Essays on the Logistics of International Trade

Abiy Ayalew Teshome
Addis Ababa, Ethiopia

M.A. Economics, University of Virginia, 2012
B.A. Economics, Stanford University, 2011

A Dissertation presented to the Graduate Faculty of the University of Virginia
in Candidacy for the Degree of Doctor of Philosophy

Department of Economics

University of Virginia
December 2018

Abstract

This dissertation tackles two complementary issues in the logistics of international trade. First, what determines transportation costs? Second, how do parties to international transactions execute the costly and often unpredictable tasks associated with cross-border manufacturing and distribution? Our understanding of these topics informs modern practices like offshoring and just-in-time production, which feature increasingly complex supply chains, and are therefore susceptible to distortions arising from market power in international shipping, and other institutional barriers to the free movement of goods across borders.

In Chapter 1, I endogenize transportation costs in a heterogeneous firm model of trade by introducing oligopolistic maritime carriers that move goods from manufacturers to final consumers. Within each destination, competition among manufacturers from various source countries generates a system, across source countries, of interdependent demands for transport. I then test the model using data on Ecuadorian auto imports from 2007 to 2012 in Chapter 2. Specifically, I determine whether the prevailing freight rates and the number of vehicles transported by a given number of carriers over a given length of time are jointly rationalizable by time-varying transport demands and convex, time-invariant cost functions in shipping. As expected, it is easier to rationalize shipping activity among smaller groups of carriers over shorter horizons as Cournot outcomes. I then bound carrier marginal costs using the set of rationalizable observations, and find evidence of dwindling profit margins since the beginning of the Great Recession, thus easing fears of distortions due to market power.

In Chapter 3, I study optimal contracting in international shipping, offering the first breakdown of the delivery process into its various components. I present stylized facts using detailed Colombian transaction-level data, showing that the allocation of delivery- related tasks within buyer-seller pairs constitutes an important margin of trade. I then model the allocation of control over such tasks. The model describes a sequential production process – consisting of manufacturing and distribution – in an incomplete-contracting environment. Contracts between exporters and importers specify shipping volumes and assign responsibility for delivery to one of the parties. The pair sequentially bargain over the value added by unverifiable efforts at each production stage. Bargaining power initially resides with the exporter, but transfers to the party in charge of distribution at the factory gate, owing to the latter's residual control rights over the output from delivery related activities. Trading partners thus allocate delivery rights to minimize the distortionary effects of these bargaining externalities. In contrast to the vast holdup literature, I find that the exporter has a strong motive to over-invest in quality, and should thus be deprived of consignment rights unless its effort is particularly important in the delivery process.

Acknowledgements

I am deeply indebted to my advisors John McLaren and Kerem Coşar for their guidance, encouraging nudges, and remarkable patience. I am grateful to Simon Anderson and Peter Debaere for feedback at various stages of my dissertation. Special thanks to Marc Santugini for reigniting my interest in teaching, and to Sage Bradburn, Sheetal Sekhri, and Richard Tanson for helping me navigate countless issues during my graduate studies. I would like to thank my friends, Samora Kariuki, Scott Hong, İlhan Güner, Guillermo Hausmann, Nate Pattison, Selcen Çakır, Ramiro Burga, Miguel Mascarua, Hundanol Kebede, and Myunghwan Yoo for their companionship and help over the years. The Bankard Fund for Political Economy provided financial support for the final chapter of this dissertation.

I dedicate this thesis to my sister, Addis, my parents, Ayalew Teshome and Snafkish Desta, and to Audrey. Your love and tireless support saw me through this.

Contents

Chapter 1

International trade with an oligopolistic transport sector

1.1 Introduction

Given the purported gains from international trade, there is immense interest in understanding the barriers to free trade. Anderson and van Wincoop (2004) list transportation and distribution costs, policy barriers, information and contract enforcement costs, and costs arising from conflicting institutions as potential impediments to trade. These barriers reduce welfare by limiting the gains from specialization and scale, and reducing competition in final goods markets.

However, as Anderson and van Wincoop (2004) and Mesquita Moreira, Volpe Martincus, and Blyde (2008) point out, transport costs stand out for three main reasons. First, they are significantly higher, in ad valorem terms, than traditional barriers like tariffs. Second, freight rates are often less predictable than trade policy barriers, which are subject to drawn-out negotiations among trading partners. Finally, freight costs are not fixed by fiat. Instead, they are equilibrium outcomes in transport markets often dominated by a handful of firms. For example, the five largest liner shipping companies in 2017 controlled nearly half of the global fleet of containerships, accounting for close to 60 percent of total capacity (UNCTAD, 2017), with (Hummels, Lugovskyy, and Skiba, 2009) reporting that more than half of all country pairs were served by at most 3 ships in 2006.

This cursory view of the transport sector suggests that potential gains from trade are vulnerable to the whims of firms in a highly concentrated industry. I explore this distinguishing feature transport costs by modeling the economic activities of three groups of agents along international value chains. At one end, utility maximizing consumers in each country desire as wide a range of differentiated goods. At the other end, heterogeneous, monopolistically competitive manufacturers based in each country seek as many profitable export opportunities for their varieties. Finally, oligopolistic maritime carriers facilitate trade between consumers and manufacturers by providing a homogenous transportation service.

In the model, tougher competition in each country – in the form of a larger set of rivals – eats into the residual demand curve of any given manufacturer, lowering demand for transport to the country in question among active manufacturers, and choking off demand for transport from the least efficient firms. Taken together, these intensive and extensive margin effects of aggregate transport use imply that individual manufacturer willingness–to–pay for transport to a given destination declines in aggregate transport use along all routes bound for that destination. Specifically, I find that inverse demand is additively separable in a term that captures technological differences across countries and consumer love–of–variety, and a term that captures competition over final goods, mediating the negative spillovers of transport use. The second component generates a system of interdependent inverse demands for transport, exacerbating concerns that carriers may suppress their output even further if they internalize the negative effects of capacity along a given route on their profitability in all routes headed for the same destination.

These theoretical insights join a recent series of papers that endogenize transport costs. Like Brancaccio, Kalouptsidi, and Papageorgiou (2017), who model trade in dry bulk commodities, I endogenize (one component of) trade costs, showing that equilibrium freight rates depend on more than just the bilateral distance between a pair of countries. Given the destination–based competition among manufacturers, my model relates trade costs along all routes bound for a particular country. In contrast, Brancaccio et al. (2017) derive a more complex set of freight rate dependencies, relating a country's trade costs to its entire network of trading partners. However, I micro-found manufacturer willingness–to–pay for transport from profit-maximizing behaviour, while Brancaccio et al. (2017) assume an exogenous distribution of valuations for transport services.

My demand–driven links between shipping markets stand in contrast to a vast body of work on supply–side constraints in shipping. Behrens and Picard (2011), Friedt and Wilson (2017), Wong (2017), and Ishikawa and Tarui (2018) study the "backhaul problem," which forces carriers to commit to the maximum capacity required for a round trip, and therefore imposes opportunity costs to shipping along routes linking countries with unbalanced trade. Wong (2017) highlights the implications of such constraints for trade liberalization when trade is facilitated by a perfectly competitive transport sector. For example, French exports to Ecuador rise in response to a reduction in Ecuadorian import tariffs on French products. In addition, Ecuadorian exports to France also increase because the incoming vessels would rather avoid an empty return voyage to France. In contrast, my (demand–side) restriction links various source countries exporting to a given destination. To return to Wong's example, a reduction in Ecuadorian tariffs on imports of French products increases French exports to Ecuador at the expense of exports from other countries, and is accompanied by a fall price of shipping from any given destination to Ecuador.

Finally, this paper assumes the same market structure in shipping as Behrens, Gaigné, and Thisse (2009) and (Hummels, Lugovskyy, and Skiba, 2009), who consider oligopolistic, quantity–setting carriers. However, Behrens, Gaigné, and Thisse (2009) model a single shipping market, with manufacturers hiring transport services at the same freight rate regardless of their location or desired final market. Admittedly, this is less stark of an assumption in their two–region model, which focuses on agglomeration forces. In contrast, I define shipping markets as unidirectional trips between two countries, allowing freight rates and capacities to differ depending on the direction of travel. Likewise, (Hummels, Lugovskyy, and Skiba, 2009) allow distinct shipping markets, but assume that each origin–destination pair is served by a distinct set of carriers, thus sidestepping multi-market contact, a key theme in this paper.

Section 1.2 presents the theoretical framework of trade with an oligopolistic transport sector. Specifically, I derive demand for shipping in Section 1.2.1, before characterizing equilibrium in the shipping sector in Section 1.2.2.

1.2 Theory

The economy consists of a finite set of countries, each populated by workers/consumers with preferences over a homogenous good and a continuum of differentiated varieties. Manufacturers in the differentiated sector vary in their productivity, and reach consumers by hiring homogenous transport services from carriers. Figure 1.1 illustrates the timing. I first derive consumer demand for final goods, taking goods prices as given, then obtain manufacturer demand for transport, taking consumer demand and freight rates as given, before characterizing equilibrium trade and freight rates as functions of the underlying consumer preferences and manufacturing technologies.

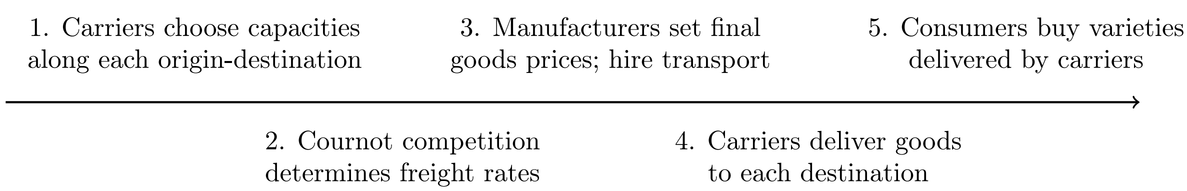

Figure 1.1: Timing

Preferences

As in Melitz and Ottaviano (2008), consumers in destination d have identical quasilinear preferences over a homogenous good, q_0, and a continuum of horizontally differentiated varieties. Specifically, each d–based consumer derives utility

$$U_d^c = q_0^c + \int_{\Omega_d} \left(\alpha q^c(\omega) - \frac{\gamma}{2} q^c(\omega)^2\right) \mathrm{d}\omega - \frac{\eta}{2}\left(\int_{\Omega_{od}} q^c(\omega)\, \mathrm{d}\omega\right)^2, \tag{1.1}$$

from the bundle $(q_0^c, (q^c(\omega))_{\omega\in\Omega_d})$, where the equilibrium set of varieties consumed in d, $\Omega_d = \cup_o \Omega_{od}$, likely contains of imported varieties. Let N_{od} denote the mass of varieties produced in o and consumed in d. The parameters α and η govern substitution between the homogeneous good and the differentiated varieties, while γ measures the love-of-variety in the differentiated goods sector.

Consumer c maximizes utility U_d^c subject to the budget constraint

$$q_0^c + \sum_o \int_{\Omega_{od}} p_{od}(\omega)\, q_{od}^c(\omega)\ \mathrm{d}\omega \leq y_d^c,$$

where per-capita income,

$$y_d^c \equiv \overline{q}_0^c + w_d + \frac{1}{\sum_o L_o}\Pi, \tag{1.2}$$

consists of the value of an exogenous endowment $\overline{q}_0^c > 0$ of the homogenous good, payments to an inelastically supplied unit of labour, and dividends from a global mutual fund. The endowment is large enough to guarantee positive demand for the homogenous good, eliminating any income effects in demand for the differentiated good. Each consumer owns a single share of the mutual fund, which evenly distributes profits from all manufacturers, Π. Carrier operators retain profits and do not demand any final goods.

Although consumers do not distinguish between varieties, it will prove helpful to associate each variety with its country of origin. With this in mind, utility maximization yields a linear inverse demand from d–based consumers for variety ω from o,

$$p_{od}(\omega) = p_d^{\max}(Q_{1d}, \ldots, Q_{Od}) - \frac{\gamma}{L_d} q_{od}(\omega), \tag{1.3}$$

where

$$Q_d = \sum_o Q_{od} \equiv \sum_o \int_{\Omega_{od}} q_{od}(\omega)\ \mathrm{d}\omega, \qquad p_d^{\max}(Q_{1d}, \ldots, Q_{Od}) \equiv \alpha - \frac{\eta}{L_d} Q_d \tag{1.4}$$

are the endogenous aggregate consumption of differentiated varieties, and the choke price in the market for differentiated goods in d. The aggregate Q_d is exogenous from the perspective of individual manufacturers, instead shifting demand for individual varieties by changing the choke price. By definition, competition among differentiated goods producers is tougher when Q_d is large.[1]

The choke price, which defines the set of varieties consumed in d according to

$$\Omega_{od}(Q_{1d}, \ldots, Q_{Od}) = \{\omega : p_{od}(\omega) \leq p_d^{\max}(Q_{1d}, \ldots, Q_{Od})\},$$

cannot exceed α, the marginal utility of the very first units of differentiated varieties. Fur-

[1] While the aggregate Q_d is sufficient for the degree of competition in d, I retain the entire profile of varieties across sources in the argument of the choke price. This simplifies the upcoming derivation of manufacturer demand for transport.

ther, large values of α or small values of η, which reflect greater demand for differentiated goods relative to the homogenous good, shift demand for any given variety outward, dampening competition among varieties. Finally, demand becomes more elastic as consumer love–of–variety, indexed by γ, increases.

Having had their fill of differentiated goods, consumers spend residual income on the homogenous good

$$q_0^c = y_d^c - \sum_o \int_{\Omega_{od}} p_{od}(\omega)\, q_{od}^c(\omega)\, \mathrm{d}\omega,$$

and always purchase a positive amount because endowments, $\overline{q}_0^c > 0$, are sufficiently large.

Production and shipping technologies

The homogenous good is costlessly tradable and sold in a perfectly competitive market, so that its price does not vary across countries. I therefore use this outside good as a numéraire. Production in the differentiated sector is subject to constant returns to scale, with a denoting unit labour requirements, drawn, at no cost, by each of $\widetilde{N}_o$ potential manufacturers from a distribution $G_o(\cdot)$. After drawing its unit labour requirement, a potential manufacturer may produce a single variety, allowing us to index manufacturers and their varieties by the unit labour requirement a. Further, the mass of varieties produced in o and consumed in d, N_{od}, is simply the number of $o-$based manufacturers selling in d.

Consumers value "goods on the shelf" rather than "goods at manufacturing plants." Subsequently, manufacturers from source o must effect a two–step process in order to supply consumers in d. They first transform labour into intermediate products resting at their o–based factories, and then transport these goods to store shelves in market d by hiring carrier services. Even manufacturers selling in their local market (setting $d = o$) require transportation services. Specifically, manufacturers combine the intermediate output with transport services in fixed proportions, so that manufacturer a produces

$$q(l, s; a) = \min\left\{\frac{l}{a}, s\right\} \tag{1.5}$$

"on-the-shelf" units of its variety by hiring l units of labour and s units of shipping services.[2]

The corresponding unit cost function depends on wages paid to factory workers, w_o, and per-unit freight rates. In principle, the relevant freight rate may vary across manufacturers. For simplicity, I assume that there is only one transport mode, abstracting from the literature on selection into various transport modes (Cosar and Demir, 2018). Carriers set capacities along origin-destination routes but cannot segment transport demand within an (o, d) route. Further, manufacturers find it prohibitively costly to ship from o to d through any transit countries and must pay the constant freight rate t_{od} per unit of transport services to d.

