-1} - w_i \sum_{k \in \mathcal{J}_i(\varphi)} f_{ik} \}. \quad (2.29)$$

As is clear from equation (2.29), when deciding whether to add a new country l to the set $\mathcal{J}_i(\varphi)$, the firm trades off the reduction in costs associated with the inclusion of that country in the set $\mathcal{J}_i(\varphi)$ against the payment of the additional fixed cost $w_i f_{il}$.

The problem in (2.29) is not straightforward to solve because the decision to include a country j in the set $\mathcal{J}_i(\varphi)$ naturally interacts with the decision to add any other country j' . For this reason, although the larger is the core productivity level φ , the higher will the marginal benefit of adding a location to any given set $\mathcal{J}_i(\varphi)$, it is not necessarily the case that the set $\mathcal{J}_i(\varphi)$ is ‘increasing’ in φ . Or, more precisely, the choice of locations $\mathcal{J}_i(\varphi_0)$ of a firm with productivity φ_0 is not necessarily a strict subset of the set of locations $\mathcal{J}_i(\varphi_1)$ chosen by a firm with a higher productivity level $\varphi_1 > \varphi_0$. For example, a highly productive firm from i might pay a large fixed cost to be able to offshore to a country l with a particularly high value of $T_l (\tau_{il} w_l)^{-\theta}$, after which the marginal incentive to add further locations might be greatly diminished whenever $(\sigma - 1)(1 - \eta) < \theta$.²¹

As we show in the Theoretical Appendix, however, these complications do not arise whenever $(\sigma - 1)(1 - \eta) \geq \theta$, in which case the addition of a

²⁰Hence, the addition of an input location decreases costs and increases revenue-based productivity for reasons quite distinct than in the love-for-variety frameworks in Halpern, Koren and Szeidl (2011), Goldberg, Khandelwal, Pavcnik and Topalova (2010), and Gopinath and Neiman (2013).

²¹The difficulties in solving for $\mathcal{J}_i(\varphi)$ are nicely discussed in Blaum, Lelarge and Peters (2013) in a model of input trade with very different features. It is worth pointing out, however, that one can easily show that the endogenous sourcing potential $\Theta_i(\varphi)$ is necessarily increasing in φ regardless of parameter values.

location to the set of active locations does not decrease the marginal benefit of adding further locations.²² As a result, one can show that the number of locations to which a firm offshores is a monotonically increasing function of productivity φ , and even more strongly, that $\mathcal{J}_i(\varphi_0) \subseteq \mathcal{J}_i(\varphi_1)$ for $\varphi_1 \geq \varphi_0$. The model thus delivers a ‘pecking order’ in the extensive margin of offshoring that is reminiscent to the one typically obtained in models of exporting with heterogeneous firms, such as in Eaton, Kortum and Kramarz (2011). Furthermore, for a sufficiently low value of core productivity φ , the only profitable location of input production might be one associated with a low fixed cost of sourcing. Under the maintained assumption that $f_{ij} > f_{ii}$ for all $j \neq i$ – so domestic sourcing is the sourcing strategy associated with the lowest fixed costs –, the model thus continues to deliver selection into offshoring based on firm core productivity.

The more tractable case with $(\sigma - 1)(1 - \eta) \geq \theta$ is more likely to apply whenever demand is elastic and thus profits are particularly responsive to variable cost reductions (high σ), and whenever input efficiency levels are relatively heterogeneous across markets (low θ), so that the expected reduction in costs achieved by adding an extra country into the set of active locations is relatively high. Naturally, this scenario is also more likely whenever head-quarter intensity η is low, and thus changes in the cost of the input bundle cost have a relatively high impact on profits.

We can obtain sharper characterizations of the solution to the sourcing strategy problem in (2.29) by making further specific assumptions. For instance, when the fixed cost of offshoring is common for all foreign countries, so $f_{ij} = f_{iO}$ for all $j \neq i$, then regardless of the value of $(\sigma - 1)(1 - \eta)/\theta$, it is clear that locations j associated with a high value of $T_j(\tau_{ij}w_j)^{-\theta}$ will necessarily be more attractive than locations associated with low values of this term. In such a case, and regardless of the value of $(\sigma - 1)(1 - \eta)/\theta$, one could then rank foreign locations $j \neq i$ according to their value of $T_j(\tau_{ij}w_j)^{-\theta}$, and denote by $i_r = \{i_1, i_2, \dots, i_{J-1}\}$ the country with the r -th

²²It is not evident which of these cases is the most relevant empirically. The share of intermediate inputs in gross output of tradeable goods is roughly 50% (see, for instance, Alvarez and Lucas, 2007), indicating $\eta \approx 0.5$. Many estimates for the elasticity of substitution exist, but the consensus is that σ is somewhere between 3 and 6 (see, for instance, Broda and Weinstein, 2006). Eaton and Kortum (2002) estimated a value of $\theta = 8.28$, which would clearly suggest $(\sigma - 1)(1 - \eta) < \theta$. Nevertheless, Simonovska and Waugh (2011) have recently estimated lower values of θ , somewhere between 2.5 and 4.5, which could imply $(\sigma - 1)(1 - \eta) > \theta$ for a value of σ in the high range of typical estimates.

highest value of $T_j (\tau_{ij} w_j)^{-\theta}$. Having constructed i_r , it then follows that for any firm with productivity φ from i that offshores to at least one country, $i_1 \in \mathcal{J}_i(\varphi)$; for any firm that offshores to at least two countries, we have $i_2 \in \mathcal{J}_i(\varphi)$, and so on. In other words, not only does the extensive margin increase monotonically with firm productivity, but it does so in a manner uniquely determined by the ranking of the $T_j (\tau_{ij} w_j)^{-\theta}$ terms.

Even with variation of fixed costs of offshoring, a similar sharp result emerges in the knife-edge case in which $(\sigma - 1)(1 - \eta) = \theta$. In that case, the addition of an element to the set $\mathcal{J}_i(\varphi)$ has no effect on the decision to add any other element to the set, and the same pecking order pattern described in the previous paragraph applies, but when one ranks foreign locations according to the ratio $T_j (\tau_{ij} w_j)^{-\theta} / f_{ij}$. This result is analogous to the one obtained in standard models of selection into exporting featuring constant marginal costs, in which the decision to service a given market is independent of that same decision in other markets.

After having solved the sourcing strategy problem in (2.29), it is straightforward to compute the aggregate volume of intermediate inputs from any country j in the industry under consideration. These imports are given by

$$M_{ij} = (\sigma - 1)(1 - \eta) \tilde{B} N_i \int_{\tilde{\varphi}_{ij}}^{\infty} \chi_{ij}(\varphi) \Theta_i(\varphi)^{(\sigma-1)(1-\eta)/\theta} \varphi^{\sigma-1} dG(\varphi), \quad (2.30)$$

where N_i is the measure of final-good entrants in country i , $\chi_{ij}(\varphi)$ is given in (2.26), $\tilde{B} = (a_{hi} w_i)^{-\eta(\sigma-1)} \gamma^{(\sigma-1)(1-\eta)} B$, and $\tilde{\varphi}_{ij}$ is the productivity of the least productive firm from i offshoring to j . As long as a higher value of $T_j (\tau_{ij} w_j)^{-\theta}$ is associated with a (weakly) higher probability that country j belongs to the set $\mathcal{J}_i(\varphi)$, it is then clear from (2.30) that a high value of $T_j (\tau_{ij} w_j)^{-\theta}$ leads to a large volume of imports from that country j on account of both the intensive and extensive margins of trade.

Interestingly, in the special case in which the fixed costs of offshoring are low enough to ensure that all firms acquire the capability to source inputs from *all* countries, equation (2.30) reduces to a modified version of the gravity equation, analogous to that in Eaton and Kortum (2002). To see this, note that whenever $\mathcal{J}_i(\varphi) = \{1, 2, \dots, J\}$ for all φ and i , (2.30) can be written as

$$M_{ij} = (\sigma - 1)(1 - \eta) \tilde{B} N_i (\bar{\Theta}_i)^{(\sigma-1)(1-\eta)/\theta} \chi_{ij} \int_{\underline{\varphi}_i}^{\infty} \varphi^{\sigma-1} dG(\varphi), \quad (2.31)$$

where

$$\bar{\Theta}_i \equiv \sum_{k=1}^J T_k (\tau_{ik} w_k)^{-\theta}$$

and

$$\chi_{ij} = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\bar{\Theta}_i}.$$

Defining $A_i = \sum_j M_{ij}$ as the total absorption of intermediate inputs by firms in i , and by $Q_j = \sum_l M_{lj}$ as the total production of intermediates in country j , it is straightforward to verify that (2.31) in fact reduces to

$$M_{ij} = \frac{(\tau_{ij})^{-\theta} \frac{A_i}{\bar{\Theta}_i}}{\sum_l (\tau_{lj})^{-\theta} \frac{A_l}{\bar{\Theta}_l}} Q_j,$$

which is analogous to equation (11) in Eaton and Kortum (2002).

When the extensive margin (or sourcing potential $\Theta_i(\varphi)$) varies across firms, such a neat expression no longer applies. This suggests that one might be able to infer the importance of the firm-level extensive margin (and of cross-country variation in the fixed costs of offshoring) from observed deviations from the traditional gravity equation. This is one of the approaches currently being explored by Antràs, Fort and Tintelnot (2014) in their study of the extensive margin of offshoring of U.S. firms.

Further Reading

This concludes my overview of the key benchmark models of international trade I will be building on in future chapters. Although most of the papers I have discussed are quite recent, there exist already a number of useful reviews of this literature. For instance, three of the chapters in the forthcoming fourth volume of the *Handbook of International Economics*, namely Melitz and Redding (2013a), Antràs and Yeaple (2013), and Costinot and Rodríguez-Clare (2013), cover these models in significant detail. The multi-country model of global sourcing is novel to this book, but it is being further developed in Antràs, Fort and Tintelnot (2014), where it is used to interpret and structurally estimate the extensive margin of U.S. intermediate input imports.

With this machinery at hand, we are now ready to begin our theoretical exploration of the implications of contractual imperfections for the global organization of production.

Chapter 3

Contracts and Export Behavior

Noble Group Limited is a global supply chain manager of agricultural and energy products, metals and minerals.¹ In January of 2004, the firm had arranged to export Brazilian soybeans to soybean crushers in China. The contracts signed in January fixed a price for the transaction, even though the delivery was only scheduled to occur in April of that same year. Unfortunately for the buyers in China, prospects for a bumper soybean crop led to a 20 percent decline in soybean prices between the months of January and April. The associated drop in the price of crushed soybeans implied that the Chinese crushers would be operating at substantial losses were they to honor the high prices fixed in their January contracts with Noble. As a result, Chinese buyers began searching for ways to nullify their January contract with Noble. Perhaps not coincidentally, that same month Chinese port authorities discovered a discoloration among a handful of red beans on a 60,000 ton soybean shipment from Brazil, which they claimed indicated the presence of carboxin, a slightly toxic fungicide. Although such discoloration (at least in small quantities) is not unusual in traded soybeans, the Chinese government proceeded to institute a ban on *all* soybean shipments from Brazil, thereby effectively voiding the contract that Noble had signed with the Chinese soybean crushers. As a result, Noble was left with millions of dollars worth of stranded cargo. Noble eventually found other buyers for its shipments, but the incident cost the company around \$25 million in demurrage losses.

This unfortunate incident of Noble Group in China exemplifies the contractual insecurity that producers face in their international transactions, the

¹The following discussion builds on Foley, Chen, Johnson and Meyer (2009).

sources of which were explained earlier in Chapter 1.² In the three chapters of this Part II of the book, I will discuss the implications of introducing contractual imperfections in the benchmark models developed in Chapter 2. In this chapter, I will develop simple imperfect-contracting variants of the Melitz (2003) model of exporting and will also discuss empirical evidence suggestive of the role of these frictions as determinants of the structure of international trade flows. In Chapter 4, I will introduce contractual frictions into the two-country model of global sourcing developed in Chapter 2. Finally, in Chapter 5, a multi-country version of this global sourcing model will be developed to guide an empirical analysis of the relevance of contractual factors for the global sourcing decisions of U.S. firms.

Contracting in the Melitz Model

As derived in Chapter 2, in the Melitz (2003) model firms set the volume of output sold and the price charged in each market in a profit-maximizing manner, and as a result, the profits that a firm from country i with productivity φ anticipates obtaining in country j are given by

$$\pi_{ij}(\varphi) = (\tau_{ij}w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_{ij}, \quad (3.1)$$

where

$$B_j = \frac{1}{\sigma} \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} P_j^{\sigma-1} \beta E_j, \quad (3.2)$$

and E_j is aggregate spending in country j .

It is worth pausing to discuss some key and often overlooked assumptions needed for a firm from i with productivity φ to *actually* realize the profit flow in equation (3.1) when choosing to export in country j . First, it is necessary for the firm to have complete information regarding all variables relevant for profits, including its own productivity level φ and the level of (residual) demand implicit in the term B_j . Second, equation (3.1) implicitly assumes that the firm can expand its production in order to meet foreign demand by costlessly hiring additional labor (or the composite factor of production) at a market wage rate w_i which is independent of the firm's operational decisions. Third, the firm is assumed to be able to costlessly contract with a local

²Interestingly, and in line with the internalization response to contractual insecurity highlighted in Chapter 1, in 2005 Noble Group acquired four soybean processing plants in China.

distributor or importer (an agent, an employee, or a firm) that will collect the sales revenue in country j and will hand them over to the exporter in i .³

Some interesting recent work in the field of International Trade has been devoted to studying the implications of relaxing the first two assumptions mentioned above. On the one hand, Segura-Cayuela and Vilarrubia (2008), Albornoz, Calvo Pardo, Corcos and Ornelas (2012), and Nguyen (2012) have all fruitfully incorporated foreign demand uncertainty in heterogeneous firm frameworks.⁴ On the other hand, a voluminous recent literature, which includes the work among others of Helpman, Itskhoki and Redding (2010) and Amiti and Davis (2012), has studied the implications of imperfect labor markets for the exporting decision, the structure of international trade and the effect of trade liberalization on labor markets, wage inequality and unemployment. As interesting as these contributions are, a treatment of these topics is beyond the scope of the current book. Instead, I will hereafter focus on relaxing the third of the assumptions mentioned above, namely that the contracting between exporters and local distributors or importers is frictionless and allows the exporter to capture the full surplus from the transaction.

Before discussing the implications of contractual imperfections in the Melitz (2003) framework it is necessary to introduce contracting into the framework and this requires us to be a bit more explicit about the agents involved in the model. For simplicity, in this chapter I will restrict attention to situations in which each export transaction involves only two agents, the exporting firm F in country i and the importer M in country j . One can think of the fixed cost of exporting $w_i f_{ij}$ as partly capturing the cost incurred by the exporter in order to be able to contract with importers from j . For the time being, I will also focus on discussing simple contracts taking the following form: at some initial date t_0 , the exporting firm F agrees to ship an amount of goods equal to q_{ij} , and in exchange the importer simultaneously agrees to pay the exporter an amount s_{ij} at some later date t_1 , perhaps corresponding to the time at which the good is received or perhaps when it has been sold and revenue has been collected. In order to avoid introducing non-essential parameters, I set the discount rate between dates t_0 and t_1 to

³Although I will abstract from such a possibility below, one could imagine that the fixed cost of exporting $w_i f_{ij}$ partly reflects the remuneration of the importer for his or her services.

⁴Conversely, models in which firms learn their productivity level φ over time, as in the seminal work of Jovanovic (1982), have not been extensively used in international trade environments.

0. Contracts with alternative timings of payments will be discussed below.

It simplifies the exposition to assume that the opportunity cost of the importer's time is 0, so that the net surplus associated with firm F with productivity φ exporting in country j continues to be given by

$$\pi_{ij}(\varphi) = \left(p_{ij}(\varphi) - \frac{\tau_{ij}}{\varphi} w_i \right) q_{ij}(\varphi) - w_i f_{ij}, \quad (3.3)$$

with $q_{ij}(\varphi) = \beta E_j P_j^{\sigma-1} p_j(\varphi)^{-\sigma}$ as dictated by the demand schedule faced by the exporting firm. In the absence of contractual frictions, the contract will set the quantity of goods $q_{ij}(\varphi)$ shipped to country j and the associated price $p_{ij}(\varphi)$ to maximize the joint surplus in (3.3), thereby leading to the joint profit flow given by

$$\pi_{ij}(\varphi) = (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_{ij}, \quad (3.4)$$

which coincides with (3.1). Only when this joint profit flow is expected to be positive, will the exporter decide to invest in being able to export to j .

Even if contracting is frictionless, whether the exporter F is able to realize that entire profit flow in (3.4) will depend on the relative bargaining power of the exporter and the importer. Given the zero reservation value of importers, the equilibrium in the Melitz (2003) framework corresponds to the case in which exporters have all the bargaining power, in the sense that they are assumed to be able to credibly make a take-it-or-leave-it offer to importers when contracting with them. To see this more formally, notice that the optimal contract from the point of view of the exporter will solve the exporter's profit subject to the importer's participation constraint, or

$$\begin{aligned} \max_{q_{ij}(\varphi), s_{ij}(\varphi)} \quad & s_{ij}(\varphi) - \frac{\tau_{ij}}{\varphi} w_i q_{ij}(\varphi) - w_i f_{ij} \\ \text{s.t.} \quad & p_{ij}(q_{ij}(\varphi)) q_{ij}(\varphi) - s_{ij}(\varphi) \geq 0, \end{aligned} \quad (3.5)$$

with $p_{ij}(q_{ij}(\varphi)) = (\beta E_j P_j^{\sigma-1})^{1/\sigma} q_{ij}(\varphi)^{-1/\sigma}$. Quite naturally, the exporter will find it optimal to make the importer's participation constraint bind, thus implying that $q_{ij}(\varphi)$ will maximize joint profits, and the exporter will end up capturing the profit flow in (3.4), as assumed in the Melitz framework.

The assumption that exporters have all the bargaining power is perhaps a natural one to make given that the model is not explicit about the role of importers in facilitating trade. If these agents have a zero opportunity

cost and add no value to exports, why should they be remunerated? In the real world, however, intermediaries serve a central role in linking demand and supply by, among others, alleviating search frictions (see Antràs and Costinot, 2011) and providing quality assurance (see Bardhan, Mookherjee and Tsumagari, 2013 or Tang and Zhang, 2012). It is therefore natural that they capture a share of the gains from international trade. Although important, a treatment of international trade intermediation is beyond the scope of this book.

Contractual Frictions in the Melitz Model

As simple as the contract discussed above is, our discussion of international contract enforcement in Chapter 1 and the above account of Noble Group's soybean misadventures in China suggests that even those simple contracts are not fully enforceable in the real world. To fix ideas, I will next develop a simple model featuring one such source of contractual insecurity, namely a limited commitment problem on the part of the importer along the lines of the seminal work of Hart and Moore (1994) and Thomas and Worrall (1994).

The lack of commitment on the part of the importer is captured by assuming that at t_1 , and before he transfers the collected sale revenue to the exporter, this importer is presented with an opportunity to divert some cash flows away from the exporter. In an extreme case, this might reflect the possibility of the importer absconding with the exporter's goods and attempting to sell them on the side, perhaps at a discount. More generally, the assumption reflects the notion that the initial contract might not compel the parties to honor its terms, thereby tempting the importer to deviate from the contract by underreporting the amount of revenues actually collected, perhaps claiming that those lower revenues were due to the low quality of the goods the exporter shipped. To simplify matters, I will let the share of diverted revenues be a common constant $1 - \mu_{ij} \in [0, 1]$ for all pairs shipping goods from i to j , but I will later briefly discuss the case in which this parameter might vary with productivity.

The parameter μ_{ij} captures the extent to which the importer feels constrained in defaulting on its contractual commitments with exporters from i and thus it is natural to treat this parameter as a measure of the degree of contract enforcement in country j . The fact that the share μ_{ij} also depends on the exporting country i implies that the level of international contract enforcement is allowed to potentially be a function of the nationality of the

two agents in the transaction, reflecting perhaps the effects of legal similarity (e.g., common versus civil law countries), a common language, or proximity (cultural or geographical).

I realize that the above modeling of contractual institutions is exceedingly simplistic, with the great complexities and nuances of this type of institutions being reduced to a single parameter μ_{ij} capturing the ‘stealing’ possibilities of agents residing in the importing country. I will stick to this simple framework for most of this chapter, but let me briefly expand on different mechanisms that might jointly contribute to a country offering a low level of contractual security to firms exporting to it. First, in some institutional environments, agents might face more opportunities to deviate from the initial contract than in other environments. This might be partly due to social norms, but is also explained by the legal environment which might determine how complete formal contracts tend to be. Let us denote by $1 - \rho_{ij}$ the probability with which a ‘default’ opportunity arises for a j -importer transacting with an i -exporter (before we assumed $\rho_{ij} = 0$). When such an opportunity to default does not arise, the importer will necessarily honor the initial contract and deliver all sale revenue to the exporter. When a default opportunity arises, however, the importer will assess the legal ramifications of a contractual breach, and will optimally decide whether to default or not. The legal consequences of a default are in turn shaped by both the probability with which a court of law will rule against a misbehaving importer (denoted by λ_{ij}) and by the amount of damages that it will be required to pay in such an eventuality. It is convenient to model these damages as a multiple d_{ij} of the sale revenues the importer had diverted from the exporter.

Notice that if d_{ij} or λ_{ij} are high enough such that $d_{ij}\lambda_{ij} > 1$, the importer will never default on the exporter and thus the exporter will be able to achieve the same profit flow as in the case with no contractual frictions (see equation (3.4)). Conversely, when $d_{ij}\lambda_{ij} < 1$, if the exporter insisted on demanding the entire sale revenue, the importer would optimally choose to fully default on the exporter because by doing so, it could obtain an expected payoff equal to a multiple $(1 - \rho_{ij})(1 - \lambda_{ij}d_{ij}) > 0$ of revenue.

Below, I will focus on the more interesting scenario in which $d_{ij}\lambda_{ij} < 1$. In such a case, the exporter is left, *in expectation*, with a share

$$\mu_{ij} \equiv \rho_{ij} + (1 - \rho_{ij}) \lambda_{ij} d_{ij} \quad (3.6)$$

of sale revenue. This expression for μ_{ij} summarizes how the prevalence of default opportunities, the competence of courts in ruling against deviating

parties, and the size and enforceability of damages jointly shape the perceived contractual security associated with different countries. Equation (3.6) also illustrates how even in situations in which contracts include choice-of-law and forum-of-law clauses (see Chapter 1), thus potentially making ρ_{ij} and λ_{ij} insensitive to j , the importing country institutions may still matter by shaping the extent to which damages set by international courts of law or arbitrators are enforced.

Later in this chapter, I will return to the general formulation of μ_{ij} in (3.6), but for the time being I will focus on the reduced form interpretation of $1 - \mu_{ij}$ as capturing the share of sale revenues that importers from j are able to divert from exporters from i .

Implications of Contractual Insecurity

How does the lack of commitment affect contracting between the exporter and the importer? The key new constraint facing the exporter when designing the initial contract, is that any remuneration to the importer lower than $(1 - \mu_{ij}) p_{ij}(q_{ij}(\varphi)) q_{ij}(\varphi)$ would necessarily lead the importer to divert cash flows. As a result, the optimal contracting problem now incorporates a new incentive compatibility (IC) constraint which is necessarily tighter than the participation constraint in the previous optimal contracting program in (3.5).

Formally, and maintaining the assumption that the exporter makes a take-it-or-leave-it offer to the importer, we now have that the date-0 quantity shipped $q_{ij}(\varphi)$ and the date-1 payment $s_{ij}(\varphi)$ solve

$$\begin{aligned} \max_{q_{ij}(\varphi), s_{ij}(\varphi)} \quad & s_{ij}(\varphi) - \frac{\tau_{ij}}{\varphi} w_i q_{ij}(\varphi) - w_i f_{ij} \\ \text{s.t.} \quad & p_{ij}(q_{ij}(\varphi)) q_{ij}(\varphi) - s_{ij}(\varphi) \geq 0 \\ & p_{ij}(q_{ij}(\varphi)) q_{ij}(\varphi) - s_{ij}(\varphi) \geq (1 - \mu_{ij}) p_{ij}(q_{ij}(\varphi)) q_{ij}(\varphi), \end{aligned}$$

with $p_{ij}(q_{ij}(\varphi)) = (\beta E_j P_j^{\sigma-1})^{1/\sigma} q_{ij}(\varphi)^{-1/\sigma}$. It is straightforward to see that $s_{ij}(\varphi)$ will now be set to exactly satisfy the (tighter) incentive compatibility constraint, thus implying that the exporter will now only capture a share μ_{ij} of revenues, and will choose $q_{ij}(\varphi)$ such that

$$\pi_{ij}(\varphi) = \max_{q_{ij}(\varphi)} \left\{ \mu_{ij} (\beta E_j P_j^{\sigma-1})^{1/\sigma} q_{ij}(\varphi)^{(\sigma-1)/\sigma} - \frac{\tau_{ij}}{\varphi} w_i q_{ij}(\varphi) - w_i f_{ij} \right\}.$$

Solving this problem, the profit function for the exporter can be written as

$$\pi_{ij}(\varphi) = \mu_{ij}^\sigma (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_{ij}. \quad (3.7)$$

Comparing equations (3.1) and (3.7), it is clear that imperfect contracting reduces the profitability of selling in country j , and the more so the lower is μ_{ij} . The reason for this is twofold: first, the exporter now shares part of the profits obtained in country j with an importer there, and second, the exporter naturally responds to this rent dissipation by reducing the desired amount of goods to ship to country j .

The Margins of Trade, Gravity and Welfare

We now turn to a more formal study of the effects of contract enforcement on the intensive and extensive margins of trade and on aggregate bilateral trade flows across countries. In analogy to the benchmark model with perfect contracting, from equation (3.7) we now have that only firms from i with productivity $\varphi > \tilde{\varphi}_{ij}$, where

$$\tilde{\varphi}_{ij} \equiv \tau_{ij} w_i \left(\frac{w_i f_{ij}}{\mu_{ij}^\sigma B_j} \right)^{1/(\sigma-1)}, \quad (3.8)$$

will find it optimal to export to country j . Clearly, for fixed w_i and B_j , the lower is μ_{ij} , the lower will be the measure of firms exporting to country j , and thus the extensive margin of trade is negatively affected by weak contract enforcement. This is illustrated in Figure 3.1 with the shift in the export productivity threshold from $\tilde{\varphi}_{ij}$ to $\tilde{\varphi}_{ij}'$.

Next, aggregating across all firms from i , we find

$$X_{ij} = N_i \int_{\tilde{\varphi}_{ij}}^{\infty} \sigma \mu_{ij}^\sigma (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} dG_i(\varphi), \quad (3.9)$$

which, as in Chapter 2, we can decompose into an extensive margin N_{ij} , namely the measure of firms from i that export in j , and an average intensive margin \bar{x}_{ij} , corresponding to the average export volume across the active exporters:

$$X_{ij} = N_{ij} \cdot \bar{x}_{ij}. \quad (3.10)$$

Whenever firm-level productivity is drawn from a Pareto distribution, we can integrate equation (3.9) and use (3.8) to express \bar{x}_{ij} as

$$\bar{x}_{ij} = \frac{\kappa}{\kappa - \sigma + 1} \sigma w_i f_{ij}, \quad (3.11)$$




Figure 3.1: Selection into Exporting with Contractual Frictions

just as in Chapter 2. Equation (3.11) then indicates that, in the particular Pareto case, the average intensive margin turns out to be unaffected by the degree of contractual enforcement. It is important to emphasize, however, that this does not imply that the intensive margin of trade *at the firm level* is unaffected by the quality of contracting institutions μ_{ij} . In fact, if a firm were to sell in two markets that differed only in their level μ_{ij} , then the firm would necessarily sell more in the market with better contract enforcement (a higher μ_{ij}), a prediction consistent with the empirical results of Araujo, Mion and Ornelas (2012), who study the effect of the quality of contracting institutions on the cross-section of firm-level export volumes of firms based in Belgium. The insensitivity of \bar{x}_{ij} to μ_{ij} in equation (3.11) is explained by the fact that countries with better contract enforcement attract a disproportionately larger set of relatively small exporters.

We can next follow the steps suggested in Melitz and Redding (2013a) to express aggregate sectoral exports from i to j in (3.10) in a slightly more familiar way. In particular, note first that the measure of active exporters in j is given by $N_{ij} = N_i (1 - G_i(\tilde{\varphi}_{ij}))$. Thus plugging (3.11) and the value of the threshold in (3.8) into (3.10), and invoking the Pareto distribution, we

have

$$X_{ij} = N_i \left(\frac{\varphi_i}{\tau_{ij} w_i} \right)^\kappa \left(\frac{\mu_{ij}^\sigma B_j}{w_i f_{ij}} \right)^{\kappa/(\sigma-1)} \frac{\kappa}{\kappa - \sigma + 1} \sigma w_i f_{ij}. \quad (3.12)$$

Next, aggregating over all markets in which firms from i sell (including their domestic market), we can express the aggregate sale revenue obtained by firms in country i as

$$Y_i = \sum_j X_{ij} = N_i \left(\frac{\varphi_i}{w_i} \right)^\kappa \left(\frac{1}{w_i} \right)^{\kappa/(\sigma-1)} \frac{\kappa}{\kappa - \sigma + 1} \sigma w_i \Theta_i \quad (3.13)$$

where

$$\Theta_i \equiv \sum_j B_j^{\frac{\kappa}{\sigma-1}} \tau_{ij}^{-\kappa} f_{ij}^{-\frac{\kappa-(\sigma-1)}{(\sigma-1)}} \mu_{ij}^{\sigma\kappa/(\sigma-1)} \quad (3.14)$$

is a measure of country i 's market potential (see Redding and Venables, 2004). Plugging equation (3.13) back into (3.12) finally delivers

$$X_{ij} = \frac{Y_i}{\Theta_i} B_j^{\frac{\kappa}{\sigma-1}} \tau_{ij}^{-\kappa} f_{ij}^{-\frac{\kappa-(\sigma-1)}{(\sigma-1)}} \mu_{ij}^{\sigma\kappa/(\sigma-1)}. \quad (3.15)$$

This expression is analogous to equation (2.14) in Chapter 2 except for the last term involving the parameter μ_{ij} (remember that B_j is given by eq. (3.2)).

Equation (3.15) demonstrates that even after introducing contractual frictions, the model continues to deliver a modified sectoral version of the gravity equation for trade flows. This feature can again serve to motivate the widespread use of empirical log-linear specifications of trade flows with exporter and importer fixed effects and measures of bilateral trade frictions. The main new lesson one derives from (3.15) is that such log-linear specifications should include a bilateral measure of the level of contractual security in transactions between producers in country i and country j . We will shortly refer back to equation (3.15) when we review the empirical literature on the effects of institutional quality on trade flows. But before doing so, it is worth addressing one more theoretical matter.

It may appear that the effect of contractual insecurity in the current model is isomorphic to the effect of standard variable trade costs τ_{ij} in the benchmark model without contractual frictions. In particular, if one were to define a broad, contract-inclusive measure of trade frictions as

$$\tilde{\tau}_{ij} \equiv \frac{\tau_{ij}}{\mu_{ij}^{\sigma/(\sigma-1)}} > \tau_{ij},$$

it is straightforward to verify that all the equations above exactly correspond to those derived in the benchmark model in Chapter 2 with $\tilde{\tau}_{ij}$ replacing τ_{ij} . One is thus tempted to conclude that, apart from serving to motivate the inclusion of contracting institutions in standard empirical models of export participation, the explicit modeling of contractual frictions has little bearing on the workings of the benchmark model. Such a conclusion is however not warranted because, as mentioned above, contractual frictions not only reduce the profitability of exporting for producers in country i but also transfer exporting surplus to importers from country j . In other words, contractual insecurity not only reduces the overall gains from international trade, but also shapes how those gains are distributed across countries.

This distinction has important bearings for the relationship between contract enforcement and welfare. For instance, in the special case in which there is only one sector in the economy, Demidova and Rodríguez-Clare (2013) have shown that unilateral reductions in variable trade frictions by a small open economy are always welfare enhancing, while the same would not always be true for an increase in μ_{ij} in the framework developed in this chapter. The reason for this is that, although country j would become a more attractive location for foreign exporters if it instituted a higher μ_{ij} , the share of sale proceeds accruing to producers in j would also diminish in that event. For a high enough value of μ_{ij} , the balance of these two effects can be shown to be necessarily negative. This result is analogous to that in Demidova and Rodríguez-Clare (2009) and Felbermayr, Jung and Larch (2013), who show that in a one-sector Melitz (2003) model, each country's unilaterally optimal import tariff is positive. An implication of this result is that countries have a *unilateral* incentive to create some amount of contractual insecurity for producers attempting to sell in their markets. Naturally, however, and as in the case of tariff wars, the unilateral optimality of contractual insecurity is associated with a globally inefficiently low level of contract enforcement.

Preliminary Empirical Evidence

The gravity equation has been one of the most widely-used empirical models of international trade since being introduced by Tinbergen (1962). It is thus not surprising that it has been employed to study the effect of contracting institutions on bilateral international trade flows. The work of Anderson and Marcouiller (2002) is a pioneering study in this literature. Anderson and Marcouiller (2002) start by imposing a model of bilateral trade flows very

similar to that in equation (3.15), although not derived from a theoretical model, as we have done above.⁵ Following their approach, we next note that if one takes a country, say the United States, as a reference country, one can use equation (3.15), together with the definition of B_j in (3.2), to derive:

$$\frac{X_{ij}}{X_{iUS}} = \left(\frac{P_j^{\sigma-1} E_j}{P_{US}^{\sigma-1} E_{US}} \right)^{\frac{\kappa}{\sigma-1}} \left(\frac{\tau_{ij}}{\tau_{iUS}} \right)^{-\kappa} \left(\frac{f_{ij}}{f_{iUS}} \right)^{-\frac{\kappa-(\sigma-1)}{(\sigma-1)}} \left(\frac{\mu_{ij}}{\mu_{iUS}} \right)^{\sigma\kappa/(\sigma-1)}. \quad (3.16)$$

Equation (3.16) shows that the ratio of exports of country i to market j relative to the exports of this same country i to the United States is a function of the relative demand or absorption in the two importing countries, as well as different terms capturing the ratio of trade barriers associated with shipping goods from i to j relative to shipping them from i to the United States. The main advantage of this approach is that the ratio X_{ij}/X_{iUS} nets out the effect of the exporter country's term Y_i/Θ_i in (3.15) that is common for all destinations j .

Anderson and Marcouiller (2002) estimate a log-linear version of equation (3.16) in which relative bilateral traditional trade barriers (variable and fixed) are proxied by a common border ratio, a common language ratio, a distance ratio and a tariff ratio. The key contract enforcement ratio is proxied by a “composite security” index that corresponds to the average score obtained by each importing country in survey-based measures of transparency and contract enforcement relative to the average score obtained by the United States in those same measures. Note in particular, that Anderson and Marcouiller (2002) assume that the contractual security experienced by exporters in country j relative to that experienced by exporters to the U.S., is common for all exporters, regardless of their country of origin i . This seems to be a restrictive assumption given that one would imagine that differences in legal proximity could make this ratio vary with i , as allowed by (3.15), and so we will revisit this assumption shortly in this chapter. A last important hurdle in estimating equation (3.16) is finding suitable proxies for the first term involving the aggregate spending ratio E_j/E_{US} and the price index ratio P_j/P_{US} . The former ratio is proxied with relative measures of GDP, while the latter is approximated with weighted sums of the physical trade cost ra-

⁵Anderson and Marcouiller (2005) do study the theoretical links between contractual insecurity and trade flows, but their framework does not predict a gravity equation in trade flows.

tios, in analogy with the “remoteness” variable often present in gravity-style estimations. This is perhaps the least satisfactory element of their empirical design because mismeasurement of these importer-specific terms could lead to important biases in the estimates of the effects of the quality of the importer’s contracting institutions on trade flows. We will return to this issue below.

Table 3.1. Importer Contracting Institutions and Relative Exports

	(1)	(2)	(3)
Log GDP ratio	0.855** (0.042)	0.866** (0.038)	0.911** (0.040)
Relative composite security		0.285** (0.073)	0.279** (0.081)
Log common border ratio	0.794** (0.155)	0.747** (0.163)	0.665** (0.186)
Log common language ratio	0.327** (0.080)	0.336** (0.082)	0.358** (0.109)
Log distance ratio	-1.109** (0.058)	-1.095** (0.056)	-1.133** (0.056)
Log adjusted tariff ratio	-2.973 (1.992)	-4.814* (2.343)	-4.699* (2.327)
Number observations	2135	2135	2159
R^2	.69	.70	
Log likelihood			-3865

Table reproduced from Table 5 in Anderson and Marcouiller (2002).

Robust standard errors with clustering by importer in parentheses.

The regressions also include a log of GDP per capita ratio and remoteness variables for language, border and distance.

+, *, ** denote 10, 5, 1 % significance.

With these caveats in mind, the key results of Anderson and Marcouiller (2002), which use 1996 data for 48 importing countries, are reproduced in Table 3.1. Column (1) presents the results of a benchmark gravity equation without institutional variables. As expected, higher relative GDP levels and lower relative traditional trade barriers of any sort are all associated with higher relative export volumes into these countries. When introducing the relative “composite security” index in column (2), this variable has a large

and statistically significant effect on relative bilateral trade flows. In column (3), Anderson and Marcouiller confirm the robustness of their results to the use of a Tobit to deal with the large number of zeros in their sample. Although it is not obvious from the nonstandardized point estimates in Table 3.1, the estimates of Anderson and Marcouiller (2002) imply that the effect of weak contracting institutions on trade flows is of a similar order of magnitude as the effect of import tariffs.

As hinted above, two obvious limitations of the Anderson and Marcouiller (2002) study is the insensitivity of their contractual insecurity measure to characteristics of the exporting country and the econometric treatment of the demand terms B_j . These limitations are in fact related to each other. The standard way to address the second concern is to control for these importer-country absorption terms with importer-specific fixed effects. This approach is, however, not feasible when the key explanatory variable of interest varies only across importing countries and thus would be subsumed in the importer fixed effect. This problem clearly applies to Anderson and Marcouiller's (2002) proxy for contractual insecurity. Notice, conversely, that this is not an issue for standard measures of trade barriers, which are defined at the exporter-importer level. A potential way to address both of these limitations is thus to construct a measure of contractual security that is a function of both the exporter and importer country.

What factors might result in particularly high contractual enforcement in transactions between two specific countries i and j ? A natural candidate might be a simple measure of whether the exporter and importer country share a common legal origin or not, because that legal relatedness might facilitate the resolution of contractual disputes. A few papers in the literature have explored the role of a common legal origin in affecting bilateral trade flows across countries. Below I focus on the particular contribution of Helpman, Melitz and Rubinstein (2008) because their estimation equation is derived from the Melitz (2003) model in a similar manner as we derived equation (3.15) above, and because their estimation technique allows one to disentangle the effects of particular explanatory variables on both the intensive and extensive margins of trade.

From a theoretical perspective, the only new feature in the framework developed by Helpman, Melitz and Rubinstein (2008) is the introduction of an upper bound in the distribution from which firms draw their productivity

level. In specific Pareto case, we now have that

$$G_i(\varphi) = \frac{1 - \left(\underline{\varphi}_i/\varphi\right)^\kappa}{1 - \left(\underline{\varphi}_i/\bar{\varphi}_i\right)^\kappa}, \quad \text{for } \bar{\varphi}_i \geq \varphi \geq \underline{\varphi}_i > 0. \quad (3.17)$$

The immediate implication of this assumption for our contracting model is that if for some exporting country i , no exporting firm draws a productivity level higher than the threshold $\tilde{\varphi}_{ij}$ defined in (3.8), then bilateral exports from i to j will be zero. This is also easy to see when writing these aggregate bilateral exports as in (3.9), but with the upper limit on the probability distribution:

$$X_{ij} = N_i \int_{\tilde{\varphi}_{ij}}^{\bar{\varphi}_i} \sigma \mu_{ij}^\sigma (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} dG_i(\varphi).$$

Defining aggregate (sectoral) output in i as $Y_i = \sum_j X_{ij}$ and also

$$V_{ij}(\tilde{\varphi}_{ij}) \equiv \int_{\tilde{\varphi}_{ij}}^{\bar{\varphi}_i} \varphi^{\sigma-1} dG_i(\varphi),$$

bilateral exports from i to j can be expressed as

$$X_{ij} = \frac{Y_i}{\tilde{\Theta}_i} B_j \tau_{ij}^{1-\sigma} \mu_{ij}^\sigma V_{ij}(\tilde{\varphi}_{ij}) \quad (3.18)$$

where

$$\tilde{\Theta}_i = \sum_j \mu_{ij}^\sigma (\tau_{ij})^{1-\sigma} B_j V_{ij}(\tilde{\varphi}_{ij}).$$

Equation (3.18) is again a modified version of the gravity equation and it is easily verified that when $\bar{\varphi}_i \rightarrow \infty$, (3.18) coincides with (3.15) after plugging the value of $\tilde{\varphi}_{ij}$ in (3.8) into $V_{ij}(\tilde{\varphi}_{ij})$.

Equation (3.18) nicely illustrates the existence of an omitted-variable bias in standard gravity-style estimation methods. Even when one partials out the terms $Y_i/\tilde{\Theta}_i$ and B_j with exporter and importer fixed effects, respectively, standard techniques do not take into account the term $V_{ij}(\tilde{\varphi}_{ij})$ capturing the extensive margin of trade from i to j .⁶ This omission is likely associated with

⁶The relevance of these biases is clear from the fact that the *direct* elasticity of trade flows to the index of contractual security is lower in equation (3.18) than in equation (3.15).

an upward bias in the elasticity of trade flows to institutional quality because as indicated by the threshold equation (3.8), μ_{ij} has a negative effect on the threshold $\tilde{\varphi}_{ij}$, thus reducing $V_{ij}(\tilde{\varphi}_{ij})$.

Helpman et al. (2008) develop a two-step estimation procedure to deal with these biases. In a first stage, a Probit selection equation is derived from the model, and the estimates of this equation are used to structurally construct a control variable for the second stage, which is a log-linear model with exporter and importer fixed effects, and various measures of bilateral trade barriers.⁷ For the procedure to work, one needs an explanatory variable that enters the first stage (the extensive margin of trade), but not the second stage. Helpman et al. (2008) argue that the cost of creating a business in a particular country satisfies this condition (they also suggest a common religion variable that allows estimation on a larger sample of countries in their sensitivity analysis).

For our purposes, the most relevant feature of Helpman et al.'s (2008) results is that their first and second stages include a variable which is equal to 1 whenever the exporter and importer share a common legal origin as defined by LaPorta, Lopez-de Silanes, Shleifer and Vishny (1999). These authors classify the legal origin of a large cross-section of countries as being either German, Scandinavian, British, French, or Socialist. As argued above, it seems natural that, other things being equal, producers located in countries sharing a common legal origin will perceive a higher degree of contractual security when transacting with each other than with producers located in countries with different legal origins. Despite the obvious coarseness of this variable, I will interpret the effect of this variable as reflecting the effect of the contracting institutions term μ_{ij} in specification (3.18).

Table 3.2 reproduces some of the main results in the Helpman et al. (2008) study. The trade data is for 1986 and covers 158 countries. The first column in the table reports the result of the first stage, in which a Probit model is used to predict the probability of positive trade flows from i to j . Not surprisingly, a lower distance and sharing a common language are both positively correlated with the probability that two countries trade with each other.⁸ Interestingly, the same is true for the common legal origin variable,

⁷The first stage estimates are also used to include a more standard Heckman-correction term for selection in the second stage.

⁸The original regressions in Helpman et al. (2008) include six additional controls: whether both countries are islands, whether they are both landlocked, whether they have colonial ties, whether they are members of the same currency union, whether they belong

and the standardized coefficients indicate that the effect of legal institutions is almost half as large as that of a common language. Perhaps surprisingly, zero trade flows appear to be more prevalent for countries that share a land border, perhaps reflecting the incidence of wars. Regulation costs, in turn, appear to have a negative effect on the extensive margin of trade.

Table 3.2: Legal Origin and Bilateral Trade Flows

	(1) Probit	(2) Benchmark	(3) NLS
Distance	-0.213** (0.016)	-1.167** (0.040)	-0.813** (0.049)
Share a land border	-0.087 (0.072)	0.627** (0.165)	0.871** (0.170)
Share a common legal origin	0.049** (0.019)	0.535** (0.064)	0.431** (0.065)
Share a common language	0.101** (0.021)	0.147+ (0.075)	-0.030 (0.087)
Regulation costs (\$ amount)	-0.108** (0.036)	-0.146 (0.100)	
Regulation costs (days & procedures)	-0.061* (0.031)	-0.216+ (0.124)	
Firm heterogeneity correction term			0.840** (0.043)
Sample selection correction			0.240* (0.099)
Observations	12,198	6,602	6,602
R^2	0.573	0.693	

Table reproduced from Table II in Helpman et al. (2008).

Robust standard errors clustered by country pair (Bootstrapped for NLS).

Regressions also include exporter and importer fixed effects as well as six other controls (island, landlocked, colonial ties, currency union, FTA, religion).

Marginal effects at sample means and pseudo R^2 reported for Probit.

+, *, ** denote 10, 5, 1 % significance.

to the same FTA, and a measure of religious proximity. I do not report these coefficients to save space.

Column (2) of Table 3.2 presents the results of a benchmark gravity estimation of the intensive margin of trade that does not correct for the biases identified above. With the exception of the common border variable (which now impacts positively trade flows), the remaining coefficients have the same sign as in the Probit regressions. The effect of a common legal origin on trade flows is positive, large and highly statistically significant. One might worry, however, that the omitted-variable bias discussed above would lead us to overestimate the effects of contractual security on the intensive margin of trade. The results in column (3), which present the second stage in the Helpman et al. (2008) procedure, confirm that such biases exist but the coefficient on legal origins is reduced by only about 20% and remains large and highly significant. We can conclude from these results that, consistently with the simple model we have developed above, contractual insecurity has a significant negative effect on bilateral trade flows and that such effect operates through both an extensive margin as well as an intensive margin.

Responses to Contractual Insecurity

Our theoretical and empirical results so far illustrate that exporters will respond to the perceived contractual insecurity associated with servicing certain foreign markets by reducing their sales or by simply opting out from selling in those markets. As explained in Chapter 1, in practice firms can resort to alternative means to alleviate such contractual insecurity. We will next discuss three of these mechanisms: investing in contract enforcement, repeated interactions with importers, and demanding prepayment from importers.

Consider first the possibility of firms investing in enhancing the contractibility of their transactions. This might involve hiring legal counsel to design the initial contract in a way that makes it more likely to be enforced, or it might be associated with resorting to international arbitration, which would typically also provide the exporter with more contractual security. Without delving into the details of these different legal mechanisms, let us suppose that if a firm from i were to invest a fixed amount $w_i f_c$ of resources in improving contractibility, the share of revenues that an importer from j would be able to divert would be reduced from $1 - \mu_{ij}$ to $1 - \bar{\mu}_{ij}$ with $\bar{\mu}_{ij} > \mu_{ij}$. The assumption that legal expenses are independent of the volume of sales is a strong one, but the results below will continue to go through as long as there is a fixed cost component to these costs, which seems a plausible

assumption to make in this setting.

Following analogous derivations as those in the previous model, it is then straightforward to verify that firms from i will optimally choose to invest in contractibility whenever

$$\bar{\mu}_{ij}^{\sigma} (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_c > \mu_{ij}^{\sigma} (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1},$$

which can alternatively be expressed as

$$\varphi^{\sigma-1} > (\varphi_{ij^c})^{\sigma-1} \equiv \frac{w_i f_c}{(\bar{\mu}_{ij}^{\sigma} - \mu_{ij}^{\sigma}) (\tau_{ij} w_i)^{1-\sigma} B_j}.$$

In words, only the largest, most productive exporters will find it optimal to incur additional legal expenses to reduce their contractual insecurity. This might explain, for instance, why arbitration cases at the International Chamber of Commerce rarely involve disputes over amounts lower than under one million U.S. dollars (see footnote 13 in Chapter 1). The selection of exporters into enhanced contractibility is depicted in Figure 3.2. The figure also shows that the endogenously higher contractibility of large exporters will tend to lead to a more skewed distribution of exports than in the version of the model in which the parameter μ_{ij} is common for all firms within an industry.

In the previous setup in which the exporter and the importer transact only once, it is optimal for importers to divert revenue from the exporter if the contract is not perceived to be enforced. I next briefly explore how the incentives of importers to misbehave might be affected by repeated interactions with a given exporter, and how this affects the dynamics of exporting volumes. In doing so, I build on the work of Araujo, Mion and Ornelas (2012) and Antràs and Foley (2013). To emphasize the differences with the static model, let us assume that the exporter and importer perceive their business relationship to be infinitely repeated. Assume also that importers come in two types: some of them are very patient and discount the future at a very low rate, while the rest are myopic and care only about current payoffs. At the beginning of each period t , the exporter and importer sign an agreement that binds the exporter to ship an amount $q_{ijt}(\varphi)$ of output to the importer in j in exchange for a payment from the importer once the goods have been sold.

With probability $1 - \rho_{ij}$, however, the importer is presented with an opportunity to divert all sale revenue and pay nothing to the exporter at the end of the period. We thus adopt here the probabilistic version of contract




Figure 3.2: Selection into Exporting with Endogenous Contractibility

enforcement discussed earlier in the chapter and that led to expression (3.6) for μ_{ij} , but for simplicity I set $d_{ij}\lambda_{ij} = 0$, so that the exporter is left with a zero payoff in case of default. Provided that the discount rate of patient importers is sufficiently low, the folk theorem implies that a trigger-strategy equilibrium exists in which patient importers never choose to default, while myopic importers always do so when an opportunity arises.⁹

To generate nontrivial dynamics, we assume that whether an agent is patient or myopic is private information to that agent. The exporter forms beliefs on the type of the particular importer they are dealing with based on the bilateral transaction history with that importer and on a public signal reflecting the prevalence of patient importers in the population. I will denote this public signal by ξ_0 and I will treat it as an exogenous parameter, although in Araujo, Mion and Ornelas (2012) it is endogenized by specifying a process of matching between exporters and importers.

Notice that a new importer will initially be assigned a probability ξ_0 of being patient, which is associated with a probability $\xi_0 + (1 - \xi_0)\rho_{ij}$ of the

⁹Of course, this requires that the importer obtains some positive payoff when it chooses to honor the contract. Still, for a discount factor close enough to 1, this required payoff can be made arbitrarily close to 0. This limiting case is considered here for simplicity.

contract being enforced in this initial period. With a history of no defaults, the exporter's belief on the importer's type will improve over time, while an incidence of a default will immediately reveal the importer to be a myopic type. Denoting by ξ_t the particular posterior probability assigned to the importer being patient, repeated application of Bayes' rule delivers

$$\xi_t = \frac{\xi_0}{\xi_0 + (1 - \xi_0) (\rho_{ij})^t} \quad (3.19)$$

when there have been no defaults up to length t , and $\xi_t = 0$ otherwise. As a result, the perceived probability of the contract signed at t being enforced is given by $\xi_t + (1 - \xi_t) \rho_{ij}$, and naturally rises with a history of no defaults. Having determined this time-varying level of contractual insecurity, the rest of the equilibrium is analogous to that of the static model with the share of revenue accruing to exporters given by

$$\mu_{ij}(t) = \xi_t + (1 - \xi_t) \rho_{ij}.$$

Hence, all firm-level equilibrium expressions continue to hold with $\mu_{ij}(t)$ replacing μ_{ij} throughout.

This extension of the model delivers several empirical predictions for the effects of weak contracting institutions on firm-level exports. As in the static model developed above, the extensive margin of trade continues to be negatively impacted by low institutional quality (low ρ_{ij}). This is both because firms are less inclined to begin selling in weak institutions countries, but also because the probability of an export relationship being discontinued is higher the lower is the probability of contracts being enforced. The effects of low formal contract enforcement on the intensive margin of trade are richer. The perceived initial probability of default is given by $(1 - \xi_0) (1 - \rho_{ij})$ and thus export relationships in weak contracting environments (countries with low ρ_{ij}) will tend to begin at low volumes. Nevertheless, the negative effect of weak contracting on the intensive margin of trade is predicted to be attenuated over time, resulting in firm-level export volumes that should rise over time.

Araujo, Mion and Ornelas (2012) study the effects of importer-country characteristics on the cross-section of firm-level exports of Belgian firms over the period 1995-2008 and find broad support for these predictions. Other things being equal, export entry is higher and export exit is lower in countries with better contracting institutions. Initial firm-level export volumes

are increasing in contract enforcement, while firm-level export growth is on average positive. Interestingly, however, this positive growth in exports appears to be faster in countries with weak contracting institutions. The simple model developed above provides a simple rationale for this fact: in countries with high default rates contractual insecurity will be lower, but the exporter will be able to learn the type of the importer at a faster rate than in an environment in which a very low default rate prevents myopic importers from defaulting. Formally, differentiation of (3.19) indicates that for a low enough t , the growth in ξ_t over time t is necessarily decreasing in ρ_{ij} .¹⁰

Choice of Payment Method: Exporter Institutions Matter

So far, we have discussed the role of investments in contractibility and repeated interactions in reducing the extent of contractual insecurity faced by exporters. If the *only* contractual friction in international transactions was the risk of importer default, then a simple solution to this problem would be for the exporter to demand pre-payment from the importer before shipping the goods. Formally, a simple modification of the static contract we have considered so far would suffice to resolve the inefficiencies associated with contractual insecurity: instead of the payment s_{ij} occurring at t_1 , the exporter could insist that it was made at t_0 . With this simple modification, the exporter would not need to worry about the payment s_{ij} satisfying an incentive compatibility constraint for the importer, and could thus choose s_{ij} to satisfy exactly the importer's participation constraint. The problem would thus reduce to that in (3.5), which we have shown above delivers payoffs identical to those in the Melitz (2003) model without contractual frictions.

Although 'cash-in-advance' transactions are not infrequent in international trade (see Antràs and Foley, 2013, for evidence from a U.S.-based exporter), the available evidence suggests that the majority of international transactions are conducted on open account (or post-shipment payment) terms.¹¹ These type of transactions roughly correspond to the timing of payments we have assumed so far. A natural question is then: why are cash-

¹⁰Araujo, Mion and Ornelas (2012) endogenize the prior ξ_0 , and show that the differentially lower growth of exports in high contract enforcement countries actually holds true for all values of t in that case.

¹¹For instance, using the World Bank Enterprise Survey database, Hoefele, Schmidt-Eisenlohr and Yu (2013) find that the average share of sales on open account terms for the firms in the sample is in excess of 80%.

in-advance terms not used more often if they effectively eliminate the risk of importer default?

The key for answering this question is that not only exporters but also importers are exposed to the risk of counterparty misbehavior in international transactions. In particular, a standard concern for importers in cash-in-advance transactions is that, after being paid, the exporter might no longer have the incentive to ship goods in the most advantageous manner for importers, thus intentionally or unintentionally reducing the amount of sale revenues that the importer would obtain when selling the goods in their local market. I next briefly develop a simple model of exporter misbehavior along the lines of the model of limited commitment by importers developed above.¹² The model will serve to illustrate the role of exporter-country institutions in shaping the different margins of international trade.

Suppose that the exporter and importer sign the following simple cash-in-advance contract. At t_0 , the exporting firm F agrees to ship an amount of goods equal to q_{ij} in exchange for an amount s_{ij} to be paid upon signing the contract at t_0 . After receiving the goods, the importer sells them in her local market and she keeps the collected sale revenue. As argued above, without any type of frictions, the exporter could set an initial payment $s_{ij}(\varphi)$ equal to the sale revenue collected by the importer at time t_1 , i.e., $p_{ij}(q_{ij}(\varphi))q_{ij}(\varphi)$, thus attaining the frictionless profit flow in (3.4).

Imagine, however, that shortly after signing the contract at t_0 , the exporter is presented with an opportunity to deviate from the initial contract in a way that would reduce its costs of production but would also reduce the expected revenues collected by the importer at t_1 . Such a deviation might entail shirking in quality-enhancing investments or in the use of shipping methods that best ensure the quality of goods when they reach the importer's market. For the time being, consider the case in which exporter misbehavior takes the extreme form of the exporter incurring no variable production costs and the importer not receiving goods or receiving totally worthless goods. Below, I will consider much less extreme cases. Faced with this opportunity to misbehave, the exporter will consider the legal implications of such a deviation before cheating on the importer. Suppose that when deviating, the importer could sue the exporter in a latter's court of law and win the case with prob-

¹²Financial constraints faced by the importer might be another factor limiting the use of cash-in-advance contracts. Although not the focus of her study, Manova (2012) has found indeed that bilateral trade flows are depressed by low quality of financial institutions in importing countries.

ability λ_{ij}^{exp} , in which case the exporter would be asked to pay the importer an amount in damages equal to a multiple d_{ij}^{exp} of the payment stipulated in the initial contract, i.e., s_{ij} . When $\lambda_{ij}^{exp} d_{ij}^{exp} > 1$, the exporter would never be tempted to cheat on the importer, so we will focus below on the more interesting case in which $\lambda_{ij}^{exp} d_{ij}^{exp} < 1$. In the latter case, in order for the exporter not be tempted to misbehave, the payment stipulated in the initial contract needs to satisfy the following incentive compatibility (IC) constraint

$$s_{ij}(\varphi) - \frac{\tau_{ij}}{\varphi} w_i q_{ij}(\varphi) \geq (1 - \lambda_{ij}^{exp} d_{ij}^{exp}) s_{ij}(\varphi). \quad (3.20)$$

Note that it is in the exporter's own interest to ensure that the initial contract satisfies this IC constraint because otherwise the importer would anticipate misbehavior with probability one, and he or she would not be willing to pay *any* amount of money to the exporter in the initial period, thus leaving both agents with a zero payoff.

Using equation (3.4) and the fact that revenues are a multiple σ of operating profits, it is straightforward to verify that provided that $\lambda_{ij}^{exp} d_{ij}^{exp} \geq (\sigma - 1)/\sigma$, the constraint in (3.20) will be slack when evaluated at the unconstrained profit-maximizing output level $q_{ij}(\varphi)$, and thus the exporter will still be able to achieve the unconstrained profit flow in (3.4). Conversely, when courts punish deviating agents with a low enough probability or when damages are low enough or unenforceable, so that $\lambda_{ij}^{exp} d_{ij}^{exp} < (\sigma - 1)/\sigma$, the exporter will no longer be able to achieve the unconstrained profit flow in (3.4). Instead, the quantity of output being shipped will need to adjust to ensure that equation (3.20) holds, and the exporter will find it optimal to make that constraint exactly bind. Furthermore, one can show that the exporter will never find it optimal to demand an ex-ante payment lower than the total sale revenues collected by the importer at t_1 , so $s_{ij}(\varphi) = p_j(q_{ij}(\varphi)) q_{ij}(\varphi)$, and from equations (3.20) we can infer that

$$p_j(q_{ij}(\varphi)) = \frac{\tau_{ij} w_i}{\lambda_{ij}^{exp} d_{ij}^{exp} \varphi}. \quad (3.21)$$

Using (3.21) together with $q_{ij}(\varphi) = \beta E_j P_j^{\sigma-1} p_j^{-\sigma}$ and the definition of B_j in (3.2), we can then express the profits of the exporter as

$$\pi_{ij}(\varphi) = \mu_{ij}^{exp} (\tau_{ij} w_i)^{1-\sigma} B_j \varphi^{\sigma-1} - w_i f_{ij}, \quad (3.22)$$

where

$$\mu_{ij}^{exp} = \begin{cases} 1 & \text{if } \lambda_{ij}^{exp} d_{ij}^{exp} \geq \frac{(\sigma-1)}{\sigma} \\ \sigma (1 - \lambda_{ij}^{exp} d_{ij}^{exp}) \left(\frac{\sigma \lambda_{ij}^{exp} d_{ij}^{exp}}{\sigma-1} \right)^{\sigma-1} < 1 & \text{if } \lambda_{ij}^{exp} d_{ij}^{exp} < \frac{(\sigma-1)}{\sigma}. \end{cases} \quad (3.23)$$

Note that μ_{ij}^{exp} is (weakly) increasing in $\lambda_{ij}^{exp} d_{ij}^{exp}$, and that $\mu_{ij}^{exp} = 1$ only when $\lambda_{ij}^{exp} d_{ij}^{exp} \geq (\sigma - 1) / \sigma$. Hence, in the range of parameter values in which the exporter is tempted to misbehave, the profits the exporter will end up obtaining will necessarily be lower than in the unconstrained problem.

Equation (3.22) illustrates that limited commitment problems on the part of the exporter end up affecting the profitability of exporting in a similar manner as limited commitment problems on the importer side. In fact, equation (3.21) is identical to (3.7) except for the term μ_{ij}^{exp} in (3.21) instead of μ_{ij}^{σ} in (3.7). The superscript *exp* in the contracting term in (3.21) serves to emphasize that the quality of the *exporter* country's contracting institutions is now key in shaping the profitability of exporting, the intensive and extensive margins of trade, and bilateral trade flows across countries. Of course, one could argue, as we did in the importer limited commitment case, that agents could resort to choice-of-law or choice-of-forum contractual clauses to partly isolate the security of a transaction from weak contracting institutions in the exporting country. Nevertheless, and as explained in Chapter 1, even when disputes are adjudicated by foreign courts, the enforceability of damages is ultimately an issue related to the local legal environment in the exporting country, and particularly whether that country has signed the New York convention.

I will next provide an overview of the empirical work linking bilateral trade flows to the quality of the exporter country's contracting institutions, but before doing so I should briefly address two further theoretical points. First, and although it is obvious to see that all equilibrium conditions with exporter limited commitment will be identical to those in (3.7) through (3.15) with μ_{ij}^{exp} replacing μ_{ij}^{σ} , there is one important, subtle difference in the general equilibrium implications of the two models. Because in this second model, importers always end up with a net payoff of zero, the effects of a low μ_{ij}^{exp} are not isomorphic to an import tariff in the importing country, but instead are analogous to those of an iceberg trade cost. An implication of this difference is that, at least in the one-sector version of the model, improvements in the quality of contracting institutions in a small exporting country will always

be beneficial for the importing country (see Demidova and Rodríguez-Clare, 2013). A second point worth making is that our model of exporter misbehavior can easily be extended to the case in which the exporter's temptation to deviate from the contract entails reducing marginal costs by a certain fraction ν_{ij}^{exp} where we now allow $\nu_{ij}^{exp} < 1$. In such a case, the exporter's incentive compatibility constraint becomes

$$\lambda_{ij}^{exp} d_{ij}^{exp} s_{ij}(\varphi) \geq \nu_{ij}^{exp} \frac{\tau_{ij}}{\varphi} w_i q_{ij}(\varphi),$$

and the same expressions (3.21) through (3.23) apply, but with $\lambda_{ij}^{exp} d_{ij}^{exp} / \nu_{ij}^{exp}$ replacing $\lambda_{ij}^{exp} d_{ij}^{exp}$ throughout. Clearly, exporter profits will be higher in that case, but as long as $\lambda_{ij}^{exp} d_{ij}^{exp} / \nu_{ij}^{exp} < (\sigma - 1) / \sigma$, contractual frictions continue to reduce the profitability of exporting.

Exporter-Country Institutions: Empirical Evidence

Earlier in the chapter we discussed the empirical work of Anderson and Marcouiller (2002) establishing a link between bilateral trade flows and the quality of the importer's contracting institutions. It is clear that the empirical strategy in that paper is not applicable to the study of the effects of exporter-country institutions since the relative exports specification in (3.16) effectively partials out exporter-specific variables. Berkowitz, Moenius and Pistor (2006) propose instead a more traditional log-linear gravity specification which can be motivated by a simple variant of equation (3.15) with μ_{ij}^{exp} replacing μ_{ij}^{σ} :

$$X_{ij} = \frac{Y_i}{\Theta_i} B_j^{\frac{\kappa}{\sigma-1}} \tau_{ij}^{-\kappa} f_{ij}^{-\frac{\kappa(\sigma-1)}{(\sigma-1)}} (\mu_{ij}^{exp})^{\kappa/(\sigma-1)}. \quad (3.24)$$

In order to control for the unobserved multilateral resistance term Θ_i and the price index implicit in B_j , Berkowitz, Moenius and Pistor (2006) introduce exporter and importer fixed effects. The author's measure of the quality of contracting institutions is an average of a country's index of rule of law, expropriation risk, corruption in government, and bureaucratic quality as computed by International Country Risk Guide. This variable is computed for the exporter and the importer in each pair of trading partners and both variables are introduced in the regression, thus allowing for both importer and exporter country institutions to affect bilateral trade flows. In order for the country fixed effects not to absorb these institutional variables, Berkowitz,

Moenius and Pistor (2006) use data from 1982 to 1992 and exploit time-series variation in both bilateral trade flows and the perceived quality of contracting institutions. Their specifications also include time fixed effects and controls for GDP, GDP per capita, and various measures of proximity between the exporter and the importer, including a measure of the remoteness related to whether a pair of countries are close to each other but distant from the rest of the world.

Table 3.3: Exporter and Importer Contracting Institutions and Bilateral Exports

Type of Goods Included	(1) Overall	(2) Overall	(3) Complex	(4) Simple
GDP importer	0.81** (0.02)	-0.15 (0.29)	0.08 (0.30)	-1.06* (0.42)
GDP exporter	0.76** (0.02)	-0.19 (0.29)	0.32 (0.30)	-1.38** (0.42)
Distance	-1.16** (0.04)	-1.03** (0.04)	-0.98** (0.04)	-1.26** (0.06)
Adjacent	0.35* (0.14)	0.40** (0.15)	0.44** (0.17)	0.27 (0.18)
Links	0.42** (0.10)	0.45** (0.10)	0.54** (0.11)	0.18 (0.15)
Language similarities	0.09 (0.18)	1.00** (0.17)	1.28** (0.19)	0.11 (0.28)
Remoteness	0.58** (0.10)	1.79* (0.78)	0.74 (0.77)	6.69** (1.22)
Quality of importer legal institutions	0.61** (0.11)	0.05 (0.10)	-0.44** (0.10)	0.66** (0.15)
Quality of exporter legal institutions	0.91** (0.13)	0.36** (0.11)	0.93** (0.11)	-0.53** (0.15)
Country dummies	No	Yes	Yes	Yes
Time dummies	No	Yes	Yes	Yes
Number of clusters (country pairs)	2792	2792	2755	2550
R^2	0.70	0.77	0.79	0.38
Number observations	23,564	23,564	22,669	18,948

Table reproduced from Table 2 and 3 in Berkowitz et al. (2006). Robust standard errors (within-group clustering) in parentheses. Regressions also include exporter and importer GDP per capita and a constant. +, *, ** denote 10, 5, 1 % significance.