[2]Stepping back from the particular application in the shipping market, the analysis may apply to any non-competitively supplied subset of manufacturer inputs.

Summarizing, manufacturer a faces a *delivered* unit cost

$$C_{od}(a, t_{od}) \equiv w_o \tau_{od}\, a + t_{od} \tag{1.6}$$

of supplying its goods to market d. Delivered costs are additive in production costs, which depend on labour expenditures $w_o a$, and tariffs levied on the cost of the goods, $\tau_{od} \geq 1$, and the freight charges. This transportation cost, t_{od}, represents the key departure from Melitz and Ottaviano (2008), which instead focuses on differences in manufacturing costs, $w_o \tau_{od} a$. Note that given the fixed proportions technology, delivered costs increase one–for–one with freight rates, regardless of manufacturer efficiency.

1.2.1 Demand for transport services

Manufacturers do not inherently value transportation. Instead transport is valuable to the extent that it grants them access to final consumers. Assuming product markets are segmented across countries, firm a chooses destination–specific outputs to maximize profits, taking as given the endogenous aggregate consumption of differentiated goods, which determines its residual demand through the choke price, and the relevant freight rate, which determines its delivered cost. Since manufacturers combine output and transport services in fixed proportions, I recast the profit–maximization problem so that firms choose the optimal level of transport services. That is, firm a solves

$$\max_{(s_{od})} \sum_d \left[p_d^{\max}(S_{1d}, \ldots, S_{Od}) - \frac{\gamma}{L_d} s_{od} - C_{od}(a, t_{od}) \right] L_d\, s_{od}. \tag{1.7}$$

Conditional on the transport demand across source countries and the relevant freight rate, an o–based manufacturer with unit labour requirement a hires

$$s^*_{od}(a |\, S_{1d}, \ldots, S_{Od}, t_{od}) \equiv \frac{L_d}{2\gamma} \left[p_d^{\max}(S_{1d}, \ldots, S_{Od}) - C_{od}(a, t_{od}) \right]^+ \tag{1.8}$$

units of transport to market d, where $x^+ \equiv \max\{0, x\}$ is the positive part of x. Firms anticipating delivered costs above the destination choke price opt to stay out of market d. Specifically, the marginal o–based manufacturer active in d, denoted $\hat{a}_{od}$, is the largest unit labour requirement consistent with nonnegative profits,

$$\begin{aligned} &C_{od}(a, t_{od}) \equiv w_o \tau_{od}\, a + t_{od} \leq p_d^{\max}(S_{1d}, \ldots, S_{Od}) \\ &\qquad \iff a \leq \frac{1}{w_o \tau_{od}} \left[p_d^{\max}(S_{1d}, \ldots, S_{Od}) - t_{od} \right] \equiv \hat{a}_{od}(S_{1d}, \ldots, S_{Od}, t_{od}). \end{aligned} \tag{1.9}$$

All else equal, the marginal o–based manufacturer is more efficient ($\hat{a}_{od}$ is low) when tariff–adjusted wages are low, and when firms expect high freight rates or fierce competition.

Figure 1.2: Profit-maximizing prices of final goods, given choke price and freight rate

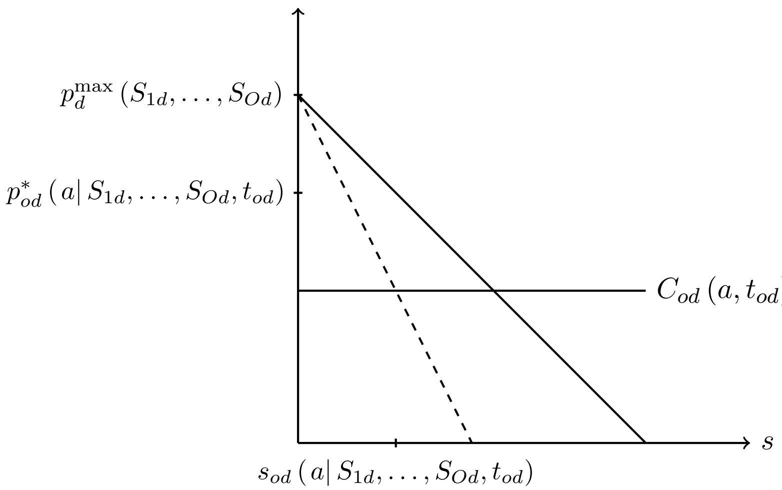

Notes: Profit–maximizing price, p_{od}^*, and output, s_{od}^*, of a manufacturer with delivered (inclusive of transport cost) unit cost function $C_{od}(a,t_{od})$. Manufacturers, indexed by their unit labour requirements a, face identical residual demands, $p_d(q) = p_d^{\max}(S_{1d},\ldots,S_{Od}) - (\gamma/L_d)\,q$, for their varieties. Marginal revenue, $p_d^{\max}(S_{1d},\ldots,S_{Od}) - (2\gamma/L_d)\,q$, is broken line.

Figure 1.3 plots $\hat{a}_{od}(S_{1d},\ldots,S_{Od},\cdot)$. Holding the choke price fixed at $p_d^{\max} \equiv p_d^{\max}(S_{1d},\ldots,S_{Od})$, selection into exporting to d from o depends on the prevailing freight rate, t_{od}, and the distribution of manufacturer costs. At one extreme, all $o-$based manufacturers remain inactive if $t_{od} > p_d^{\max}(S_{1d},\ldots,S_{Od})$; if the freight rate exceeds the choke price, then no manufacturer with positive factory costs can survive. At the other extreme, all firms serve the market if freight rates are sufficiently low ($t_{od} \leq p_d^{\max}(S_{1d},\ldots,S_{Od}) - w_o\tau_{od}\,a_o^M$), which occurs if a_o^M is relatively low, so that origin o has an absolute advantage in the sector. I consider equilibria featuring non-trivial selection into exporting, which requires intermediate levels of t_{od}, or unbounded production costs, $a_o^M = \infty$.

Returning to the set of active firms, many of the individual-level comparative statics mirror the results in Melitz and Ottaviano (2008). All else equal, manufacturers sell more output – and therefore demand more transport services – in larger destinations. Further, fierce competition in destination d lowers sales among active manufacturers, regardless of their country of origin.

As for delivered costs (see (1.6)), more efficient (low a) manufacturers in countries facing low trade barriers, τ_{od}, demand more transport services. After all, delivered costs are increasing in manufacturer unit costs, and are particularly sensitive to productivity in countries with high tariff barriers. The key departure from Melitz and Ottaviano (2008) lies in the freight component of delivered costs. All else equal, manufacturers facing low freight rates sell more output and demand more transport. Finally, note that marginal reductions in competitiveness, $p_d^{\max}(S_{1d},\ldots,S_{Od})$, exactly offset marginal increases in delivered costs, $C_{od}(a,t_{od})$.

Figure 1.3: Selection of o-based manufacturers into serving market d

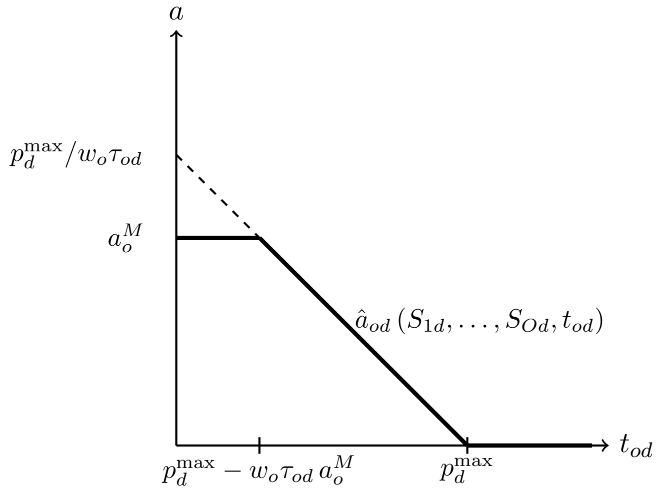

Notes: Conditional on the the choke price, $p_d^{\max} \equiv p_d^{\max}(S_{1d}, \ldots, S_{Od})$, and the prevailing freight rate, t_{od}, only manufacturers with unit-input-requirements below the solid line self-select into serving a given market. The parameter a_o^M defines the support of the distribution of unit-labour requirements, $G_o(a)$. All manufacturers export if a_o^M is sufficiently low.

Combining (1.8) and (1.9), the route–(o, d) aggregator,

$$\sigma_{od}(S_{1d}, \ldots, S_{Od}, t_{od}) \equiv \widetilde{N}_o \int_0^{\hat{a}_{od}(S_{1d},\ldots,S_{Od},t_{od})} s^*_{od}(a | S_{1d}, \ldots, S_{Od}, t_{od}) \, \mathrm{d}G_o(a), \tag{1.10}$$

returns the level of aggregate demand for transport from o to d when all o–based manufacturers *anticipate* a d–bound aggregate shipping profile, $(S_{1d}, \ldots, S_{Od})$, and incur freight rates, t_{od}. Proposition 1.1 and Figure 1.4 characterize the route–(o, d) aggregator as a function of the anticipated profile of global d–bound transport demand, the price of shipping from o to d, and all exogenous country characteristics.

Proposition 1.1. (Demand aggregator) *Between them, o–based manufacturers demand more transport to larger, more accessible markets, whose consumers (i) value differentiated goods more than the homogenous good (α large and/or η small); and (ii) freely substitute among varieties (γ small). Further, aggregate demand is increasing in the number of potential manufacturers in country o, and decreasing in the prevailing wage.*
Crucially, aggregate demand for transport is higher when manufacturers anticipate (i) low shipping prices; and (ii) low levels of aggregate transport use from any given country.

For a fixed freight rate t_{od}, the aggregator is unlikely to return the o–th component of an arbitrary profile, $(S_{1d}, \ldots, S_{od}, \ldots, S_{Od})$, of d–bound transport demand levels. However, we can guarantee a solution to this fixed point problem by adjusting the freight rate based on the aggregate demand profile.

Figure 1.4: Properties of the demand aggregator

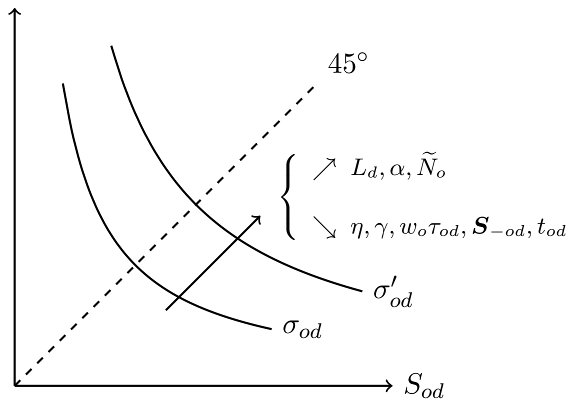

Notes: The aggregator (1.10) is decreasing in the anticipated freight rate and transport use along any route destined for d. Inverse demand reconciles individual manufacturer demands and aggregate demand along a route according to the fixed point problem (1.11). The route (o, d) demand aggregator, σ_{od}, is decreasing in all shipping levels both own, S_{od}, and cross, $S_{o'd}$, capacities.

Definition 1.1. Manufacturer willingness–to–pay (inverse demand) for transport from o to d, denoted $t_{od}(S_{1d}, \ldots, S_{Od})$, implicitly reconciles aggregate and individual demand for transport from o to d. Specifically,

$$\sigma_{od}\left(S_{1d}, \ldots, S_{od}, \ldots, S_{Od}\ ,\ t_{od}\left(S_{1d}, \ldots, S_{od}, \ldots, S_{Od}\right)\right) = S_{od} \tag{1.11}$$

for all aggregate demand profiles $(S_{1d}, \ldots, S_{od}, \ldots, S_{Od})$, where the route–$(o, d)$ aggregator, $\sigma_{od}(\cdot)$, is given by (1.10).

The function $t_{od}(S_{1d}, \ldots, S_{Od})$ reflects consumer preferences for final goods; after all manufacturers value shipping services as a means of reaching consumers. In addition, exporter willingness–to–pay for transport depends on the supply side of the market for differentiated goods. Specifically, demand for transport varies with the number of potential manufacturers, the returns to labour, and the trade barriers with the final goods market.[3]

Finally, driven by their aversion to competition over final goods consumers, manufacturers are willing to pay more for transport when they anticipate low transport use along any d–bound lane. Figure 1.4 illustrates the negative cross–capacity effects. Starting at a profile $(S_{1d}, \ldots, S_{od}, \ldots, S_{Od}, t_{od})$ that solves the fixed point problem (1.11), consider an increase in aggregate transport use along route (o', d), for $o' \neq o$. By Proposition 1.1, this shifts the aggregator downward, lowering aggregate demand along (o, d) beneath its initial value. Proposition 1.1 implies that, all else equal, the freight rate from o must fall to return demand from o to return its initial value. That is, inverse demand along a given route

[3]See Section 1.C for an alternative approach that obtains manufacturer willingness–to–pay for transport by inverting a system of demands for shipping.

is decreasing in capacities elsewhere. Consider now an increase in capacity along (o, d), a movement down the curve reflecting a decline in demand for shipping among o–based firms, each *anticipating* fiercer competition from its compatriots. Yet again, freight rates must fall to encourage o–based to serve market d. In other words, the law of demand holds.

The remainder of this section characterizes $t_{od}(S_{1d}, \ldots, S_{Od})$ more formally by applying Proposition 1.1. See Section 1.A for all proofs.

Proposition 1.2. *Inverse demand along route (o, d) takes the form*

$$t_{od}(S_{1d}, \ldots S_{od}, \ldots, S_{Od}) = p_d^{\max}(S_{1d}, \ldots S_{od}, \ldots, S_{Od}) - \tilde{t}_{od}(S_{od}), \tag{1.12}$$

where $\tilde{t}_{od}(\cdot)$ is an increasing, concave function passing through the origin. Specifically, inverse demand along (o, d) satisfies the differential equations,

$$\frac{\partial t_{od}}{\partial S_{od}} = \frac{\partial p_d^{\max}}{\partial S_{od}} - \frac{\partial \tilde{t}_{od}}{\partial S_{od}} = -\left(\frac{\eta}{L_d} + \frac{2\gamma}{L_d}\frac{1}{N_{od}}\right), \qquad \frac{\partial t_{od}}{\partial S_{o'd}} = \frac{\partial p_d^{\max}}{\partial S_{o'd}} = -\frac{\eta}{L_d} \quad \text{for } o' \neq o, \tag{1.13}$$

subject to the initial condition $t_{od}(0, (S_{jd})_{j\neq o}) = p_d^{\max}$ for all $(S_{jd})_{j\neq o}$.

Inverse demand for shipping bears a striking resemblance to consumer demand for final goods. It consists of a *competitive effect* that operates through the choke price for differentiated varieties, and a *compatriot effect* operating through $\tilde{t}_{od}(\cdot)$.[4]

1.2.1.1 Competitive effect of transport use

This final goods choke price, $p_d^{\max}(S_{1d}, \ldots S_{od}, \ldots, S_{Od})$, contributes to exporter demand for transport along any d–bound route, and is common across all source countries. The competitive channel therefore inherits the properties of the choke price, discussed after equation (1.4). In particular, inverse demand for shipping is decreasing in transport use along any route destined for d, reflecting manufacturer aversion to competition.