The first two columns of Table 3.3 reproduce the results obtained by Berkowitz et al. (2006) when running their specification with and without the country and year fixed effects, respectively. As is clear from column (1), when ignoring these fixed effects, all variables affect bilateral trade flows in the expected way and the institutional quality variables related to both the exporter and the importer are highly statistically significant, with exporter institutions appearing to matter more than importer institutions. When introducing the exporter, importer and year fixed effects in column (2), a first noteworthy fact is that the effect of GDP on bilateral trade flows vanishes. This is not entirely surprising since the fixed effects were supposed to control for terms in the gravity equation involving GDP. More relevant for the current discussion is the fact that the variable capturing the quality of contracting institutions in the exporting country remains both positive and highly statistically significant, while the importer country institutional quality variable remains positive but loses its statistical significance.

I will shortly discuss a set of additional results in the Berkowitz et al. (2006) paper that anticipated the voluminous literature on the institutional determinants of comparative advantage. Before doing so, however, I will draw on the work of Waugh (2010) and briefly outline an alternative way to identify the potential role of contracting institutions in shaping bilateral trade flows across countries. Let us return to the modified gravity equation in (3.24), and note that we can use it to express the ratio of exports from i to j to the domestic absorption of the importing country j as:

$$\frac{X_{ij}}{X_{jj}} = \frac{Y_i/\Theta_i}{Y_j/\Theta_j} \frac{\tau_{ij}^{-\kappa}}{\tau_{jj}^{-\kappa}} \frac{f_{ij}^{-\frac{\kappa-(\sigma-1)}{(\sigma-1)}}}{f_{jj}^{-\frac{\kappa-(\sigma-1)}{(\sigma-1)}}} \frac{(\mu_{ij}^{exp})^{\kappa/(\sigma-1)}}{(\mu_{jj}^{exp})^{\kappa/(\sigma-1)}}. \quad (3.25)$$

Next suppose that within-country or domestic barriers to trade – technological and contractual – do not vary significantly across countries, so we can set $\tau_{jj} = \tau_d$, $f_{jj} = f_d$, and $\mu_{jj}^{exp} = \mu_d$ for all j . This is a strong assumption, so I will return to it below. Assume also that transportation barriers are symmetric across countries, so $\tau_{ij} = \tau_{ji}$ and $f_{ij} = f_{ji}$ for any two countries i and j . Conversely, and as long as $i \neq j$, let μ_{ij}^{exp} be *only* a function of the quality of the exporter-country institutions, in the spirit of our discussion above. Let us thus simply denote $\mu_{ij}^{exp} = \mu_i$. Taking logs of (3.25) then delivers

$$\ln \left(\frac{X_{ij}}{X_{jj}} \right) = \alpha + \Psi_i - \Psi_j - \kappa \ln \tau_{ij} - \frac{\kappa - (\sigma - 1)}{\sigma - 1} \ln f_{ij} + \frac{\kappa}{\sigma - 1} \ln \mu_i, \quad (3.26)$$

where $\Psi_i = \ln(Y_i/\Theta_i)$. A key feature of equation (3.26) is that when regressing the left-hand-side on exporter and importer fixed effects, and empirical proxies for the bilateral trade costs between i and j (distance, language, and so on), the only reason for a country's fixed effect as an exporter to be different from that as an importer is for μ_i to be less than 1, i.e., for contract enforcement to be imperfect.




Figure 3.3: Inferred Exporter-Specific Ease of Exporting and the Rule of Law

Waugh (2010) runs the specification in (3.26) with data on bilateral trade flows and domestic absorption for 77 countries in 1996. His findings indicate the existence of very significant asymmetries in a country's fixed effect as an exporter and as an importer, and he shows that these asymmetries are correlated with income per capita. He interprets his results as suggesting that poor countries face much larger trade barriers when exporting than rich countries do. Waugh (2010) associates this asymmetry to an exporter-specific term in iceberg trade barriers (so $\tau_{ij} \neq \tau_{ji}$), while above I have ascribed these asymmetries to differences in export-country contract enforcement across countries. Admittedly, this is quite arbitrary but as Figure 3.3 indicates there exists a very significant positive correlation between the implied measure of

$\frac{\kappa}{\sigma-1} \ln \mu_i$ one backs out from the data and the ‘Rule of Law’ measure from the Governance Matters III Database, which is a standard proxy for contract enforcement (see Nunn, 2007, for details). Furthermore, the implied effect of contract enforcement on trade flows is remarkably large. For example, if the quality of contracting institutions in Guatemala were to increase to the level of those in Ecuador - which corresponds to an increase of 0.65 standard deviations in the rule of law measure – Guatemalan exports relative to the domestic absorption of an average importing country would increase by 118%.

In order to extract the exporter-specific trade impediment from data on trade flows and domestic absorption we have made strong assumptions. For instance, it seems reasonable that countries with poor institutions will also feature particularly weak contract enforcement in domestic transactions. It is apparent from (3.25), however, that this would tend to restore the symmetry between the fixed effect of a country as an exporter and as an importer, and thus this would presumably work to attenuate the strong positive correlation observed in Figure 3.3, rather than provide an alternative explanation for it.

Hopefully, the reader will view the evidence reviewed so far as suggestive of the relevance of contracting institutions as a determinant of bilateral trade flows across countries. It would be quite a stretch, however, to claim that the results in Table 3.3 or Figure 3.3 come anywhere close to identifying a causal effect of contracting institutions on trade flows. A particularly important concern is that we have ascribed to contractual institutions an effect that might in reality be caused by other country characteristics that happen to be correlated with the quality of this type of institutions.¹³

Berkowitz et al. (2006) acknowledge the potential existence of omitted variable biases in their estimates and suggest an ingenious identification strategy based on the notion that contracting institutions are likely to have a differential effect on different types of goods. More specifically, it seems natural to suppose that the type of contractual difficulties highlighted in this chapter are more likely to apply to complex goods than to simple goods. In fact, it is rather simple to extend the above model of exporter misbehavior to formalize this insight. For that purpose, assume that whether the exporter will be presented with an opportunity to misbehave or not occurs with a probability ρ which is a function of the type of good being traded. In par-

¹³For example, the positive relationship in Figure 3.3 remains positive but loses its statistical significance when controlling for income per capita.

ticular, assume that ρ is higher for complex, less contractible goods than for simple, homogeneous goods. Provided that both producers know whether such misbehavior is possible or not before they sign the initial contract (but not before the fixed exporting cost is incurred), it is straightforward to show that the equilibrium of the model will be identical to that above, but with

$$\tilde{\mu}_{ij}^{exp} \equiv \rho + (1 - \rho) \mu_{ij}^{exp}$$

replacing μ_{ij}^{exp} throughout. It is then clear that the effect of better contract enforcement on profitability, firm-level exports, bilateral exports, and so on is lower, the higher is ρ , i.e., the more complex goods are.

In order to test this prediction, Berkowitz et al. (2006) employ the Rauch (1999) classification of goods into differentiated and homogeneous and run their specification on each set of goods separately. Their results are reproduced in columns (3) and (4) of Table 3.3. A striking feature of their findings is that high levels of contract enforcement in the exporting country are shown to increase exports of complex goods but at the same time they *reduce* exports of simple goods. Conversely, and somewhat puzzlingly, good contracting institutions in the importing country enhance imports of simple goods, but reduce those of complex goods!

In order to rationalize their findings, Berkowitz et al. (2006) argue that the quality of contracting institutions will not only affect the security with which international transactions are conducted, but will also shape the efficiency with which the traded goods are produced, thus becoming a source of comparative advantage. Viewed from that perspective, their results in columns (3) and (4) are less surprising. They simply might reflect that countries with strong contracting institutions gain comparative advantage in complex (contract-intensive) goods, and as a result they tend to feature disproportionately high levels of exports of these complex goods and disproportionately low levels of imports of simple goods.

Domestic Institutions and Comparative Advantage

The idea that the quality of domestic institutions may constitute a source of comparative advantage has featured prominently in the trade literature in recent years. The vast literature on the topic is reviewed in Nunn and Trefler (2013a), so I will only sketch a few key contributions here. The earlier papers in that literature were closest in spirit to the work of Berkowitz et al. (2006).

Nunn (2007), Levchenko (2007), and Costinot (2009) all explored how domestic contracting institutions shape productivity differentially across sectors depending on characteristics of those sectors. They each proposed a measure of contract intensity different from the dichotomous one used by Berkowitz et al. (2006), and each showed that the effect of contracting institutions on trade flows was disproportionately higher in the industries identified to be relatively contract-intensive.

I will next briefly overview Nunn's paper because it has been the most influential one in this literature. In Chapter 5, I will perform empirical tests closely related to those in Levchenko (2007), so I will provide more details on his work at that point. Nunn (2007) proposes as a proxy for contract intensity a measure of the proportion of an industry's intermediate inputs that are relationship-specific. To construct that proportion, he builds on the classification of goods developed by Rauch (1999), which distinguishes between goods sold on organized exchanges, those with reference prices in trade publications, and all residual goods, which are assumed to be differentiated or customized. More specifically, Nunn uses U.S. Input-Output Use tables to construct an industry's use of intermediate inputs provided by other industries, and then infers the extent to which these inputs are customized from Rauch's classification of goods.

Table 3.4. The Ten Least and Ten Most Contract Intensive Industries

10 Least contract intensive: lowest z_i^{rs1}		10 Most contract intensive: highest z_i^{rs1}	
0.024	Poultry processing	0.810	Photogr. & photoc. equip. manuf.
0.024	Flour milling	0.819	Air & gas compressor manuf.
0.036	Petroleum refineries	0.822	Analytic laboratory instr. manuf.
0.036	Wet corn milling	0.824	Other engine equipment manuf.
0.053	Aluminum sheet, plate & foil manuf.	0.826	Oth. electronic component manuf.
0.058	Primary aluminum production	0.831	Packaging machinery manuf.
0.087	Nitrogenous fertilizer manuf.	0.840	Book publishers
0.099	Rice milling	0.851	Breweries
0.111	Prim. nonferrous metal	0.854	Musical instrument manuf.
0.132	Tobacco stemming & redrying	0.872	Aircraft engine & parts manuf.

Note: Table reproduced from Table II in Nunn (2007)

The resulting least and most contract intensive industries according to Nunn's (2007) definition are reproduced in Table 3.4. The ordering of indus-

tries appears sensible. For instance, the two least contract-intensive industries are ‘poultry processing’ and ‘flour milling’, which indeed use highly homogeneous inputs (chickens and wheat, respectively), while the most contract-intensive industry is aircraft manufacturing, which requires the use of highly customized inputs.

With this industry measure of contract intensity at hand, Nunn (2007) then uses international trade data for 146 countries and 222 industries in 1997 to explore whether countries with better contract enforcement appear to feature disproportionately large levels of exports in contract intensive sectors. As a proxy for the level of contract enforcement in a particular country, Nunn (2007) uses the ‘Rule of Law’ variable from the Governance Matters III Database, which consists of a weighted average of 17 measures of judicial quality and contract enforcement. Nunn’s specifications are of the form

$$\ln(X_{si}) = \alpha_s + \alpha_i + \beta_1 z_s \mu_i + \beta_2 h_s H_i + \beta_3 k_s K_i + \gamma c_s C_i + \varepsilon_{si},$$

where X_{si} denotes total exports in industry s from country i to all other countries in the world, z_s is contract intensity in industry s , μ_i is a measure of the quality of contract enforcement in (the exporting) country i , H_i and K_i denote country i ’s endowments of skilled labor and capital, and h_s and k_s are the skill and capital intensities of production in industry s . The term $c_s C_i$ represents a vector of control interactions of industry and country characteristics, while α_s and α_i denote industry fixed effects and country fixed effects, respectively.

Table 3.5 reproduces the benchmark results in Nunn (2007). The first two columns demonstrate that the interaction of contract intensity and judicial quality has a positive and statistically significant effect on exports, which is suggestive of the importance of contracting variables for the structure of international trade flows. Column (2) features a lower number of observations than column (1) because the sample is restricted to those countries and industries for which suitable proxies for factor abundance and intensities are available. The addition of these Heckscher-Ohlin interactions in column (3) has a negligible impact on the estimate of β_1 , while the standardized beta coefficients in that column indicate that the effect that judicial quality has on the pattern of trade is greater than the combined effects of both capital and skilled labor. The inclusion of additional controls in column (4) has little impact on these conclusions. Nunn (2007) presents several robustness tests and also attempts to deal with endogeneity concerns by using legal origin as an instrument for judicial quality and by using propensity score techniques.

Table 3.5: The Determinants of Comparative Advantage

	(1)	(2)	(3)	(4)
Judicial quality interaction	0.289** (0.013)	0.318** (0.020)	0.326** (0.023)	0.296** (0.024)
Skill interaction			0.085** (0.017)	0.063** (0.017)
Capital interaction			0.105** (0.031)	0.074+ (0.041)
Log income \times value added				-0.137* (0.067)
Log income \times intra-industry trade				0.546** (0.056)
Log income \times TFP growth				-0.010 (0.049)
Log income \times capital				0.021 (0.018)
Log income \times input variety				0.522** (0.103)
R^2	0.72	0.76	0.76	0.76
Number of observations	22,598	10,976	10,976	10,816

Table reproduced from Table IV in Nunn (2007). Regressions also include country and industry fixed effects. Standardized beta coefficients reported. Standard errors in parentheses. +, *, ** denote 10, 5, 1 % significance.

Building on the insights of this empirical literature on the effects of contracting institutions, other researchers have explored the role of other types of institutions in shaping comparative advantage across sectors. Manova (2008, 2012), for instance, explores the role of the quality of financial institutions in shaping the extensive and intensive margin of trade. Her empirical strategy builds on the seminal work of Rajan and Zingales (1998), who categorized sectors into more or less financially dependent depending on their external finance requirements. Relatedly, Cuñat and Melitz (2012) study how differences in the flexibility of labor market institutions across countries affect comparative advantage by building an industry-level measure of the importance of within-sectoral reallocations of labor as a response to shocks. In a very nicely executed study, Chor (2010) attempts to disentangle the partial

effect of each of these institutional determinants of comparative advantage in a unified empirical model.

Despite the recent focus in the literature on the role of *domestic* contracting institutions in shaping trade flows, it is not a warranted conclusion from these studies that *international* contract enforcement is irrelevant for explaining trade flows across countries. First, the aforementioned findings of Helpman et al. (2008) regarding the effect of having a common legal origin on aggregate bilateral trade flows are hard to rationalize in models in which international contract enforcement is perfect. Second, apart from their results discussed above, Berkowitz et al. (2006) also found that the effects of exporter and importer legal quality appear to be significantly affected by whether countries have ratified the New York convention or not. For instance, their estimates indicate that for the case of complex goods, the quality of exporter institutions matters disproportionately more when the export partner has not yet signed the New York convention and thus international enforcement of damages is more doubtful. For a third illustration of the importance of imperfect international contract enforcement, I next briefly return to the choice of payment-method decision faced by exporters and importers when negotiating their initial contracts.

Back to Trade Finance

So far we have illustrated how the quality of importer country institutions shapes the profitability and structure of exports whenever contracts are associated with post-shipment payment (or simply, open account) terms, while the quality of exporter country institutions plays a similar role in cash-in-advance transactions. Obviously, this constitutes a simplistic description of the effect of institutions on exporting. It seems natural, for instance, that exporter country institutions will matter even in open account terms to the extent that the consequences of exporter misbehavior might manifest themselves long after the goods have been received by the importer, or even after these goods have been sold to local consumers. Similarly, importer country institutions might affect the profitability of cash-in-advance transactions to the extent that they shape the financing costs faced by exporters. Intuitively, in countries where defaults are not sufficiently punished, not only exporters but also banks will shy away from extending credit to importers.

An active literature in international trade has explored the determinants of the choice of payment mode in international transactions, with a special

emphasis on the role of weak contracting institutions. This literature includes among others, the work of Amiti and Weinstein (2011), Antràs and Foley (2013), Ahn (2011), Hoefele, Schmidt-Eisenlohr and Yu (2013), Olsen (2013), and Schmidt-Eisenlohr (2013). Antràs and Foley (2013), in particular, focus on the role of importer country institutions, while allowing these to affect the profitability of both open account transactions (via default risk) as well as cash-in-advance transactions (via financing costs). Their key theoretical finding is that in the plausible case in which local banks in the importing country are better able than exporters to pursue financial claims against importers, one would expect exports to locations characterized by weak contractual enforcement to be more likely to occur on cash-in-advance as opposed to open account terms.¹⁴

One of the main challenges in studying the financing arrangements used to support international trade is that detailed data on how different types of transactions are financed are not readily available. Antràs and Foley (2013) overcome this dearth of data by analyzing detailed transaction-level data from a single U.S.-based firm that exports frozen and refrigerated food products, primarily poultry. The data cover roughly \$7 billion in sales to more than 140 countries over the 1996-2009 period and contain comprehensive information on the financing terms used in each transaction. A key advantage of the dataset is that by focusing on the sales of a single exporter based in the U.S., any institution-driven variation in the choice of payment mode must be ascribed to importer-country institutions or, following our broader interpretation of the parameter μ_{ij} in (3.6), to legal proximity between the U.S. and the importing country.

Antràs and Foley (2013) find robust evidence that variation in importer country contract enforcement has a strong effect on the method of payment offered to importers. Figure 3.4 reproduces the results in Figure 3 of their paper. For each of the proxies of contractual enforcement in the figure, the share of transactions occurring on cash-in-advance share is strikingly lower in strong contract enforcement countries than in weak contract enforcement countries. For instance, in common law countries, 4.0% of sales occur on cash-in-advance terms and 79.8% of sales occur on open account terms, while in civil law countries these shares are 63.8% and 20.4%. Similarly stark differ-

¹⁴Antràs and Foley (2013) also consider the possibility of exporters and importers resorting to letters of credit, but these financial instruments mediate a small share of world trade in modern times (see also Olsen, 2013, for more on letters of credit).




Figure 3.4: Financing Terms and the Enforcement of Contracts

ences appear when the sample is split using measures of contract viability from International Country Risk Guide (ICRG), payment delay (also from ICRG), and the enforceability of contracts (from Knack and Keefer, 1995). Antràs and Foley (2013) show that these patterns persist after controlling for several country-level variables as well as product fixed effects. Consistently with the results of the dynamic model of repeated interactions developed earlier in this chapter, they also find that first-time buyers are disproportionately more likely to be demanded to prepay for their purchases, but that as the exporter establishes a relationship with an importer, the share of cash-in-advance transactions falls smoothly over time.

In a recent paper, Hoefele, Schmidt-Eisenlohr and Yu (2013) have employed information from the World Bank Enterprise Survey to study the effects of variation in *exporting* country institutions on the choice of trade finance by firms. Consistently with the models developed in this chapter, Hoefele, Schmidt-Eisenlohr and Yu (2013) find that the use of cash-in-advance terms is more prevalent in exporting countries with strong contracting institutions, in which exporter misbehavior is less of a concern.

The Road Ahead

This chapter has explored both theoretically as well as empirically the significance of weak contract enforcement for the export decisions of firms and, more broadly, for the structure of international trade flows. The focus, however, has been on how contractual frictions affect the international *exchange* of goods. As explained in Chapter 1, the rapid growth in intermediate input trade has been one of the most prominent developments in the world economy in recent years. At the same time, the contractual relationships that support the phenomenon of offshoring are much more intricate than those that support the mere shipment of goods across countries. Thus, weak contract enforcement has the potential to affect the global organization of production in more profound ways than we have studied so far. In the next chapter, we will begin to explore these more complex contractual aspects of global sourcing.

Chapter 4

Contracts and Global Sourcing

In his highly entertaining book *Poorly Made in China*, Paul Midler (2009) describes his experiences as an offshoring consultant in China, where his command of Chinese made him a particularly valuable asset for American companies eager to outsource in that country. The bulk of Midler's book is centered around his work in assisting the chief operating officer of Johnson Carter, a U.S. personal care import company, in his negotiations with King Chemical, a Chinese manufacturer located outside Guangzhou (the names in the book are fictitious but the stories are genuine). Johnson Carter is in the business of supplying large chain stores in the U.S. with house-brand shampoo and soaps.

Midler's book vividly illustrates just how irrelevant formal outsourcing contracts can prove to be in China, reflecting the old Chinese adage that 'signing a contract is just the first step in the real negotiations'. Chinese manufacturers – and King Chemical, in particular – appear to attract clients by offering very low prices in order to secure contracts. After locking in U.S. buyers – in this case, Johnson Carter –, they work to devise numerous ways to increase their profit margins. These maneuvers often involve reductions of material costs, last-minute price increases, or even the use of the client's designs to sell the same products to alternative buyers at higher prices. In one particular instance, and after a contract order had been signed, King Chemical sought to reduce the money they spent on plastic by secretly switching to thinner bottles which were much more likely to collapse when squeezed. In the case of another order, Midler discovered that King Chemical was filling the bottles with significantly less soap than the 850 milliliters they had agreed on. Arbitrary price increases were particularly frustrating to Johnson

Carter, as King Chemical's price hikes were typically poorly justified and were invariably demanded right after the U.S. company had secured a large order from a U.S. retailer, at which point the U.S. company had no time to turn to an alternative Chinese manufacturer. This last-minute price spikes were hardly specific to King Chemical. As Midler writes, " 'Price go up!' was the resounding chorus heard across the manufacturing sector" in China.

The goal of this chapter is to initiate the analysis of the implications of weak contract enforcement for the international organization of production. With that goal in mind, I will start by going back to the benchmark two-country model of global sourcing developed in Chapter 2, and I will highlight the contracting assumptions underlying the results of that model. I will then introduce a model of contractual frictions that will shed light on the effect of contracting institutions on the intensive and extensive margins of intermediate input trade. Next, I will develop a series of extensions of the model that will open the door for model-based empirical tests of the effects of contracting considerations on the offshoring decisions of firms. In Chapter 5, I will present suggestive evidence based on U.S. import data of the empirical relevance of the concepts highlighted in this chapter.

A Brief Recap of the Global Sourcing Model

Let us now go back to the model of global sourcing with heterogeneous firms introduced in Chapter 2. To build intuition, I will initially focus on the simple two-stage, two-country offshoring model inspired by the work of Antràs and Helpman (2004). Before transitioning to the empirics, it will prove useful to extend the model to a multi-country environment as outlined towards the end of Chapter 2, but I defer any discussion of multi-country environments until Chapter 5.

I will not review the benchmark global sourcing model in great detail, but it may be useful to remind the reader a few of its key features. We are largely concerned with the behavior of firms in a differentiated good sector in which final goods are produced by combining two stages, headquarter services and manufacturing production, under a Cobb-Douglas technology. Firms differ in their productivity level, which as in Melitz (2003), is only revealed upon paying a fixed cost of entry. There are two countries, the North and the South. The North has comparative advantage in entry and headquarter services, so these stages always occur there. Conversely, the South has comparative advantage in manufacturing production, and absent any fixed costs of off-

shoring, all firms would want to fragment production and combine Northern headquarter services and Southern manufacturing production. Offshoring is however costly and entails a fixed cost of f_O units of Northern labor, while Northern production of manufacturing goods entails a lower fixed cost equal to $f_D < f_O$ units of Northern labor.

Overall, the total cost of production associated with *Domestic* sourcing in the North and *Offshoring* to South are given, respectively, by

$$C_D(q, \varphi) = \left(f_D + \frac{q}{\varphi} \right) w_N,$$

and

$$C_O(q, \varphi) = f_O w_N + \frac{q}{\varphi} (w_N)^\eta (\tau w_S)^{1-\eta},$$

where φ is a firm-specific productivity parameter throughout the book, w_i denotes the wage rate in country $i = N, S$, τ reflects variable costs of offshoring (input trade costs, communications costs,...), and η represents the headquarter intensity of production.

Assuming further, and for simplicity, that trade in final goods is free, we then used the demand equation (2.3) to derive the operating profits associated with each of the two sourcing strategies available to firms in the North. In particular, a Northern firm with productivity φ would obtain a profit flow equal to

$$\pi_D(\varphi) = (w_N)^{1-\sigma} B \varphi^{\sigma-1} - f_D w_N \quad (4.1)$$

when sourcing domestically, and a profit flow equal to

$$\pi_O(\varphi) = \left((w_N)^\eta (\tau w_S)^{1-\eta} \right)^{1-\sigma} B \varphi^{\sigma-1} - f_O w_N, \quad (4.2)$$

when offshoring the manufacturing stage of production to the South. In these equations, the demand shifter B is given by

$$B = \frac{1}{\sigma} \left(\frac{\sigma}{(\sigma-1)P} \right)^{1-\sigma} \beta (w_N L_N + w_S L_S).$$

As in our discussion of the export profit functions in the Melitz (2003) model, the realization of the entire profit flow (4.2) by the Northern final-good producer rests on strong assumptions, including having information on all parameters of the model, and frictionless contracting between producers (and between producers and workers). Below, I will maintain the

assumption that all the parameters of the model are deterministic and common knowledge to all producers, and that labor markets work competitively and efficiently. This will focus our attention on the implications of weak contracting enforcement between Northern headquarters and Southern manufacturing plants. A discussion of contracting issues requires a more detailed discussion of the agents involved in production and the timing of events, a task to which I turn next.

Microeconomic Structure and Contracting

We shall begin by assuming that there are only two agents relevant for contracting in the production of each final-good differentiated variety. On the one hand, there is the final-good producer – agent F – which is in charge of incurring all fixed costs of production (including entry), and who also controls the provision of headquarter services. On the other hand, manufacturing production is controlled by a manager – agent M – in that production facility. Both agents have an outside opportunity that delivers them an income level which for simplicity I normalize to 0.

I next describe the timing of events, which is also illustrated in Figure 4.1. At some initial date t_0 , the final-good producer F incurs the fixed cost of entry $w_N f_e$, upon which the productivity level φ is revealed, and F decides whether to have the manufacturing stage of production controlled by a Northern or a Southern manager. At the end of this same period, F approaches a manager M in the chosen location and offers him or her a formal sourcing contract (details below). This initial stage is followed by an investment stage t_1 , at which F produces headquarter services and M undertakes manufacturing production. I assume for now that these investments occur simultaneously, but we will contemplate models with sequential production below. Once the investments have been incurred and before the manager M hands over the manufactured goods to the final-good producer F , we shall consider the possibility that the terms in the initial contract are renegotiated and bargained over at stage t_2 . Finally, the terms of this renegotiation (or of the initial contract in the absence of renegotiation) are executed at a final stage t_3 , when the final good is also produced and sold.

I will first illustrate that a seemingly simple initial contract might suffice for F to be able to attain the ‘frictionless’ levels of operating profits in (4.1) and (4.2). To fix ideas, let us consider the case in which, at t_0 , F has decided to approach a Southern manufacturer (we will later study the loca-




Figure 4.1: Timing of Events

tion choice). Suppose then that in the initial contract, F offers the Southern manager a contract that stipulates a quantity m^c of manufacturing production to be provided by M in exchange for a fee s^c received by M . Assume also that the initial contract includes a clause such that if any party deviates from this initial contract, the other party is entitled to arbitrarily large damages. With that clause, F can be assured that any level of manufacturing services stipulated in the initial contract will be honored by the manager. As a result, F can safely choose $h^c(\varphi)$, $m^c(\varphi)$ and $s^c(\varphi)$ to solve the following problem, where the constraint reflects the participation constraint of the Southern manager:

$$\begin{aligned} \max_{h(\varphi), m(\varphi), s(\varphi)} \quad & p(q(\varphi))q(\varphi) - w_N h(\varphi) - w_N f_O - s(\varphi) \\ \text{s.t.} \quad & s(\varphi) - \tau w_S m(\varphi) \geq 0 \end{aligned} \tag{4.3}$$

Naturally, F will set $s^c(\varphi)$ in the initial contract to make the participation constraint of the manager M exactly bind, and as a result, F will choose $h^c(\varphi)$ and $m^c(\varphi)$ to maximize the overall surplus from the relationship, thus resulting in profit-maximizing investments and a profit flow identical to that in equation (4.2) above. Note that the timing of events, investments or payments is quite irrelevant for this result.

Incomplete Contracts and Weak Contract Enforcement

As simple as the above contract may seem, it is hard to imagine that agents will in fact be able to (1) write such a type of contract, and (2) find a court of law to enforce it. More specifically, note that the initial contract needs to specify the level of m , which corresponds to the value or services

obtained from the manufacturing stage of production. It does not merely reflect, in particular, the number of physical units of manufacturing goods. This distinction is crucial and carries important consequences. Consider our above anecdotal evidence from Midler's book *Poorly Made in China*. Johnson Carter presumably signed a contract with King Chemical for the provision of a certain number of bottles of soap in exchange for a certain amount of money (or price per bottle). It is hard to imagine, however, that the formal contract specified a variety of characteristics of the product which were clearly relevant for how the combination of Johnson Carter's 'headquarter services' and King Chemical's manufactured goods was to translate into sale revenues in U.S. retail chains. For instance, the contract almost surely did not indicate the plastic content of the bottles or the chemical composition of the soap. Flimsy bottles or abrasive soap would each lead to a low value of m in the model, no matter how large the number of bottles actually manufactured.

The general lesson here is that while certain aspects of manufacturing, such as the number of units of goods to be produced, the price per unit, or the date of the delivery are relatively easy to incorporate into a formal written contract, many other aspects of production are not. And, crucially, these noncontractible elements of production are often key in shaping the quality of goods or their compatibility with other parts of the production process.

Beyond the obvious incompleteness of real-life commercial contracts, there still remains the issue of contract enforcement. Even if two expert parties were able to design a highly comprehensive contract specifying what each party is supposed to contribute to production and making explicitly what defines high or low quality or compatibility, it is still questionable that a court of law will be able to understand such a contract and enforce it properly. And even when it does, in international transactions there still remains the uncertainty over the cross-border enforcement of damages (see Chapters 1 and 3 for a discussion of this issue). A natural response of the contracting parties is thus to shy away from specifying too harsh penalties for deviations from highly detailed clauses in written contracts. In Midler's soap bottles examples, Johnson Carter could have presumably insisted on the initial contract specifying the plastic content of bottles or that each bottle contained exactly 850 milliliters of soap, but it is somewhat hard to believe that a court of law would be able to verify whether that contract had been honored or not.

In sum, as simple as it might have seemed, the type of contract that we

have considered so far is in fact too complicated to realistically discipline the behavior of agents. If such a contract is not feasible or enforceable, then what type of contracts are? One could envision the possibility that the initial contract would at least specify all the characteristics of the contractible aspects of production, while stipulating large penalties for deviations from that contract, thus ruling out the possibility of any ex-post renegotiation. It is not hard to see, however, that these type of contracts will typically deliver unappealing outcomes. For instance, imagine that the initial contract were to stipulate the number of physical units of m to be traded as well as their price in a binding manner, and without allowing M or F to renegotiate those terms upon observing the quality of the manufactured goods. In such a case, M would have every incentive to produce the inputs m in the least-cost possible manner, which would typically result in a low quality level of those goods. Foreseeing this debasements in quality, F would not be willing to offer a particularly large price for those goods in the initial contract, and the overall surplus of such a contractual relationship would end up being quite low and possibly zero, if M were to produce useless goods. In these circumstances, F might be better off by offering a less complete contract at the initial stage, as we will show below.

Another possibility would be for the initial contract to specify a simple sharing rule for the sale revenue obtained by F when selling the final good to consumers. Even if the initial contract did not specify the quality characteristics of the manufactured inputs in a binding manner, it seems natural that the incentives of M to skim on quality will be attenuated when his or her profits are a function of the the willingness of consumers to pay for goods that embody those manufactured inputs. Whenever certain aspects of the investments in h and m remain noncontractible, however, these type of revenue-sharing contracts will not lead to the frictionless levels of investment and profits of the complete-contracting environment, a result reminiscent of Holmstrom's (1982) moral-hazard-in-teams problem. As a result, even when revenue-sharing arrangements are feasible and enforceable, we will see that in some cases it may well be in the interest of the parties to opt out of them in the initial contract. Needless to say, the appeal of these contracts is further diminished whenever they are not perfectly enforceable, perhaps due to manipulation by the agent collecting revenues, who may be tempted to underreport them. Gennaioli (2013) has studied optimal contracting in a model with potentially biased judges, and has shown that the enforcement risk generated by these biases, often leads parties to write simple contracts

that are not contingent on revenues.