Proposition 1.3. (Competition and transport use) *Regardless of their location, all manufacturers are willing to pay more for transport to large destinations (L_d large), where consumers value differentiated goods more than the homogenous good (α large and/or η small).*

Intuitively, incumbent manufacturers scale up production to larger markets or markets with high relative demand for differentiated sector goods. Furthermore, marginally inactive firms now opt to produce for the market in question. Combining these intensive and extensive margin effects delivers an unambiguous increase in the level of aggregate demand for shipping

[4] Strictly speaking, these terms are interdependent because S_{od} appears in both expressions. Nonetheless, this distinction is helpful as several comparative statics operate solely through the choke price. They would be independent if each country country was "small" relative to the global economy, in which case the choke price would be $\alpha - \frac{\eta}{L_d}\int_{o'} S_{o'd}\,\mathrm{d}o'$.

from o to d. By Proposition 1.1, the freight rate along (o, d) must rise so as to restore initial demand for shipping.[5]

1.2.1.2 Compatriot use of transport services

The competitive component of demand for transport depends solely on destination characteristics. What of the manufacturers who demand these services? The compatriot effect mediates the effects of country–of–origin characteristics like number of potential producers, $\widetilde{N}_o$, the productivity distribution, $G_o(a)$, and bilateral characteristics like tariffs on exports from o to d, τ_{od}.

Proposition 1.4. (Source country characteristics and transport demand) *Willingness-to-pay for transport from o to d is (i) decreasing in the wage w_o; (ii) increasing in $\widetilde{N}_o$, the number of potential manufacturers in country o; and (iii) and independent of conditions in other source countries.*

$$\partial t_{od}/\partial \widetilde{N}_o > 0, \qquad \partial t_{od}/\partial w_o < 0, \qquad \partial t_{od}/\partial \widetilde{N}_{o'} = \partial t_{od}/\partial w_{o'} = 0, \quad o' \neq o.$$

Conditions elsewhere have no direct effect on demand for transport along a given route. After all, conditional on market toughness, profit–maximizing manufacturers in any given country are indifferent towards rivals located elsewhere. As for the own–country effects, while transport demand among active firms is insensitive to the number of potential manufacturers, a larger value of $\widetilde{N}_o$ encourages a larger mass of entrants, all else equal. By Proposition 1.1, the freight rate along (o, d) must fall to restore equilibrium. Finally, wages and free–on–board tariffs have similar effects on transport demand since they are complementary in delivered costs.

Proposition 1.5. (Bilateral trade barriers and transport demand) *Willingness-to-pay for transport from o to d is decreasing in the (free–on–board) tariff τ_{od}, and independent of conditions in other source countries,*

$$\partial t_{od}/\partial \tau_{od} < 0, \qquad \partial t_{od}/\partial \tau_{o'd} = 0 \qquad \text{for } o' \neq o.$$

Higher wages in o or greater tariffs on imports from o curtail production for market d by o–based firms. Specifically, large values of $w_o \tau_{od}$ shrink the set of active firms, and reduce exports among surviving firms. By Proposition 1.1, the freight rate along (o, d) must fall to encourage enough manufacturer activity, thereby keeping exports from o to d constant.

Proposition 1.6. (Love–of–variety and transport demand) *All else equal, manufacturers in any given source country are willing to pay less for transport to destinations where*

[5] Note that, unlike the preference parameters α and η, market size also operates through the compatriot channel (where it also raises manufacturer willingness–to–pay). As a result, the overall effect of an increase in L_d on inverse demand is positive, which does not contradict Proposition 1.3.

consumers freely substitute among varieties. That is,

$$\partial t_{od}/\partial\gamma < 0.$$

Intuitively, manufacturer residual demand becomes more elastic as love–of–variety increases. Reassuringly, the degree of product differentiation has no extensive margin effect, only lowering demand among the set of exporters. By Proposition 1.1, the freight rate along (o, d) must fall to generate enough exports to counteract the initial effect.

The extensive margin of trade Armed with a better understanding of the compatriot effect, consider the number of o–based manufacturers selling to d, given by the fraction of the exogenous number of potential manufacturers with sufficiently low unit labour requirements. Substituting (1.12) into the marginal exporter's unit labour requirement in (1.9),

$$\hat{a}_{od}\left(S_{od}\right) \equiv \frac{1}{w_o\tau_{od}}\tilde{t}_{od}\left(S_{od}\right), \tag{1.9$'$}$$

so that conditional on shipping activity along route (o, d), the unit labour requirement of the marginal o–based manufacturer active in d is independent of shipping activity along all other routes. Subsequently, the number of manufacturers serving d from o is also determined solely by S_{od},

$$N_{od}\left(S_{od}\right) = G_o\left(\hat{a}_{od}\left(S_{od}\right)\right) \times \widetilde{N}_o. \tag{1.14}$$

Proposition 1.7. *The marginal exporter's unit labour requirement (and hence number of o–based firms) selling S_{od} units of the differentiated good to d, denoted $\hat{a}_{od}\left(S_{od}\right)$ and $N_{od}\left(S_{od}\right)$, are increasing in the traded volume S_{od}, and consumer love–of–variety, γ, and decreasing in trade costs between o and d, and the size of the market. Finally, a larger number of potential manufacturers implies a lower exporting threshold but more firms.*

The next section characterizes equilibrium in the carrier game, describing the sensitivity of aggregate capacities and the corresponding freight rates to origin, destination, and bilateral characteristics. Although this task does not require a closed form expression for the compatriot effect, $\tilde{t}_{od}(\cdot)$, I wrap up this section by solving for inverse demand when manufacturer costs are Pareto-distributed, easing comparison with the benchmark Melitz and Ottaviano (2008).

1.2.1.3 Pareto distributed manufacturing costs: Willingness–to–pay

Suppose unit labour requirements in source o are distributed according to

$$G_o\left(a\right) \equiv \left(a/a_o^M\right)^{\theta}, \qquad 0 < a \leq a_o^M,$$

where $a_o^M > 0$ is the least efficient $o-$based manufacturer, and $\theta \geq 1$ determines the dispersion of cost draws in any given country. In this case, the compatriot effect takes the form

$$\tilde{t}_{od}\left(S_{od}\right) = \kappa_{od}\, S_{od}^{\frac{1}{\theta+1}}, \tag{1.15}$$

where

$$\kappa_{od} \equiv \left((\theta+1)\, \frac{2\gamma}{L_d} \frac{1}{\left(a_o^M\right)^{-\theta} \widetilde{N}_o} \frac{1}{\left(w_o \tau_{od}\right)^{-\theta}} \right)^{\frac{1}{\theta+1}} > 0 \tag{1.16}$$

is independent of the final goods demand shifters α and η, as claimed in Proposition 1.3. Instead, the compatriot effect is stronger among unequal (large θ) producers of highly differentiated (large γ) goods, and weaker when producers compete in large markets, or reside in countries enjoying an absolute advantage, $\left(a^M\right)^{-\theta} \widetilde{N}_o$, or relatively free trade with the final goods market, (high $\left(w_o \tau_{od}\right)^{-\theta}$).

Substituting (1.15) into (1.9′), the marginal $o-$based exporter's unit input requirement and the mass of varieties simplify to

$$\begin{aligned} \hat{a}_{od}\left(S_{od}\right) &= \frac{\kappa_{od}}{w_o \tau_{od}} S_{od}^{\frac{1}{\theta+1}} \\ N_{od}\left(S_{od}\right) &= \widetilde{N}_o \times \left(\frac{1}{a_o^M} \frac{\kappa_{od}}{w_o \tau_{od}} S_{od}^{\frac{1}{\theta+1}} \right)^{\theta} \end{aligned} \tag{1.17}$$

which, in line with Proposition 1.7, depend on the transport sector only through own transport demand, S_{od}. Conditional on destination and origin characteristics, higher shipping capacity in a given market is associated with a larger mass of exporters from that market, and a less efficient marginal exporter.

1.2.2 Equilibrium in the shipping sector

This section characterizes equilibrium outcomes in the shipping sector by combining manufacturer demand for shipping, derived in the previous section, with an oligopolistically structured transport sector, consisting of $F \geq 1$ carriers, each transporting goods along routes defined by unidirectional trips between countries.[6]

Content with the intricacies of an interdependent system of demands for transport, I consider a relatively simple supply side. Specifically, the cost of providing transport services between any pair of countries is independent of activity elsewhere.[7] Subsequently, carrier competition over the provision of transport services, like manufacturer competition for final consumers, is separable across destination markets. In particular, carrier f faces a constant marginal cost of ψ_{od}^f per container (a convenient unit of transport) along route (o, d). In principle, part of this cost may reflect payments to the same type of labour used in manufacturing. I assume that carriers incur shipping costs in units of the homogenous good, which

[6] This is consistent with zero costs of entry into any given shipping route. While it would be more realistic to consider asymmetric market configurations in the shipping sector, the analysis becomes intractable, depending on the set of carriers available in each route.

[7] I reconsider this point in the empirical section.

sidesteps potential competition between the manufacturing and transport sectors for jointly used factors provided sufficiently large consumer endowments of the homogenous good, $\overline{q}_0^c$.

Letting S_{od}^f denote the level of transport service offered by carrier f along route (o, d), and $S_{od} = \sum_f S_{od}^f$ the corresponding aggregate capacity along the same route, carrier f solves

$$\max_{\left(S_{od}^f\right)_{od}} \sum_{od} \left[t_{od}\left(S_{1d}, \ldots, S_{Od}\right) - \psi_{od}^f \right] S_{od}^f, \tag{1.18}$$

taking rival outputs, and manufacturer inverse demands as given. Unlike Wong (2017) and the related backhaul literature, carriers may commit to different levels of transport between any given country pair; S_{od}^f may differ from S_{do}^f. Since carrier decisions are separable across destinations, the remainder of this section characterizes equilibria along routes destined for an arbitrary market d.

I consider interior equilibria, where each carrier is active along each d–bound route; $S_{od}^f > 0$ for all f and all (o, d). Increasing S_{od}^f affects carrier profits through two channels. First, f extracts the markup $t_{od} - \psi_{od}^f$ from marginal o–based manufacturers wishing to sell in d. Second, the carrier loses $S_{o'd}^f \left|\partial t_{od} / \partial S_{o'd}\right|$ on the infra-marginal units of shipping services offered in all d–bound shipping lanes $(o', d)_{o'}$. Using (1.12), the first-order condition for an interior optimum in route (o, d) sets f's marginal profitability of shipping from o to d to zero:

$$\begin{aligned}
\mathtt{m}_{od}\pi^f &\equiv \left(t_{od}\left(\cdot\right) - \psi_{od}^f\right) + \sum_j S_{jd}^f \frac{\partial t_{jd}}{\partial S_{od}} \\
&= p_d^{\max}\left(\cdot\right) - \tilde{t}_{od}\left(S_{od}\right) - \psi_{od}^f + \sum_j \left(\frac{\partial p_d^{\max}}{\partial S_{od}} - \frac{\partial \tilde{t}_{jd}}{\partial S_{od}}\right) S_{jd}^f \\
&= \underbrace{p_d^{\max}\left(\cdot\right) - \tilde{t}_{od}\left(S_{od}\right) - \psi_{od}^f}_{\text{Marginal}_{od}} + \underbrace{\left(\frac{\partial p_d^{\max}\left(\cdot\right)}{\partial S_{od}} - \frac{\partial \tilde{t}_{od}\left(S_{od}\right)}{\partial S_{od}}\right) S_{od}^f}_{\text{Inframarginal}_{od}} + \underbrace{\frac{\partial p_d^{\max}\left(\cdot\right)}{\partial S_{od}} \sum_{j \neq o} S_{jd}^f}_{\text{Inframarginal}_{-od}}.
\end{aligned} \tag{1.19}$$

The second line separates inverse demand along each route into its competitive and compatriot components, as per (1.12), while the last line follows from the fact that the compatriot component of inverse demand elsewhere is, by definition, insensitive to extra capacity along route (o, d). Identifying different shipping lanes with different goods, carrier competition is therefore analogous to the multi-product oligopoly problem (see, for example, Bulow et al. (1985), and Eckel and Neary (2010) in an international trade setting). Specifically, carriers internalize the negative spillovers ($\partial p_d^{\max} / \partial S_d < 0$) of extra capacity in one transport market on demand to every other d–bound shipping lane. This "cannibalization" effect lowers marginal revenue along each route, suppressing carrier output.

The system of marginal profitabilities across markets, which determines optimal carrier capacities and the corresponding freight rates, depends on source, destination, and bilateral characteristics through manufacturer demand for transport. In particular, these characteristics affect marginal profitability along (o, d) by either shifting demand for transport from o to d (allowing carriers to extract a larger markup on marginal consumers), or by exacerbating or mitigating the decline, following an increase in capacity along (o, d), in willingness–to–pay along all d–bound routes. Proposition 1.8 gathers these comparative statics, relegating the proof to Section 1.B.2.

Proposition 1.8. *Carrier f's marginal returns to capacity along route (o, d) are larger (i) when shipping less differentiated varieties of goods in high relative demand; (ii) when f incurs low shipping costs along route (o, d); (iii) when market d has a large consumer base, and is open to exports from o; and (iv) when source o has a large pool of potential exporters facing low manufacturing wages.*

For example, by Proposition 1.3, exporters are willing to pay more to send their goods to larger consumer markets ($\partial t_{od}/\partial L_d > 0$). In addition, larger markets attenuate the decline in willingess–to-pay for transport in all markets induced by extra capacity along (o, d) (that is, $\partial^2 t_{od}/\partial S_{o'}\partial L_d > 0$). Combining these effects, $\partial \mathtt{m}_{od}\pi^f/\partial L_d > 0$.

While Proposition 1.8 summarizes individual carrier incentives at the margin of capacity along any given route, I am primarily interested in country-level predictions. Reassuringly, this result carries over once we consider equilibrium conditions for *aggregate* capacities on the various d–bound routes. In particular, note that a carrier f's marginal returns along a given route, (1.19), depend on its capacities in the various markets, $(S^f_{1d}, \ldots, S^f_{Od})$, and the profile of capacities on d–bound routes, $(S_{1d}, \ldots, S_{Od})$. Thus, like optimal manufacturer demand for transport, equilibrium in shipping supply is fully characterized by the set of country–level aggregates. Since each carrier sets marginal profits to zero, equilibrium *aggregate* capacities necessarily set the (across carrier) average marginal returns to zero. Put differently, equilibrium aggregate capacities are the (unique) roots of the system $\{\mathtt{m}_{od}\pi\,(S_{1d}, \ldots, S_{Od}) = 0\}_o$, where

$$\mathtt{m}_{od}\pi \equiv p^{\max}_d\,(S_{1d}, \ldots, S_{Od}) - \tilde{t}_{od}\,(S_{od}) - \overline{\psi}_{od} + \frac{\partial p^{\max}_d}{\partial S_d}\textstyle\sum_j S_{jd}/F - \dfrac{\partial \tilde{t}_{od}}{\partial S_{od}} S_{od}/F, \qquad (1.20)$$

and $\overline{\psi}_{od} \equiv \sum_f \psi^f_{od}/F$ and $S_{od}/F \equiv \sum_f S^f_{od}/F$ are the average unit cost, and level of shipping along route (o, d). As the average of individual marginal profitabilities, the function $\mathtt{m}_{od}\pi\,(\cdot)$ inherits all the properties outlined in Proposition 1.8, simplifying comparative statics on aggregate capacities and the corresponding freight rates.[8]

Optimal aggregate capacities reflect the negative externalities from transport use among manufacturers. Consider an increase in capacity along some d–bound route. All else equal, this lowers marginal profitability along all routes destined for d, with a stronger effect in the source of the extra capacity.[9] Capacities in each market must therefore fall to restore zero marginal profitability. Panel (a) of Figure 1.5 illustrates this capacity tradeoff by plotting the

[8] Carrier–level comparative statics then follow from combining aggregate comparative statics with (1.19).
[9] See Section 1.B.1 for details.