Borrowing tools from the mechanism design literature, an important body of theoretical work has proposed a variety of ingenious ‘mechanisms’ to restore efficiency in environments in which the contracting parties (i) have symmetric information, (ii) can commit not to renegotiate an initial contract, and (iii) can resort to a third party (presumably a court of law or arbitrator) to enforce off-the-equilibrium-path penalties (see, for instance, Aghion, Dewatripont and Rey, 1994, or Maskin and Tirole, 1999). This literature is often criticized for suggesting somewhat convoluted contracts that are not observed in the real world. I find that criticism unconvincing: after all, one could have similarly criticized some key contributions to the auction theory literature, and yet they have subsequently had an enormous impact in real-world auctions (see, for instance, Milgrom, 2000). My main reservation with mechanism-design resolutions to incomplete contracting is that they rely on the ability of a third party to enforce contracts, and as argued repeatedly in this book, this is a real sticking point when studying the international organization of production.

‘Totally Incomplete’ Contracts

Given the discussion above and for pedagogical reasons, I will begin by considering environments in which, when contracts are incomplete, they are so in a rather extreme way. With that in mind, consider the following definition:

Definition *A contract is said to be ‘totally incomplete’ whenever no aspect of the contract is perceived to be enforceable, with the possible exception of a lump-sum transfer exchanged at the time of the agreement.*

For reasons discussed in Chapter 1, it seems natural to assume that certain contracts that are feasible or enforceable in domestic transactions might not be feasible or enforceable in international transactions. To illustrate the implications of this asymmetry I will this consider first the case in which contracting is complete or perfect in the domestic sourcing relationships, while contracting is *totally incomplete* in offshoring relationships. Later in the chapter, I will consider environments with partial contractibility in each of these two sourcing options.

With complete contracting in domestic sourcing, we have demonstrated above that F will be able to design a contract such that the levels headquarter services h and manufacturing production m are set at their joint-profit max-

imizing level, thus resulting in the frictionless profit flow $\pi_D(\varphi)$ in equation (4.2).

Let us next consider the more interesting implications of incomplete contracting in offshoring relationships. What happens when the initial contract does not stipulate the levels of h or m nor a payment to be paid to the manufacturer M contingent on the volume of m produced or contingent on sale revenues? In that case, the only option left for the parties is to decide on the terms of exchange at t_2 (remember the timing of events in Figure 4.1). The next question is then: how should one model this bargaining/contracting stage? Given that we have assumed above that the final good producer makes a take-it-or-leave-it offer to the manager at t_0 , it might seem natural to maintain that assumption for stage t_3 . This would however ignore what Oliver Williamson famously termed the “fundamental transformation” (see, for instance, Williamson, 1985). This transformation refers to the fact that even though the final good producer might have chosen a particular manager M from a competitive fringe of managers, once the investments h and m have been incurred, a contractual separation is likely to prove costly to both parties. To the extent that parties feel ‘locked-in’ with each other, the initial competitive environment at t_0 has thus been fundamentally transformed into a bilateral monopoly scenario.

As explained in Chapter 1, in the global sourcing environments that we are considering in this book, there are several natural sources of *lock in* between final-good producers and suppliers. First, manufacturing inputs are often customized to their intended buyers and cannot easily be resold at full price to alternative buyers. Second, certain types of headquarter services are also designed with particular suppliers in mind, and it would prove costly to reuse these services with alternative suppliers. Third, even though we have abstracted from modeling them in the initial period, search frictions are particularly relevant in international environments and they are likely make ex-post separations particularly costly for both final-good producers and suppliers, who would at the very least suffer from delays in obtaining a return on their investments.

In sum, in the presence of lock-in effects, incomplete contracting leads to a situation of bilateral monopoly in which the terms of exchange between F and M will only be determined ex-post (at t_2), after these agents have incurred investments that are by then sunk and have a relatively lower value outside that particular business relationship. The combination of incomplete contracting and lock-in effects, leads to what is often referred to as a *hold-up*

problem, which in our particular context is two-sided. More specifically, on one end, F will try to push down the price paid for M 's manufacturing input, realizing that M might be inclined to accept a reduced price due to the lower value of those inputs for alternative buyers. At the same time, however, M will try to raise the price of m as much as possible, knowing that it might also be in F 's best interest to accept a relatively high price if that avoids having to search for a new supplier.

In deciding how assertively to bargain, each party takes into account that a too aggressive offer might lead the other party to refuse to trade, an outcome that is not appealing given the lower value of the sunk investments outside the relationship. As a result, even when bargaining is efficient and trade takes place in equilibrium, the *possibility* of a disagreement and associated failure to trade implies that F and M will tend to have lower incentives to invest in h and m than in the complete contracting case, in which case the initial contract ensures that trade will occur. In more technical terms, with incomplete contracting, the payoff obtained by each party in the ex-post negotiations will put a positive weight on *off-the-equilibrium path* situations in which the return to each party's investments is lower than on the equilibrium path.

In the literature, it is common to characterize the ex-post bargaining at t_2 using the Nash Bargaining solution and assuming symmetric information between F and M with regards to all parameters of the model. In such a case, each party ends up with a payoff equal to the value of their outside option (their payoff under no trade) plus a share of the ex-post gains from trade, which correspond to the difference between the sum of the agent's payoffs under trade and their sum under no trade. For the time being, I will assume that the outside option for each party is equal to 0. In other words, I am assuming that the manufactured input m is fully specialized to F and thus useless to other producers, while headquarters h are also fully tailored to the supplier M and could not be productively combined with inputs provided by other manufacturers. I will also consider the case of *symmetric* Nash bargaining, which implies that F and M share *equally* the ex-post gains, which with zero outside options equal sale revenues. Obviously, these are restrictive assumptions, but I will consider more general environments below.

In sum, with symmetric Nash bargaining, each party will anticipate obtaining a payoff equal to one-half of sale revenue at t_2 , and thus the levels of

$h(\varphi)$ and $m(\varphi)$ will be set at t_1 to solve

$$\max_h \frac{1}{2}p(q(\varphi))q(\varphi) - w_N h \quad (4.4)$$

and

$$\max_m \frac{1}{2}p(q(\varphi))q(\varphi) - \tau w_S m, \quad (4.5)$$

respectively, where remember that

$$q(\varphi) = \varphi \left(\frac{h}{\eta}\right)^\eta \left(\frac{m}{1-\eta}\right)^{1-\eta} \quad (4.6)$$

and $p(q(\varphi)) = B^{1/\sigma} \sigma (\sigma - 1)^{(\sigma-1)/\sigma} q(\varphi)^{-1/\sigma}$.¹

For comparability with the complete-contracting case, I will assume again that at t_0 there is a competitive fringe of potential managers M willing to work for each final-good producer F at a reservation wage equal to 0. Each F then decides the terms of the initial contract and makes a take-it-or-leave-it offer to one of those managers. Because the initial contract is allowed to include a lump-sum transfer between parties, F can set the transfer such that the participation constraint of M exactly binds. So, as with complete contracts, F ends up with a profit level equal to

$$\pi_O(\varphi) = p(q(\varphi))q(\varphi) - w_N h(\varphi) - \tau w_S m(\varphi) - w_N f_O. \quad (4.7)$$

The key difference, however, is that the levels of $h(\varphi)$ and $m(\varphi)$ can no longer be set in the initial contract in an enforceable manner. As a result, $h(\varphi)$ and $m(\varphi)$ are no longer set to maximize $\pi_O(\varphi)$ in (4.7), but instead are chosen simultaneously and non-cooperatively by F and M to solve programs (4.4) and (4.5), respectively.

In analogy to the way we studied contractual frictions in Chapter 3 and to our formulation of the complete-contracting program in (4.3), a compact way to represent the ex-ante problem faced by the final good producer under ‘totally incomplete’ contracts is:

¹It should be noted that I am assuming that the agreement in t_2 is always enforced. Enforcement might become an issue if the parties exchange the goods at t_2 and the payment occurs at t_3 , but not if the payment occurs at t_2 , simultaneously with the exchange of goods.

$$\begin{aligned}
& \max_{h(\varphi), m(\varphi), s(\varphi)} \quad \frac{1}{2} p(q(\varphi)) q(\varphi) - w_N h(\varphi) - w_N f_O - s(\varphi) \\
& \text{s.t.} \quad s(\varphi) + \frac{1}{2} p(q(\varphi)) q(\varphi) - \tau w_S m(\varphi) \geq 0 \\
& \quad h(\varphi) = \arg \max_h \left\{ \frac{1}{2} p(q(\varphi)) q(\varphi) - w_N h(\varphi) \right\} \\
& \quad m(\varphi) = \arg \max_m \left\{ \frac{1}{2} p(q(\varphi)) q(\varphi) - \tau w_S m(\varphi) \right\}.
\end{aligned} \tag{4.8}$$

A simple comparison of (4.3) and (4.8) shows that the effect of incomplete contracting is captured by the addition of the last two constraints, which represent incentive compatibility constraints faced by both F and M . Because the bargaining payoffs of each agent put a positive weight (in this case, $1/2$) on an off-the-equilibrium-path zero return to producing in h and m , it is naturally the case that the hold up problem discussed above leads to inefficiently low investment levels at t_1 , and consequently, results in depressed overall profits as well.

Solving formally program (4.8) delivers a volume of profits that F anticipates obtaining when choosing foreign sourcing at t_0 equal to

$$\pi_O = ((w_N)^\eta (\tau w_S)^{1-\eta})^{1-\sigma} B \Gamma_O \varphi^{\sigma-1} - w_N f_O, \tag{4.9}$$

where

$$\Gamma_O = (\sigma + 1) \left(\frac{1}{2} \right)^\sigma < 1 \text{ for } \sigma > 1. \tag{4.10}$$

Note that (4.9) is identical to the complete-contracting expression (4.2) except for the term Γ_O , which is necessarily lower than 1. Hence, this term reflects the loss of efficiency due to incomplete contracting. Furthermore, for the relevant range $\sigma > 1$, Γ_O is a decreasing function of σ , and thus in environments with tougher competition (i.e., lower markups), the profit losses from incomplete-contracting frictions are relatively larger (see the Theoretical Appendix for a formal proof). Intuitively, the last two constraints in (4.8) suggest that the effect of incomplete contracting on the choice of h and m is analogous to a doubling of the marginal cost of each of these stages, and this will tend to reduce profits more, the more price elastic demand is. This marginal cost inflation only becomes irrelevant in the limiting case $\sigma \rightarrow 1$, since in that case a firm's market is independent of its cost.

Choice of Location and Prevalence of Offshoring

Having computed the anticipated profits associated with domestic sourcing and offshoring, we can next study the choice of location of the final-good producer in the initial period t_0 . Note from equations (4.1) and (4.9) that we can write these profits functions succinctly as

$$\pi_\ell(\varphi) = \psi_\ell B \varphi^{\sigma-1} - w_N f_\ell \quad \text{for } \ell = D, O,$$

with

$$\frac{\psi_D}{\psi_O} = \frac{1}{\Gamma_O} \left(\frac{w_N}{\tau w_S} \right)^{-(1-\eta)(\sigma-1)}. \quad (4.11)$$

As in the benchmark models reviewed in Chapter 2, the profit functions $\pi_D(\varphi)$ and $\pi_O(\varphi)$ are linearly increasing in the transformation of productivity $\varphi^{\sigma-1}$, with the relative slope of the two functions now being governed by the ratio ψ_D/ψ_O . Figure 4.2 depicts these functions for the case in which wage differences across countries are high so that $\psi_D < \psi_O$ (see the line π_O^h), and the case in which wage differences are sufficiently low to imply $\psi_D > \psi_O$ (see the line π_O^l). In the latter case, no firm finds it optimal to offshore in the South, while in the former case, only the most productive firms will. The key difference with the complete-contracting case is that, other things equal, firms will now find it less profitable to offshore to the South due to the term Γ_O in (4.11).

Let us now aggregate the decisions of the various firms in a sector and study the determinants of the relative importance of offshoring in different industries. As in Chapter 2 and with an eye to the empirical applications in Chapter 5, it seems natural to focus on the share of spending on *imported* manufacturing inputs over total manufacturing input purchases in a particular industry as a measure of the prevalence of offshoring. Because contracts are complete in domestic transactions, for all firms sourcing domestically input purchases constitute a share $(1-\eta)(\sigma-1)/\sigma$ of revenue and a multiple $(1-\eta)(\sigma-1)$ of operating profits (defined as revenue minus variable costs)

Matters are trickier in the case of foreign intermediate inputs, since one needs to take a stance on how these inputs are priced. One possibility is to assume that the headquarter H 's spending on inputs corresponds to the payment obtained by the manufacturing manager M in the ex-post bargaining at t_2 , which in the model above is simply $\frac{1}{2}$ of revenue and a share $\sigma/(\sigma+1)$




Figure 4.2: Equilibrium Sorting and Contractual Frictions

of operating profits. Alternatively, one could appeal to the existence of the ex-ante lump-sum transfer $s(\varphi)$ at t_0 to argue that H is *effectively* purchasing inputs at marginal cost and thus total spending on foreign inputs is instead given by $\tau w_S m(\varphi)$. From the first-order-condition of the program in the last constraint of program (4.8), this corresponds to a share $\frac{1}{2}(1-\eta)(\sigma-1)/\sigma$ of sale revenue or a share $(1-\eta)(\sigma-1)/(\sigma+1)$ of operating profits.

I do not want to take a strong stance on a particular pricing practice so I will instead take the more agnostic approach of assuming that foreign inputs are priced such that these input expenditures constitute the same multiple $(1-\eta)(\sigma-1)$ of operating profits as in the case of domestic input purchases. I do not believe that this assumption is crucial for the results derived below, but it will significantly simplify the derivations.

Assuming a Pareto distribution of productivity, we can then follow similar steps as in the derivation of equations (2.22) and (2.23) to solve for the share Υ_O of imported manufacturing input purchases in a given industry as

$$\Upsilon_O = \frac{\Gamma_O \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)}}{\left(\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} \right)^{\kappa-(\sigma-1)} - 1 + \Gamma_O \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)}}. \quad (4.12)$$

where

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\Gamma_O \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)}. \quad (4.13)$$

These expressions are analogous to equations (2.22) and (2.23) in Chapter 2 except for the presence of the term $\Gamma_O < 1$.

As is clear from these equations, the prevalence of offshoring is naturally increasing in the term Γ_O , and is thus reduced by the presence of contractual frictions. The remaining comparative statics are analogous to those in the complete-contracting case. The share of imported inputs is increasing in wage differences (w_N/w_S) and productivity dispersion ($1/\kappa$) and decreasing in (relative) fragmentation barriers (f_O/f_D , τ) and headquarter intensity (η). Conversely, the overall effect of the elasticity of substitution on the share Υ_O is now ambiguous. As in the complete-contracting case, a higher σ makes offshoring more prevalent by reducing the Pareto shape parameter $\kappa/(\sigma - 1)$ for firm sales and by increasing the percentage gain from a reduction in marginal costs of labor from w_N down to τw_S . Nevertheless, as mentioned above in our discussion of equation (4.10), a larger σ now also aggravates the inefficiencies caused by incomplete contracting, hence reducing Υ_O .

Extensions of the Basic Model

So far, our analysis has illustrated that weak contract enforcement across countries will tend to reduce the profitability associated with firms engaging in global sourcing strategies and will thus lead to a larger reliance on domestic intermediate inputs. We have illustrated this insight with a highly stylized framework with many simplifying assumptions, and as a result the model has not delivered any particularly valuable empirical predictions other than the intuitive negative effect of weak contracting on offshoring. I will next turn to studying more general environments that relax some of the strong assumptions above. This will serve to verify the robustness of the key comparative statics emphasized so far, and also to develop a richer set of comparative statics that are better suited to guide empirical work on the contractual determinants of the global sourcing decisions of firms.

I will consider six basic extensions of the model: (i) a generalization of the bargaining process at t_2 , (ii) the possibility of restrictions on ex-ante transfers at t_0 , (iii) environments with partial contractibility at t_0 , (iv) investments at

t_1 that are only partially relationship-specific, (v) global sourcing decisions with multiple suppliers, and (vi) sequential production in which the investments of different suppliers occurs at different points in time. In each of these cases, we will confirm that the basic insights obtained so far continue to apply to those more general environments, but we will obtain sharper predictions about the differential effect of contractual institutions on firm profitability across firms, sectors and countries.

To be as succinct as possible, I will focus on outlining how each of these extensions affects and enriches the determination of the term Γ_O in (4.9) capturing the inefficiencies associated with incomplete contracting in offshore relationships. I will also develop variants of the model in which contractual frictions reduce the profitability of domestic sourcing via an analogous term Γ_D which is also shaped by the parameters of the model in ways to be discussed below. In the next chapter, I will return to the aggregation of firms' decisions to illustrate how Γ_O and Γ_D jointly affect the share of imported manufacturing input purchases in a given industry as in equation (4.12). This will serve to motivate my discussion of empirical evidence based on U.S. import data.

Generalized Nash Bargaining and Revenue-Sharing Contracts

In our basic model above, we have assumed that F and M share the ex-post gains from trade equally at t_2 . In some circumstances, it may make sense to assume that the primitive bargaining power of final-good producers might be higher (or perhaps lower) than that of supplying managers. The large literature on non-cooperative models of bargaining emanating from the seminal work of Rubinstein (1982), has uncovered several potential determinants of primitive bargaining power. It is well-known, for instance, that relatively impatient or risk averse agents will tend to have relatively low bargaining power, and the same will be true about agents for which a bargaining delay might be particularly costly for reasons other than impatience, such as credit constraints (see, for instance, Rubinstein, 1982, or Roth, 1985).

Rather than developing any of these microfoundations in great detail, let me just assume that the final-good producer F obtains a share β of the ex-post gains from trade, with the manager M obtaining the remaining share $1 - \beta$. This is often referred to as the *generalized* Nash bargaining solution. With the maintained assumption that the investments incurred to produce h and m are fully relationship-specific, β and $1 - \beta$ will also correspond to

the shares of revenue obtained by F and M , respectively, in their ex-post negotiations at t_2 . Below, we will develop variants of the model in which the division of revenue will be shaped by factors other than primitive bargaining power.

It should be clear by now that with complete contracting, the particular bargaining solution adopted to characterize the t_2 stage is irrelevant, and thus the profits associated with domestic sourcing are still given by (4.1). Conversely, when solving for the equilibrium associated with offshoring, one needs to replace the term $1/2$ in (4.4) and the second constraint in (4.8) with β , and the term $1/2$ in (4.5) and the third constraint in (4.8) with $1 - \beta$. Naturally, the larger is β , the lower will tend to be the underinvestment in the provision of headquarter services, but the larger will be the underinvestment in the provision of manufacturing services.

Solving for the equilibrium profits obtained by F under generalized Nash bargaining, delivers a profit flow identical to that in equation (4.9), but with

$$\Gamma_O = \Gamma_\beta \equiv (\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta))) (\beta^\eta (1 - \beta)^{1-\eta})^{\sigma-1}. \quad (4.14)$$

It is straightforward (though somewhat tedious) to verify that, regardless of the value of the primitive bargaining power β , incomplete contracting still necessarily reduces the profitability of offshoring, i.e., $\Gamma_\beta < 1$ for $\sigma > 1$.² Furthermore, in the Technical Appendix I show that the size of these contractual distortions continues to be increasing in the elasticity of demand (or $\partial\Gamma_\beta/\partial\sigma < 0$). The main novelty in (4.14) is that the level of contractual frictions is no longer only shaped by the elasticity of substitution σ , but now also depends on headquarter intensity η and on the bargaining power parameter β . The effects of these parameters on the level of Γ_β are non-monotonic and interact closely with each other. More specifically, it can be shown that Γ_β is decreasing in η when $\beta < 1/2$, while it is increasing in η when $\beta > 1/2$.³ Notice also that when $\eta \rightarrow 0$, $\Gamma_\beta \rightarrow (\sigma - (\sigma - 1)(1 - \beta)) ((1 - \beta))^{\sigma-1}$ and thus Γ_β is a decreasing function of β , while when $\eta \rightarrow 1$, $\Gamma_\beta \rightarrow (\sigma - (\sigma - 1)\beta) \beta^{\sigma-1}$, and Γ_β is instead an increasing function of β . In other words, whether increases or decreases in the bargaining power of the final good producer increase or

²This is a special case of the proof of Proposition 1 in Antràs and Helpman (2008). The result can also be proven more directly by noting that Γ_β was derived from an optimization problem that is more constrained than the one that delivered the profit flow in (4.2), and the latter profit flow can be obtained from (4.9) by setting $\Gamma_\beta = 1$.

³The proof is straightforward but cumbersome, so I relegate it to the Theoretical Appendix.

decrease profits also crucially depends on the level of headquarter intensity. Intuitively, and as argued above, when β increases, the underinvestment in headquarter services is alleviated, while the underinvestment in manufacturing production is aggravated. Whether the net effect is positive or negative naturally depends on the intensity with which these two stages are combined in production. In fact, straightforward calculations show that there exists a unique value β^* that maximizes Γ_β , and it satisfies

$$\frac{\beta^*}{1 - \beta^*} = \sqrt{\frac{\eta}{1 - \eta} \frac{\sigma - (\sigma - 1)(1 - \eta)}{\sigma - (\sigma - 1)\eta}}. \quad (4.15)$$

In line with our intuition above, this profit-maximizing level of β^* is an increasing function of headquarter intensity η .

This result also shed lights on the implications of the model when allowing for revenue-sharing contracts to be signed at t_0 . In particular, imagine a situation in which the ex-ante contract was not ‘totally incomplete’ but rather was allowed to include a division rule contingent on the volume of revenue generated at t_3 . Denoting revenue by $R = p(q)q$ and the sharing rule by $\beta(R)$, the optimal initial contract would now solve (I am omitting the argument φ in all functions for simplicity):

$$\begin{aligned} \max_{h,m,s,\beta(R)} \quad & \beta(R)R - w_N h - w_N f_O - s \\ \text{s.t.} \quad & s + (1 - \beta(R))R - \tau w_S m \geq 0 \\ & h = \arg \max_h \{\beta(R)R - w_N h\} \\ & m = \arg \max_m \{(1 - \beta(R))R - \tau w_S m\}. \end{aligned}$$

If one restricts attention to linear sharing rules where $\beta(R)$ is independent of R , then our above discussion indicates that the optimal contract will set $\beta(R) = \beta^*$, where β^* is given in (4.15). Even with this more complete initial contract, the frictionless profit flow in (4.2) cannot possibly be attained because remember that Γ_β in (4.14) is less than 1 for any $\beta \in (0, 1)$. Importantly, this conclusion is not specific to the case of linear sharing rules. As shown by Holmstrom (1982), for general sharing rules $\beta(R)$ satisfying budget-balance, the resulting investment levels h and m will continue to differ from the efficient ones, and the equilibrium profits associated with offshoring will necessarily fall short of those under complete contracts.

Limitations on Ex-Ante Transfers: Financial Constraints

So far, I have assumed that F and M are allowed to freely exchange lump-sum transfers when signing the initial contract at t_0 . This transfer can be inferred from the participation constraint of the manager M , which implies

$$s(\varphi) = \tau w_S m(\varphi) - \frac{1}{2} p(q(\varphi)) q(\varphi).$$

Plugging the equilibrium values of $m(\varphi)$ and $q(\varphi)$ it is straightforward to show that $s(\varphi) \leq 0$ and thus the optimal contract calls for the manager M to post a bond in order to be able to transact with the final good producer F . In practice, it is not obvious that supplying firms will be willing or able to make that initial transfer. This is due to at least two reasons. First, managers might worry that the final good producer will disappear after period t_0 without incurring the fixed cost of offshoring or investing in headquarter services.⁴ Second, depending on the financial environment in the manager's country, it may be hard for M to raise from financiers the full amount of cash $s(\varphi)$ stipulated in the contract we have considered so far. We next explore the implications of the existence of constraints on these ex-ante transfers.

To fix ideas, consider the case in which M can pledge to external financiers in his domestic economy at most a share ϕ of the net income it receives from transacting with F . I will not specify the source of these financial frictions, but they could stem from a limited commitment friction on the part of M along the lines of the models we have explored in Chapter 3. The equilibrium under financial constraints can again be reduced to a program analogous to that in (4.8), but with the additional constraint

$$-s(\varphi) \leq \phi \left[\frac{1}{2} p(q(\varphi)) q(\varphi) - \tau w_S m(\varphi) \right].$$

This constraint is tighter than the original participation constraint in (4.8) and will bind in equilibrium. Solving the program, we then find that F again obtains a payoff equal to that in equation (4.9), but now with

$$\Gamma_O = \Gamma_\phi \equiv (\sigma + \phi - (\sigma - 1)(1 - \phi)\eta) \left(\frac{1}{2} \right)^\sigma. \quad (4.16)$$

⁴It can be shown, however, that if the manager insists that the fixed cost of offshoring be incurred prior to the payment $s(\varphi)$, the final-good producer would no longer have an incentive to abscond with the transfer.

It is straightforward to see that Γ_ϕ is increasing in ϕ and attains the symmetric Nash bargaining value of $(\sigma + 1)(1/2)^\sigma$ when financial constraints disappear, i.e., $\phi = 1$. A first implication is thus that the profitability of offshoring will be increasing in the quality of financial contracting in the manager's foreign country, as summarized by the parameter ϕ . Intuitively, offshoring is now not only associated with distorted investments, but it also entails a loss of rents for the final good producer F . It is interesting to note that Γ_ϕ is also decreasing in headquarter intensity. This is due to the fact that headquarter services are complementary to manufacturing production, and the higher is η , the larger are the rents that M obtains in the ex-post bargaining at t_2 relative to the costs of production her or she incurs at t_1 . Consequently, the larger is η , the larger is the loss of rents for F associated with a low ϕ . For the same reason, we observe in (4.16) that an increase in improvement in the quality of financial contracting will have a differentially large positive effect on the profitability of offshoring in production processes with high headquarter intensity (i.e., $\partial(\partial\Gamma_\phi/\partial\phi)/\partial\eta > 0$).⁵

Partial Contractibility

It is obviously unrealistic to assume, as we have done so far, that contracts in international transactions are 'totally incomplete'. It seems natural that some aspects of production can be specified in a contract in a manner that contracting parties feel confident that those aspects of the contract will be enforced. Moreover, it is also unrealistic to assume, as we have done so far, that contracts in domestic transactions are complete. Surely some aspects of production are nonverifiable or certain contracts are perceived to be hard to enforce in domestic transactions. I next incorporate partial contractibility into our global sourcing model following the approach in Antràs and Helpman (2008).

It will prove useful to assume that the production of headquarter services and manufacturing inputs now entails a continuum of processes or activities all of them carried out at t_1 . A fraction of these processes is assumed to be ex-ante contractible, in the sense that contracts specifying how those processes should be carried out can be designed in a way that a court of law can verify their fulfillment and penalize any deviation from what was stipulated in

⁵It is also worth pointing out that, under plausible parametric restrictions, in this extension of the model too, higher demand elasticities are associated with larger incomplete-contracting distortions (see the Technical Appendix for details).

the contract. Conversely, the complementary fraction of processes is non-contractible and contracts specifying these activities would fail to discipline their provision at t_1 .

There are two main determinants of the degree to which the overall production process is contractible or not. First, a low fraction of contractible activities could reflect technological factors that make it particularly hard to write down enforceable contracts disciplining the behavior of the agents engaged in production. For instance, the production of new and high-tech goods is more contractually demanding than that of more traditional and standardized goods. Second, even focusing on the same production process, it seems reasonable to assume that the fraction of contractible activities will vary across countries reflecting international variation in the quality of contracting institutions. In other words, certain types of contracts are perceived to be enforceable in some environments but perhaps not in others.

To further illustrate these two determinants, it may be useful to consider the following analogy to refereeing in (European) football. There are some rules in football that are almost trivial to enforce, such as ensuring that each team has no more than eleven players on the pitch or preventing a team from performing more than three substitutions. On the other hand, other rules are much trickier to enforce, such as calling a close offside infraction or deciding whether the entire ball crossed the goal line in a ghost-goal situation. It is thus not surprising that players anticipate the former rules to be properly enforced with a much higher likelihood than the latter ones. At the same time, the quality of the referees is also obviously critical in predicting whether the rules will be correctly applied or not. Skilled referees (or linesmen) make fewer mistakes in enforcing these rules than incompetent ones.

In order to formally introduce partial contractibility into the framework, we now let headquarter services h and manufacturing production m be a Cobb-Douglas aggregate of the services of a continuum of measure one of activities, so

$$h = \exp \left[\int_0^1 \log h(i) di \right] \quad (4.17)$$

and

$$m = \exp \left[\int_0^1 \log m(i) di \right]. \quad (4.18)$$

Our key new assumption is that activities related to input $k = h, m$ in the range $[0, \mu_{kj}]$ (with $0 \leq \mu_{kj} \leq 1$) are contractible in country $j = N, S$, in

the sense that the characteristics of these activities can be fully specified in advance in an enforceable ex-ante contract involving a manufacturer M from country j . Hence, the initial contract is no longer ‘totally incomplete’ because in addition to a lump-sum transfer between F and M , it also specifies the level of contractible activities to be carried out at t_1 . The remaining activities in the range $(\mu_{kj}, 1]$ continue to be noncontractible as in our benchmark model and F and M decide on the terms of exchange for those activities only after they have been produced. Because the initial contract does not compel any of the two parties to provide a positive amount of these noncontractible tasks, the threat point for each party in the negotiations at t_2 is to withhold the services from those activities, which in light of the Cobb-Douglas production technologies (4.6), (4.17), and (4.18), and our maintained assumption that all investments are fully relationship-specific, would lead to a zero payoff for both parties.⁶ Thus each agent ends up capturing a constant share of sale revenues, and for simplicity, I will now revert back to the assumption of symmetric Nash bargaining, so that the two parties end up sharing evenly total sale revenues.

The symmetry assumptions on technology built into (4.17) and (4.18) allows us to simplify the problem of the firm conditional on having selected a location $j = N, S$, to the choice of an ex-ante transfer s , a common value h_c for all contractible headquarter activities, a common value h_n for all noncontractible headquarter services, and analogous values m_c and m_n for contractible and noncontractible manufacturing tasks, respectively. Formally, we can now write the problem (ignoring fixed costs) as

$$\begin{aligned}
& \max_{h_c, h_n, m_c, m_n, s} && \frac{1}{2}R - w_N (\mu_{hj} h_c + (1 - \mu_{hj}) h_n) - s \\
& \text{s.t.} && s + \frac{1}{2}R - c_j (\mu_{mj} m_c + (1 - \mu_{mj}) m_n) \geq 0 \\
& && h_n = \arg \max_h \left\{ \frac{1}{2}R - w_N (1 - \mu_{hj}) h \right\} \\
& && m_n = \arg \max_m \left\{ \frac{1}{2}R - c_j (1 - \mu_{mj}) m \right\}.
\end{aligned} \tag{4.19}$$

⁶An implicit assumption in the analysis is that these noncontractible tasks are not yet fully embodied in into the manufactured inputs at the time of bargaining.

where revenue is given by

$$R = B^{1/\sigma} \sigma (\sigma - 1)^{-(\sigma-1)/\sigma} \varphi^{(\sigma-1)/\sigma} \times \left(\frac{(h_c)^{\mu_{hj}} (h_n)^{1-\mu_{hj}}}{\eta} \right)^{(\sigma-1)\eta/\sigma} \left(\frac{(m_c)^{\mu_{mj}} (m_n)^{1-\mu_{mj}}}{1-\eta} \right)^{(\sigma-1)(1-\eta)/\sigma} \quad (4.20)$$

and where $c_j = w_N$ when $j = N$ and $c_j = \tau w_S$ when $j = S$.

This problem is somewhat tedious to solve, so I will not go over the derivations here. The interested reader can find the details in the Theoretical Appendix, where I reproduce the derivations in Antràs and Helpman (2008), where we solved the same problem for a general division of revenue (β_h, β_m) , rather than $(1/2, 1/2)$. Using these results, we find that, in the case of domestic sourcing, the profits obtained by F are now given by

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

where

$$\Gamma_D(\mu_N) = \left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_N)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_N)} \left(\frac{1}{2} \right)^\sigma \quad (4.21)$$

and

$$\mu_N \equiv \eta \mu_{hN} + (1 - \eta) \mu_{mN}.$$

The derived parameter μ_N measures the average contractibility associated with domestic sourcing and is a weighted sum of the contractibility of head-quarter services and manufacturing.

F 's profits under foreign sourcing can be similarly computed, resulting in

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

with

$$\Gamma_O(\mu_S) = \left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_S)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_S)} \left(\frac{1}{2} \right)^\sigma \quad (4.22)$$

and

$$\mu_S \equiv \eta \mu_{hS} + (1 - \eta) \mu_{mS}.$$

As in the simpler models developed above, the term $\Gamma_\ell(\mu_j)$ captures the contractual frictions associated with the sourcing options $\ell = D$ and

$\ell = O$, which entail manufacturing in country $j = N$ and country $j = S$, respectively. Differentiation of (4.21) and (4.22) demonstrates that each of these terms is increasing in their associated index of contractibility (see the Theoretical Appendix). Hence, as in our simpler model above, contract incompleteness reduces the profitability of production but the effect is now smoothly shaped by the partial contractibility parameters μ_{hj} and μ_{mj} for $j = N, S$. In fact, our initial model is a special case of the current one, with complete contracting in domestic sourcing (so $\mu_{hN} = \mu_{mN} = 1$), and ‘totally incomplete contracts’ in foreign sourcing (or $\mu_{hS} = \mu_{mS} = 0$).⁷

Our notation associates the relevant degree of contractibility in foreign sourcing with the quality of contractual institutions in the South. In particular, agents engaged in this type of sourcing strategy perceive that the quality of Southern institutions will be the key one determining the extent to which contracts specifying certain aspects of production, including headquarter service provision, will be enforced. This is a strong assumption to make. One would expect that the contractual insecurity of offshoring relationships would be a function of both the Northern and Southern institutions and perhaps their legal similarity, as argued in Chapter 3 when studying the exporting decision. Nevertheless, by using this notation, I seek to stress the notion that the quality of contracting institutions in the country where manufacturing takes place will be an important determinant of the profitability of offshore transactions. I will later appeal to this result when discussing the empirical evidence in Chapter 5. Of course, as I discussed in Chapter 3, parties can seek to insulate a given transaction from weak contract enforcement in ‘the South’ by including choice-of-law and forum-of-law clauses (see Chapter 1). Still, Southern institutions will likely remain crucial in determining the degree to which damages set by international courts of law or arbitrators are enforced.

Notice that equations (4.21) and (4.22) not only illustrate the positive effect of better contract enforcement on profitability, but they also shed light on the differential effect of such an improvement on institutions depending on other features of the environment. For instance, tedious differentiation of these expressions delivers the intuitive result that an increase in the reduced-form aggregate contractibility μ_j for $j = N, S$, will have a disproportionately larger effect on profitability whenever σ is high, that is when

⁷This is easily verified by plugging these values of μ_{hj} and μ_{mj} for $j = N, S$ into the profit functions above and comparing these with (4.1) and (4.9).

the final-good producer faces a particularly competitive environment.⁸ Intuitively, the higher the price elasticity faced by the final-good producer, the costlier will the investment inefficiencies associated with weak contracting prove to be. Apart from this interaction effect, in the Theoretical Appendix, we also show that the elasticity of demand continues to have an unambiguous negative effect on $\Gamma_\ell(\mu_j)$ in this more general framework.

When inspecting how the terms μ_N and μ_S are shaped by the contractibility of the different process of production, it is also evident that improvements in contractibility will interact with the headquarter intensity of production depending on the source of these changes in contractibility. For instance, if improvements in Southern institutions affect disproportionately the contractibility of manufacturing, then this version of the model predicts that these improvements will have a disproportionate effect on profitability in sectors with low headquarter intensity. Conversely, if Southern institutions disproportionately affect the extent to which F will capture the full marginal return from his or her investments in headquarter services, then the model predicts a larger impact of improved contracting on headquarter intensive sectors.

Partial Relationship Specificity

Although relationship-specific investments are pervasive in economic transactions, the assumption of *full* relationship-specificity in our basic model is extreme. Even when particular transactions end up not occurring, suppliers can generally recoup part of the cost of their investment, perhaps by reselling their goods to alternative buyers. Similarly, contractual breaches by suppliers may reduce the overall profitability of headquarter services, but will generally not render them useless. A proper modeling of partial-relationship-specificity would require the introduction of a secondary market for inputs as well as of the negotiations between final-good producers and suppliers in that market, which in turn might depend on the outside options of agents in a tertiary market, and so on (see, for instance, Grossman and Helpman, 2002). The main idea would then be that the lower is the degree of specificity, the larger is the value of inputs in the secondary markets and thus the lower should be the incentive of agents to underinvest. I will next consider a reduced-form version of such a model.

⁸More precisely, in the Theoretical Appendix we establish that $\partial(\partial \ln \Gamma_\ell / \partial \mu_j) / \partial \sigma > 0$.

In particular, assume that there indeed exists a second market for inputs in which the manager M can obtain a price $p_m^s(\varphi)$ for each unit of m , while the final-good producer F anticipates obtaining a monetary return $p_h^s(\varphi)$ per unit of headquarter services. These constitute the outside options for each party at the bargaining stage t_2 we have studied above. I will assume that the agents perceive these secondary market transaction prices $p_h^s(\varphi)$ and $p_m^s(\varphi)$ as unaffected by their actions and in particular their investment levels (more on this below).

Assuming again symmetric Nash bargaining, the payoff obtained by final-good producer will now be given $p_h^s(\varphi)h + \frac{1}{2}(R - p_h^s(\varphi)h - p_m^s(\varphi)m)$, while the supplier will obtain $p_m^s(\varphi)m + \frac{1}{2}(R(\varphi) - p_h^s(\varphi)h - p_m^s(\varphi)m)$. As a result, the levels of investments $h(\varphi)$ and $m(\varphi)$ at t_1 will satisfy the following first-order conditions:

$$\begin{aligned}\frac{1}{2}\left(\frac{\partial R(\varphi)}{\partial h} + p_h^s(\varphi)\right) &= w_N \\ \frac{1}{2}\left(\frac{\partial R(\varphi)}{\partial m} + p_m^s(\varphi)\right) &= \tau w_S.\end{aligned}\tag{4.23}$$

Consider next the determination of the prices $p_h^s(\varphi)$ and $p_m^s(\varphi)$. In a frictionless environment *without* any relationship-specificity, one would expect that this secondary market would provide a thick market for each input and that, in equilibrium, the price commanded by these inputs would correspond to the monetary value of their marginal product. In that case, we would have $p_k^s(\varphi) = \partial R(\varphi) / \partial k$ for $k = h, m$, and the corresponding investments in (4.23) would coincide with the efficient ones under complete contracts. In other words, in the absence of relationship-specificity of investments, weak contract enforcement is irrelevant as the hold-up problem disappears. Conversely, in the other extreme case with full relationship specificity, we instead have $p_h^s(\varphi) = p_m^s(\varphi) = 0$, and the model collapses back to our basic model.⁹

In order to consider environments with partial relationship specificity assume then that the secondary market price commanded by each input is a share $1 - \epsilon$ of the actual value of the marginal product of this input, so that

⁹The setup I am developing is admittedly special in that I am allowing the value of the marginal product of the manufacturer's investment m in the secondary market to be a function of the productivity level φ of the final-good producer with whom it initially contracted. This might reflect the fact that the secondary market is thick for any level of φ or perhaps that the supplier is able to assimilate F 's technology while producing m .

larger values of ϵ are associated with larger degrees of customization or relationship specificity. As is clear from equation (4.23), this then corresponds to the case in which F and M choose investments h and m while anticipating obtaining a share $\beta_h = \beta_m = 1 - \epsilon/2$ of the actual value of the marginal return to these investments.

The rest of the equilibrium of this variant of the model is as in our basic model. Notice in particular that the parties will still find it efficient to reach an agreement at t_2 and thus the secondary market is never used in equilibrium.¹⁰ Overall, the equilibrium can be solved in a manner analogous to the program in (4.19), but with $1 - \epsilon/2$ replacing $1/2$ in the second and third constraints.¹¹

As mentioned before, Antràs and Helpman (2008) solved this program for a general division of revenue (β_h, β_m) , so we can just plug $\beta_h = \beta_m = 1 - \epsilon/2$ into their equilibrium equations (see the Theoretical Appendix). This yields a level of profits associated with domestic sourcing ($\ell = D$ and $j = N$) and foreign sourcing ($\ell = O$ and $j = S$) equal to

$$\pi_\ell = (c_j)^{1-\sigma} B\Gamma_\ell(\mu_j, \epsilon) \varphi^{\sigma-1} - w_N f_\ell,$$

with

$$\Gamma_\ell(\mu_j, \epsilon) = \left(1 + \frac{\frac{\epsilon}{2}}{1 - \frac{\epsilon}{2}} \frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_j)}\right)^{\sigma - (\sigma - 1)(1 - \mu_j)} \left(1 - \frac{\epsilon}{2}\right)^\sigma, \quad (4.24)$$

and where remember that $\mu_j = \eta\mu_{hj} + (1 - \eta)\mu_{mj}$, $c_j = w_N$ when $j = N$ and $c_j = \tau w_S$ when $j = S$. As in the model with full relationship specificity, it continues to be the case that improvements in contractibility are associated with larger values of $\Gamma_\ell(\mu_j, \epsilon)$. Similarly, the negative effect of σ on $\Gamma_\ell(\mu_j, \epsilon)$ and the positive “interaction” effect $\partial(\partial \ln \Gamma_\ell / \partial \mu_j) / \partial \sigma > 0$ continue to apply in this more general environment (see the Theoretical Appendix).

The main new feature of expression (4.24) is that the inefficiencies derived from incomplete contracting are now increasing in the degree of specificity ϵ in the sense that $\Gamma_\ell(\mu_j, \epsilon)$ decreases in ϵ . This intuitive result is not immediate from inspection of equation (4.24), but it can be verified by analyzing the

¹⁰More specifically, given the concavity of the revenue function, we necessarily have that $R(\varphi) > (1 - \epsilon/2) \left(\frac{\partial R(\varphi)}{\partial h} h + \frac{\partial R(\varphi)}{\partial m} m\right)$.

¹¹Despite the fact that F and M do not each receive a share $1 - \epsilon/2$ of revenue, their investments are determined *as if they did*.

partial derivative $\partial \ln \Gamma_\ell(\mu_j, \epsilon) / \partial \epsilon$. The reader is referred to the Theoretical Appendix for the mathematical derivations, which also demonstrate that the cross-partial derivative $\partial (\ln \partial \Gamma_\ell(\mu_j, \epsilon) / \partial \epsilon) / \partial \mu_j$ is positive. In words, the positive effect of higher quality of contracting institutions on firm profitability is predicted to be disproportionately higher in production processes with high degrees of specificity ϵ . Or, put differently, the model seems to be consistent with the fact that countries with weak contracting environments appear to export manufactured goods featuring relatively low levels of specificity, as empirically shown by Nunn (2007). I will further illustrate this result in Chapter 5, when I develop a multi-country version of the model.

Multiple Inputs and Multilateral Contracting

So far, I have focused on situations in which F is concerned only with the provision of one input. In modern manufacturing processes final-good producers instead combine intermediate inputs provided by various suppliers. I will next return to the version of the global sourcing model introduced in Chapter 2, in which the manufacturing stage of production entails the procurement of a continuum of measure one of inputs indexed by v , all produced simultaneously at t_1 . Assuming that the services from these stages are imperfectly substitutable with each other with a constant and symmetric elasticity of substitution equal to $\sigma_\rho \equiv 1/(1-\rho)$, we can now write the production function as

$$q(\varphi) = \varphi \left(\frac{h}{\eta} \right)^\eta \left(\frac{\left[\int_0^1 m(v)^\rho dv \right]^{1/\rho}}{1-\eta} \right)^{1-\eta}. \quad (4.25)$$

Note that if one interprets $q(\varphi)$ as the *quality-adjusted* volume of output, this formulation is perfectly consistent with the notion that, from an engineering point of view, all stages might be essential. For example, producing a car requires four wheels, two headlights, one steering wheel, and so on, but the value of this car for consumers will typically depend on the services obtained from these different components, with a high quality in certain parts potentially making up for inferior quality in others.

I will assume that the continuum of inputs are not only symmetric in technology but are also produced with the same marginal cost in a given location. Manufacturing also continues to entail fixed costs that depend on the location of this activity but I assume that these fixed costs are indepen-

dent of the number of inputs produced in a location. For this reason, it is natural to focus on symmetric equilibria in which all manufacturing inputs are produced in the same location. In Chapter 5 we will consider a more interesting and realistic environment in which firms source inputs from various locations.

Headquarter service provision continues to be controlled by the final good producer, agent F . To obtain the various intermediate inputs, F now needs to contract with a continuum of managers $M(v)$ each controlling one input. If all the aspects associated with the production of the different inputs could be specified in an enforceable manner in an initial contract, then it is straightforward to show that the resulting profit functions for the final-good producer associated with domestic sourcing and offshoring would be exactly identical to those of the single manufacturing input model. These profits flows are given by equations (4.1) and (4.2). Note that, given our symmetry assumptions and complete contracting, these profit flows are independent of the value of the input substitution parameter ρ . As we will next demonstrate, this parameter will play a much more relevant role in the presence of contractual frictions.

Consider now the case of partial contractibility introduced above, in which some of the characteristics of production are contractible, while others are not. Specifically, headquarter services h and each manufacturing input $m(v)$ are a Cobb-Douglas aggregate of the services of a continuum of measure one of activities, as in equations (4.17) and (4.18), and only a share μ_{hj} and μ_{mj} of those activities are contractible when manufacturing takes place in country $j = N, S$. Note that, for simplicity, the share μ_{mj} is common for all inputs v . The terms of exchange related to the noncontractible activities are only decided at t_2 , after they have been performed but not yet embodied into production. The threat point for each party in the negotiations at t_2 is to withhold the services from those activities.

The key novel feature of this richer environment is that the ex-post negotiations at t_2 are now multilateral, rather than bilateral. How should one model these ex-post bargaining? One possible way would be to apply Nash bargaining to our multilateral setup with each agent obtaining their outside option plus a share of the difference between joint surplus under cooperation and the sum of outside options (see, for instance, Osborne and Rubinstein, 1990, p. 23). With zero outside options, and hence full relationship specificity, this would amount to all agents obtaining a constant share of revenues. This would lead, however, to a situation analogous to a moral hazard in teams

problem (see Holmstrom, 1982) with an arbitrarily large number of agents. In such a case, the agents would have no incentive to invest in noncontractible tasks and revenue would be zero. In sum, a minimal amount of contractual frictions would be sufficient to drive production efficiency to zero.

This extreme result is in part due to our Cobb-Douglas assumptions in (4.17) and (4.18) but it also reflects the limitations of the Nash bargaining solution in multi-agent environments. In particular, this solution does not allow for situations of *partial cooperation* in which even if one supplier rejects an agreement, the other agents are still allowed to cooperate with each other and obtain some surplus. For this reason, in multilateral bargaining setups it is customary to adopt the Shapley value as the solution concept characterizing the equilibrium of these negotiations. In a bargaining game with a finite number of players, each player's Shapley value is the average of her contributions to all coalitions that consist of players ordered below her in all feasible permutations.¹²

A complication arises from the fact that, in our environment, we have a continuum of agents bargaining over surplus. Acemoglu, Antràs and Helpman (2007) resolve this issue by considering a discrete-player version of the game and computing the *asymptotic* Shapley value of Aumann and Shapley (1974). I will next develop an alternative, heuristic derivation of this Shapley value.

First note that agent F is an *essential* player in the bargaining game and thus a supplier $M(v)$'s marginal contribution is equal to zero when being added to a coalition that does not include the firm. When that coalition does include the firm and a measure n of suppliers, the marginal contribution of supplier v is equal to $\Delta R(v, n) = \partial R(\varphi, n) / \partial n$, where $R(\varphi, n) = p(q(\varphi, n)) q(\varphi, n)$ and $q(\varphi, n)$ is as in (4.25) but with the integral running up to n rather than 1. Using Leibniz' rule and invoking symmetry,

¹²More formally, in a game with M players, let $g = \{g(1), \dots, g(M)\}$ be a permutation of $1, 2, \dots, M$, and let $z_g^j = \{j' \mid g(j) > g(j')\}$ be the set of players ordered below j in the permutation g . Denoting by G the set of feasible permutations, and by $v : G \rightarrow \mathbb{R}$ the value (or surplus generated) of the coalition consisting of any subset of the M players, the Shapley value of player j is then

$$s_j = \frac{1}{(M+1)!} \sum_{g \in G} [v(z_g^j \cup j) - v(z_g^j)].$$

this marginal contribution can be succinctly written as

$$\Delta R(v, n) = \frac{(\sigma - 1)(1 - \eta)}{\sigma \rho} R(\varphi) \left(\frac{m_n(v)}{m_n(-v)} \right)^\rho n^{\frac{(\sigma-1)}{\sigma\rho} - 1},$$

where $m(-v)$ represents the (symmetric) investments of all suppliers other than v . Note that in deriving this expression we also have imposed, without loss of generality, a symmetric choice for contractible manufacturing tasks, or $m_c(v) = m_c$ for all v .

The Shapley value of $M(v)$ is the average of her marginal contributions to coalitions that consist of players ordered below her in all feasible orderings. A supplier that has a measure n of players ordered below her has a marginal contribution of $\Delta R(v, n)$ if the firm is ordered below her – which occurs with probability n –, and 0 otherwise. Averaging over all possible orderings of the players and using the above formula for $\Delta R(v, n)$ we obtain the following payoff for supplier $M(v)$:

$$P_m(v) = \int_0^1 n \Delta R(v', n) dv' = \frac{(\sigma - 1)(1 - \eta)}{(\sigma - 1)(1 - \eta) + \sigma \rho} R(\varphi) \left(\frac{m_n(v)}{m_n(-v)} \right)^\rho. \quad (4.26)$$

A number of features of (4.26) are worth noting. First, in equilibrium, all suppliers invest equally in all the noncontractible activities, and thus each receives a share $(\sigma - 1)(1 - \eta) / (\sigma - 1 + \sigma \rho)$ of revenue, leaving F with the residual share $\sigma \rho / ((\sigma - 1)(1 - \eta) + \sigma \rho)$. The bargaining power of the firm is thus naturally increasing in the substitutability of inputs as governed by ρ , since the suppliers' bargaining threats are less effective in that case. Second, and although in equilibrium suppliers end up with an equal share of sale revenue, equation (4.26) indicates that suppliers perceive their non-contractible investments to have a non-negligible (i.e., measurable) effect on their payoffs, and thus the moral-hazard-in-teams, zero-investment result mentioned above does not apply here. Third, the degree of substitutability ρ crucially impacts the marginal return to suppliers' investments by shaping the degree to which increases in the investments of a given supplier affect output. Intuitively, when inputs are highly complementary (low ρ), the marginal return to increasing the production of one input v while holding the rest fixed is particularly low.

Having solved for the division of surplus at t_2 , the rest of the equilibrium is as in previous models. In particular, the program is analogous to that in (4.19), but with $\beta_h = \sigma \rho / ((\sigma - 1)(1 - \eta) + \sigma \rho)$ replacing $1/2$ in

F 's choice of h and with each supplier $M(v)$ choosing $m_n(v)$ to maximize $P_m(v) - c_j(1 - \mu_{mj})m_n(v)$. In equilibrium, the latter choice is isomorphic to that of a single supplier choosing a *common* m_n for all v to maximize $\beta_m R(\varphi) - c_j(1 - \mu_{mj})m_n$ with $\beta_m = \rho\sigma / ((\sigma - 1)(1 - \eta) + \sigma\rho)$.¹³ Thus we can apply the general formula in Antràs and Helpman (2008) (see the Theoretical Appendix) to express the level of contractual frictions $\Gamma_\ell(\mu_j, \rho)$ associated with manufacturing taking place in country j as

$$\Gamma_\ell(\mu_j, \rho) = \left(1 + \frac{1}{\rho} \frac{(\sigma - 1)(1 - \eta)}{\sigma - (\sigma - 1)(1 - \mu_j)}\right)^{\sigma - (\sigma - 1)(1 - \mu_j)} \left(\frac{\rho\sigma}{\rho\sigma + (\sigma - 1)(1 - \eta)}\right)^\sigma, \quad (4.27)$$

with again $\mu_j = \eta\mu_{hj} + (1 - \eta)\mu_{mj}$ (where $j = N$ when $\ell = D$ and $j = S$ when $\ell = O$).

The generality of the results in Antràs and Helpman (2008) allows us to conclude, without having to differentiate this expression, that $\Gamma_\ell(\mu_j, \rho)$ is again increasing in the degree of contractibility μ_j and decreasing in the elasticity of demand σ .¹⁴ In addition, the positive effect of contract enforcement affects high-substitutability sectors disproportionately, or $\partial(\partial \ln \Gamma_\ell / \partial \mu_j) / \partial \sigma > 0$.

The main novelty of equation (4.27) is that the degree of input substitutability is now a key determinant of the extent to which contractual frictions depress the profitability of production. Straightforward differentiation demonstrates that $\Gamma_\ell(\mu_j, \rho)$ is increasing in ρ , and thus contractual frictions are lower, the more substitutable inputs are. Intuitively, investments tend to be less distorted in that case because a higher level of ρ (i) provides more ex-post surplus to F thus enhancing the investments in headquarter services by F , and (ii) increases the sensitivity of suppliers' ex-post payoffs to their own investments. Naturally, a high ρ also reduces the share of ex-post surplus accruing to suppliers, but given the functional forms in the model, this is a dominated effect.

Differentiation of (4.27) also demonstrates (see the Theoretical Appendix) that $\partial(\partial \ln \Gamma_\ell(\mu_j, \rho) / \partial \mu_j) / \partial \rho < 0$ and thus the effect of an improvement

¹³This follows from noting that β_m must be such that $\rho P_m(v) = \beta_m \frac{(\sigma - 1)(1 - \eta)}{\sigma} R(\varphi)$ whenever $m_n(v) = m_n$ for all v .

¹⁴This latter comparative static result would appear to be complicated by the fact that the bargaining weights β_h and β_m are now endogenous and a function of σ . But since $\Gamma_\ell(\mu_j, \rho)$ in (8.9) is increasing in β_h and β_m , and each of these two shares is decreasing in σ , this does not affect the sign of the derivative $\partial \ln \Gamma_\ell / \partial \sigma$.

in contractual institutions has a differentially larger effect in sectors featuring higher input complementarities. The model thus suggests that, other things equal, foreign sourcing to countries with particularly weak contract enforcement should be more prevalent in sectors with higher substitutability between inputs. We will further formalize and test this result in Chapter 5.

Sequential Production

The variant of the model with multiple suppliers that I have developed above assumes that all stages of production are performed simultaneously. In real-life manufacturing processes, there is often a natural sequencing of stages. First, raw materials are converted into basic components, which are next combined with other components to produce more complicated inputs, before themselves being assembled into final goods. Antràs and Chor (2013) develop a sequential production variant of the model with a continuum of inputs we have just studied. The key new feature of their analysis is that the relationship-specific investments made by suppliers in upstream stages can affect the incentives of suppliers involved in downstream stages thereby generating investment inefficiencies that vary systematically along the value chain.

The model developed by Antràs and Chor (2013) turns out to be very tractable but some of the details of the analysis are somewhat intricate, so I refer the reader to the paper and its Supplemental Appendix for many details. Antràs and Chor (2013) assume a production technology analogous to (4.25) but with $v \in [0, 1]$ indexing the position of an input in the value chain, with a larger v corresponding to stages further downstream (closer to the final end product). Although they develop extensions with headquarter services and partial contractibility, I will focus below on their benchmark model in which $\eta = 0$ and in which all investments are noncontractible.

The final-good producer F plays two roles in the model. On the one hand, it is in charge of assembling the measure one of sequentially produced inputs into a final good valued by consumers. Second, it sequentially negotiates with suppliers once their stage input has been produced and the firm has had a chance to inspect it. It is simplest to consider the case in which this negotiation at stage v is treated independently from the bilateral negotiations that take place at other stages (see Antràs and Chor, 2013, for alternative formulations). Because each intermediate input v is assumed compatible only with the firm's output, the supplier's outside option at the bargaining stage

is 0. Hence, the quasi-rents over which the firm and the supplier negotiate are given by the incremental contribution to total revenue generated by supplier v at that stage. In light of (4.20) and (4.25), this incremental contribution is given by

$$\begin{aligned} \Delta R(v) &= \frac{(\sigma - 1)}{\sigma \rho} B^{1/\sigma} \sigma (\sigma - 1)^{(\sigma-1)/\sigma} \varphi^{(\sigma-1)/\sigma} \\ &\quad \times \left(\int_0^v m(u)^{(\sigma_\rho-1)/\sigma_\rho} du \right)^{(1-\sigma_\rho/\sigma)/(\sigma_\rho-1)} m(v)^\rho, \end{aligned} \quad (4.28)$$

where remember that $\sigma_\rho = 1/(1 - \rho)$. Assume that the share of these quasi-rents accruing to F are given by $\beta(v)$. Below, I will allow this share to be affected by the location of manufacturing production.

Notice that if $\sigma > \sigma_\rho$, then the investment choices of suppliers are *sequential complements* in the sense that higher investment levels by prior suppliers increase the marginal return of supplier v 's own investment $m(v)$. Conversely, if $\sigma < \sigma_\rho$, investment choices are *sequential substitutes* because high values of upstream investments reduce the marginal return to investing in $m(v)$. Because the supplier at position v chooses $m(v)$ to maximize $(1 - \beta(v)) \Delta R(v) - c_j m(v)$, equation (4.28) illustrates the trickle down effect that upstream investment inefficiencies can have on downstream stages.

Exploiting the recursive structure of the model, Antràs and Chor (2013) show that if agent F is able to use ex-ante transfers to extract all surplus from suppliers, then the overall profits obtained by a final-good producer with productivity φ when all inputs are produced under a marginal cost equal to j are given by

$$\pi_j = (c_j)^{1-\sigma} B \Gamma_j (\{\beta(v)\}_{v=0}^1) \varphi^{\sigma-1}$$

where

$$\begin{aligned} \Gamma_\ell (\{\beta(v)\}_{v=0}^1) &= \frac{(\sigma - 1)}{(\sigma_\rho - 1)} \left(\frac{\sigma_\rho}{\sigma} \right)^{\frac{\sigma - \sigma_\rho}{\sigma_\rho - 1}} \int_0^1 \left\{ \left(\frac{\sigma_\rho}{1 - \beta(v)} - (\sigma_\rho - 1) \right) \right. \\ &\quad \left. \times (1 - \beta(v))^{\sigma_\rho} \left[\int_0^v (1 - \beta(u))^{\sigma_\rho - 1} du \right]^{\frac{\sigma - \sigma_\rho}{\sigma_\rho - 1}} \right\} dv. \end{aligned} \quad (4.29)$$

In the case of a symmetric bargaining power at all stages, so $\beta(v) = \beta$ for all v , equation (4.29) reduces to

$$\Gamma_\ell (\{\beta(v)\}_{v=0}^1) = \left(\frac{\sigma_\rho}{\sigma} \right)^{\frac{\sigma - \sigma_\rho}{\sigma_\rho - 1}} \left(\frac{\sigma_\rho}{1 - \beta} - (\sigma_\rho - 1) \right) (1 - \beta)^\sigma.$$

This expression in turns collapses to the single-supplier index of contractual frictions in equation (4.14) when $\sigma = \sigma_\rho$ (and $\eta = 0$, of course). This is intuitive since in that knife-edge case, the payoff to a supplier is independent of other suppliers' investments, and the trickle-down effects mentioned above become irrelevant.

More interesting implications from the modeling of sequential production can be obtained when allowing the bargaining share $\beta(v)$ to vary along the value chain and across manufacturing locations. To build intuition, it is instructive to consider first the case in which the (infinite-dimensional) vector of $\beta(v)$'s is chosen to maximize F 's profits, and thus $\Gamma_j(\{\beta(v)\}_{v=0}^1)$. Antràs and Chor (2013) show that this seemingly complicated problem can be reduced to a standard calculus of variation problem which delivers the surprisingly simple Euler-Lagrange condition

$$\frac{\partial \beta^*(v)}{\partial v} = \frac{1 - \sigma_\rho / \sigma}{\sigma_\rho - 1} v^{\frac{\sigma_\rho(\sigma-1)}{(\sigma_\rho-1)\sigma}}.$$

The key implication of this expression is that the relative size of the input and final-good elasticities of substitution σ_ρ and σ governs whether the incentive for F to retain a larger surplus share increases or decreases along the value chain. Intuitively, when σ is high relative to σ_ρ , investments are sequential complements, and high upstream values of $\beta(v)$ are particularly costly since they reduce the incentives to invest not only of these early suppliers but also of all suppliers downstream. Conversely, when σ is small relative to σ_ρ , investments are sequential substitutes, and low values of $\beta(v)$ in upstream stages are now relatively detrimental, since they reduce the incentives to invest for downstream suppliers, who are already underinvesting to begin with.¹⁵

This result has interesting implications for the choice between domestic and foreign sourcing whenever these sourcing strategies are associated with

¹⁵Imposing two boundary conditions on the Euler-Lagrange equation, the optimal stage- v bargaining share can in fact be solved in closed form, and is given by $\beta^*(v) = 1 - v^{\frac{\sigma_\rho/\sigma-1}{\sigma_\rho-1}}$. In the sequential complements case ($\sigma > \sigma_\rho$) this implies $\beta^*(v) < 0$ for all v , so F has an incentive to allocate to suppliers more than their entire incremental contribution. This extreme result does not apply when F cannot extract all surplus from suppliers via ex-ante lump-sum transfer or when the model includes headquarter service provision (see Antràs and Chor, 2013, for details). Importantly, in those cases, it continues to be the case that the sign of $\partial \beta^*(v) / \partial v$ is determined by the relative size of σ and σ_ρ , and $\partial \beta^*(v) / \partial v > 0$ whenever $\sigma > \sigma_\rho$.

different levels of contract enforcement or with different bargaining shares for F in its negotiations with suppliers. To see this, consider first the case in which contracting in domestic Northern transactions is complete, while foreign sourcing is associated with totally incomplete contracting. Our results above then suggest that, in the sequential complements case ($\sigma > \sigma_\rho$), foreign sourcing is particularly unappealing in upstream stages. Thus, if domestic and foreign sourcing coexist along the value chain, then only relatively downstream inputs will be offshored.¹⁶ Conversely, in the sequential substitutes case, ($\sigma < \sigma_\rho$) one would expect relatively upstream stages to be offshored. In sum, the model predicts that the ‘upstreamness’ of an input should be a relevant determinant of the extent to which it is procured from foreign suppliers, with the sign of that dependence being crucially shaped by the relative size of σ and σ_ρ .

Note, however, that very different results might arise if domestic and foreign sourcing do not differ significantly in their contractibility, but are associated with F obtaining a higher share of surplus under domestic sourcing than under offshoring, i.e., $\beta_D(v) > \beta_O(v)$, as suggested for instance by Antràs and Helpman (2008). In such a case, offshoring would be relatively more appealing in upstream stages in the sequential complements case, and relatively more appealing in downstream stages in the sequential substitutes case. We will explore the empirical relevance of these different scenarios in the next chapter.

Summary and Next Steps

This chapter has explored the determinants of the global sourcing decisions of firms in the presence of incomplete contracting frictions in vertical relationships. The different variants of our global sourcing model have delivered a rich set of comparative statics and have also provided tools for testing these predictions with data on intermediate input trade. In the next chapter, I will test the empirical success of the model with detailed data on U.S. imports by product and source country. In the process, and given the cross-country dimension of the data, it will prove necessary to develop a multi-country version of the model.

¹⁶This result can be proved in a manner analogous to Proposition 2 in Antràs and Chor (2013).

Chapter 5

Contracts and Sourcing: Evidence

Earlier 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 the last chapter, 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 observed 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 (2.22) in Chapter 4, which solves for the share of spending on imported manufacturing inputs over total manufacturing input purchases in a particular industry. The equation relates 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 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 from 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

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.

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 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 those 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.

As an example, consider two manufacturing industries with very different levels of Numm’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’), Numm (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




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

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

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.

States. It is important to emphasize that these tests are well grounded on 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.

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 7, 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 is 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 result 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

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

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

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

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

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

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.

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 reproduce 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 is 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

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

observations when computing the averages in Table 5.1.

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

As is clear from Table 5.1, most sectors on the left-panel of the table produce goods that are relatively difficult or expensive to ship across borders. Conversely, most sectors on 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 the cross-section and time-series variation in 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

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

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 HS ten-digit products computed by Broda and Weinstein (2006). All 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 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 7 are available for download at <http://scholar.harvard.edu/antras/scholar.harvard.edu/antras/pages/books>.

Having discussed the variables included and their sources, we can now turn to describe 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.

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.

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, in a recent paper, Mutreja (2013) documents that cross-country dispersion in capital equipment is much larger than the dispersion in capital structures and that the ratio of equipment to structures is much higher in rich than in poor countries.⁵ The effects of productivity dispersion and the elasticity of demand on offshoring shares are analyzed in Column (4) of Table 5.2. The coefficient on these variables are small in magnitude and are imprecisely estimated.

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 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

⁵Using 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 autos 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.

at producing goods that happen to be cheap to transport or feature a high equipment capital intensity. By introducing source country 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, but in my (weak) defense, this approach is fairly standard in the literature (more on this below).

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 identifies only the sec-

tor to which the good being transacted belongs to. 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 this '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. Column (1) of Table 5.3 reports cross-industry results analogous to those in column (4) of Table 5.2 but with the IO2002 classifica-

tion 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.⁶

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.⁷

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 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.

⁶The 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 $[0, 1]$.

⁷The 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.

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 Treffer (2013) 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.⁸

⁸I follow Wright (2014) in also removing industries that purely process raw materials

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.⁹ 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.

so we can more comfortably treat inputs as differentiated.

⁹A 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 Refined 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) for which 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.¹⁰

¹⁰The full list of 18 countries includes Iceland, Italy, Finland, Liechtenstein, Switzer-

As is clear from column (4) of Table 5.3, this Nunn-Trefler correction has a qualitatively similar effect than 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 on purely cross-industry variation, it is advisable to fit the data through this straighter jacket. As in 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.¹¹ 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).¹² 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

land, Sweden, Taiwan, Belgium, Bermuda, Norway, Denmark, Korea, Japan, Spain, Israel, Austria, France, and Germany.