Figure 1.5: Equilibrium aggregate shipping: comparative statics

a. Destination characteristics **b.** Source–country characteristics

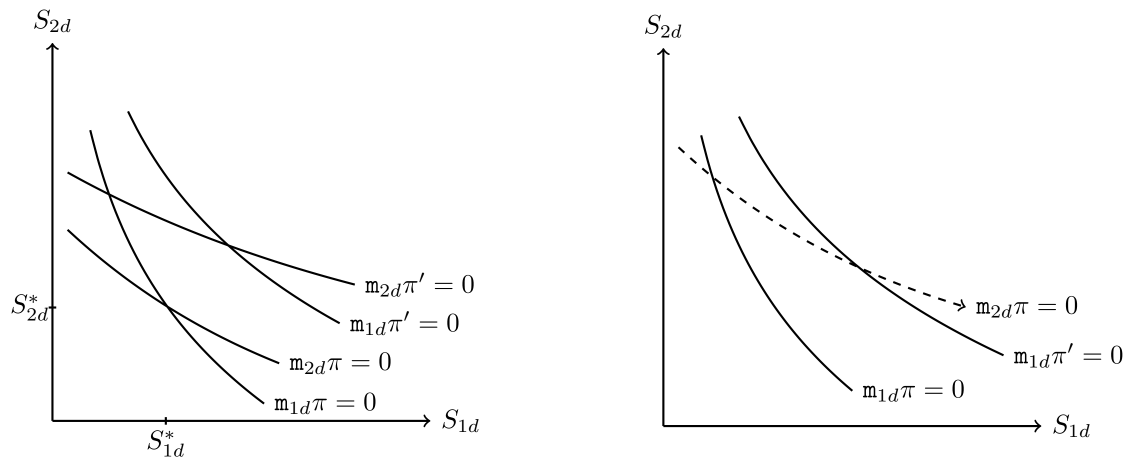

Notes: Equilibrium aggregate shipping volumes in countries 1 and 2 occurs at the intersection of the zero-level sets of $\mathtt{m}_{od}\pi$, the (across carrier average of) marginal profitability of shipping from o to d, given in (1.20). **Panel a** The volume of trade along all routes increases as the destination market expands $L_d \nearrow$; relative demand for differentiated goods increases, $\alpha \nearrow$ or $\eta \searrow$; or product differentiation falls, $\gamma \searrow$. **Panel b** The volume of trade from country 1 increases, while that from country 2 falls, as the number of potential manufacturers in 1 increases $\widetilde{N}_1 \nearrow$, the average cost of shipping from 1 falls, $\overline{\psi}_{1d} \searrow$, or the tariff–adjusted manufacturing cost falls in country 1, $w_1\tau_{1d} \searrow$.

level sets of marginal profitability along routes emanating from countries 1 and 2. Because of the negative spillovers of capacity on marginal profitability between any pair of countries, it follows that $\mathrm{d}S_{-od}/\mathrm{d}S_{od} < 0$ along the levels set of $\mathtt{m}_{od}\pi$. Remarkably, I obtain this result without resorting to decreasing returns to scale in the provision of shipping services, or constraining capacities along different routes (as in the backhaul literature). Instead, it follows purely from self–cannibalization concerns among carriers serving multiple markets.

Equilibrium aggregate capacities along the two d–bound routes, (S^*_{1d}, S^*_{2d}), occurs at the intersection of the zero–marginal–profitabiltiy level sets. I show in Section 1.B.1 that level sets of $\mathtt{m}_{1d}\pi$ cross those of $\mathtt{m}_{2d}\pi$ once, and from above. Apart from guaranteeing a unique solution, this single–crossing property greatly simplifies (qualitative) comparative statics with respect to all source–specific parameters, which, by definition, only affect marginal profitability in the country in question. In particular, capacity along any given route increases whenever such a factor raises marginal profitability. The remainder of this section characterizes equilibrium capacities and freight rates, appealing, where possible, to the effects on marginal profitabilities in the two–country case. See Section 1.B.3 for the magnitudes of these changes.

Proposition 1.9. (Equilibrium trade & Destination market characteristics)

1. *Aggregate capacities are higher along routes destined for larger markets. However, freight rates are insensitive to market size.*

2. *Aggregate capacities and freight rates are everywhere higher when consumers value sector goods (α large).*

3. *Aggregate capacities are everywhere higher, and freight rates everywhere lower, when η is small.*

4. *With Pareto and Fréchet distributed costs, aggregate capacities and freight rates are everywhere lower in sectors with greater love-of-variety.*

Changes in destination characteristics have the same qualitative effects on the marginal profitability of shipping along all routes, although the magnitudes may differ (see section 1.B.3 for details). For example, larger markets encourage additional capacity along each route by shifting each marginal profitability curve outwards, resulting in higher capacities along each route, and higher aggregate d–bound volumes. All else equal, this increase in volume lowers freight rates along each route. However, markets size also has a positive *direct* effect on manufacturer willingness–to–pay for shipping, which happens to exactly offset the negative effects operating through equilibrium capacities. The rest of the results follow by analogy, appealing to the comparative statics of marginal profitability in Proposition 1.8.

Proposition 1.10. (Equilibrium trade & Source country characteristics)

1. *A larger set of potential manufacturers in a given country (i) raises shipping activity from that country; (ii) lowers capacities in other source countries; and (iii) raises freight rates in all routes.*

2. *High wages in any source country lower shipping flows from that country and raise those from other countries. With Pareto distributed costs, such high costs lower freight rates along all shipping routes destined for the market in question.*

A large set of potential manufacturers, or lower wages in a source country lowers the marginal returns of serving that country by dampening demand for shipping (thus lowering markups on marginal manufacturers) and exacerbating the decline in manufacturer willingness–to–pay for transport.

Proposition 1.11. (Equilibrium trade & bilateral charactersitics)

1. *When the average (across carriers) cost of serving a route increases, total shipping activity in that route falls, while capacity elsewhere increases. Subsequently, freight rates rise in the route in question, and fall elsewhere.*

2. *High tariffs on a country's imports lower shipping flows from that country and increase shipping activity from other source countries. With Pareto distributed costs, these tariffs lower freight rates along all shipping routes destined for the market in question.*

Proof. See Section 1.B. □

The first result is hardly surprising. Carriers cut back on capacity in costly markets, passing on some of the cost to manufacturers in the form of higher freight rates. The spillover into other markets should not be interpreted as carriers somehow "diverting" a fixed level of capacity to other markets. After all, there is no explicit capacity constraint to speak of in the Cournot model. As for trade costs, Proposition 1.11 shows that changes in wages and tariffs have the same effect on Cournot equilibrium outcomes. This follows from the complementarity between wages, w_o, and tariffs, τ_{od}, in manufacturer delivered costs (1.6). Picking up from Section 1.2.1.3, the remainder of this chapter analyzes trade policy when manufacturing costs are Pareto–distributed.

1.2.2.1 Pareto distributed manufacturing costs: Trade policy

It will be helpful to define

$$\phi_{od}\left(S_{od}\right) = \lambda_2 \times \left(\frac{\widetilde{N}_o}{w_o \tau_{od} a_o^M}\right)^{\frac{\theta}{\theta+1}} S_{od}^{\frac{\theta}{\theta+1}},$$

a monotonic transformation of the level of shipping activity that the average carrier is willing to give up along lane (o, d) for additional activity along some other route while keeping marginal profitability along (o, d) constant.[10] All else equal, carriers are willing to give up more activity ($\phi_{od}\left(S_{od}\right)$ is large) from countries endowed with a large mass (large $\widetilde{N}$) of relatively productive (small a^M) potential manufacturers facing low tariffs (low τ).

Consider an increase in tariffs on imports from country o. I show in Section 1.B.3 that the volume of trade from o and $o' \neq o$ respond as

$$\begin{aligned} \partial S_{od}^{*} / \partial w_o \tau_{od} &= -\left(1 - \frac{\phi_{od}\left(S_{od}\right)}{1 + \sum_j \phi_{jd}\left(S_{jd}\right)}\right) \frac{\theta}{w_o \tau_{od}} S_{od} < 0 \\ \partial S_{o'd}^{*} / \partial w_o \tau_{od} &= \frac{\phi_{o'd}\left(S_{o'd}\right)}{1 + \sum_j \phi_{jd}\left(S_{jd}\right)} \frac{\theta}{w_o \tau_{od}} S_{od} > 0. \end{aligned}$$

In keeping with Proposition 1.11, higher tariffs on a country's exports reduce trade activity from that country, and increase exports from all other countries. Further, aggregate d–bound

[10] Put differently, $\phi_{od}\left(S_{od}\right)$ is the marginal rate of substitution of shipping along any other route for activity along lane (o, d).

As for the positive constant,

$$\lambda_2 \equiv \frac{\eta}{2\gamma} \frac{\left(F+1\right)\left(\theta+1\right)^{\frac{2\theta+1}{\theta+1}}}{F\left(\theta+1\right)+1} \left(\frac{2\gamma}{L_d}\right)^{\frac{\theta}{\theta+1}} > 0.$$

trade falls:

$$\partial S_d^* / \partial w_o \tau_{od} = -\frac{1}{1+\sum_j \phi_{jd}\left(S_{jd}\right)} \frac{\theta}{w_o \tau_{od}} S_{od} < 0.$$

Trade liberalization therefore diverts trade towards the liberalizing country, and results in larger aggregate trade volumes.

1.3 Conclusion

This chapter presents a model of international trade in an oligopolistic transport sector. It derives demand for transport services from underlying manufacturer behaviour before characterizing Cournot equilibrium freight rates and the corresponding levels of trade between countries in a given sector. Given the interdependent nature of demand for shipping, carriers active in shipping lanes destined for a given destination internalize the (negative) effects of additional capacity along one lane on their profitability elsewhere. This self–cannibalization concern discourages carrier activity beyond what we would expect in a set of independent oligopolistic markets. Finally, optimal trade policy must take into account the transport sector's role as a trade facilitator.

Appendix

1.A Inverse demand for transport

In this section, I prove various comparative statics of manufacturer willingess-to-pay for transport (Propositions 1.1–1.5). Recall that inverse demand along a given route is *defined* to guarantee consistent aggregate demand along said route, taking activity elsewhere as given. That is, inverse demand guarantees

$$\sigma_{od}\left(S_{1d}, \ldots S_{od}, \ldots, S_{Od}\ ,\ t_{od}\left(S_{1d}, \ldots S_{od}, \ldots, S_{Od}\right)\right) = S_{od}$$

for all d–bound profiles $(S_{1d}, \ldots, S_{Od})$, where

$$\begin{aligned} \sigma_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) &\equiv \widetilde{N}_o \int_0^{\hat{a}_{od}(S_{1d},\ldots,S_{Od},t_{od})} s^*_{od}\left(a \mid S_{1d}, \ldots, S_{Od}, t_{od}\right) \mathrm{d}G_o\left(a\right) \\ &= N_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) \times \overline{s}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) \end{aligned} \tag{1.21}$$

is the route–(o, d) aggregator. The second line allows us to track changes in the extensive, $\overline{s}_{od}\,\mathrm{d}N_{od}$, and intensive, $N_{od}\,\mathrm{d}\overline{s}_{od}$, demand margins where

$$N_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) = G_o\left(\hat{a}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right)\right) \times \widetilde{N}_o \tag{1.22}$$

is the number of o–based manufacturers active in d, and

$$\overline{s}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) \equiv \mathbb{E}\left[s^*_{od}\left(a \mid S_{1d}, \ldots, S_{Od}, t_{od}\right) \mid a \leq \hat{a}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right)\right],$$

is average demand for transport among exporting firms. Given the linearity of demand for shipping in unit labour requirements, average demand for shipping is equal to demand for shipping from the firm with the average unit labour requirement among the set of active firms, denoted

$$\overline{a}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right) \equiv \mathbb{E}\left[a \mid a \leq \hat{a}_{od}\left(S_{1d}, \ldots, S_{Od}, t_{od}\right)\right].$$

It is relatively simple to apply the implicit function theorem since the freight rate t_{od}

only appears in aggregate demand along (o, d). It will be helpful to define

$$\zeta_{od}\left(S_{1d},\ldots,S_{Od},t_{od}\right)\equiv\frac{\partial\overline{a}_{od}\left(S_{1d},\ldots,S_{Od},t_{od}\right)/\partial x}{\partial\hat{a}_{od}\left(S_{1d},\ldots,S_{Od},t_{od}\right)/\partial x}>0 \tag{1.23}$$

as the relative sensitivity of the average unit labour requirement to some variable x, which is positive because the truncated mean, $\overline{a}_{od}$, is increasing in the truncation point, $\hat{a}_{od}$. I assume $\zeta_{od}\leq 1$, which holds, for example, when costs are Pareto or Fréchet distributed.

1.A.1 Properties of the demand aggregator σ_{od}

Applying Leibniz rule, and using the fact that the threshold manufacturer demands no transport services,

1. Freight rate, t_{od}

$$\partial\sigma_{od}/\partial t_{od}=\overbrace{-\frac{L_d}{2\gamma}\left(1-\zeta_{od}\right)N_{od}}^{\text{intensive}<0}+\overbrace{-\frac{L_d}{2\gamma}\zeta_{od}N_{od}}^{\text{extensive}<0}=-\frac{L_d}{2\gamma}N_{od}, \tag{1.24}$$

so that a higher expected freight rate lowers the number of active manufacturers *and* demand for shipping among exporters.

2. Aggregate level of global d-bound transport demand, $S_d=\sum_i S_{id}$

$$\partial\sigma_{od}/\partial S_d=\overbrace{\frac{\partial p_d^{\max}}{\partial S_d}\frac{L_d}{2\gamma}\left(1-\zeta_{od}\right)N_{od}}^{\text{intensive}<0}+\overbrace{\frac{\partial p_d^{\max}}{\partial S_d}\frac{L_d}{2\gamma}\zeta_{od}N_{od}}^{\text{extensive}<0}=\frac{\partial p_d^{\max}}{\partial S_d}\frac{L_d}{2\gamma}N_{od}=-\frac{\eta}{2\gamma}N_{od}, \tag{1.25}$$

so that a higher expected d–bound shipping lowers the number of active manufacturers *and* demand for shipping among exporters, regardless of the source of the additional activity.

3. Market size, L_d

$$\begin{aligned}\partial\sigma_{od}/\partial L_d &= \overbrace{\left(\frac{\partial p_d^{\max}}{\partial L_d}\frac{L_d}{2\gamma}\left(1-\zeta_{od}\right)+\frac{1}{L_d}\overline{s}_{od}\right)N_{od}}^{\text{intensive}>0}+\overbrace{\frac{\partial p_d^{\max}}{\partial L_d}\frac{L_d}{2\gamma}\zeta_{od}N_{od}}^{\text{extensive}>0} \\ &= \frac{1}{L_d}S_{od}+\frac{\partial p_d^{\max}}{\partial L_d}\frac{L_d}{2\gamma},\end{aligned} \tag{1.26}$$

so that larger markets attract new varieties and motivate exporters to expand output, and hence demand for transport.