¹¹To 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.

¹²It is worth stressing that the standard errors in these regressions are clustered conservatively at the industry level.

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 evidence 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, and symmetric bargaining. 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

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

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.¹³ 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

¹³To 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.

contractibility and their ratio is ambiguous.¹⁴ 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.

This 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 purchases 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 de-

¹⁴For 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 .

pendence 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 firm's 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. Curiously, 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 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 (2007)'s 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 an 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 association.¹⁵

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

¹⁵As 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.

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 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.¹⁶

¹⁶The interested reader can consult the whole set of regression coefficients by accessing the data and programs available online at <http://scholar.harvard.edu/antras/scholar.harvard.edu/antras/pages/books>.

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	$\simeq 0.15$	$\simeq 0.19$	$\simeq 0.20$	0.168	0.200	0.199

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

In the first three columns of Table 5.6, I report the results of adding these new nine 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 of them 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 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.

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 an 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 better 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 the product 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 degree 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 these are allowed to vary across countries. Furthermore, while a_{hi} is a technological parameter common across firms based in i , the manufacturing productivity parameters 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

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

These Fréchet draws are assumed to be independent across firms and locations, and notice also that they are assumed to be 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 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

$$\chi_{ij} = \frac{T_j (\tau_{ij} w_j)^{-\theta}}{\sum_{k \in J} T_k (\tau_{ik} w_k)^{-\theta}}. \quad (5.6)$$

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

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

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 (4.27))

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

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

$$\chi_{ij} = \frac{T_j \left(\tau_{ij} w_j \Gamma_{ij}^{1/(1-\eta)(1-\sigma)} \right)^{-\theta}}{\sum_{j \in J} T_k \left(\tau_{ik} w_k \Gamma_{ik}^{1/(1-\eta)(1-\sigma)} \right)^{-\theta}}. \quad (5.9)$$

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 relative lower share of the intermediate inputs purchased by firms from country i .¹⁷

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

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

where I will write Γ_{ijv} as

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

Notice 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. On the other hand, 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}) .

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_{ijv} as country (or country-pair) characteristics. In some specifications, I will however allow the effect of financial

¹⁷As 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.

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 fixed effects helped isolate the effect of sectoral-level characteristics on offshoring shares. Indeed, if the inclusion of country fixed completely partialled out the effect of any country level variable on offshoring shares χ_{ijv} , then indeed 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 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

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

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 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 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

$$\ln \chi_{ijv} = \alpha_{iv} + \alpha_{ij} + \beta Z_{ij} z_v + \gamma \ln \tau_{ijv} + \delta \ln \Gamma_{ijv} + \varepsilon_{ijv}. \quad (5.13)$$

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 fixed effects. The effect of technological or factor endowments differences as sources of comparative advantage is captured by the interaction terms $Z_{ij}z_v$, 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 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.¹⁸

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).¹⁹ 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 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

¹⁸See Nunn and Trefler (2013a) for a nice overview of the literature on trade and institutions.

¹⁹In 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.

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 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.

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

Dep. Var.: log(Imports)	(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^2	$\simeq 0.610$	$\simeq 0.607$	0.622	0.623	0.637

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

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.

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 Trefler, 2013*a*). 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 of 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 interac-

tions, 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 , ϵ_v , and η_v in the same manner that I did so 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. 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 5.8 report the results of running 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 almost twice as large in absolute terms and statistically significant at the 5% level.

Table 5.8. Testing the Global Sourcing Model

Dep. Var.: log(Imports)	(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. Int. \times Credit/GDP			0.074**	0.045**	0.044**	0.045**
			(0.007)	(0.009)	(0.012)	(0.012)
Headq. Inten. \times Rule			0.093**	0.050**	0.050**	0.047**
			(0.010)	(0.012)	(0.013)	(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^2	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.

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.²⁰ Interestingly, each of the five institutional interactions appears to be highly significant and with a sign consistent

²⁰This 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

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 ($\rho_v\mu_j$), which remains negative but is now statistically indistinguishable from zero.

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 ($\sigma_v\mu_j$), 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

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.

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/scholar.harvard.edu/antras/pages/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 country and industry 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 key challenge in the empirical analyses performed in this chapter is that the key industry characteristics (specificity, input substitutability, demand 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, so I presume that 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.

Chapter 6

The Internalization Decision

Chapter 7

Internalization: Empirical Evidence

Chapter 8

Some Implications for Policy

A Theoretical Appendix

The proof of a few theoretical results in the main text of this book were only sketched or simply omitted to enhance the flow of text. In this Appendix, I provide the details of these proofs or refer the reader to the relevant papers where the proofs can be found.

A.1 Optimal Sourcing Strategy in the Multi-Country Global Sourcing Model in Chapter 2

In this Appendix, I formally prove some statements related to the characterization of the optimal sourcing strategy in the multi-country global sourcing Model in Chapter 2.

Remember that the problem of choosing an optimal sourcing strategy is given by

$$\begin{aligned} \max_{\mathcal{J}_i(\varphi)} \pi_i(\varphi, \mathcal{J}_i(\varphi)) &= (a_{hi}w_i)^{-\eta(\sigma-1)} \left(\gamma \sum_{k \in \mathcal{J}_i(\varphi)} T_k (\tau_{ik}w_k)^{-\theta} \right)^{(\sigma-1)(1-\eta)/\theta} B\varphi^{\sigma-1} \\ &\quad - w_i \sum_{k \in \mathcal{J}_i(\varphi)} f_{ik}. \end{aligned}$$

With a discrete number of locations, we can rewrite the problem as follows:

$$\begin{aligned} \max_{I_1, I_2, \dots, I_J \in \{0,1\}^J} \pi_i(\varphi, I_1, I_2, \dots, I_J) &= \left(\sum_{k=1}^J I_k T_k (\tau_{ik}w_k)^{-\theta} \right)^{(\sigma-1)(1-\eta)/\theta} \tilde{B}\varphi^{\sigma-1} \\ &\quad - w_i \sum_{k=1}^J I_k f_{ik}, \end{aligned} \tag{8.1}$$

where $\tilde{B} = (a_{hi}w_i)^{-\eta(\sigma-1)} \gamma^{(\sigma-1)(1-\eta)} B$. The dummy variable I_j thus takes a value of 1 when $j \in \mathcal{J}_i(\varphi)$ and 0 otherwise.

The key thing to note is that, provided that $(\sigma - 1)(1 - \eta) > \theta$, the modified objective function in (8.1) features increasing differences in (I_j, I_k) for $j, k \in \{1, \dots, J\}$ such that $j \neq k$, and also features increasing differences in (I_j, φ) for any $j \in \{1, \dots, J\}$. Invoking standard results in monotone comparative statics, we can then conclude that for $\varphi_1 \geq \varphi_0$, we must have $(I_1^*(\varphi_1), I_2^*(\varphi_1), \dots, I_J^*(\varphi_1)) \geq (I_1^*(\varphi_0), I_2^*(\varphi_0), \dots, I_J^*(\varphi_0))$. Naturally, this rules out a situation in which $I_j^*(\varphi_1) = 0$ but $I_j^*(\varphi_0) = 1$, and thus we can conclude that $\mathcal{J}_i(\varphi_0) \subseteq \mathcal{J}_i(\varphi_1)$ for $\varphi_1 \geq \varphi_0$.

A.2 Comparative Statics of the Global Sourcing Model in Chapter 4

In this Appendix, I will provide formal proofs for some comparative statics mentioned in Chapters 4 and 5. When a result has been proved in an existing paper, I will simply refer the reader to that paper.

Derivation of a General Formula for the Offshoring Share

In the first part of Chapter 5, I studied the determinants of the cross-section of offshoring shares. In Chapter 4 I derived a formula for this share but under the strong assumptions of complete contracting in the North, ‘totally’ incomplete contracting in the South, a single input, and symmetric bargaining. In Chapter 5, I appealed to a general formula that applies to all the extensions of the two-country model developed in Chapter 4. Let me now provide more details on that derivation.

As explained in Chapter 4, the share of foreign input purchases in total input purchases typically depends on how these inputs are priced in the presence of incomplete contracting and renegotiation. Below, I stick to the assumption in the main text that the ratio of input expenditures to sale revenue is common for firms sourcing domestically and offshoring. As a result, the offshoring share is identical to the fraction of industry sales captured by firms offshoring intermediate inputs. With a constant price elasticity of demand $\sigma > 1$, firm revenues are in turn a multiple σ of operating profits. Operating profits are in turn equal to overall profits plus fixed costs, or

$$\begin{aligned}\pi_D(\varphi) + f_D w_N &= (w_N)^{1-\sigma} B \Gamma_D \varphi^{\sigma-1} \\ \pi_O(\varphi) + f_O w_N &= \left((w_N)^\eta (\tau w_S)^{1-\eta} \right)^{1-\sigma} B \Gamma_O \varphi^{\sigma-1}.\end{aligned}$$

Assuming selection into offshoring – i.e., condition (2.21) in Chapter 2, we can define the thresholds $\tilde{\varphi}_O > \tilde{\varphi}_D$ satisfying $\pi_D(\tilde{\varphi}_D) = 0$ and $\pi_O(\tilde{\varphi}_O) = \pi_D(\tilde{\varphi}_O)$. It

is straightforward to verify that

$$\frac{\tilde{\varphi}_O}{\tilde{\varphi}_D} = \left[\frac{f_O/f_D - 1}{\frac{\Gamma_O}{\Gamma_D} \left(\frac{w_N}{\tau w_S} \right)^{(1-\eta)(\sigma-1)} - 1} \right]^{1/(\sigma-1)}. \quad (8.2)$$

The share of revenues (and of input purchases) accounted for by offshoring firms is then given by

$$\Upsilon_O = \frac{\int_{\tilde{\varphi}_O}^{\infty} \left((w_N)^\eta (\tau w_S)^{1-\eta} \right)^{1-\sigma} B \Gamma_O \varphi^{\sigma-1} dG(\varphi)}{\int_{\tilde{\varphi}_D}^{\tilde{\varphi}_O} (w_N)^{1-\sigma} B \Gamma_D \varphi^{\sigma-1} dG(\varphi) + \int_{\tilde{\varphi}_O}^{\infty} \left((w_N)^\eta (\tau w_S)^{1-\eta} \right)^{1-\sigma} B \Gamma_O \varphi^{\sigma-1} dG(\varphi)}.$$

Assuming a Pareto distribution of productivity – i.e., $G(\varphi) = 1 - (\underline{\varphi}/\varphi)^\kappa$ for $\varphi \geq \underline{\varphi} > 0$ –, this expression further simplifies to

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

where $\tilde{\varphi}_O/\tilde{\varphi}_D$ is given in 8.2. This corresponds to the general offshoring share equation (5.2) in Chapter 5. This formula is identical to the one applying in the complete-contracting case except for the term Γ_O/Γ_D .

With this expression in hand, we can next turn to the study of comparative statics in the different variants of the global sourcing model developed in Chapter 4. Below I will focus on how the different parameters of the model shape the ratio Γ_O/Γ_D , which differs across variants of the model. As argued in the main text, leaving aside this term Γ_O/Γ_D , the share Υ_O is increasing in w_N/w_S and σ , and decreasing in τ , f_O/f_D , κ and η (these results are straightforward to prove by simple differentiation (making use of $\kappa \geq \sigma - 1$)).

Symmetric Nash Bargaining Model

Consider first the basic model with complete contracting in the North, ‘totally’ incomplete contracting in the South, a single input, and symmetric bargaining. This implies $\Gamma_D = 1$ and thus (see eq. (4.10)).

$$\frac{\Gamma_O}{\Gamma_D} = (\sigma + 1) \left(\frac{1}{2} \right)^\sigma.$$

But note that

$$\frac{\partial (\Gamma_O/\Gamma_D)}{\partial \sigma} = - \left(\frac{1}{2} \right)^\sigma ((1 + \sigma) \ln 2 - 1) < 0,$$

and thus the offshoring share is lower in higher elasticity sectors on account of the effect of contractual frictions. This effect is of the opposite sign to the ‘standard’ one operating in the complete-contracting case, and thus the overall effect of σ on the offshoring share Υ_O is ambiguous.

Generalized Nash Bargaining Model

Let us now turn to the basic model with generalized Nash bargaining. Again we have $\Gamma_D = 1$ and thus

$$\frac{\Gamma_O}{\Gamma_D} = \Gamma_\beta \equiv (\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta))) \left(\beta^\eta (1 - \beta)^{1-\eta} \right)^{\sigma-1},$$

as indicated by equation (4.14).

As mentioned in the main text, the effects of β and η on Γ_β are ambiguous and interact with each other. More specifically, we next show that Γ_β is decreasing in η when $\beta < 1/2$, while it is increasing in η when $\beta > 1/2$. To see this, first note that

$$\left. \frac{\partial \ln \Gamma_\beta}{\partial \eta} \right|_{\eta=0} = (\sigma - 1) \left(\frac{1 - 2\beta}{1 - \beta + \sigma\beta} + \ln \left(\frac{\beta}{1 - \beta} \right) \right)$$

and

$$\left. \frac{\partial \ln \Gamma_\beta}{\partial \eta} \right|_{\eta=1} = (\sigma - 1) \left(\frac{1 - 2\beta}{\sigma(1 - \beta) + \beta} + \ln \left(\frac{\beta}{1 - \beta} \right) \right).$$

It is not hard to show that each of these two expressions is negative for $\beta < 1/2$ and positive for $\beta > 1/2$. In particular, one can use $1 - x + \ln x \leq 0$ and $\ln 1/x - (1 - x) \geq 0$, with $x = \beta/(1 - \beta)$ to rewrite these expressions in a way that makes this obvious by inspection.¹ Next, notice that

$$\frac{\partial^2 \ln \Gamma_\beta}{\partial \eta^2} = - \frac{(\sigma - 1)^2 (1 - 2\beta)^2}{(\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta)))^2} < 0.$$

¹For completeness, note that $\frac{1-2\beta}{1-\beta+\sigma\beta} + \ln\left(\frac{\beta}{1-\beta}\right) = (2\beta-1) \frac{\sigma\beta}{(1-\beta)(1-\beta+\sigma\beta)} + \left(1 - \frac{\beta}{1-\beta}\right) + \ln\left(\frac{\beta}{1-\beta}\right) = (2\beta-1) \frac{(\sigma-2)\beta+1}{\beta(\sigma-1)\beta+1} + \ln\left(\frac{\beta}{1-\beta}\right) - \left(1 - \frac{1-\beta}{\beta}\right)$ and $\frac{1-2\beta}{\sigma(1-\beta)+\beta} + \ln\left(\frac{\beta}{1-\beta}\right) = (2\beta-1) \frac{\sigma-1-(\sigma-2)\beta}{(1-\beta)(\sigma+\beta-\sigma\beta)} + \left(1 - \frac{\beta}{1-\beta}\right) + \ln\left(\frac{\beta}{1-\beta}\right) = \sigma(2\beta-1) \frac{1-\beta}{\beta(\sigma+\beta-\sigma\beta)} + \ln\left(\frac{\beta}{1-\beta}\right) - \left(1 - \frac{1-\beta}{\beta}\right)$.

In sum, we have that when $\beta < 1/2$, $\partial \ln \Gamma_\beta / \partial \eta < 0$ when evaluated at $\eta = 0$, while for $\beta > 1/2$, $\partial \Gamma_\beta / \partial \eta > 0$ when evaluated at $\eta = 1$. Together with the concavity of Γ_β , we can then conclude that $(\partial \ln \Gamma_\beta / \partial \eta)(\beta - 1/2) \geq 0$ for all η , with strict inequality for $\beta \neq 1/2$. The practical relevance of this result is that it complicates the overall comparative static of the offshoring share Υ_O with respect to η (remember that under complete contracting, Υ_O is unambiguously decreasing in η).

We next consider how the ambiguous effect of changes in β interacts with η . I begin by noting that simple differentiation delivers

$$\frac{\partial \ln(\Gamma_O/\Gamma_D)}{\partial \beta} = (\sigma - 1) \frac{\eta(1 - \eta) + \sigma\eta^2 - ((\sigma - 1)\eta + 1)2\eta\beta + \sigma(2\eta - 1)\beta^2}{\beta(1 - \beta)(\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta)))}$$

and

$$\frac{\partial^2 \ln(\Gamma_O/\Gamma_D)}{\partial \beta^2} = -(\sigma - 1)^2 \left(\frac{(2\eta - 1)^2}{(\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta)))^2} + \frac{\eta(1 - \beta) + \beta(\beta - \eta)}{(\sigma - 1)\beta^2(1 - \beta)^2} \right) < 0.$$

Thus, Γ_O/Γ_D is maximized for the value(s) of β that solve the quadratic equation in the numerator of $\partial \ln(\Gamma_O/\Gamma_D)/\partial \beta$. It turns out that there is only one solution β^* of this quadratic equation satisfying $\beta^* \in [0, 1]$. Rearranging this solution, we find equation (4.15) in 4, which makes it clear that β^* is increasing in η .

Consider finally how the elasticity of demand affects the ratio Γ_O/Γ_D . Simple (though tedious) differentiation confirms first that

$$\frac{\partial^2 (\ln(\Gamma_O/\Gamma_D))}{\partial \sigma^2} = -\frac{(\beta + \eta - 2\beta\eta)^2}{(\sigma - (\sigma - 1)(\beta\eta + (1 - \beta)(1 - \eta)))^2} < 0.$$

Hence $\partial \ln(\Gamma_O/\Gamma_D)/\partial \sigma$ is bounded above by the value of this derivative when evaluated at the lowest possible value of σ , namely $\sigma = 1$. But note that

$$\left. \frac{\partial \ln(\Gamma_O/\Gamma_D)}{\partial \sigma} \right|_{\sigma=1} = \beta + \eta - 2\beta\eta + \ln(\beta^\eta(1 - \beta)^{1-\eta}).$$

To evaluate this expression, notice that it increases in η when $\beta > 1/2$, while it decreases in η when $\beta < 1/2$.² Furthermore, the expression equals $\beta + \ln(1 - \beta) \leq 0$ when $\eta = 0$, $1 - \beta + \ln(\beta) \leq 0$ when $\eta = 1$, and $\frac{1}{2} + \ln(\frac{1}{2}) < 0$ when $\beta = 1/2$. We can thus conclude that $\partial \ln(\Gamma_O/\Gamma_D)/\partial \sigma < 0$ for $\sigma > 1$.

²This in turn can be shown again by applying the inequalities $1 - x + \ln x \leq 0$ and $\ln 1/x - (1 - x) \geq 0$ with $x = \beta/(1 - \beta)$, and decomposing $1 - 2\beta + \ln \frac{\beta}{1-\beta} = \beta \frac{2\beta-1}{1-\beta} + 1 - \frac{\beta}{1-\beta} + \ln \frac{\beta}{1-\beta} = (2\beta - 1) \frac{1-\beta}{\beta} - \left(1 - \frac{1-\beta}{\beta}\right) + \ln \frac{\beta}{1-\beta}$.

Limitations on Ex-Ante Transfers: Financial Constraints

Remember that the case in which M cannot transfer to F ex-ante more than a share ϕ of his or her ex-post rents, delivered

$$\frac{\Gamma_O}{\Gamma_D} = \Gamma_\phi \equiv (\sigma + \phi - (\sigma - 1)(1 - \phi)\eta) \left(\frac{1}{2}\right)^\sigma,$$

since again we assumed $\Gamma_D = 1$. It is obvious from this expression that Γ_ϕ increases in ϕ and η , and these effects interact in a positive matter, or $\partial^2 \Gamma_\phi / (\partial \phi \partial \eta) > 0$. The positive effect of η on Γ_ϕ again renders ambiguous the overall effect of headquarter intensity on the offshoring share Υ_O (with complete contracting, Υ_O is unambiguously decreasing in η).

Consider next the effect of the demand elasticity σ . Straightforward differentiation delivers

$$\frac{\partial \ln \Gamma_\phi}{\partial \sigma} = \frac{(1 - \eta + \phi\eta)}{(\sigma + \phi - (\sigma - 1)(1 - \phi)\eta)} - \ln 2,$$

as well as $\partial^2 \ln \Gamma_\phi / \partial \sigma^2 < 0$. It is then straightforward to show that for a sufficiently high σ , we necessarily have $\partial \ln \Gamma_\phi / \partial \sigma < 0$. In fact, the weak condition $\sigma + \phi > (\ln 2)^{-1} = 1.4427$ is sufficient for this inequality to hold, regardless of the value of η .

Partial Contractibility

In the extension of the model with partial contractibility in both countries, I alluded to the results in Antràs and Helpman (2008) to motivate the following expressions for the index of contracting distortions under domestic sourcing and offshoring:

$$\Gamma_D = \left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_N)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_N)} \left(\frac{1}{2} \right)^\sigma; \quad (8.3)$$

$$\Gamma_O = \left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_S)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_S)} \left(\frac{1}{2} \right)^\sigma, \quad (8.4)$$

where

$$\begin{aligned} \mu_N &\equiv \eta \mu_{hN} + (1 - \eta) \mu_{mN}; \\ \mu_S &\equiv \eta \mu_{hS} + (1 - \eta) \mu_{mS}. \end{aligned}$$

In fact, these expressions are a special case of those that apply in the framework in Antràs and Helpman (2008). Because I will be referring to these more general

results repeatedly in the derivations below, it might be useful to sketch here the steps that lead to that more general formula.

With that in mind, consider the following generalization of the problem in (4.19) after substitution of the participation constraint pinning down the ex-ante transfer:

$$\begin{aligned} \max_{h_c, h_n, m_c, m_n, s} \quad & R - w_N (\mu_{hj} h_c + (1 - \mu_{hj}) h_n + \mu_{mj} m_c + (1 - \mu_{mj}) m_n) - s \\ \text{s.t.} \quad & h_n = \arg \max_h \{ \beta_h R - w_N (1 - \mu_{hj}) h_n \} \\ & m_n = \arg \max_m \{ \beta_m R - c_j (1 - \mu_{mj}) m_n \}, \end{aligned}$$

where revenue is given by

$$\begin{aligned} R = & B^{1/\sigma} \sigma (\sigma - 1)^{-(\sigma-1)/\sigma} \varphi^{(\sigma-1)/\sigma} \\ & \times \left(\frac{(h_c)^{\mu_{hj}} (h_n)^{1-\mu_{hj}}}{\eta} \right)^{(\sigma-1)\eta/\sigma} \left(\frac{(m_c)^{\mu_{mj}} (m_n)^{1-\mu_{mj}}}{1-\eta} \right)^{(\sigma-1)(1-\eta)/\sigma} \end{aligned} \quad (8.5)$$

and where $c_j = w_N$ when $j = N$ and $c_j = \tau w_S$ when $j = S$. The problem above thus cover the cases of symmetric and generalized Nash bargaining, but it also encompasses environments with partial relationship-specificity in which F and M only bargain over a fraction of revenue ex-post, and thus $\beta_h + \beta_m < 1$. And, as discussed below, this formulation will also prove useful in the characterization of the equilibrium under multiple suppliers.

In order to derive the formula for profits associated with this more general problem, notice first that from the two constraints of the problem, we have

$$\begin{aligned} h_n &= \frac{\beta_h (\sigma - 1) \eta}{\sigma w_N} R \\ m_n &= \frac{\beta_m (\sigma - 1) \eta}{\sigma w_j} R. \end{aligned}$$

Plugging these expressions into (8.5), delivers

$$\begin{aligned} R = & \left(B^{1/\sigma} \sigma (\sigma - 1)^{-(\sigma-1)/\sigma} \varphi^{(\sigma-1)/\sigma} \right)^{\frac{\sigma}{\sigma - (\sigma-1)(1-\mu_j)}} \left(\frac{h_c}{\eta} \right)^{\frac{(\sigma-1)\eta\mu_{hj}}{\sigma - (\sigma-1)(1-\mu_j)}} \left(\frac{m_c}{1-\eta} \right)^{\frac{(\sigma-1)(1-\eta)\mu_{mj}}{\sigma - (\sigma-1)(1-\mu_j)}} \\ & \times \left(\frac{\beta_h (\sigma - 1)}{\sigma w_N} \right)^{\frac{(\sigma-1)\eta(1-\mu_{hj})}{\sigma - (\sigma-1)(1-\mu_j)}} \left(\frac{\beta_m (\sigma - 1)}{\sigma w_j} \right)^{\frac{(\sigma-1)(1-\eta)(1-\mu_{mj})}{\sigma - (\sigma-1)(1-\mu_j)}}. \end{aligned} \quad (8.6)$$

Given the Cobb-Douglas structure, we can then characterize the choice of con-

tractible investments as satisfying

$$h_c = \frac{(\sigma - 1)\eta \left(1 - \frac{(\sigma-1)}{\sigma} (\beta_h \eta (1 - \mu_{hj}) + \beta_m (1 - \eta) (1 - \mu_{mj}))\right)}{(\sigma - (\sigma - 1) (1 - \mu_j)) w_N} R \quad (8.7)$$

$$m_m = \frac{(\sigma - 1) (1 - \eta) \left(1 - \frac{(\sigma-1)}{\sigma} (\beta_h \eta (1 - \mu_{hj}) + \beta_m (1 - \eta) (1 - \mu_{mj}))\right)}{(\sigma - (\sigma - 1) (1 - \mu_j)) w_j} \quad (8.8)$$

As a result, operating profits are given by

$$\left(\frac{\sigma - (\sigma - 1) (\beta_h (1 - \mu_{hj}) + \beta_m (1 - \mu_{mj}))}{\sigma - (\sigma - 1) (1 - \mu_j)} \right) \frac{R}{\sigma},$$

where R can be solved by plugging the above expression (8.7) and (8.8) into (8.6). This delivers, after some manipulations

$$\begin{aligned} R &= \sigma B \left((w_N)^\eta (w_j)^{1-\eta} \right)^{1-\sigma} \varphi^{\sigma-1} (\beta_h)^{(\sigma-1)\eta(1-\mu_{hj})} (\beta_m)^{(\sigma-1)(1-\eta)(1-\mu_{mj})} \\ &\quad \times \left(\frac{\sigma - (\sigma - 1) (\beta_h \eta (1 - \mu_{hj}) + \beta_m (1 - \eta) (1 - \mu_{mj}))}{\sigma - (\sigma - 1) (1 - \mu_j)} \right)^{(\sigma-1)\mu_j}, \end{aligned}$$

and thus

$$\begin{aligned} \pi_D(\varphi) + f_D w_N &= (w_N)^{1-\sigma} B \Gamma_D \varphi^{\sigma-1} \\ \pi_O(\varphi) + f_O w_N &= \left((w_N)^\eta (\tau w_S)^{1-\eta} \right)^{1-\sigma} B \Gamma_O \varphi^{\sigma-1}, \end{aligned}$$

where

$$\begin{aligned} \Gamma_\ell &= \left(\frac{\sigma - (\sigma - 1) (\beta_h \eta (1 - \mu_{hj}) + \beta_m (1 - \eta) (1 - \mu_{mj}))}{\sigma - (\sigma - 1) (1 - \mu_j)} \right)^{\sigma - (\sigma - 1) (1 - \mu_j)} \\ &\quad \times (\beta_h)^{(\sigma-1)\eta(1-\mu_{hj})} (\beta_m)^{(\sigma-1)(1-\eta)(1-\mu_{mj})} \end{aligned} \quad (8.9)$$

captures the contractual frictions associated with the sourcing options $\ell = D$ and $\ell = O$, which entail manufacturing in country $j = N$ and country $j = S$, respectively. Setting $\beta_h = \beta_m = 1/2$, it is straightforward to verify that equation (8.9) reduces to equations (8.3) and (8.4) above.

Having derived these equations, we can next turn to discussing some key comparative statics. Below, I will focus on an analysis of the general formula (8.9), with the understanding that the results obtained below also apply to the particular

case with $\beta_h = \beta_m = 1/2$. Consider first the effect of the indices of contractibility μ_{hj} and μ_{mj} , and its weighted average μ_j . As shown in Antràs and Helpman (2008) (see the proof of their Proposition 1), Γ_ℓ is necessarily non-decreasing in each of these parameters. The proof in that paper is rather cumbersome, so it may be worth offering a much simpler proof here. Consider the case of an increase in μ_{hj} (the derivations associated with a change in μ_{mj} are analogous). Taking logs of (8.9), differentiating and rearranging terms, we can write

$$\frac{\partial \ln \Gamma_\ell}{\partial \mu_{hj}} = \eta(\sigma - 1)(-\ln \mathcal{Q} - (1 - \mathcal{Q}) - \ln \beta_h - (1 - \beta_h)) + \mathcal{W}, \quad (8.10)$$

where

$$\mathcal{Q} = \frac{\sigma - (\sigma - 1)(1 - \mu_j)}{\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_{hj}) + \beta_m(1 - \eta)(1 - \mu_{mj}))}$$

and

$$\mathcal{W} = \eta(\sigma - 1)^2(1 - \beta_h) \frac{(1 - \eta)(1 - \mu_{mj})(1 - \beta_m) + \eta(1 - \mu_{hj})(1 - \beta_h)}{\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_{hj}) + \beta_m(1 - \eta)(1 - \mu_{mj}))}.$$

It is clear that the second term is positive, while the first one is non-negative too because $-\ln x - (1 - x) \geq 0$ for all x . Thus, $\partial \ln \Gamma_\ell / \partial \mu_{hj} \geq 0$.

It is also clear from inspection of equation (8.9), that as stated in the main text, the effect of improvements in contractibility interacts with the headquarter intensity of production depending on the source of these changes in contractibility. Increases in μ_{hj} will be particularly beneficial when η is high, while the converse is true for μ_{mj} . For the same reason, and as in the model with totally incomplete contracting and generalized Nash bargaining, the effect of changes in headquarter intensity on Γ_ℓ is ambiguous.

Let us now turn to the effect of the elasticity of demand σ on Γ_ℓ . Tedious differentiation of Γ_ℓ delivers

$$\frac{\partial^2 \ln \Gamma_\ell}{\partial \sigma^2} = - \frac{(1 - \mu - \eta(1 - \mu_h)\beta_h - (1 - \eta)(1 - \mu_m)\beta_m)^2}{(1 - \mu + \sigma\mu)(\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_h) + \beta_m(1 - \eta)(1 - \mu_m)))^2} < 0$$

and

$$\left. \frac{\partial \ln \Gamma}{\partial \sigma} \right|_{\sigma=1} = 1 - \mu + \eta(1 - \mu_h)(\ln \beta_h - \beta_h) + (1 - \eta)(1 - \mu_m)(\ln \beta_m - \beta_m) \leq 0.$$

To prove the negative sign in the second equation, note that this expression is maximized when $\beta_h = \beta_m = 1$, and at that level $\partial \ln \Gamma / \partial \sigma|_{\sigma=1} = 1 - \mu - \eta(1 - \mu_h) - (1 - \eta)(1 - \mu_m) = 0$. In light of these results, we can conclude that $\partial \ln \Gamma / \partial \sigma < 0$

for all $\sigma > 1$, and thus contractual frictions are again aggravated by high demand elasticities in this variant of the model.

We next show how the effects of contractibility and the elasticity of demand interact with each other. In particular, differentiating $\partial \ln \Gamma_\ell / \partial \mu_{hj}$ in (8.10) with respect to σ , we find:

$$\frac{\partial^2 \ln \Gamma_\ell}{\partial \mu_{hj} \partial \sigma} = \frac{1}{(\sigma - 1)} \frac{\partial \ln \Gamma_\ell}{\partial \mu_{hj}} + \eta (\sigma - 1) \frac{\partial (-\ln \mathcal{Q} - (1 - \mathcal{Q}))}{\partial \sigma} + \frac{\partial \mathcal{W}}{\partial \sigma}.$$

We have established before that the first term is non-negative. Differentiating the second and third terms, we find

$$\frac{\partial (\ln \mathcal{Q}^{-1} - 1 + \mathcal{Q})}{\partial \sigma} = \frac{(\sigma - 1) \left(\frac{(1-\eta)(1-\mu_m)(1-\beta_m) + \eta(1-\mu_h)(1-\beta_h)}{\sigma - (\sigma-1)(\beta_h \eta(1-\mu_h) + \beta_m(1-\eta)(1-\mu_m))} \right)^2}{(\sigma - (\sigma - 1)(1 - \eta \mu_h - (1 - \eta) \mu_m))}$$

and

$$\frac{\partial \mathcal{W}}{\partial \sigma} = (1 - \beta_h) \frac{(1 - \eta)(1 - \mu_m)(1 - \beta_m) + \eta(1 - \mu_h)(1 - \beta_h)}{(\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_h) + \beta_m(1 - \eta)(1 - \mu_m)))^2},$$

and thus these terms are non-negative as well. In sum, we can conclude that $\frac{\partial^2 \ln \Gamma_\ell}{\partial \mu_{hj} \partial \sigma} \geq 0$, as stated in the main text. Notice, that the result is *not* particular to the special case $\beta_h = \beta_m = 1/2$, nor does it require $\beta_h + \beta_m = 1$. It is worth pointing out that it is important that we are considering the partial derivative of the logarithm of Γ_ℓ . Computing $\frac{\partial^2 \Gamma_\ell}{\partial \mu_{hj} \partial \sigma}$, we find that this expression may take negative values for some parameter values. This justifies the use of logarithms of import flows in certain empirical specifications in Chapter 5, as discussed in the main text.

I have thus far focused on providing formal proofs of the results mentioned in Chapter 4, which are key for interpreting the cross-country, cross-industry results in the second part of Chapter 5. The first part of that chapter focuses on studying the determinants of the offshoring share Υ_O , which in turn depend on the ratio Γ_O/Γ_D :

$$\begin{aligned} \frac{\Gamma_O}{\Gamma_D} &= \frac{\left(\frac{\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_{hS}) + \beta_m(1 - \eta)(1 - \mu_{mS}))}{\sigma - (\sigma - 1)(1 - \mu_S)} \right)^{\sigma - (\sigma - 1)(1 - \mu_S)}}{\left(\frac{\sigma - (\sigma - 1)(\beta_h \eta(1 - \mu_{hN}) + \beta_m(1 - \eta)(1 - \mu_{mN}))}{\sigma - (\sigma - 1)(1 - \mu_N)} \right)^{\sigma - (\sigma - 1)(1 - \mu_N)}} \\ &\quad \times (\beta_h)^{(\sigma - 1)\eta(\mu_{hN} - \mu_{hS})} (\beta_m)^{(\sigma - 1)(1 - \eta)(\mu_{mN} - \mu_{mS})}. \end{aligned}$$

From the results above, it is immediate that Γ_O/Γ_D is increasing in μ_S and its components μ_{hS} and μ_{mS} , and decreasing in μ_N and its components μ_{hN} and μ_{mN} .

Less trivially, we can also use the results above to show that Γ_O/Γ_D is decreasing in the elasticity of demand σ provided that contract enforcement is higher in domestic transactions vis a vis offshoring transactions. In particular, notice that

$$\frac{\partial \ln(\Gamma_O/\Gamma_D)}{\partial \sigma} = \frac{\partial \ln(\Gamma_O)}{\partial \sigma} - \frac{\partial \ln(\Gamma_D)}{\partial \sigma},$$

and provided that $\mu_{hN} \geq \mu_{hS}$ and $\mu_{mN} \geq \mu_{mS}$, we can appeal to the above results $\frac{\partial^2 \ln \Gamma_\ell}{\partial \mu_{hj} \partial \sigma} \geq 0$ and $\frac{\partial^2 \ln \Gamma_\ell}{\partial \mu_{mj} \partial \sigma} \geq 0$ to conclude that $\partial \ln(\Gamma_O/\Gamma_D) / \partial \sigma \leq 0$.

Finally, it is important to emphasize that our results above do not suggest that offshoring shares will be higher for more “contractible” goods. To see this, suppose that contractibility in the South is always a fraction $\delta < 1$ of the one in the North, so we can write $\mu_{hS}/\mu_{hN} = \mu_{mS}/\mu_{mN} = \delta$. For the special case, $\beta_h = \beta_m = 1/2$, we then have

$$\frac{\Gamma_O}{\Gamma_D} = \frac{\left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \delta \mu_N)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \delta \mu_N)}}{\left(\frac{\sigma}{\sigma - (\sigma - 1)(1 - \mu_N)} + 1 \right)^{\sigma - (\sigma - 1)(1 - \mu_N)}}.$$

Increases in μ_N can then be interpreted as overall increases in the contractibility of goods, since they affect their contractibility proportionately, regardless of the country where production takes place. It is not hard to confirm, however, that the effect of μ_N on the expression above is non-monotonic. For instance, if one sets $\sigma = 10$ and $\delta = 0.9$, Γ_O/Γ_D is lower when $\mu_N = 0.7$ than when either $\mu_N = 0.5$ or $\mu_N = 0.9$.

Relationship-Specificity

As discussed in the main text, this is a special case of the more general Antràs-Helpman (2008) framework with $\beta_h = \beta_m = 1 - \epsilon/2$, with $\epsilon \in [0, 1]$. The results derived above for the case of partial contractibility thus continue to apply. Improvements in contractibility are associated with larger values of Γ_ℓ , the elasticity of demand σ affects Γ_ℓ negatively, and the positive “interaction” effect $\partial(\partial \ln \Gamma_\ell / \partial \mu_j) / \partial \sigma > 0$ continue to apply. Similarly, we have that the offshoring share is negatively impacted by the elasticity of demand σ on account of the term Γ_O/Γ_D (remember though that there is a positive counterbalancing effect that applies even in the complete-contracting case).

Let us then focus on the new comparative statics that emerge when introducing relationship-specificity. Consider first the direct effect of the specificity parameter ϵ . Simple differentiation of (4.24) delivers

$$\frac{\partial \ln \Gamma_\ell(\mu_j, \epsilon)}{\partial \epsilon} = - \frac{\sigma \epsilon (\sigma - 1) (1 - \mu_j)}{(2 - \epsilon) (2(1 - \mu_j) + (2 - \epsilon) \sigma \mu_j + (\sigma - 1 + \mu_j) \epsilon)} < 0$$

and

$$\frac{\partial^2 \ln \Gamma_\ell(\mu_j, \epsilon)}{\partial \epsilon \partial \mu_j} = \frac{2\sigma^2 \epsilon (\sigma - 1)}{(2 - \epsilon) (2(1 - \mu_j) + (2 - \epsilon) \sigma \mu_j + (\sigma - 1 + \mu_j) \epsilon)^2} > 0,$$

as stated in the main text. Hence, profitability is decreasing in specificity, and improvements in contractibility are particularly profitability-enhancing at high levels of specificity. Furthermore, we can use the latter result to conclude that

$$\frac{\partial \ln(\Gamma_O/\Gamma_D)}{\partial \epsilon} = \frac{\partial \ln(\Gamma_O)}{\partial \epsilon} - \frac{\partial \ln(\Gamma_D)}{\partial \epsilon} \leq 0$$

$\mu_{hN} \geq \mu_{hS}$ and $\mu_{mN} \geq \mu_{mS}$. In words, whenever contract enforcement is higher in domestic transactions relative to offshore transactions, higher levels of specificity tend to be associated with lower offshoring shares Υ_O .

Multiple Inputs and Multilateral Contracting

As mentioned in the main text, the equilibrium expressions of this variant of the model are analogous to those in Antràs and Helpman (2008) whenever $\beta_h = \beta_m = \sigma\rho / ((\sigma - 1)(1 - \eta) + \sigma\rho)$. Plugging these values into (8.9) delivers equation (4.27). Because (4.27) is a special case of (8.9), we can conclude once again that $\partial \Gamma_\ell(\mu_j, \rho) / \partial \mu_j \geq 0$. Furthermore, we can also appeal to previous results to establish that $\partial \Gamma_\ell(\mu_j, \rho) / \partial \sigma < 0$. This latter comparative static result would appear to be complicated by the fact that β_h and β_m are now a function of σ . But since $\Gamma_\ell(\mu_j, \rho)$ in (8.9) is increasing in β_h and β_m , and each of these two shares is decreasing in σ , we can again conclude that $\partial \Gamma_\ell(\mu_j, \rho, \beta_h(\sigma), \beta_m(\sigma)) / \partial \sigma < 0$. In addition, the cross-partial derivative $\partial (\partial \ln \Gamma_\ell(\mu_j, \rho) / \partial \mu_j) / \partial \sigma$ continues to be positive, despite the dependence of β_h and β_m on σ . To see this, we can just appeal to equation (8.10) and note that each of the terms in that expression is decreasing in β_h , which in turn decreases in σ . More precisely, we have that (i) $-\ln \mathcal{Q} - (1 - \mathcal{Q})$ is decreasing in \mathcal{Q} whenever $\mathcal{Q} < 1$, (ii) \mathcal{Q} is indeed lower than 1 and is increasing in β_h , (iii) $-\ln \beta_h - (1 - \beta_h)$ is decreasing in β_h for $\beta_h < 1$, and (iv) \mathcal{W} is decreasing in β_h .

We can next turn to the effects of ρ which is the new parameter introduced in this variant of the model. Simple differentiation of equation (4.27) indicates

$$\frac{\partial \ln \Gamma_\ell(\mu_j, \rho)}{\partial \rho} = \frac{(\sigma - 1)^3 (1 - \eta)^2 (1 - \mu_j)}{\rho(\rho\sigma + (\sigma - 1)(1 - \eta)) ((\sigma - (\sigma - 1)(1 - \mu_j)) \rho + (\sigma - 1)(1 - \eta))} > 0$$

and

$$\frac{\partial^2 \ln \Gamma_\ell(\mu_j, \rho)}{\partial \rho \partial \mu_j} = -(\sigma - 1)^3 \frac{(1 - \eta)^2}{\rho((\sigma - (\sigma - 1)(1 - \mu_j)) \rho + (\sigma - 1)(1 - \eta))^2} < 0,$$

which are the two key novel comparative statics highlighted in the main text of that section. Again, this last cross-partial derivative is useful in deriving predictions for the offshoring share Υ_O since for $\mu_{hN} \geq \mu_{hS}$ and $\mu_{mN} \geq \mu_{mS}$, this result implies $\frac{\partial \ln(\Gamma_O/\Gamma_D)}{\partial \rho} = \frac{\partial \ln(\Gamma_O)}{\partial \rho} - \frac{\partial \ln(\Gamma_D)}{\partial \rho} \geq 0$. In sum, whenever contract enforcement is higher in domestic transactions relative to offshore transactions, higher degrees of input substitutability tend to be associated with higher offshoring shares Υ_O .

B Data Appendix

In this Appendix, I provide more details on the data sources, refinements of the data, and some details on the estimation techniques associated with the empirical work in Chapters 5 and 7. The entire dataset and Stata program codes are available at <http://scholar.harvard.edu/antras/scholar.harvard.edu/antras/pages/books>. I often refer to this url as the ‘book’s website’.

B.1 Raw U.S. Import and Export Data

The basis for the empirical work conducted in this book is the *U.S. Related-Party Trade* database collected by the U.S. Bureau of Customs and Border Protection and managed by the U.S. Census Bureau. This dataset can be downloaded from the following U.S. census website <http://sasweb.ssd.census.gov/relatedparty/>. The data is available at different levels of industry aggregation, but the most disaggregated level available online is the six-digit North American Industry Classification System (NAICS). At the time of writing this Appendix, the data is available for the period 2002-2012. In the empirical work I instead used data for the period 2000-2011. Data for 2000 and 2001 is available as part of the entire dataset on the book’s website.

Throughout the data construction, non-manufacturing sectors were dropped. In addition, industries that are not reported in the Related Party Trade data were also dropped. In the U.S. Related Party Trade data, five manufacturing industries are reported at the 5-digit NAICS: 31131X (Alumina and Aluminum Production and Processing), 31181X (Bread and Bakery Product Manufacturing), 31511X (Hosiery and Sock Mills), 33631X (Motor Vehicle Gasoline Engine and Engine Parts Manufacturing), and 33641X (Aerospace Product and Parts Manufacturing). This last industry 33641X is somewhat different from the rest. First, it is the only five-digit industry with zero import flows. Furthermore, it is the only one of these industries for which six-digit industries with the same initial five-digits (in particular, 336411, 336412, 336413, 336414, 336415, and 336419) are reported in the dataset. For these reasons, this five-digit sector 33641X was dropped from

the dataset. All other NAICS-level industry variables described below were also constructed for the four surviving synthetic five-digit industries.

The U.S. Related-Party Trade database reports imports and exports associated with each of 233 foreign countries. The data for South Sudan is however only available for 2011, so this country was dropped from the sample. The raw data was rectangularized by treating any missing values as zero import or export flows. Overall, we work with data 390 industries, 232 countries, and 12 years, for a total of 1,085,760 observations.

The data define related-party import transactions involving parties "with various types of relationships including any person directly or indirectly, owning, controlling or holding power to vote, 6 percent of the outstanding voting stock or shares of any organization." On the other hand, a related-party export transaction is one "between a U.S. exporter and a foreign consignee, where either party owns, directly or indirectly, 10 percent or more of the other party." Although this ownership requirements are very low, I argued in Chapter 1 that BEA data suggests that intrafirm trade is generally associated with one of the entities having a controlling stake in the other entity. The dataset also contains data on non-related imports and exports, which involve parties that "have no affiliation with each other or who do not meet the relevant equity requirements" for related-party trade. Although in principle an indicator of whether a transaction involves or not related parties is required for *all* import or export transactions recorded by the U.S. Bureau of Customs and Border Protection, in practice that information is missing in some cases. The dataset labels those volumes of trade as "not reported."

Table B.1 presents descriptive statistics for the key variables in the U.S. Related-Party Trade database. A few features of the table are worth highlighting. First, related-party imports account for 51.6% of overall U.S. manufacturing imports over 2000-11, and for 31.8% of overall U.S. manufacturing exports. Second, the average (unweighted) related-party import and export shares are however much lower (23.8% for imports and 10.4% for exports). Third, the large number of zeros in the data implies that the number of observations with a well-defined intrafirm import share is less than one third the total number of observations (i.e., country and six-digit NAICS combinations in the data). Conversely, on the export side, more than fifty percent of the observations feature positive exports and well-defined intrafirm export shares. Fourth, non-reported import transactions account for a negligible 0.04% of U.S. imports (that percentage goes up to 3.71% for exports, but I do not employ that export share in the empirical analyses in the book).

Table B.1. Some Descriptive Statistics from the NAICS Related-Party Trade Database

Variable (in \$ except shares)	Mean	Std. Dev	Min	Max	N
Total Imports	14,712,628	282,669,054	0	44,917,394,621	1,085,760
a) Related Party	7,587,177	220,242,612	0	44,134,184,241	1,085,760
b) Non-Related Party	7,119,158	112,745,454	0	20,981,735,046	1,085,760
c) Not Reported	6,294	567,136	0	269,396,613	1,085,760
Related-Party Share a/(a+b)	0.2380	0.3265	0	1	312,884
Total Exports	8,594,996	113,419,707	0	19,996,871,796	1,085,760
Related	2,736,289	61,075,666	0	13,174,432,899	1,085,760
Non-Related Party	5,539,536	67,292,534	0	14,757,989,972	1,085,760
Not Reported	319,172	7,736,500	0	1,574,623,834	1,085,760
Related-Party Share a/(a+b)	0.1039	0.2150	0	1	565,145

For some of the tests performed in the book, I make use of this same U.S. Related-Party Trade database, but at the finer six-digit Harmonized System (HS) industrial disaggregation. This dataset is not publicly available but can be purchased from the U.S. Census. Over the period 2000-11, this more detailed dataset contains information on U.S. imports (related, non-related and non-reported) for 5,705 products and for 238 countries and territories. Despite the fact that this finer disaggregation generates more than 16 million potential observations on U.S. imports, only for about 10 percent of these cases (1,572,949 to be precise) are U.S. imports positive. This in turn leads to an overall number of 1,568,711 intrafirm trade shares in the data, exceeding by a factor of five the 312,884 available shares when using the publicly available NAICS dataset.³ Still, the mean and variance of intrafirm trade shares are very similar to those reported in Table B.1 for the six-digit NAICS data.

I will employ this dataset in regressions of the intrafirm import share on product (HS6) fixed effects and country-level variables. Because I do not have data on these country-level variables for all 238 countries and territories in the sample, the typical regression using these data will have less than 1,568,711 observations.

B.2 U.S. Import and Export Data at IO2002 Level

In many of the empirical tests presented in this book, offshoring is correlated with variables that because of their nature and characteristics can only be computed

³For 4,238 observations, the data indicates positive imports that are entirely recorded as “Non-Reported”. Because I define the share of intrafirm trade as related-party imports divided by the sum of related-party imports and non-related party imports, I cannot compute a well-defined intrafirm trade share in those cases.

with Input-Output data (more on this below). For this reason, the natural industry classification to work with in those cases is the I-O commodity code classification. More specifically, I use the 2002 Input-Output industrial classification, or IO2002 for short.

Following Antràs and Chor (2013), the raw data in NAICS industry codes were mapped to six-digit IO2002 industries using a correspondence provided by the Bureau of Economic Analysis (BEA) as a supplement to the 2002 U.S. Input-Output Tables.⁴ This concordance is a straightforward many-to-one mapping for the manufacturing industries (NAICS first digit = 3). Two industries required a separate treatment, as the NAICS data were at a coarser level of aggregation than could be mapped into six-digit IO2002 codes. A synthetic code 31131X was created to merge IO 311313 (Beet sugar manufacturing) and 31131A (Sugar cane mills and refining), while a separate code 33641X merged IO 336411, 336412, 336413, 336414, 33641A (all related to the manufacture of aircraft and related components). This approach is somewhat distinct from the one I used to handle these aircraft subsectors in the NAICS dataset, where I simply dropped 33641X rather than merging it into a single category with all other five digit sectors starting with 33641. Because the NAICS sector 33641X features no U.S. imports, this small divergence should have little impact on the results. All other industry variables described below were also constructed for these two synthetic IO2002 codes 31131X and 33641X.

Overall, the IO2002 identifies U.S. imports and exports (related, non-related, and non-reported) for 253 sectors, 232 countries and 12 years of data, for a total of 704,352 observations, of which again many are zeroes. Table B.2 provides some basic descriptive statistics from this related-party IO2002 trade data.

Table B.2. Some Descriptive Statistics from the IO2002 Related-Party Trade Database

Variable (in \$ except shares)	Mean	Std. Dev	Min	Max	N
Total Imports	22,679,545	376,502,066	0	44,917,395,456	704,352
a) Related Party	11,695,648	281,267,307	0	44,134,184,241	704,352
b) Non-Related Party	10,974,196	166,941,030	0	20,981,735,046	704,352
c) Not Reported	9,702	749,284	0	269,396,613	704,352
Related-Party Share a/(a+b)	0.2438	0.3265	0	1	227,829
Total Exports	13,549,842	175,761,983	0	27,862,790,144	704,352
Related	4,259,789	86,715,148	0	13,174,432,899	704,352
Non-Related Party	8,797,599	103,436,155	0	15,297,237,562	704,352
Not Reported	492,455	12,228,037	0	2,585,156,012	704,352
Related-Party Share a/(a+b)	0.1068	0.2121	0	1	416,933

⁴See, for instance, <http://www.bea.gov/industry/xls/2002DetailedItemOutput.xls>.

B.3 Isolating the Intermediate Input Component of U.S. Imports and Exports

In this section, I provide more details on the Wright (2014) methodology for isolating the intermediate input component of trade flows. The key input for this data correction is a list of End-Use industrial categories available from the U.S. Bureau of Economic Analysis. The BEA uses these end-use codes to allocate goods to their final use, within the National Income and Product Accounts. Importantly, U.S. imports and exports at the ten-digit Harmonized System level are similarly allocated to end-use codes. Foreign Trade Statistics distinguish six one-digit end-use categories: (0) Foods, feeds, and beverages; (1) Industrial supplies and materials; (2) Capital goods, except automotives; (3) Automotive vehicles, parts and engines; (4) Consumer goods (nonfood), except auto; and (5) Other merchandise. Apart from these six principal end-use categories, the classification is further subdivided into about 140 broad commodity groupings. Wright (2014) advocates dropping all products with an end-use code equal to 0, 4 or 5, as well as a subset of the commodity groupings in the other three end-codes. The full list of dropped BEA commodity groupings can be found in Table 7 of Wright (2014).

The practical implementation of this correction consists of four steps:

Step 1. We begin by mapping each ten-digit HS product to a BEA end-use code. In order to maximize such a mapping, we put together multiple years of concordance tables published by the Census Bureau. The Census website provides tables for recent years from 2008 to 2013. We also downloaded older tables for years from 1993 to 1997 from Jon Haveman's trade data website (<http://goo.gl/5pyijB>). Technically, import and export HS codes are administered by different federal agencies (Export codes, known as Schedule B, is administered by the Census, and import codes, known as Harmonized Tariff System (HTS), is administered by the U.S. International Trade Commissions (USITC)). As such, a complete mapping requires putting together concordance tables for both import codes and export codes. We first map each HS product to end-use code, using the most recent concordance table published in 2013, and move to the previous year's table, if the mapping is still incomplete.

Step 2. We then use detailed ten-digit HS U.S. import and export data by foreign country for 2000-11 available from Peter Schott's website at Yale University and drop all ten-digit flows consisting of finished goods, as dictated by the concordance constructed in step 1. See Schott (2008) for more details on the ten-digit U.S. import and export data.

Step 3. Next, we aggregate the ten-digit HS U.S. import and export data back to the IO2002 level using a concordance between ten-digit HS codes and IO2002 codes available from the U.S. Bureau of Economic Analysis website at <http://www.bea.gov/industry/xls/HSConcord.xls>. We do so both for total flows as well as for only the intermediate input component of these flows (after dropping final goods). Comparing these two flows we obtain an IO2002-country-year specific ‘discount factor’ by which overall imports and exports need to be multiplied to obtain intermediate input imports and exports.

Step 4. We then apply these discount factors to the Related-Party IO2002 trade data constructed based on the NAICS Related-Party trade data. Note that the applied discount factor varies across IO2002 codes, countries and years, and is also distinct for imports and exports. When the discount factor is 0, this implies that particular IO2002-country-year observation does not contain any intermediate input flows. Whenever an IO2002 sector features zero aggregate imports of intermediate inputs in each year we treat that sector as a final-good sector and drop it from the sample in all Wright-adjusted regressions. We also drop IO2002 sectors 325411 (‘Medicinal and botanical manufacturing’) and 33299B (‘Arms, ordinance, and accessories’) because they feature tiny input import flows only for 2000 and 2001, and their recorded offshoring shares for those years are negative. Overall we entirely drop 39 industries, which are listed in Table B.3.

Apart from these dropped sectors, many other industries feature small aggregate volumes of intermediate input flows, and many zero flows for imports from particular countries. Even though these observations are not dropped, their associated discount factor will be tiny or equal to zero (and in the latter case they will be dropped from log-linear specifications). Table B.3 presents basic descriptive statistics comparing total and intermediate input U.S. imports and exports. The table indicates that, overall, intermediate input flows account for 53.1% of total U.S. imports and 67.8% of total U.S. exports. Although not reported in the table, one can also compute the share of intermediate inputs in related-party imports and exports. These turn out to be only marginally higher than for overall trade, equalling on aggregate for 54.4% of intrafirm imports and 68.4% of intrafirm exports.

Apart from this Wright (2014) correction, in some of the empirical tests in the book I follow Nunn and Treffer (2013*b*) in restricting the sample to the set of countries for which it is more plausible that U.S. intermediate input purchases are associated with U.S. headquarters purchasing inputs from abroad (rather than foreign headquarters exporting inputs to U.S. suppliers). The details of this correction appear in the main text of Chapter 5, so I will not repeat them here. Quantitatively, this correction removes five countries from the sample accounting for a mere 3.18% of U.S. imports. As mentioned in the main text, I have also

experimented with a more extensive correction based on Nunn and Treffer (2013*b*) that drops 18 countries accounting for 32.52% of U.S. imports. These corrected flows are available from the files in the book's website.

Table B.3. IO2002 Sectors Excluded from the Sample by the Wright (2014) Correction

311111	Dog and cat food manufacturing	314110	Carpet and rug mills
311119	Other animal food manufacturing	315100	Apparel knitting mills
311210	Flour milling and malt manuf.	315230	Women's and girls' cut & sew apparel
311230	Breakfast cereal manufacturing	321991	Manufactured mobile home manuf.
31131X	Sugar Manufacturing	322291	Sanitary paper product manufacturing
311320	Chocolate & confectionery manuf	325411	Medicinal and botanical manufacturing
311340	Nonchocolate confectionery ma.	325620	Toilet preparation manufacturing
311410	Frozen food manufacturing	331314	Secondary smelting/alloying of alum.
311513	Cheese manufacturing	33299B	Arms, ordinance, and accessories
31151A	Fluid milk and butter manuf.	33461A	Software, audio, & video media reprod.
311520	Ice cream and frozen dessert ma.	335221	Household cooking appliances
311615	Poultry processing	335222	Household refrigerator & freezers
311820	Cookie, cracker, and pasta ma.	336213	Motor home manufacturing
311910	Snack food manufacturing	336991	Motorcycle, bicycle, & parts
311920	Coffee and tea manufacturing	336992	Military armored vehicle & tanks
311930	Flavoring syrup & concentrate ma.	337110	Wood kitchen cabinet & countertops
312110	Soft drink and ice manufacturing	337121	Upholstered household furniture
312120	Breweries	33721A	Office furniture manufacturing
312130	Wineries	337910	Mattress manufacturing
312140	Distilleries		

Table B.4. Descriptive Statistics Illustrating the Effects of the Wright (2014) Correction

Variable (in \$)	Mean	Std. Dev	Min	Max	N
Total Imports	22,679,545	376,502,066	0	44,917,395,456	704,352
Wright-Adjusted Input Imports	12,042,991	216,171,921	0	44,206,174,208	704,352
Total Exports	13,549,842	175,761,983	0	27,862,790,144	704,352
Wright-Adjusted Input Exports	9,180,273	140,980,583	0	27,860,740,096	704,352

B.4 Computing Offshoring Shares

The construction of offshoring shares requires data not only on U.S. imports and exports, but also on U.S. domestic shipments. Ignoring the Wright adjustment for

intermediate inputs, industry-level offshoring shares are simply computed as the ratio of U.S. imports to the sum of U.S. shipments plus U.S. imports minus U.S. exports in that sector. The analogous offshoring share for a given foreign country is computed as U.S. imports from that particular country divided by the same industry-level denominator, so the sum of country-industry-level offshoring shares corresponds to the aggregate industry-level offshoring share.

Data on U.S. shipments was obtained from the NBER-CES Manufacturing database for the period 2000-09 and from the Annual Survey of Manufacturing for the last two years 2010 and 2011. Both of these data sources are available at the six-digit NAICS level, which facilitates its merging with the six-digit NAICS U.S. import and export data. The mapping between the trade data and the NBER-CES database is quite clean, but merging the ASM data required a series of small adjustments to deal with the fact that 88 industries are reported at the more aggregated five-digit NAICS level in the ASM dataset. In order to minimize the loss of industry categories associated with adding those two last years of data, we imputed shipment values for all six-digit sectors available in the trade and NBER data for which only the more aggregated five-digit industry was available in the ASM data. This imputation was based on breaking up the 2010 and 2011 ASM total values of shipments at the five-digit level into six-digit values based on the relative weights of the different six-digit segments in the NBER-CES data over the period 2005-09.

As mentioned in Chapter 5, for a small percentage of industries and years the recorded value of total shipments falls short of the value of U.S. exports. This is true for years prior to 2009 so it is not explained by the adjustments described in the previous paragraph. These observations are typically dropped in the empirical tests in this book.

So far, I have described the construction of offshoring shares for the NAICS case and without any adjustment for intermediate input trade. In order to compute offshoring shares at the IO2002 level, we simply repeated the steps above but using as the basis U.S. imports and exports at the IO2002 level, as well as U.S. shipments at the same industry classification. The latter series was obtained by filtering the constructed NAICS shipment series described above through the same BEA concordance table used to transition from NAICS to IO2002 trade flows. Wright-adjusted offshoring shares were computed in an analogous manner based on Wright-adjusted U.S. imports, exports and shipments. To isolate the intermediate input component of U.S. shipments we applied discount factor to overall shipments equal to the average of the import and export ‘Wright’ discount factors applied to trade flows in that industry over the period 2000-11. Again, in some cases, the resulting value of input shipments fell short of the value of input exports, and these observations are dropped from most regressions.

This concludes our discussion of the construction of the main dependent variables in the empirical tests in the book. I now turn to describing the explanatory variables used in those tests.

B.5 Industry-Country-Level Covariates

Freight Costs. Sectoral and exporter-specific measures of freight costs associated with U.S. imports were downloaded from Peter Schott’s website (see Schott, 2010, for further documentation). More specifically, freight costs are computed as the ratio of CIF imports to FOB imports for a given product and origin country for the period 2000-05. Although this variable varies year by year, to salvage observations for 2006 through 2011, we construct a time-invariant measure of freight costs equal to the average of the ratio of CIF imports to FOB imports for a given country and product over the period 2000-05. We then assign this average measure to all 12 observations associated with a given exporting country and sector. The data are originally available at the six-digit NAICS level but we also constructed them at the IO2002 level using the same BEA concordance table employed throughout the book.

U.S. Tariffs. U.S. tariffs corresponds to U.S. applied tariffs from the World Integrated Trade Solution (WITS) database maintained by the World Bank. We again construct a measure at the exporter-sector level based on the average of this variable over multiple years, and in this case we do so for 2000-10 given data availability. The data is originally available at the six-digit HS level. We used the concordance in Pierce and Schott (2009) to transition from six-digit HS codes to six-digit NAICS codes. To construct IO2002 tariff levels, we used the same BEA concordance from HS6 to IO2002 employed when applying the Wright intermediate input correction, as described in step 3 of section B.3 above.

B.6 Industry-Level Covariates

Trade Costs. Industry-level freight costs and tariffs were computed based on the industry-country series we have just described in section B.5 but averaging them over all exporting countries. In both the NAICS and IO2002 cases, data on freight costs and tariffs was missing for some industry codes. For those sectors, we imputed a value equal to the weighted freight costs and tariffs of the sectors with which the industry shared the same first four digits, or (if the value was still missing) the same first three digits, using industry shipment values as weights.

R&D Intensity. We build on Nunn and Treffer (2013*b*), who calculated R&D expenditures to total sales on an annual basis for the period for 1998-2006 using

the U.S. firms in the Bureau van Dijk's Orbis dataset. Their original data is reported for IO1997 industries. To obtain IO2002 values, we follow Antràs and Chor (2013) and construct a crosswalk from IO1997 to IO2002 through the NAICS industry codes. More specifically, the R&D intensity for each IO2002 industry was calculated as the weighted average value of $\log(0.001 + R\&D/Sales)$ over that of its constituent IO1997 industries over the years 2000-2005, using the industry output values in the 1997 U.S. I-O Tables as weights. There remained 13 IO2002 industries without R&D intensity values after the above procedure. A similar procedure to that described above for trade costs was used to obtain the R&D intensity for the remaining 13 IO2002 codes (based on the R&D intensity of the IO2002 codes with which the industry shared the same first four or three digits). This yielded a complete series for R&D intensity of the 'selling' sector. In many specifications we instead use a measure of the R&D intensity of the 'average buyer' of an industrial good. This buyer version of the variable was computed as a weighted average of the R&D intensity of the industries that purchase the good in question (call it good v), with weights equal to these buying sector's input purchase values of good v as reported in the 2002 U.S. I-O Tables. The construction of R&D intensity at the six-digit NAICS level was analogous although in imputing values for missing observations we also made use of four-digit and three-digit measures of R&D intensity kindly provided by Heiwai Tang.

Capital and Skill Intensities. These were obtained from the NBER-CES Manufacturing Industry Database (Becker et al., ?). Skill intensity is the log of the number of non-production workers divided by total employment. Physical capital intensity is the log of the real capital stock per worker. Equipment capital intensity and plant capital intensity are respectively the log of the equipment and plant capital stock per worker. The NBER-CES data is originally available at the six-digit NAICS level, so matching it to the Related-Party Trade dataset only required minor adjustments related to trade data for five manufacturing industries being reported at the 5-digit NAICS (as explained in B.1 above). These industries were in turn mapped to IO2002 codes using the same procedure described in section B.2 above for the related party trade data. For each factor intensity variable, a simple average of the annual values from 2000-2005 was taken to obtain the seller industry measures. The factor intensities for the average buyer were then calculated using the same procedure as described for the average buyer R&D intensity.

Detailed Capital Equipment Intensities. In some specifications we report results that break capital equipment intensity into the separate effects of expenditures on (i) automobiles and trucks for highway use, (ii) computers and peripheral data processing equipment, and (iii) all other machinery and equipment computers. These were obtained from the Annual Survey of Manufactures (2002-2010)

which reports the data at the six-digit NAICS level. As mentioned before when discussing data on shipments from 2010 and 2011 (which also originate in the ASM), for 88 industries capital expenditures are reported at the more aggregated five-digit NAICS level. In order to impute those values to the six-digit sectors within each of these 88 industries we followed the same approach as in the case for industry shipments, but using overall equipment expenditures from the NBER-CES dataset as weights. The final measures of auto, computer, and other equipment intensity were obtained by dividing by these types of capital expenditures by the wage bill and taking logarithms. Values for these variables at the IO2002 level were obtained by filtering those variables through the BEA concordance described above, while ‘average buyer’ versions of these variables were also computed using the same approach as with R&D intensity and capital and skill intensity.

Productivity Dispersion. As in Antràs and Chor (2013), we build on Nunn and Treffer (2008), who constructed dispersion for each HS6 code as the standard deviation of log exports for its HS10 sub-codes across U.S. port locations and destination countries in the year 2000, from U.S. Department of Commerce data. We associated the dispersion value of each HS6 code to each of its HS10 sub-codes. These were mapped into IO2002 industries using the IO-HS concordance, taking a trade-weighted average of the dispersion value over HS10 constituent codes; the weights used were the total value of U.S. imports for each HS10 code from 1989-2006, from Feenstra, Romalis and Schott (2002). A similar procedure to that described above for trade costs and R&D intensity was used to obtain the dispersion measure for the remaining 13 IO2002 codes. The construction a productivity or size dispersion measure at the six-digit NAICS level was analogous to that of the R&D intensity and required using data from encompassing five- or four-digit sectors to impute values to industries that otherwise would have been left missing values.