4. Relative demand for differentiated goods, $(\alpha, -\eta)$

$$\begin{aligned} \partial\sigma_{od}/\partial\alpha &= \overbrace{\frac{\partial p_d^{\max}}{\partial\alpha}\frac{L_d}{2\gamma}\left(1-\zeta_{od}\right)N_{od}}^{\text{intensive}>0} + \overbrace{\frac{\partial p_d^{\max}}{\partial\alpha}\frac{L_d}{2\gamma}\zeta_{od}N_{od}}^{\text{extensive}>0} &= \frac{L_d}{2\gamma}N_{od} \\ \partial\sigma_{od}/\partial\eta &= \overbrace{\frac{\partial p_d^{\max}}{\partial\eta}\frac{L_d}{2\gamma}\left(1-\zeta_{od}\right)N_{od}}^{\text{intensive}<0} + \overbrace{\frac{\partial p_d^{\max}}{\partial\eta}\frac{L_d}{2\gamma}\zeta_{od}N_{od}}^{\text{extensive}<0} &= -\frac{1}{2\gamma}N_{od}S_d, \end{aligned} \tag{1.27}$$

so that (i) previously inactive firms choose to produce; and (ii) incumbent exporter expand output, demanding more transport following an outward shift in consumer demand for differentiated varieties.

5. Love–of–variety, γ

$$\partial\sigma_{od}/\partial\gamma = \overbrace{-\frac{1}{\gamma}S_{od}}^{\text{intensive}<0} + \overbrace{0}^{\text{extensive}}. \tag{1.28}$$

Facing a more elastic demand curve, exporters scale back output. Inactive firms are unaffected.

6. Number of potential producers, $\widetilde{N}_o$

$$\partial\sigma_{od}/\partial\widetilde{N}_o = \overbrace{0}^{\text{intensive}} + \overbrace{\frac{1}{\widetilde{N}_o}S_{od}}^{\text{extensive}}. \tag{1.29}$$

All else equal, more firms choose to serve market d.

7. Wages and tariffs, $w_o\tau_{od}$

$$\partial\sigma_{od}/\partial w_o\tau_{od} = \overbrace{\frac{L_d}{2\gamma}\left(\zeta_{od}\hat{a}_{od} - \overline{a}_{od}\right)N_{od}}^{\text{intensive}\lesseqgtr 0} + \overbrace{-\frac{L_d}{2\gamma}\hat{a}_{od}\zeta_{od}N_{od}}^{\text{extensive}<0} = -\frac{L_d}{2\gamma}\overline{a}_{od}N_{od}, \tag{1.30}$$

so that demand aggregate demand falls if firms face higher wages or trade barriers.

1.A.2 Shipping activity and exporter willingness–to–pay for transport

Proof. Applying (1.25),

$$\partial\left(\sigma_{od}-S_{od}\right)/\partial S_{o'd} = -\frac{\eta}{2\gamma}N_{od}, \qquad \partial\left(\sigma_{od}-S_{od}\right)/\partial S_{od} = -\left(\frac{\eta}{2\gamma}N_{od}+1\right).$$

Substituting this, and (1.24), into (1.12) delivers (1.13), the system of differential equations in Proposition 1.2.

The initial condition

$$t_{od}\left(S_{1d},\ldots,S_{Od}\right)|_{S_{od}=0}=p_d^{\max}\qquad\text{for all } S_{-od}$$

guarantees an o–based threshold manufacturer, barely making ends meet in d, regardless of activity elsewhere.

To get inverse demand in the form (1.12), stop short of the last step of (1.25). By the implicit function theorem,

$$\partial t_{od}\left(S_{1d},\ldots,S_{Od}\right)/\partial S_d=\partial p_d^{\max}/\partial S_d, \tag{1.31}$$

so that inverse demand along (o,d) additively separable in the destination choke price. Integrating with respect to the aggregate level of global d-bound transport demand,

$$t_{od}\left(S_{1d},\ldots,S_{Od}\right)=p_d^{\max}\left(S_{1d},\ldots,S_{Od}\right)-\tilde{t}_{od}\left(S_{od}\right),$$

where $\tilde{t}_{od}\left(S_{od}\right)$ accounts for the fact that shipping from o contributes to S_d. Writing inverse demand as a function of own quantity demanded and total activity elsewhere yields equation (1.12) in the main text.

Turning to the shape of $\tilde{t}_{od}\left(\cdot\right)$, the boundary condition is equivalent to

$$\begin{array}{rcl} p_d^{\max}\left(0,\left(S_{o'd}\right)_{o'\neq o}\right) & = & t_{od}\left(0,\left(S_{o'd}\right)_{o'\neq o}\right)\equiv p_d^{\max}\left(0,\left(S_{o'd}\right)_{o'\neq o}\right)-\tilde{t}_{od}\left(0\right)\\ \iff\qquad \tilde{t}_{od}\left(0\right) & = & 0\end{array}$$

by the additive separability of inverse demand in the final goods choke price. From (1.13),

$$\partial t_{od}/\partial S_{od}=\partial p_d^{\max}/\partial S_d-2\gamma/L_dN_{od}\iff\partial\tilde{t}_{od}/\partial S_{od}=2\gamma/L_dN_{od},$$

so that $\tilde{t}_{od}\left(\cdot\right)$ is increasing. Further,

$$\partial^2\tilde{t}_{od}/\partial S_{od}^2=-\left(\zeta_{od}/S_{od}\right)\partial\tilde{t}_{od}/\partial S_{od}<0. \tag{1.32}$$

where (1.23) simplifies to

$$\zeta_{od}\left(S_{od}\right)\equiv\frac{\overline{a}'_{od}\left(S_{od}\right)}{\hat{a}'_{od}\left(S_{od}\right)}. \tag{1.23$'$}$$

Finally, by definition ((1.13) in the main text), the compatriot effect $\tilde{t}_{od}\left(S_{od}\right)$ solves the differential equation

$$\frac{\partial\tilde{t}_{od}}{\partial S_{od}}=\frac{2\gamma}{L_d}\frac{1}{\widetilde{N}_o\,G_o\left(\tilde{t}_{od}/w_o\tau_{od}\right)}, \tag{1.33}$$

where I use (1.9′) in place of the threshold exporter. Note that (1.33) is separable, so that it can be rearranged into integral form

$$S_{od} + \text{constant} = \frac{L_d}{2\gamma}\widetilde{N}_o \int G_o\left(\tilde{t}_{od}/w_o\tau_{od}\right)\, \mathrm{d}\tilde{t}_{od}, \tag{1.34}$$

which has a unique, but not necessarily closed–form, solution. Substituting the Pareto distribution into (1.34) and rearranging yields the family of solutions

$$\tilde{t}_{od}\left(S_{od}\right) = \left[\frac{L_d}{2\gamma}\widetilde{N}_o\frac{1}{\theta+1}\left(\frac{1}{a_o^M}\frac{1}{w_o\tau_{od}}\right)^{\theta}\right]^{-\frac{1}{1+\theta}}\left(S_{od}+\text{constant}\right)^{\frac{1}{1+\theta}}.$$

The initial condition $\tilde{t}_{od}(0) = 0$ implies that the constant of integration is zero, giving (1.15). □

1.A.3 Shifts in manufacturer willingness–to–pay for transport

Changes in the economic environment affect carrier incentives by altering manufacturer demand for transport (Propositions 1.3, 1.4, and 1.5). Propositions 1.3, 1.4, and 1.5 follow from the implicit-function theorem, using (1.24), and (1.26) to (1.30). See the main text for accompanying discussion.

$$\begin{aligned}
&\underline{\text{Market size}}\\
&\partial t_{od}/\partial L_d = \partial p_d^{\max}/\partial L_d - \partial\tilde{t}_{od}/\partial L_d = \frac{1}{L_d}\left(\left(\eta/L_d\right)S_d + \left(2\gamma/L_dN_{od}\right)S_{od}\right) > 0.\\
&\underline{\text{Relative demand for differentiated goods}}\\
&\partial t_{od}/\partial\alpha = \partial p_d^{\max}/\partial\alpha = 1,\\
&\partial t_{od}/\partial\eta = \partial p_d^{\max}/\partial\eta = -\frac{1}{L_d}S_d > 0,\\
&\underline{\text{Love-of-variety}}\\
&\partial t_{od}/\partial\gamma = -\partial\tilde{t}_{od}/\partial\gamma = -\frac{1}{\gamma}\left(2\gamma/L_dN_{od}\right)S_{od} < 0,\\
&\underline{\text{Tariffs}}\\
&\partial t_{od}/\partial w_o\tau_{od} = -\partial\tilde{t}_{od}/\partial w_o\tau_{od} = -\overline{a}_{od} < 0,\\
&\underline{\text{Potential manufacturers}}\\
&\partial t_{od}/\partial\widetilde{N}_o = -\partial\tilde{t}_{od}/\partial\widetilde{N}_o = \frac{1}{\widetilde{N}_o}\left(2\gamma/L_dN_{od}\right)S_{od} > 0.
\end{aligned} \tag{1.35}$$

Changes in the slope of inverse demand for transport

It will be helpful to collect the effects of the various model parameters on the detrimental effect of capacity along a given route on manufacturer willingness–to–pay for transport on that route, $\partial t_{od}(\cdot)/\partial S_{od}$. Picking up from (1.35),

$$\begin{aligned}
&\underline{\text{Market size}}\\
&\partial^2 t_{od}/\partial S_{od}\,\partial L_d = -\frac{1}{L_d}\left(1-\zeta_{od}\right)\partial t_{od}/\partial S_{od} = -\frac{1}{L_d}\left(2\gamma/L_d N_{od}\right)\left(1-\zeta_{od}\right) < 0.\\
&\underline{\text{Relative demand for differentiated goods}}\\
&\partial^2 t_{od}/\partial S_{od}\,\partial \alpha = 0,\\
&\partial^2 t_{od}/\partial S_{od}\,\partial \eta = -\frac{1}{L_d} < 0,\\
&\underline{\text{Love-of-variety}}\\
&\partial^2 t_{od}/\partial S_{od}\,\partial \gamma = \frac{1}{\gamma}\left(1-\zeta_{od}\right)\partial t_{od}/\partial S_{od} = \frac{1}{\gamma}\left(2\gamma/L_d N_{od}\right)\left(1-\zeta_{od}\right) > 0,\\
&\underline{\text{Tariffs}}\\
&\partial^2 t_{od}/\partial S_{od}\,\partial\left(w_o\tau_{od}\right) = \frac{1}{S_{od}}\left(\hat{a}_{od}-\overline{a}_{od}\right)\zeta_{od} > 0,\\
&\underline{\text{Potential manufacturers}}\\
&\partial^2 t_{od}/\partial S_{od}\,\partial \widetilde{N}_o = -\frac{1}{\widetilde{N}_o}\left(1-\zeta_{od}\right)\partial t_{od}/\partial S_{od} = -\frac{1}{\widetilde{N}_o}\left(2\gamma/L_d N_{od}\right)\left(1-\zeta_{od}\right) < 0.
\end{aligned} \tag{1.36}$$

1.A.4 Extensive margin of trade

This section proves Proposition 1.7 by differentiating $\hat{a}_{od}(S_{od})$, the unit–labour requirement of the threshold o–based exporter active in d, (1.9′), and $N_{od}(S_{od})$, the number of o–based exporters in d, (1.14).

$$\begin{aligned}
&\underline{\text{Volume of trade}} &&\\
&\partial\hat{a}_{od}/\partial S_{od} = \left(\hat{a}_{od}-\overline{a}_{od}\right)/S_{od} > 0 && \partial N_{od}/\partial S_{od} = \widetilde{N}_o g_o\left(\hat{a}_{od}\right)\times\partial\hat{a}_{od}/\partial S_{od} > 0\\
&\underline{\text{Market size}} &&\\
&\partial\hat{a}_{od}/\partial L_d = -\left(\hat{a}_{od}-\overline{a}_{od}\right)/L_d < 0 && \partial N_{od}/\partial L_d = \widetilde{N}_o g_o\left(\hat{a}_{od}\right)\times\partial\hat{a}_{od}/\partial L_d < 0\\
&\underline{\text{Love-of-variety}} &&\\
&\partial\hat{a}_{od}/\partial\gamma = \left(\hat{a}_{od}-\overline{a}_{od}\right)/\gamma > 0 && \partial N_{od}/\partial\gamma = \widetilde{N}_o g_o\left(\hat{a}_{od}\right)\times\partial\hat{a}_{od}/\partial\gamma > 0\\
&\underline{\text{Tariffs}} &&\\
&\partial\hat{a}_{od}/\partial w_o\tau_{od} = -\left(\hat{a}_{od}-\overline{a}_{od}\right)/w_o\tau_{od} < 0 && \partial N_{od}/\partial w_o\tau_{od} = \widetilde{N}_o g_o\left(\hat{a}_{od}\right)\times\partial\hat{a}_{od}/\partial w_o\tau_{od} < 0\\
&\underline{\text{Potential manufacturers}} &&\\
&\partial\hat{a}_{od}/\partial\widetilde{N}_o = -\left(\hat{a}_{od}-\overline{a}_{od}\right)/\widetilde{N}_o < 0 && \partial N_{od}/\partial\widetilde{N}_o = \left(1-\zeta_{od}\right)G_o\left(\hat{a}_{od}\right) > 0
\end{aligned}$$

1.B Cournot equilibrium outcomes

This section presents comparative statics on aggregate capacities and freight rates when all carriers are active in exactly the same set of markets by applying the implicit-function theorem to (1.20). Changes in volume of shipping from o to d following a change in some parameter is a linear combination of the effect of the parameter on the marginal profitability of shipping along the various d–bound lanes. In fact, it turns out to be a linear combination of the effect on the marginal profitability of shipping from o to d, and the differential effects of the parameter in marginal profitability of shipping along (o, d) against each of the other d–bound routes.

1.B.1 Marginal rate of substitution of capacities along different routes

Differentiating $\mathtt{m}_{od}\pi$ with respect to capacity from o'

$$\frac{\partial \mathtt{m}_{od}\pi}{\partial S_{o'd}} = \begin{cases} -\left(1+\dfrac{1}{F}\right)\dfrac{\eta}{L_d} & o' \neq o \\ \\ -\left[\left(1+\dfrac{1}{F}\right)\dfrac{\eta}{L_d} + \left(1+\dfrac{1}{F}\left(1-\zeta_{od}\right)\right)\dfrac{2\gamma}{L_d}\dfrac{1}{N_{od}}\right] & o' = o, \end{cases} \tag{1.37}$$

$o' = o$: diminishing marginal profitability; $o' \neq o$: unsurprising that marginal profitability decreasing in capacity elsewhere; self-cannibalization governed by competitive effect.

Single–crossing property Taking the ratio of the cross– and own–capacity effects, Equation (1.37) implies that, between them, profit-maximizing carriers along (o, d) must ship

$$0 < \left|\frac{\mathrm{d}S_{od}}{\mathrm{d}S_{o'd}}\right| = \frac{(F+1)\,\eta N_{od}}{(F+1)\,\eta N_{od} + 2\gamma\left(F+1-\overline{a}'_{od}\left(S_{od}\right)/\hat{a}'_{od}\left(S_{od}\right)\right)} < 1 \tag{1.38}$$

fewer units from o to d in exchange for an additional unit from source o'. By symmetry, profit-maximizing carriers along (o', d) must ship $0 < |\mathrm{d}S_{o'd}/\mathrm{d}S_{od}| < 1$ fewer units along (o', d) for an additional unit along (o, d). Combining these observations yields the single–crossing property in a two–country world.