Demand Elasticities. U.S. demand elasticities at the IO2002 level were computed as in Antràs and Chor (2013). We begin with the U.S. import demand elasticities for HS10 products computed by Broda and Weinstein (2006). This was merged with a comprehensive list of HS10 codes from Pierce and Schott (2009). For each HS10 code missing an elasticity value, we assigned a value equal to the trade-weighted average elasticity of the available HS10 codes with which it shared the same first nine digits. This was done successively up to codes that shared the same first two digits, to fill in as many HS10 elasticities as possible. Using the IO-HS concordance provided by the BEA with the 2002 U.S. I-O Tables, we then took the trade-weighted average of the HS10 elasticities within each IO2002 category. At each stage, the weights used were the total value of U.S. imports by HS10 code from 1989-2006, calculated from Feenstra, Romalis and Schott (2002). U.S.

demand elasticities at the six-digit NAICS level were computed in an analogous way based on the same HS10 elasticities but using the HS-NAICS concordance in Pierce and Schott (2009) to compute six-digit NAICS averages. In both cases, there remained industries without elasticity values after the above procedures. Values for these sectors were imputed following the same approach as for other variables above, using data from encompassing five- or four-digit sectors. Finally, in order to compute ‘average buyer’ demand elasticities, we took a weighted average of the elasticities of industries that purchase the input in question, with weights equal to these input purchase values as reported in the 2002 U.S. I-O Tables. This is the same approach used to construct buyer versions of R&D, capital and skill intensity.

Input Substitutability. We begin with the import demand elasticities estimated by Broda and Weinstein (2006), but this time we use their estimates at the SITC Revision 3 three-digit level (rather than ten-digit HS level). As documented in their paper (see in particular their footnote 22), these elasticities were estimated in part off the substitution seen across HS10 product codes that fall under each SITC three-digit heading. These estimates would contain information on the degree of substitution across inputs under the assumption that the constituent HS10 products in each SITC three-digit category are typically used together as inputs in production. The three-digit SITC elasticities were mapped into IO2002 codes by first assigning them to HS codes using the concordance in Feenstra, Romalis and Schott (2002) and then using the HS to IO concordance provided by the BEA.

Specificity. This measure is borrowed from Antràs and Chor (2013), who in turn build on Rauch (1999) and Nunn (2007). For each IO2002 industry, it is calculated as the fraction of HS10 constituent codes classified by Rauch (1999) as neither reference-priced nor traded on an organized exchange, under Rauch’s “liberal” classification. The original Rauch classification was for SITC Rev. 2 products; these were associated with HS10 codes using a mapping derived from U.S. imports in Feenstra, Romalis and Schott (2002). A higher value of this share is interpreted as the industry producing more differentiated goods, which in an input setting we associate with specificity.

Contractibility. We experiment with four measures of contractibility at the IO2002 industrial level. ‘Nunn contractibility’ was computed by Antràs and Chor (2013) from the 2002 U.S. I-O Tables following the methodology of Nunn (2007). We begin with the Rauch-Nunn sectoral measure of specificity above. For each IO2002 industry, we then calculate a weighted average specificity of the inputs used by that industry, where the weights correspond to each input’s share in the overall input purchases of the industry in question. We took one minus this value as a measure of the ‘Nunn contractibility’ of each IO2002 industry. Levchenko and Costinot contractibility were obtained from Chor (2010), who in turn built

them following the methodology of Levchenko (2007) and Costinot (2009), respectively. These two measures were normalized so that higher levels imply higher contractibility or lower dependence on formal contract enforcement. In particular, Levchenko contractibility is computed as the Herfindahl index of intermediate input use – rather than minus the Herfindahl as in Levchenko (2007) –, while Costinot contractibility is equal to the negative of the measure of complexity in Costinot (2009). Chor (2010) computed the Levchenko and Nunn measures at the 1987 Standard Industry Classification (SIC) level. We used a concordance from the U.S. census to map them into NAICS codes, and then the NAICS-IO2002 BEA concordance used for several measures above to obtain these variables at the IO2002 level.⁵ Finally, BJRS contractibility corresponds to the measure of intermediation in Bernard, Jensen, Redding and Schott (2010), who calculated this from U.S. establishment-level data as the weighted average of the wholesale employment share of firms in 1997, using the import share of each firm as weights. We use the IO2002 version of this variable constructed by Antràs and Chor (2013) (see their Data Appendix for more details).

We also compute ‘average buyer’ contractibility for each of these four contractibility measures using the same procedure described for computing average buyer R&D, capital and skill intensities.

Financial and Labor Contractibility. The Rajan and Zingales (1998) external dependence measure, the Braun (2002) asset tangibility measure, and the Cuñat and Melitz (2012) sales volatility measure are all borrowed from Chor (2010), who computed them at the 1987 SIC level. As with the Levchenko and Costinot contractibility variables discussed above, we used a concordance from the U.S. census to map them into NAICS codes, and then the NAICS-IO2002 BEA concordance used for several measures above to obtain these variables at the IO2002 level.

Downstreamness. Calculated from the 2002 U.S. I-O Tables, as described in Antràs and Chor (2013). It corresponds to a weighted index of the average position in the value chain at which an industry’s output is used (i.e., as final consumption, as direct input to other industries, as direct input to industries serving as direct inputs to other industries, and so on), with the weights being given by the ratio of the use of that industry’s output in that position relative to the total output of that industry. I next provide some more specific details on this measure for the interested reader. To build intuition, recall the basic input-output identity:

$$Y_i = F_i + Z_i,$$

where Y_i is total output in industry i , F_i is the output of i that goes toward final consumption and investment (“final use”), and Z_i is the use of i ’s output as inputs

⁵See <http://www.census.gov/eos/www/naics/concordances/concordances.html>.

to other industries (or its “total use” as an input). In a world with N industries, this identity can be expanded as follows:

$$Y_i = F_i + \underbrace{\sum_{j=1}^N d_{ij} F_j}_{\text{direct use of } i \text{ as input}} + \underbrace{\sum_{j=1}^N \sum_{k=1}^N d_{ik} d_{kj} F_j + \sum_{j=1}^N \sum_{k=1}^N \sum_{l=1}^N d_{il} d_{lk} d_{kj} F_j + \dots}_{\text{indirect use of } i \text{ as input}}, \quad (8.11)$$

where d_{ij} for a pair of industries (i, j) , $1 \leq i, j \leq N$, is the amount of i used as an input in producing one dollar worth of industry j 's output. Building on this identity, Antràs and Chor (2013) suggest computing the (weighted) average position of an industry's output in the value chain, by multiplying each of the terms in (??) by their distance from final use plus one and dividing by Y_i :

$$\begin{aligned} U_i &= 1 \cdot \frac{F_i}{Y_i} + 2 \cdot \frac{\sum_{j=1}^N d_{ij} F_j}{Y_i} \\ &+ 3 \cdot \frac{\sum_{j=1}^N \sum_{k=1}^N d_{ik} d_{kj} F_j}{Y_i} \\ &+ 4 \cdot \frac{\sum_{j=1}^N \sum_{k=1}^N \sum_{l=1}^N d_{il} d_{lk} d_{kj} F_j}{Y_i} + \dots \end{aligned} \quad (8.12)$$

It is clear that $U_i \geq 1$ and that larger values are associated with relatively higher levels of upstreamness of industry i 's use. Although computing (8.12) might appear to require computing an infinite power series, notice that provided that $d_{ij} < 1$ for all (i, j) (a natural assumption), the numerator of the above measure equals the i -th element of the $N \times 1$ matrix $[I - D]^{-2} F$, where D is an $N \times N$ matrix whose (i, j) -th element is d_{ij} and F is a column matrix with F_i in row i .⁶ In order to obtain a measure of downstreamness (rather than upstreamness), Antràs and Chor (2013) simply take the reciprocal of U_i , which necessarily lies in the interval $[0, 1]$. Antràs, Chor, Fally and Hillberry (2012) discuss additional appealing features of this downstreamness measure.

B.7 Country-Level Covariates

Relative Factor Abundance. Physical capital abundance corresponds to the log of the physical capital per worker averaged over 2000-2005. Physical capital was constructed by Davin Chor based on investment data from the Penn World Tables (version 7.1) using the perpetual inventory method. Skill abundance is

⁶Because $Y = [I - D]^{-1} F$, this numerator also equals the i -th element of the $N \times 1$ matrix $[I - D]^{-1} Y$, where Y is a column matrix with Y_i in row i .

measured as the average years of schooling at all levels (primary, secondary, and tertiary) averaged over 2000 and 2005 based on the Barro and Lee (2013) dataset.

Rule of Law. Country rule of law is obtained from the Worldwide Governance Indicators (see also Kaufmann et al., 2010). The annual index range of -2.5 to 2.5 and it was averaged over the period 2000-05.

Financial Development. Computed as private credit provided by banking sector is measured as the percentage of GDP, averaged over 2000-05, based on data from the World Bank's World Development Indicators.

Labor Market Flexibility. This corresponds to the country labor market flexibility index for the year 2004 used by Cuñat and Melitz (2012). It was originally constructed by the World Bank building on the work of Botero, Djankov, Porta, Lopez-de Silanes and Shleifer (2004).

GDP per capita. Computed as the log of Real GDP per capita in constant 2005 dollars from the Penn World Tables (version 7.1), averaged over the period 2000-05.

Bibliography

- Acemoglu, D., Antràs, P. and Helpman, E. (2007), ‘Contracts and technology adoption’, *American Economic Review* **97**(3), 916–943.
- Aghion, P., Dewatripont, M. and Rey, P. (1994), ‘Renegotiation design with unverifiable information’, *Econometrica* **62**(2), 257.
- Ahn, J. (2011), A theory of domestic and international trade finance, IMF Working Paper 11/262, International Monetary Fund.
- Albornoz, F., Calvo Pardo, H. F., Corcos, G. and Ornelas, E. (2012), ‘Sequential exporting’, *Journal of International Economics* **88**(1), 17–31.
- Alvarez, F. and Lucas, J. (2007), ‘General equilibrium analysis of the Eaton–Kortum model of international trade’, *Journal of Monetary Economics* **54**(6), 1726–1768.
- Amiti, M. and Davis, D. R. (2012), ‘Trade, firms, and wages: Theory and evidence’, *The Review of Economic Studies* **79**(1), 1–36.
- Amiti, M. and Weinstein, D. E. (2011), ‘Exports and financial shocks’, *The Quarterly Journal of Economics* **126**(4), 1841–1877.
- Anderson, J. E. and Marcouiller, D. (2002), ‘Insecurity and the pattern of trade: An empirical investigation’, *Review of Economics and Statistics* **84**(2), 342–352.
- Anderson, J. E. and Marcouiller, D. (2005), ‘Anarchy and autarky: Endogenous predation as a barrier to trade’, *International Economic Review* **46**(1), 189–213.
- Anderson, J. E. and van Wincoop, E. (2003), ‘Gravity with gravitas: A solution to the border puzzle’, *American Economic Review* **93**(1), 170–192.

- Antràs, P. (2003), ‘Firms, contracts, and trade structure’, *Quarterly Journal of Economics* **118**(4), 1375–1418.
- Antràs, P. and Chor, D. (2013), ‘Organizing the global value chain’, *Econometrica* **81**(6), 2127–2204.
- Antràs, P., Chor, D., Fally, T. and Hillberry, R. (2012), ‘Measuring the upstreamness of production and trade flows’, *American Economic Review* **102**(3), 412–416.
- Antràs, P. and Costinot, A. (2011), ‘Intermediated trade’, *The Quarterly Journal of Economics* **126**(3), 1319–1374.
- Antràs, P. and Foley, F. (2013), ‘Poultry in motion: A study of international trade finance practices’.
- Antràs, P., Fort, T. C. and Tintelnot, F. (2014), ‘The margins of global sourcing: Theory and evidence from U.S. firms’. Work in Progress.
- Antràs, P., Garicano, L. and Rossi-Hansberg, E. (2006), ‘Offshoring in a knowledge economy’, *The Quarterly Journal of Economics* **121**(1), 31–77.
- Antràs, P. and Helpman, E. (2004), ‘Global sourcing’, *Journal of Political Economy* **112**(3), 552–580.
- Antràs, P. and Helpman, E. (2008), Contractual frictions and global sourcing, in ‘The Organization of Firms in a Global Economy’, Harvard University Press, Cambridge, MA.
- Antràs, P. and Rossi-Hansberg, E. (2009), ‘Organizations and trade’, *Annual Review of Economics* **1**(1), 43–64.
- Antràs, P. and Staiger, R. (2012), ‘Offshoring and the role of trade agreements’, *American Economic Review* **102**(7), 3140–3183.
- Antràs, P. and Yeaple, S. (2013), Multinational firms and international trade structure, in ‘Handbook of International Economics’, Vol. 4, Elsevier.
- Araujo, L. F., Mion, G. and Ornelas, E. (2012), Institutions and export dynamics, Technical report, London School of Economics Working Paper.
- Arkolakis, C., Costinot, A. and Rodríguez-Clare, A. (2012), ‘New trade models, same old gains?’, *American Economic Review* **102**(1), 94–130.

- Arkolakis, C., Demidova, S., Klenow, P. J. and Rodríguez-Clare, A. (2008), ‘Endogenous variety and the gains from trade’, *American Economic Review* **98**(2), 444–450.
- Atalay, E., Hortacsu, A. and Syverson, C. (2013), ‘Vertical integration and input flows’. mimeo University of Chicago.
- Aumann, R. J. and Shapley, L. S. (1974), *Values of Non-Atomic Games*, Princeton University Press.
- Baker, G. P. and Hubbard, T. N. (2003), ‘Make versus buy in Trucking: Asset ownership, job design, and information’, *American Economic Review* **93**(3), 551–572.
- Baldwin, R. E. and Forslid, R. (2010), ‘Trade liberalization with heterogeneous firms’, *Review of Development Economics* **14**(2), 161–176.
- Bardhan, P., Mookherjee, D. and Tsumagari, M. (2013), ‘Middlemen margins and globalization’, *American Economic Journals: Micro* . Forthcoming.
- Barro, R. J. and Lee, J. W. (2013), ‘A new data set of educational attainment in the world, 1950–2010’, *Journal of Development Economics* **104**, 184–198.
URL: <http://www.sciencedirect.com/science/article/pii/S0304387812000855>
- Behar, A. and Freund, C. (2011), ‘Factory europe? brainier but not brawnier’, *Unpublished Manuscript* .
- Berkowitz, D., Moenius, J. and Pistor, K. (2006), ‘Trade, law, and product complexity’, *Review of Economics and Statistics* **88**(2), 363–373.
- Bernard, A. B., Eaton, J., Jensen, J. B. and Kortum, S. (2003), ‘Plants and productivity in international trade’, *American Economic Review* **93**(4), 1268–1290.
- Bernard, A. B., Jensen, J. B., Redding, S. J. and Schott, P. K. (2007), ‘Firms in international trade’, *Journal of Economic Perspectives* **21**(3), 105–130.
- Bernard, A. B., Jensen, J. B., Redding, S. J. and Schott, P. K. (2010), ‘Intrafirm trade and product contractibility’, *American Economic Review* **100**(2), 444–48.
- Bernard, A. B., Jensen, J. B. and Schott, P. K. (2009), Importers, exporters, and multinationals: A portrait of firms in the U.S. that trade goods, *in* T. Dunne, J. B. Jensen and M. J. Roberts, eds, ‘Producer Dynamics: New Evidence from Micro Data’, NBER.

- Blaum, J., Lelarge, C. and Peters, M. (2013), ‘Non-homothetic import demand: Firm productivity and quality bias’, *Unpublished paper* .
- Botero, J. C., Djankov, S., Porta, R. L., Lopez-de Silanes, F. and Shleifer, A. (2004), ‘The regulation of labor’, *The Quarterly Journal of Economics* **119**(4), 1339–1382.
- Braun, M. (2002), ‘Financial contractibility and assets’ hardness: Industrial composition and growth’. mimeo Harvard University.
- Broda, C. and Weinstein, D. E. (2006), ‘Globalization and the gains from variety’, *The Quarterly Journal of Economics* **121**(2), 541–585.
- Campa, J. M. and Goldberg, L. S. (1997), ‘The evolving external orientation of manufacturing: A profile of four countries’, *Federal Reserve Bank of New York Economic Policy Review* **3**, 53–81.
- Chaney, T. (2008), ‘Distorted gravity: The intensive and extensive margins of international trade’, *American Economic Review* **98**(4), 1707–1721.
- Chen, H., Kondratowicz, M. and Yi, K.-M. (2005), ‘Vertical specialization and three facts about U.S. international trade’, *The North American Journal of Economics and Finance* **16**(1), 35–59.
- Chor, D. (2010), ‘Unpacking sources of comparative advantage: A quantitative approach’, *Journal of International Economics* **82**(2), 152–167.
- Clermont, K. M. and Eisenberg, T. (2007), ‘Xenophilia or xenophobia in U.S. courts? before and after 9/11’, *Journal of Empirical Legal Studies* **4**(2), 441–464.
- Coase, R. H. (1937), ‘The nature of the firm’, *Economica* **4**(16), 386–405.
- Costinot, A. (2009), ‘On the origins of comparative advantage’, *Journal of International Economics* **77**(2), 255–264.
- Costinot, A. and Rodríguez-Clare, A. (2013), Trade theory with numbers: Quantifying the consequences of globalization, in ‘Handbook of International Economics’, Elsevier.
- Cuñat, A. and Melitz, M. J. (2012), ‘Volatility, labor market flexibility, and the pattern of comparative advantage’, *Journal of the European Economic Association* **10**(2), 225–254.

- Davis, D. R. and Weinstein, D. E. (2001), ‘An account of global factor trade’, *American Economic Review* **91**(5), 1423–1453.
- Deardorff, A. V. (2001), ‘Fragmentation in simple trade models’, *The North American Journal of Economics and Finance* **12**(2), 121–137.
- Demidova, S. and Rodríguez-Clare, A. (2009), ‘Trade policy under firm-level heterogeneity in a small economy’, *Journal of International Economics* **78**(1), 100–112.
- Demidova, S. and Rodríguez-Clare, A. (2013), ‘The simple analytics of the melitz model in a small economy’, *Journal of International Economics* **90**(2), 266–272.
- Dixit, A. K. and Stiglitz, J. E. (1977), ‘Monopolistic competition and optimum product diversity’, *The American Economic Review* **67**(3), 297–308.
- Djankov, S., La Porta, R., Lopez-De-Silanes, F. and Shleifer, A. (2003), ‘Courts’, *The Quarterly Journal of Economics* **118**(2), 453–517.
- Eaton, J. and Kortum, S. (2002), ‘Technology, geography, and trade’, *Econometrica* **70**(5), 1741–1779.
- Eaton, J., Kortum, S. and Kramarz, F. (2011), ‘An anatomy of international trade: Evidence from french firms’, *Econometrica* **79**(5), 1453–1498.
- Ethier, W. J. (2001), Dixit-stiglitz, trade and growth, in S. Brakman and B. J. Heijdra, eds, ‘The Monopolistic Competition Revolution in Retrospect’, Cambridge University Press.
- Fally, T. (2012), *Production Staging: Measurement and Facts*. mimeo University of Colorado.
- Feenstra, R. C. (1998), ‘Integration of trade and disintegration of production in the global economy’, *Journal of Economic Perspectives* **12**(4), 31–50.
- Feenstra, R. C. and Hanson, G. H. (1996a), Foreign investment, outsourcing and relative wages, in R. Feenstra, G. Grossman and D. Irwin, eds, ‘The Political Economy of Trade Policy: Papers in Honor of Jagdish Bhagwati’, MIT Press.
- Feenstra, R. C. and Hanson, G. H. (1996b), ‘Globalization, outsourcing, and wage inequality’, *American Economic Review* **86**, 240–245.
- Feenstra, R. C., Romalis, J. and Schott, P. K. (2002), U.S. imports, exports, and tariff data, 1989-2001, Working Paper 9387, National Bureau of Economic Research.

- Felbermayr, G., Jung, B. and Larch, M. (2013), 'Optimal tariffs, retaliation, and the welfare loss from tariff wars in the melitz model', *Journal of International Economics* **89**(1), 13–25.
- Foley, C. F., Chen, M., Johnson, M. and Meyer, L. (2009), 'Noble group'. 210-021, October 2009.
- Fort, T. C. (2013), Breaking up is hard to do: Why firms fragment production across locations, Working Paper 13-35, Center for Economic Studies, U.S. Census Bureau.
- Gans, J. S. and Shepherd, G. B. (1994), 'How are the mighty fallen: Rejected classic articles by leading economists', *Journal of Economic Perspectives* **8**(1), 165–179.
- Garetto, S. (2013), 'Input sourcing and multinational production', *American Economic Journal: Macroeconomics* **5**(2), 118–151.
- Gennaioli, N. (2013), 'Optimal contracts with enforcement risk', *Journal of the European Economic Association* **11**(1), 59–82.
- Goldberg, P. K., Khandelwal, A. K., Pavcnik, N. and Topalova, P. (2010), 'Imported intermediate inputs and domestic product growth: Evidence from india', *The Quarterly Journal of Economics* **125**(4), 1727–1767.
- Gopinath, G. and Neiman, B. (2013), 'Trade adjustment and productivity in large crises', *American Economic Review* p. Forthcoming.
- Grossman, G. M. and Helpman, E. (2002), 'Integration versus outsourcing in industry equilibrium', *The Quarterly Journal of Economics* **117**(1), 85–120.
- Grossman, G. M. and Rossi-Hansberg, E. (2008), 'Trading tasks: A simple theory of offshoring', *American Economic Review* **98**(5), 1978–97.
- Grossman, S. J. and Hart, O. D. (1986), 'The costs and benefits of ownership: A theory of vertical and lateral integration', *Journal of Political Economy* **94**(4), 691–719.
- Halpern, L., Koren, M. and Szeidl, A. (2011), 'Imported inputs and productivity', *Working Paper, CEU* **8**, 28.
- Harrison, A. and Scorse, J. (2010), 'Multinationals and anti-sweatshop activism', *American Economic Review* **100**(1), 247–273.

- Hart, O. and Moore, J. (1994), ‘A theory of debt based on the inalienability of human capital’, *The Quarterly Journal of Economics* **109**(4), 841–879.
- Helpman, E. (1984), ‘A simple theory of international trade with multinational corporations’, *Journal of Political Economy* **92**(3), 451–471.
- Helpman, E., Itskhoki, O. and Redding, S. (2010), ‘Inequality and unemployment in a global economy’, *Econometrica* **78**, 1239–1283.
- Helpman, E. and Krugman, P. (1985), *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*, The MIT Press.
- Helpman, E., Melitz, M. J. and Yeaple, S. R. (2004), ‘Export versus FDI with heterogeneous firms’, *The American Economic Review* **94**(1), 300–316.
- Helpman, E., Melitz, M. and Rubinstein, Y. (2008), ‘Estimating trade flows: Trading partners and trading volumes’, *The Quarterly Journal of Economics* **123**(2), 441–487.
- Hoefele, A., Schmidt-Eisenlohr, T. and Yu, Z. (2013), Payment choice in international trade: Theory and evidence from cross-country firm level data, CESifo Working Paper Series 4350, CESifo Group Munich.
- Holmstrom, B. (1982), ‘Moral hazard in teams’, *The Bell Journal of Economics* **13**(2), 324.
- Hummels, D. (2007), ‘Transportation costs and international trade in the second era of globalization’, *Journal of Economic Perspectives* **21**(3), 131–154.
- Hummels, D., Ishii, J. and Yi, K.-M. (2001), ‘The nature and growth of vertical specialization in world trade’, *Journal of International Economics* **54**, 75–96.
- Johnson, R. C. and Noguera, G. (2012a), ‘Accounting for intermediates: Production sharing and trade in value added’, *Journal of International Economics* **86**(2), 224–236.
- Johnson, R. C. and Noguera, G. (2012b), *Fragmentation and trade in value added over four decades*, Working Paper, Dartmouth College.
- Jones, R. W. (2000), *Globalization and the Theory of Input Trade*, MIT Press.
- Jones, R. W. and Neary, P. (1984), The positive theory of international trade, in R. W. Jones and P. B. Kenen, eds, ‘Handbook of International Economics’, Vol. Volume 1, Elsevier, p. 1–62.

- Jovanovic, B. (1982), 'Selection and the evolution of industry', *Econometrica* **50**(3), 649.
- Kaufmann, D., Kraay, A. and Mastruzzi, M. (2010), 'The worldwide governance indicators: Methodology and analytical issues'. World Bank Policy Research Working Paper 5430.
- Knack, S. and Keefer, P. (1995), 'Institutions and economic performance: Cross-country tests using alternative institutional measures', *Economics & Politics* **7**(3), 207–227.
- Koopman, R., Wang, Z. and Wei, S.-J. (2014), 'Tracing value-added and double counting in gross exports', *American Economic Review* **104**(2), 459–494.
- Kremer, M. and Maskin, E. (2006), 'Globalization and inequality', *mimeo Harvard University*.
- Krugman, P. (1980), 'Scale economies, product differentiation, and the pattern of trade', *The American Economic Review* **70**(5), 950–959.
- Krugman, P. R. (1979), 'Increasing returns, monopolistic competition, and international trade', *Journal of International Economics* **9**(4), 469–479.
- Kuhn, T. S. (1996), *The Structure of Scientific Revolutions*, University of Chicago Press.
- Lafontaine, F. and Slade, M. (2007), 'Vertical integration and firm boundaries: The evidence', *Journal of Economic Literature* **45**(3), 629–685.
- LaPorta, R., Lopez-de Silanes, F., Shleifer, A. and Vishny, R. (1999), 'The quality of government', *Journal of Law, Economics and Organization* **15**(1), 222–279.
- Leamer, E. E. (1984), *Sources of International Comparative Advantage: Theory and Evidence*, MIT Press.
- Levchenko, A. A. (2007), 'Institutional quality and international trade', *The Review of Economic Studies* **74**(3), 791–819.
- Lu, D. (2011), 'Exceptional exporter performance? evidence from chinese manufacturing firms'. mimeo Rochester.
- Manova, K. (2008), 'Credit constraints, equity market liberalizations and international trade', *Journal of International Economics* **76**(1), 33–47.
- Manova, K. (2012), 'Credit constraints, heterogeneous firms, and international trade', *The Review of Economic Studies* p. rds036.

- Maskin, E. and Tirole, J. (1999), ‘Unforeseen contingencies and incomplete contracts’, *The Review of Economic Studies* **66**(1), 83–114.
- Melitz, M. J. (2003), ‘The impact of trade on intra-industry reallocations and aggregate industry productivity’, *Econometrica* **71**(6), 1695–1725.
- Melitz, M. J. and Ottaviano, G. I. P. (2008), ‘Market size, trade, and productivity’, *Review of Economic Studies* **75**(1), 295–316.
- Melitz, M. J. and Redding, S. (2013*a*), Heterogeneous firms and trade, in ‘Handbook of International Economics’, Vol. Volume 4, Elsevier.
- Melitz, M. and Redding, S. (2013*b*), ‘Firm heterogeneity and aggregate welfare’. mimeo Harvard University.
- Midler, P. (2009), *Poorly Made in China: An Insider’s Account of the Tactics Behind China’s Production Game*, 1 edn, Wiley.
- Milgrom, P. (2000), ‘Putting auction theory to work: The simultaneous ascending auction’, *Journal of Political Economy* **108**(2), 245–272.
- Miroudot, S., Lanz, R. and Ragoussis, A. (2009), ‘Trade in intermediate goods and services’, *OECD Trade Policy Working Papers* (93).
- Mutreja, P. (2013), ‘Equipment and structures capital: Accounting for income differences’. University of Syracuse.
- Nguyen, D. X. (2012), ‘Demand uncertainty: Exporting delays and exporting failures’, *Journal of International Economics* **86**(2), 336–344.
- Novy, D. (2013), ‘International trade without CES: estimating translog gravity’, *Journal of International Economics* **89**(2), 271–282.
- Nunn, N. (2007), ‘Relationship-specificity, incomplete contracts, and the pattern of trade’, *The Quarterly Journal of Economics* **122**(2), 569–600.
- Nunn, N. and Treffer, D. (2008), The boundaries of the multinational firm: An empirical analysis, in ‘The Organization of Firms in a Global Economy’, Harvard University Press, Cambridge, pp. 55–83.
- Nunn, N. and Treffer, D. (2013*a*), Domestic institutions as a source of comparative advantage, in ‘Handbook of International Economics’, Vol. 4, North Holland.
- Nunn, N. and Treffer, D. (2013*b*), ‘Incomplete contracts and the boundaries of the multinational firm’, *Journal of Economic Behavior & Organization* **94**, 330–344.

- Olsen, M. (2013), 'How firms overcome weak international contract enforcement: Repeated interaction, collective punishment, and trade finance'.
- Osborne, M. J. and Rubinstein, A. (1990), *Bargaining and markets*, Academic Press.
- Pierce, J. R. and Schott, P. K. (2009), 'Concording U.S. harmonized system categories over time, Working Paper 14837, National Bureau of Economic Research.
- Rajan, R. G. and Zingales, L. (1998), 'Financial dependence and growth', *American Economic Review* **88**(3), 559–86.
- Ramondo, N., Rappoport, V. and Ruhl, K. J. (2013), 'Horizontal versus vertical foreign direct investment: Evidence from U.S. multinationals'. mimeo NYU Stern.
- Rauch, J. E. (1999), 'Networks versus markets in international trade', *Journal of International Economics* **48**(1), 7–35.
- Redding, S. and Venables, A. J. (2004), 'Economic geography and international inequality', *Journal of International Economics* **62**(1), 53–82.
- Rodrik, D. (2000), 'How far will international economic integration go?', *Journal of Economic Perspectives* **14**(1), 177–186.
- Romalis, J. (2004), 'Factor proportions and the structure of commodity trade', *American Economic Review* **94**(1), 67–97.
- Roth, A. E. (1985), 'A note on risk aversion in a perfect equilibrium model of bargaining', *Econometrica* **53**(1), 207–211.
- Rubinstein, A. (1982), 'Perfect equilibrium in a bargaining model', *Econometrica* **50**(1), 97–109.
- Ruhl, K. (2013), 'An overview of US intrafirm-trade data sources'. mimeo NYU Stern.
- Schmidt-Eisenlohr, T. (2013), 'Towards a theory of trade finance', *Journal of International Economics* **91**(1), 96–112.
- Schott, P. K. (2004), 'Across-product versus within-product specialization in international trade', *The Quarterly Journal of Economics* **119**(2), 647–678.
- Schott, P. K. (2008), 'The relative sophistication of chinese exports', *Economic Policy* **23**(53), 5–49.

- Schott, P. K. (2010), 'U.S. manufacturing exports and imports by SIC or NAICS category and partner country, 1972 to 2005'. Yale School of Management.
- Segura-Cayuela, R. and Vilarrubia, J. M. (2008), Uncertainty and entry into export markets, Banco de España Working Paper 0811, Banco de España.
- Simonovska, I. and Waugh, M. E. (2011), The elasticity of trade: Estimates and evidence, Working Paper 16796, National Bureau of Economic Research.
- Tang, H. and Zhang, Y. (2012), 'Quality differentiation and trade intermediation'.
- Thomas, J. and Worrall, T. (1994), 'Foreign direct investment and the risk of expropriation', *Review of Economic Studies* **61**(1), 81–108.
- Tinbergen, J. (1962), *Shaping the World Economy; Suggestions for an International Economic Policy*, New York: Twentieth Century Fund.
- Tintelnot, F. (2013), 'Global production with export platforms'. mimeo Princeton University.
- Trefler, D. (1993*a*), 'International factor price differences: Leontief was right!', *Journal of Political Economy* **101**(6), 961–987.
- Trefler, D. (1993*b*), 'Trade liberalization and the theory of endogenous protection: An econometric study of U.S. import policy', *Journal of Political Economy* **101**(1), 138–160.
- Trefler, D. (1995), 'The case of the missing trade and other mysteries', *The American Economic Review* **85**(5), 1029–1046.
- Trefler, D. and Zhu, S. C. (2010), 'The structure of factor content predictions', *Journal of International Economics* **82**(2), 195–207.
- Vanek, J. (1968), 'The factor proportions theory: The n-factor case', *Kyklos* **21**(4), 749–756.
- Waugh, M. E. (2010), 'International trade and income differences', *American Economic Review* **100**(5), 2093–2124.
- Williamson, O. E. (1985), *The Economic Institutions of Capitalism*, The Free Press, New York.
- Wright, G. C. (2014), 'Revisiting the employment impact of offshoring', *European Economic Review* **66**, 63–83.

- Yeats, A. J. (2001), Just how big is global production sharing?, *in* S. W. Arndt and H. Kierzkowski, eds, 'Fragmentation New Production Patterns in the World Economy', Oxford University Press, Oxford.