It will be convenient to proceed with a monotonic transformation of the marginal rate of substitution in (1.38). Specifically, consider

$$\begin{aligned} \phi_{od}\left(S_{od}\right) &\equiv \frac{|\mathrm{d}S_{od}/\mathrm{d}S_{o'd}|}{1-|\mathrm{d}S_{od}/\mathrm{d}S_{o'd}|} \\ &= \frac{\eta}{2\gamma}\frac{(F+1)\,N_{od}\left(S_{od}\right)}{F+1-\zeta_{od}\left(S_{od}\right)} > 0. \end{aligned} \tag{1.39}$$

Lemma. *Carriers relinquish a greater volume of capacity along route (o,d) in exchange for capacity elsewhere (ϕ_{od} is large) when*

1. *at large volumes of capacity/ capacity along (o,d) is large / already shipping large amounts;*
2. *shipping markets are more interdependent (η large or L_d small);*
3. *o–based firms face few restrictions when exporting to d ($w_o \tau_{od}$ small); and*
4. *o has a large number of potential manufacturers.*

The marginal rate of substitution is independent of the level of capacity elsewhere, $S_{-od} \equiv \sum_{o' \neq o} S_{o'd}$, and the overall level of final goods demand, α.

Proof. All else equal, ϕ_{od} is increasing in the number of o–based manufacturers exporting to d, described in Proposition 1.7. Similarly, ϕ_{od} is increasing in the relative sensitivity of average exporter to additional capacity, □

$$\zeta_{od}(S_{od}) \equiv \overline{a}'_{od}(S_{od}) / \hat{a}'_{od}(S_{od}),$$

This ratio well defined when only a subset of o–based potential manufacturers are active in d. It is positive since the average and marginal unit labour requirements move in the same direction. Assuming differentiable distributions and densities,

$$\partial \zeta_{od} / \partial S_{od} = (X_{od} / w_o \tau_{od}) \times \partial \tilde{t}_{od} / \partial S_{od},$$

where

$$X_{od}(S_{od}) \equiv (1 - \zeta_{od}(S_{od})) \frac{g_o(\hat{a}_{od})}{G_o(\hat{a}_{od})} + \zeta_{od}(S_{od}) \left(\frac{g'_o(\hat{a}_{od})}{g_o(\hat{a}_{od})} - \frac{g_o(\hat{a}_{od})}{G_o(\hat{a}_{od})} \right), \tag{1.40}$$

is a convex combination of the reverse hazard rate, g_o/G_o, and the difference, $\partial \ln g_o / \partial a - g_o/G_o$, between the growth rate of the density, and the reverse hazard rate, all evaluated at the marginal unit labour requirement, $\hat{a}_{od} = \hat{a}_{od}(S_{od})$. Likewise, the comparative statics of ζ_{od} with respect to parameter θ are

$$\partial \zeta_{od} / \partial \theta = X_{od} \times \partial \hat{a}_{od} / \partial \theta.$$

The second term in (1.40), and hence X_{od}, is everywhere positive if and only if $G_o(\cdot)$ is log–convex. However, the Fréchet, log-normal, and exponential, and Pareto distributions are log–concave (Bagnoli and Bergstrom, 2005). The sign of X_{od} therefore depends on the relative magnitude of the two terms in (1.40).

It is easy to verify that the negative effect weakly dominates for commonly used distributions in international trade (Fréchet, log-normal, exponential, and Pareto), so that

$X_{od} \leq 0$, with equality for the Pareto distribution. Assuming one of these commonly used distributions, and applying the extensive margin results from Section 1.A.4,

$$\begin{array}{lll} \text{Market size:} & \partial\zeta_{od}/\partial L_d & \geq 0 \\ \text{Love-of-variety:} & \partial\zeta_{od}/\partial\gamma & \leq 0 \\ \text{Tariffs:} & \partial\zeta_{od}/\partial w_o\tau_{od} & \leq 0 \\ \text{Potential manufacturers:} & \partial\zeta_{od}/\partial\widetilde{N}_o & \geq 0. \end{array} \tag{1.41}$$

Combining these results with comparative statics of the number of varieties from Section 1.A.4, yields the comparative statics of $\phi_{od}(S_{od})$, which, recall, indexes the marginal rate of substitution of capacity elsewhere for capacity along (o, d). In particular, define

$$Y_{od}(S_{od}) \equiv \frac{g_o \circ \hat{a}_{od}(S_{od})}{G_o \circ \hat{a}_{od}(S_{od})} + \frac{X_{od}(S_{od})}{F + 1 - \zeta_{od}(S_{od})}.$$

$Y_{od}(S_{od})$ is strictly positive when costs are distributed Pareto, log-normal or Fréchet, so that

$$\begin{array}{lll} \text{Capacity:} & \phi'_{od}(S_{od}) & > 0 \\ \text{Market size:} & \partial\phi_{od}/\partial L_d & < 0 \\ {}^*\text{Relative demand for differentiated goods:} & \partial\phi_{od}/\partial\alpha & = 0 \\ & \partial\phi_{od}/\partial\eta & > 0 \\ \text{Love-of-variety:} & \partial\phi_{od}/\partial\gamma & < 0 \\ \text{Tariffs:} & \partial\phi_{od}/\partial w_o\tau_{od} & < 0 \\ \text{Potential manufacturers:} & \partial\phi_{od}/\partial\widetilde{N}_o & > 0 \\ {}^*\text{Number of carriers:} & \partial\phi_{od}/\partial F & < 0, \end{array} \tag{1.42}$$

where an asterisk denotes that the result is independent of the particular distribution of manufacturing costs.

1.B.2 Marginal profitability of shipping

This section characterizes carrier–level marginal profitability of shipping from o to d, proving Proposition 1.8. Recall, from (1.19), the effect of additional capacity by carrier f on its profits across all d–bound markets,

$$\mathtt{m}_{od}\pi^f = p_d^{\max}(S_d) - \tilde{t}_{od}(S_{od}) - \psi_{od}^f + \left(\frac{\partial p_d^{\max}}{\partial S_d} - \frac{\partial \tilde{t}_{od}}{\partial S_{od}}\right) S_{od}^f + \frac{\partial p_d^{\max}}{\partial S_d} \sum_{j \neq o} S_{jd}^f.$$

Differentiating $\mathtt{m}_{od}\pi$ with respect to f's capacity along routes (o, d) and (o', d) gives a sense of the effects of additional capacity on marginal profits along (o, d), while the derivative with respect to model parameters measures the complementarity between capacity along (o, d) and said parameter.

Differentiating $\mathtt{m}_{od}\pi$ with respect to capacity from $o' \neq o$,

$$\partial \mathtt{m}_{od}\pi^f / \partial S^f_{o'd} = 2\partial p^{\max}_d / \partial S_d = -2\eta / L_d, \qquad o' \neq o$$

since the competitive component is linear in capacities, so that $\partial^2 p^{\max}_d / \partial S^2_d = 0$, and capacity along (o', d) has no effect on the compatriot component along (o, d). The marginal returns to capacity in one market are decreasing in f's activity elsewhere. Further,

$$\begin{aligned} \partial \mathtt{m}_{od}\pi^f / \partial S^f_{od} &= 2\left(\partial p^{\max}_d / \partial S_d - \partial \tilde{t}_{od} / \partial S_{od}\right) - S^f_{od}\, \partial^2 \tilde{t}_{od} / \partial S^2_{od} \\ &= 2\,\partial p^{\max}_d / \partial S_d + \left[\left(S^f_{od}/S_{od}\right)\zeta_{od} - 2\right] \partial \tilde{t}_{od} / \partial S_{od} \\ &= -2\eta / L_d + \left[\left(S^f_{od}/S_{od}\right)\zeta_{od} - 2\right] 2\gamma / L_d N_{od}, \end{aligned}$$

implying diminishing returns to capacity along a given route whenever $\zeta_{od}\left(S_{od}\right) \equiv \overline{a}'_{od}\left(S_{od}\right) / \hat{a}'_{od}\left(S_{od}\right) \leq 2$. I impose the stronger sufficient condition $\overline{a}'_{od}\left(S_{od}\right) \leq \hat{a}'_{od}\left(S_{od}\right)$.

As for an arbitrary source, destination, or carrier characteristic, θ, the change in marginal profitability is

$$\begin{aligned} \partial \mathtt{m}_{od}\pi^f / \partial \theta &= \partial t_{od} / \partial \theta + \textstyle\sum_j S^f_{jd}\, \partial^2 t_{jd} / \partial S_{od} \partial \theta \\ &= \partial t_{od} / \partial \theta - \partial^2 \tilde{t}_{od} / \partial S_{od} \partial \theta + \partial^2 p^{\max}_d / \partial S_d \partial \theta \textstyle\sum_j S^f_{jd}. \end{aligned}$$

If θ increases demand for transport along (o, d) (that is, if $\partial t_{od} / \partial \theta \geq 0$), *and* mitigates the decline in manufacturer willingness–to–pay in every other route ($\partial^2 t_{jd} / \partial S_{od} \partial \theta \geq 0$), then θ

unambiguously raises marginal profitability along (o,d).

Market size

$$\begin{aligned}\partial \mathtt{m}_{od}\pi^f/\partial L_d &= \partial p_d^{\max}/\partial L_d - \partial \tilde{t}_{od}/\partial L_d - S_{od}^f\,\partial^2 \tilde{t}_{od}/\partial S_{od}\partial L_d + S_d^f\,\partial^2 p_d^{\max}/\partial S_d \partial L_d \\ &= \frac{1}{L_d}\left[2\,(\eta/L_d)\;S_d^f + (2\gamma/L_d N_{od})\;(2-\zeta_{od})\,S_{od}^f\right] > 0,\end{aligned}$$

Relative demand for differentiated goods

$\partial \mathtt{m}_{od}\pi^f/\partial\alpha = \partial p_d^{\max}/\partial\alpha = 1,$

$\partial \mathtt{m}_{od}\pi^f/\partial\eta = \partial p_d^{\max}/\partial\eta + S_d^f\,\partial^2 p_d^{\max}/\partial S_d\partial\eta = -2\,(\eta/L_d)\,S_d^f/\eta < 0,$

Love-of-variety

$$\partial \mathtt{m}_{od}\pi^f/\partial\gamma = -\left(\partial\tilde{t}_{od}/\partial\gamma + S_{od}^f\,\partial^2\tilde{t}_{od}/\partial S_{od}\partial\gamma\right) = -\frac{1}{\gamma}\,(2\gamma/L_d N_{od})\;(2-\zeta_{od})\,S_{od}^f < 0,$$

Tariffss

$$\partial \mathtt{m}_{od}\pi^f/\partial w_o\tau_{od} = -\left(\partial\tilde{t}_{od}/\partial w_o\tau_{od} + S_{od}^f\,\partial^2\tilde{t}_{od}/\partial S_{od}\partial\widetilde{N}_o\right) = -\left(\overline{a}_{od} + \zeta_{od}\left(\hat{a}_{od} - \overline{a}_{od}\right)\right) < 0,$$

Potential manufacturers

$$\partial \mathtt{m}_{od}\pi^f/\partial\widetilde{N}_o = -\left(\partial\tilde{t}_{od}/\partial\widetilde{N}_o + S_{od}^f\,\partial^2\tilde{t}_{od}/\partial S_{od}\partial\widetilde{N}_o\right) = \frac{1}{\widetilde{N}_o}\,(2\gamma/L_d N_{od})\;(2-\zeta_{od})\,S_{od}^f > 0,$$

Carrier marginal cost

$\partial \mathtt{m}_{od}\pi^f/\partial\psi_{od}^f = -1.$

Averaging $\mathtt{m}_{od}\pi^f$ across carriers yields $\mathtt{m}_{od}\pi$, the object driving much of the analysis in the main text.

1.B.3 Equilibrium capacities and freight rates

The optimal capacity from o to d is part of a system equalizing marginal profitabilities across d–bound lanes. Applying the implicit function theorem with respect to parameter θ hinges on the properties of the (symmetric) matrix of cross–capacity effects from (1.37), $[\partial \mathtt{m}_{od}\pi/\partial S_{o'd}]_{o,o'}$, and the shifts in marginal profitability, $[\partial \mathtt{m}_{od}\pi/\partial\theta]_o$,

$$[\partial S_{od}^*/\partial\theta]_o\,[\partial \mathtt{m}_{od}\pi/\partial S_{o'd}]_{o,o'} = [\partial \mathtt{m}_{od}\pi/\partial\theta]_o\,. \tag{1.43}$$

Claim. The matrix of cross–capacity effects, $[\partial \mathtt{m}_{od}\pi/\partial S_{o'd}]_{o,o'}$, is invertible.

Proof. Its determinant is

$$\Delta\left(S_{1d},\ldots,S_{Od}\right) = \left(-\frac{F+1}{F}\frac{\eta}{L_d}\right)^O \left(1+\textstyle\sum_j \phi_{jd}\left(S_{jd}\right)\right)\prod_j \frac{1}{\phi_{jd}\left(S_{jd}\right)},$$

where $\phi_{jd}(S_{jd})$ is given in (1.39). The determinant may vanish for one of two reasons: (i) if $\phi_{jd}(S_{jd}) = 0$ for some country o; or (ii) if $1+\sum_j \phi_{jd}(S_{jd}) = 0$, so that $\Delta \neq 0$ if $\phi_{jd}(S_{jd}) > 0$ for all $j \in \mathcal{O}$. This sufficient condition holds whenever (i) the transport sector consists of a large number of carriers; or (ii) the threshold unit labour requirement is everywhere more sensitive to additional capacity than the average, that is, $\overline{a}'_{od}(S_{od}) \leq \hat{a}'_{od}(S_{od})$ for all S_{od}. The latter is equivalent to

$$(\hat{a} - \mathbb{E}[a|\, a \leq \hat{a}])\, g(\hat{a}) \leq G(\hat{a}) \qquad \text{for all } \hat{a}. \tag{1.44}$$

While this condition holds for all decreasing densities $g(a)$, it is more plausible that efficient (low a) manufacturers are relatively rare. In any case, this condition also holds for the Pareto and Fréchet distributions commonly used in international trade. □

Armed with comparative statics $[\partial S^*_{od}/\partial\theta]_o$, I compute the corresponding changes in freight rates as

$$\frac{\partial t^*_{od}}{\partial\theta} = \underbrace{\frac{\partial t_{od}}{\partial\theta}}_{\text{Direct effect}} + \underbrace{-\frac{\eta}{L_d}\sum_{o'}\frac{\partial S^*_{o'd}}{\partial\theta}}_{\text{Competitive channel}} + \underbrace{-\frac{2\gamma}{L_d}\frac{1}{N_{od}}\frac{\partial S^*_{od}}{\partial\theta}}_{\text{Compatriot channel}}. \tag{1.45}$$

Changing θ directly affects manufacturer willingness–to–pay for transport. The Cournot equilibrium quantities may also respond to such a change, offering an indirect mechanism to freight rates. The remainder of this section provides comparative statics with respect to origin, destination, and bilateral characteristics.

1. Market size, L_d

By the implicit function theorem, (1.43),

$$\partial S^*_{od}/\partial L_d = S_{od}/L_d > 0, \qquad \partial S^*_d/\partial L_d = S_d/L_d > 0, \tag{1.46}$$

that is, aggregate capacities are higher in all routes destined for larger markets. Combining the direct and indirect effects,

$$\partial t^*_{od}/\partial L_d = 0. \tag{1.47}$$

The direct effect exactly offsets the indirect effects.

2. Relative demand for differentiated goods, $(\alpha, -\eta)$**.**

By the implicit function theorem, (1.43),

$$\begin{aligned}
\partial S^*_{od}/\partial\alpha &= \frac{F}{F+1}\frac{L_d}{\eta}\frac{\phi_{od}}{1+\sum_o \phi_{od}} > 0, \qquad \partial S^*_{d}/\partial\alpha = \frac{F}{F+1}\frac{L_d}{\eta}\frac{\sum_o \phi_{od}}{1+\sum_o \phi_{od}} > 0, \\
\partial S^*_{od}/\partial\eta &= -\frac{1}{\eta}\frac{\phi_{od}}{1+\sum_o \phi_{od}} S_d < 0, \qquad \partial S^*_{d}/\partial\eta = -\frac{1}{\eta}\frac{\sum_o \phi_{od}}{1+\sum_o \phi_{od}} S_d < 0.
\end{aligned} \tag{1.48}$$

Capacities along each route, and hence the aggregate, destined for d are higher when shipping goods in demand. Combining the direct and indirect effects,

$$\begin{aligned}
\partial t^*_{od}/\partial\alpha &= \frac{1}{1+\sum_j \phi_{jd}}\left(1+\frac{1}{F+1}\sum_j \phi_{jd} - \frac{2\gamma}{\eta}\frac{F}{F+1}\phi_{od}\right) > 0 \\
\partial t^*_{od}/\partial\eta &= \frac{1}{1+\sum_j \phi_{jd}}\left(\frac{2\gamma}{\eta}\frac{\phi_{od}}{N_{od}} - 1\right)\frac{S_d}{L_d} > 0.
\end{aligned} \tag{1.49}$$

3. Love–of–variety, γ**.**

By the implicit function theorem, (1.43),

$$\partial S^*_{od}/\partial\gamma = -\frac{1}{\gamma}\left(S_{od} - \frac{\phi_{od}}{1+\sum_j \phi_{jd}} S_d\right), \tag{1.50}$$

whose sign depends on the value of $(S_{1d}, \ldots, S_{Od})$. In particular,

$$\partial S^*_{od}/\partial\gamma < 0 \iff \frac{S_{od}}{\sum_j S_{jd}} > \frac{\phi_{od}}{1+\sum_j \phi_{jd}}$$

Despite the ambiguity surrounding individual transport use, aggregate d–bound transport use falls

$$\partial S^*_{d}/\partial\gamma = -\frac{1}{\gamma}\frac{1}{1+\sum_j \phi_{jd}} S_d < 0.$$

As for freight rates,

$$\begin{aligned}
\partial t^*_{od}/\partial\gamma &= -\frac{1}{\gamma} S_{od}\frac{\partial\tilde{t}_{od}}{\partial S_{od}} + \frac{1}{\gamma}\left(S_{od} - \frac{\phi_{od}}{1+\sum_j \phi_{jd}} S_d\right)\frac{\partial\tilde{t}_{od}}{\partial S_{od}} - \frac{1}{\gamma} S_d \frac{1}{1+\sum_j \phi_{jd}}\frac{\partial p^{\max}}{\partial S_d} \\
&= -\frac{1}{\gamma}\frac{1}{1+\sum_j \phi_{jd}}\left(-\frac{\eta}{L_d} + \frac{2\gamma}{L_d}\frac{\phi_{od}}{N_{od}}\right) S_d < 0.
\end{aligned} \tag{1.51}$$

4. Number of potential producers, $\widetilde{N}$.

Suppose the number of potential firms in country o increases. The effects on capacity are

$$\partial S^*_{od}/\partial \widetilde{N}_o = \left(1 - \frac{\phi_{od}}{1+\sum_j \phi_{jd}}\right)\frac{S_{od}}{\widetilde{N}_o} > 0, \qquad \partial S^*_{o'd}/\partial \widetilde{N}_o = -\frac{\phi_{o'd}}{1+\sum_j \phi_{jd}}\frac{S_{od}}{\widetilde{N}_o} < 0. \quad (1.52)$$

As for aggregate d–bound trade,

$$\partial S^*_{d}/\partial \widetilde{N}_o = \frac{1}{1+\sum_j \phi_{jd}}\frac{S_{od}}{\widetilde{N}_o} > 0.$$

The rise in exports from o' more than offsets the decline from other countries. The effects on freight rates along routes (o,d) and (o',d), for $o' \neq o$, are

$$\begin{aligned} \partial t^*_{od}/\partial \widetilde{N}_o &= \frac{1}{1+\sum_j \phi_{jd}}\frac{S_{od}}{\widetilde{N}_o}\left(-\frac{\eta}{L_d} + \frac{2\gamma}{L_d}\frac{\phi_{od}}{N_{od}}\right) > 0 \\ \partial t^*_{o'd}/\partial \widetilde{N}_o &= \frac{1}{1+\sum_j \phi_{jd}}\frac{S_{od}}{\widetilde{N}_o}\left(-\frac{\eta}{L_d} + \frac{2\gamma}{L_d}\frac{\phi_{o'd}}{N_{o'd}}\right) > 0. \end{aligned} \quad (1.53)$$

5. Wages and tariffs, $w_o\tau_{od}$.

Let

$$\epsilon_{od}\left(S_{od}\right) \equiv \frac{\partial \ln \overline{a}_{od}\left(S_{od}\right)/\partial \ln S_{od}}{\partial \ln \hat{a}_{od}\left(S_{od}\right)/\partial \ln S_{od}}$$

denote the relative capacity–elasticity of the average o–based unit labour requirement. For example, $\epsilon_{od}\left(S_{od}\right) \equiv 1$ when manufacturing costs are Pareto–distributed. Following an increase in the number of potential firms in o increases

$$\begin{aligned} \partial S^*_{od}/\partial w_o\tau_{od} &= -\overline{a}_{od}\left[\frac{L_d}{2\gamma}N_{od} + \frac{1}{F+1}\frac{L_d}{\eta}\phi_{od}\left(\epsilon_{od}-1\right)\right]\left(1 - \frac{\phi_{od}}{1+\sum_j \phi_{jd}}\right), \\ \partial S^*_{o'd}/\partial w_o\tau_{od} &= \overline{a}_{od}\left[\frac{L_d}{2\gamma}N_{od} + \frac{1}{F+1}\frac{L_d}{\eta}\phi_{od}\left(\epsilon_{od}-1\right)\right]\frac{\phi_{o'd}}{1+\sum_j \phi_{jd}} > 0. \end{aligned} \quad (1.54)$$

As for aggregate d–bound trade,

$$\partial S^*_{d}/\partial w_o\tau_{od} = -\overline{a}_{od}\left[\frac{L_d}{2\gamma}N_{od} + \frac{1}{F+1}\frac{L_d}{\eta}\phi_{od}\left(\epsilon_{od}-1\right)\right]\frac{1}{1+\sum_j \phi_{jd}} < 0.$$

The decline in exports from o more than offsets the rise from other countries. The effects on freight rates along routes (o,d) and (o',d), for $o' \neq o$, are

$$\begin{aligned}
\partial t^*_{od}/\partial w_o\tau_{od} &= -\overline{a}_{od}\left[\begin{array}{c} 1-\left(\frac{L_d}{2\gamma}N_{od}+\frac{1}{F+1}\frac{L_d}{\eta}\phi_{od}\left(\epsilon_{od}-1\right)\right) \\ \times\left(\left(1-\frac{\phi_{od}}{1+\sum_j \phi_{jd}}\right)\frac{2\gamma}{L_d}\frac{1}{N_{od}}+\frac{1}{1+\sum_j \phi_{jd}}\frac{\eta}{L_d}\right)\end{array}\right] \\
\partial t^*_{o'd}/\partial w_o\tau_{od} &= \frac{\overline{a}_{od}}{1+\sum_j \phi_{jd}}\left[\begin{array}{c}\frac{L_d}{2\gamma}N_{od}+\frac{1}{F+1}\frac{L_d}{\eta}\phi_{od}\left(\epsilon_{od}-1\right) \\ \times\left(-\frac{\eta}{L_d}+\frac{2\gamma}{L_d}\frac{\phi_{o'd}}{N_{o'd}}\right)\end{array}\right] < 0.
\end{aligned} \tag{1.55}$$

6. Average carrier costs, $\overline{\psi}_{od}$.

Consider an increase in the average cost of shipping from o to d. The effects on capacity along routes (o,d) and (o',d), for $o' \neq o$, are

$$\partial S^*_{od}/\partial\overline{\psi}_{od} = -\left(1-\frac{\phi_{od}}{1+\sum_j \phi_{jd}}\right) < 0, \qquad \partial S^*_{o'd}/\partial\overline{\psi}_{od} = \frac{F}{F+1}\frac{L_d}{\eta}\frac{\phi_{od}\phi_{o'd}}{1+\sum_j \phi_{jd}} > 0. \tag{1.56}$$

Note that cross effects are symmetric. As for aggregate d–bound trade,

$$\partial S^*_{d}/\partial\overline{\psi}_{od} = -\frac{F}{F+1}\frac{L_d}{\eta}\frac{\phi_{od}}{1+\sum_j \phi_{jd}} < 0.$$

The decline in exports from o more than offsets the rise from other countries. The effects on freight rates along routes (o,d) and (o',d), for $o' \neq o$, are

$$\begin{aligned}
\partial t^*_{od}/\partial\overline{\psi}_{od} &= \frac{F}{F+1}\frac{\phi_{od}}{1+\sum_j \phi_{jd}}+\frac{2\gamma}{L_d}\frac{1}{N_{od}}\left(1-\frac{\phi_{od}}{1+\sum_j \phi_{jd}}\right) > 0 \\
\partial t^*_{o'd}/\partial\overline{\psi}_{od} &= -\frac{F}{F+1}\frac{L_d}{\eta}\frac{\phi_{o'd}}{1+\sum_j \phi_{jd}}\left(-\frac{\eta}{L_d}+\frac{2\gamma}{L_d}\frac{1}{N_{od}}\phi_{od}\right) < 0.
\end{aligned} \tag{1.57}$$

1.C Demand for transport

In contrast to the main analysis, this section describes demand for shipping as a function of freight rates, and then inverts this demand system for a well–defined Cournot game. While independent of the main analysis, I include this dual approach for the sake of completeness.

Let $\boldsymbol{t}_{\cdot d} \equiv (t_{1d}, \ldots, t_{Od})$ be the vector of freight rates along routes destined for d, and abbreviate demand for shipping along route (o, d) by a manufacturer with unit labour requirement a, given freight rates $\boldsymbol{t}_{\cdot d}$, (1.8), to

$$s_{od}^{*}\left(a \mid p_{d}^{\max}\left(\boldsymbol{t}_{\cdot d}\right), t_{od}\right)=\frac{L_{d}}{2 \gamma}\left[p_{d}^{\max}\left(\boldsymbol{t}_{\cdot d}\right)-C_{od}\left(a, t_{od}\right)\right] \equiv s_{od}^{*}\left(a \mid \boldsymbol{t}_{\cdot d}\right).$$

Similarly, let

$$\hat{a}_{od}^{*}\left(\boldsymbol{t}_{\cdot d}\right) \equiv \hat{a}_{od}\left(p^{\max}\left(\boldsymbol{t}_{\cdot d}\right), t_{od}\right), \qquad \overline{a}_{od}\left(\boldsymbol{t}_{\cdot d}\right) \equiv \mathbb{E}\left[a \mid a \leq \hat{a}_{od}^{*}\left(\boldsymbol{t}_{\cdot d}\right)\right] \tag{1.58}$$

denote unit labour requirements of the marginal and average o–based manufacturers active in d, with corresponding demands $\hat{s}_{od}^{*}\left(\boldsymbol{t}_{\cdot d}\right) \equiv s_{od}^{*}\left(\hat{a}_{od}\left(\boldsymbol{t}_{\cdot d}\right) \mid \boldsymbol{t}_{\cdot d}\right)$ and $\overline{s}_{od}^{*}\left(\boldsymbol{t}_{\cdot d}\right) \equiv s_{od}^{*}\left(\overline{a}_{od}\left(\boldsymbol{t}_{\cdot d}\right) \mid \boldsymbol{t}_{\cdot d}\right)$.

The analogue to (1.11) is

$$S_{od}\left(\boldsymbol{t}_{\cdot d}\right)=\int_{0}^{\hat{a}_{od}^{*}\left(\boldsymbol{t}_{\cdot d}\right)} s_{od}^{*}\left(a \mid \boldsymbol{t}_{\cdot d}\right) \cdot L_{o} \, \mathrm{d} G_{o}(a). \tag{1.59}$$

The choke price, written here as a function of freight rates, is the solution to the fixed point problem (1.4)

$$\begin{aligned} p_{d}^{\max} &= \alpha-\frac{\eta}{L_{d}} \sum_{o} S_{od}\left(p_{d}^{\max}, t_{od}\right) \\ &= \alpha-\frac{\eta}{2 \gamma} \sum_{o} \int_{0}^{\hat{a}_{od}\left(p_{d}^{\max}, t_{od}\right)}\left[p_{d}^{\max}-C_{od}\left(a, t_{od}\right)\right] L_{o} \, \mathrm{d} G_{o}(a) \end{aligned} \tag{1.60}$$

where the second equality follows from substituting the profit-maximizing quantities (1.8).

If $p_{d}^{\max} \leq t_{od}$, then even the most efficient o–based manufacturer prefers to stay out of the market, leading to zero sales from source o; $S_{od}\left(p_{d}^{\max}, t_{od}\right)=0$. Note that the choke price is at least as great as the lowest freight rate. Otherwise, $p_{d}^{\max} \leq \min_{o} t_{od}$ implies that no manufacturers serve the market in question; the aggregate quantity across all origins, $\sum_{o} S_{od}\left(p_{d}^{\max}, t_{od}\right)$, evaluates to zero. Therefore, *holding freight rates fixed,* an increase in the choke price weakly increases aggregate exports from any given source, so that $\sum_{o} S_{od}\left(p_{d}^{\max}, t_{od}\right)$ is weakly increasing in $p_{d}^{\max}$. The RHS thus crosses the 45-degree line exactly once, which guarantees a unique solution, $p_{d}^{\max}\left(\boldsymbol{t}_{\cdot d}\right)$, to (1.60).

Thus, the choke price is higher (and competition among manufacturers less fierce) when

Figure 1.C.1: Choke price, conditional on freight rates

a. Conditional on freight rates **b.** Level sets of $p_d^{\max}(\boldsymbol{t}_{\cdot d})$

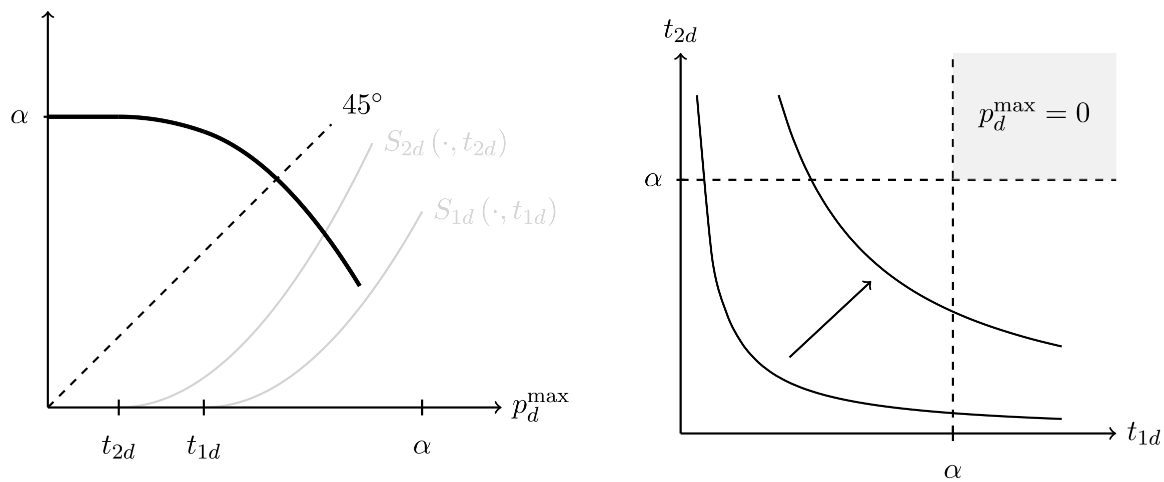

Notes: **Panel a** Equilibrium choke price, given freight rates, as determined in (1.60). The light increasing lines plot $S_{od}(p_d^{\max}, t_{od})$ for $o = 1, 2$. Exports from o to d are zero if the relevant freight rate, t_{od}, exceeds the choke price. Otherwise, $S_{od}(\cdot, t_{od})$ is strictly increasing (recall that $p_d^{\max} \leq \alpha$). The dark downward sloping line displays $\alpha - (\eta/L_d)\sum_o S_{od}(p_d^{\max}, t_{od})$, the right-hand-side of (1.60). Its intersection with the 45-degree line gives the choke price conditional on freight rates, $p_d^{\max}(\boldsymbol{t}_{\cdot d})$. **Panel b** displays equilibrium choke price for any (t_{1d}, t_{2d}) pair. The choke price is effectively zero if $\min_o\{t_{od}\} > \alpha$.

freight rates are high. When $p_d^{\max}(\boldsymbol{t}_{\cdot d})$ is differentiable,

$$\frac{\partial p_d^{\max}(\boldsymbol{t}_{\cdot d})}{\partial t_{od}} = \frac{\eta\, N_{od}(\boldsymbol{t}_{\cdot d})}{2\gamma + \sum_{o' \in \mathcal{O}} \eta\, N_{o'd}(\boldsymbol{t}_{\cdot d})} \in [0, 1), \tag{1.61}$$

where

$$N_{od}(\boldsymbol{t}_{\cdot d}) \equiv \widetilde{N}_o G_o\left(\hat{a}_{od}^*(\boldsymbol{t}_{\cdot d})\right) \tag{1.62}$$

is the mass of o–based exporters active in d. The choke price is more sensitive to freight rates in countries with a large mass of sellers in the destination in question. This, in turn, may be because there is a large number of (exogenous) potential entrants, or because the country is particularly good at producing the final good.

The insights from (1.61) allow us to complete the chain from freight rates to selection into exporting. Specifically, an increase in the freight rate $t_{o'd}$ affects the marginal o–based exporter if both o and o' export to d. If freight rates along (o, d) are high enough to thwart exports, then a marginal increase in $t_{o'd}$ does not affect entry. By the same token, if freight rates along (o', d) are so high as to choke off sales to the destination in question, then a further increase leaves the situation unchanged. I ignore these two cases by assuming that all countries export to d. Substituting (1.61) into (1.9), the extensive margin of exports in

o depend on freight rates as

$$\frac{\partial \hat{a}^*_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}} = \begin{cases} \dfrac{1}{w_o \tau_{od}} \dfrac{\partial p_d^{\max}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}} & \in \dfrac{1}{w_o \tau_{od}}[0,1) \quad \text{if } o' \neq o \\ \dfrac{1}{w_o \tau_{od}} \left(\dfrac{\partial p_d^{\max}(\boldsymbol{t}_{\cdot d})}{\partial t_{od}} - 1 \right) & \in \dfrac{1}{w_o \tau_{od}}[-1,0) \quad \text{if } o' = o. \end{cases} \tag{1.63}$$

The mass of exporters from any given origin is therefore decreasing in the the own-freight rate, and increasing in cross-rates.

1.C.1 Effects of freight rate changes on demand for shipping

Equation (1.59) suggests that aggregate demand for shipping along route (o,d) depend on the the freight rates along routes destined for d, with cross-market effects mediated by the choke price, (1.61). I explore this relationship further by writing total shipping demand along route (o,d) as the product of the mass of o–based varieties sold in d and demand from the average surviving manufacturer,

$$S_{od}(\boldsymbol{t}_{\cdot d}) = N_{od}(\boldsymbol{t}_{\cdot d}) \times \overline{s}^*_{od}(\boldsymbol{t}_{\cdot d}).$$

Following Head and Mayer (2014), I decompose changes in log demand due to changes in freight rate $t_{o'd}$ into the extensive and intensive-and-compositional margins,

$$\frac{\partial \ln S_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}} = \underbrace{\frac{\partial \ln N_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}}}_{\text{extensive}} + \underbrace{\frac{\partial \ln \overline{s}^*_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}}}_{\text{intensive + compositional}}. \tag{1.64}$$

The linearity of shipping demand in a implies that average shipping demand coincides with demand from the average manufacturer,

$$\mathbb{E}\left[s^*_{od}(a|\,\boldsymbol{t}_{\cdot d}) |\, a \leq \hat{a}^*_{od}(\boldsymbol{t}_{\cdot d}) \right] = \overline{s}^*_{od}(\boldsymbol{t}_{\cdot d}),$$

so that the second term in (1.64) simplifies to

$$\begin{aligned} \frac{\partial \ln \overline{s}^*_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}} &= \underbrace{t_{o'd} \frac{\mathbb{E}\left[\partial s^*_{od}(a|\,\boldsymbol{t}_{\cdot d}) / \partial t_{o'd} |\, a \leq \hat{a}^*_{od}(\boldsymbol{t}_{\cdot d}) \right]}{\overline{s}^*_{od}(\boldsymbol{t}_{\cdot d})}}_{\text{intensive}} \\ &\quad + \underbrace{t_{o'd} \frac{\partial \hat{a}^*_{od}(\boldsymbol{t}_{\cdot d})}{\partial t_{o'd}} \frac{g_o(\hat{a}^*_{od}(\boldsymbol{t}_{\cdot d}))}{G_o(\hat{a}^*_{od}(\boldsymbol{t}_{\cdot d}))} \left(\frac{\hat{s}^*_{od}(\boldsymbol{t}_{\cdot d})}{\overline{s}^*_{od}(\boldsymbol{t}_{\cdot d})} - 1 \right)}_{\text{compositional}}. \end{aligned} \tag{1.65}$$

Maintaining the assumption that both o and o' export to d,

$$\frac{\partial S_{od}\left(\boldsymbol{t}_{\cdot d}\right)}{\partial t_{o'd}} = \begin{cases} \dfrac{L_d}{2\gamma}\dfrac{\eta\, N_{o'd}\left(\boldsymbol{t}_{\cdot d}\right)}{2\gamma+\sum_l \eta\, N_{ld}\left(\boldsymbol{t}_{\cdot d}\right)} N_{od}\left(\boldsymbol{t}_{\cdot d}\right) & \text{if } o' \neq o \\ \\ -\dfrac{L_d}{2\gamma}\dfrac{2\gamma+\sum_{i\neq o} \eta\, N_{id}\left(\boldsymbol{t}_{\cdot d}\right)}{2\gamma+\sum_l \eta\, N_{ld}\left(\boldsymbol{t}_{\cdot d}\right)} N_{od}\left(\boldsymbol{t}_{\cdot d}\right) & \text{if } o' = o, \end{cases} \tag{1.66}$$

where I substitute (1.61) and (1.63). Demand is downward sloping in the own freight rate, and increasing in cross-rates. Further,

$$\frac{\partial S_{od}\left(\boldsymbol{t}_{\cdot d}\right)}{\partial t_{o'd}} = \frac{\partial S_{o'd}\left(\boldsymbol{t}_{\cdot d}\right)}{\partial t_{od}},$$

mirroring the symmetry of *inverse* demand highlighted in the main text.

1.C.2 Invertibility of demand system

Cournot competition requires a well-defined inverse demand system. Since it is fruitless to attempt to invert the demand system for freight rates that map to zero demand, I restrict attention to

$$\begin{aligned} T_d^* &\equiv \left\{\boldsymbol{t}_{\cdot d} :\ S_{od}\left(\boldsymbol{t}_{\cdot d}\right) > 0 \text{ for all } o \in \mathcal{O}\right\} \\ &= \left\{\boldsymbol{t}_{\cdot d} :\ \hat{a}^*_{od}\left(\boldsymbol{t}_{\cdot d}\right) > 0 \text{ for all } o \in \mathcal{O}\right\}, \\ &= \left\{\left(t_{1d}, \ldots, t_{Od}\right) :\ t_{od} < p_d^{\max}\left(t_{1d}, \ldots, t_{Od}\right) \text{ for all } o \in \mathcal{O}\right\}, \end{aligned}$$

the set of freight rates consistent with strictly positive demand from all countries. To reach the third line, note that – holding freight rates in other countries fixed – demand along (o, d) is zero, $\hat{a}^*_{od}\left(\boldsymbol{t}_{\cdot d}\right) = 0$, if and only if t_{od} exceeds the equilibrium choke price $p_d^{\max}\left(\boldsymbol{t}\right)$. In other words, demand is invertible only if freight rates are low enough manufacturers export from each country.

In the two–country case portrayed in Figure 1.C.2, T_d^* is the non-rectangular region enclosed by the $S_1 = 0$ and $S_2 = 0$ lines. This precludes applying the results in Cheng (1985) and Okuguchi (1987), which hold when T_d^* is rectangular, i.e., the Cartesian product of O intervals.

The following argument allows us to invert the demand system, at least in the two-country case. The inverse demands at the demand levels $(S_1, S_2) = \left(\overline{S}'_1, \overline{S}'_2\right)$ in Figure 1.C.2 is given by the intersection of the corresponding level sets, and are denoted by $t_o\left(\overline{S}'_1, \overline{S}'_2\right)$. It is clear from the figure that this willingness-to-pay is weakly decreasing in either argument. Note that, as a special case (when $\overline{S}'_o = 0$), the intersection gives the choke price in country o. Finally, the greatest demand for shipping from, say, country 1, conditional on demand $\overline{S}'_2$ in country 2 is given by the level set through the intersection of the $\overline{S}'_2$-level set and the $t_1 = 0$ axis. Proceeding in this manner, we obtain the system of inverse demand functions

Figure 1.C.2: Demand and inverse demand for shipping

a. Levels sets of demand in 2 countries **b.** Inverse demand in country 1

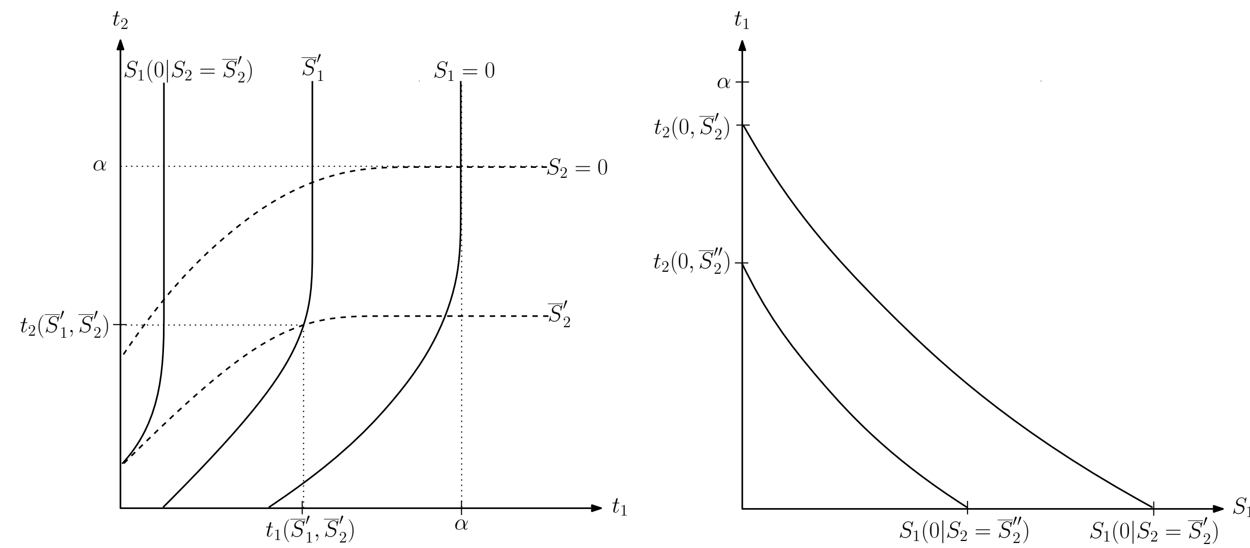

Notes: **Panel a** displays level sets of $S_o(t_1, t_2)$, demands in countries 1 (solid) and 2 (dashed). Levels are $\overline{S}_o'' > \overline{S}_o' > 0$. The $S_o = 0$ line separates regions in (t_1, t_2) space with positive demand from country $o = 1, 2$ from those with zero demand. Demand for shipping from country o is weakly decreasing in the own rate t_o, and weakly increasing in the cross rate t_{-o}. **Panel b** displays inverse demand in country 1, conditional on country 2 demand.

$t_o(\boldsymbol{S}) \equiv t_o(S_1, S_2)$ for $o = 1, 2$.

For more than two countries, I appeal to Berry, Gandhi, and Haile (2013), who provide sufficient conditions for invertibility of the demand system $\{S_{od}(\boldsymbol{t}_{\cdot d})\}_o$. As luck would have it, their results apply even when T_d^* is non-rectangular. Berry et al. (2013) define an artificial country 0 with demand

$$S_{0d}(\boldsymbol{t}_{\cdot d}) \equiv 1 - \sum_{o \in \mathcal{O}} S_{od}(\boldsymbol{t}_{\cdot d}) = 1 - \frac{1}{\eta}\left(\alpha - p_d^{\max}(\boldsymbol{t}_{\cdot d})\right), \tag{1.67}$$

where I use (1.4). The demand system yields inverse demands, $\{t_{od}(\boldsymbol{S}_{\cdot d})\}_o$, if conditions (C1) and (C2) hold.

C1. Weak substitutability $S_{od}(\boldsymbol{t}_{\cdot d})$ is weakly increasing in $t_{o'd}$ for all $o \in \mathcal{O} \cup \{0\}$, and all $o' \in \mathcal{O} \setminus \{o\}$.

Proof. For all but the artificial country, we know from (1.59) that cross-market price effects are mediated by the choke price. The choke price operates along both the extensive margin (allowing more manufacturers to export when $p_d^{\max}$ is large), and along the intensive margin (allowing active manufacturers to charge to sell more). As a result, an increase in $p_d^{\max}$

unambiguously raises aggregate demand $S_{od}(\boldsymbol{t}_{\cdot d})$. From (1.61), the choke price $p_d^{\max}(\boldsymbol{t}_{\cdot d})$ is weakly increasing in each $t_{o'd}$, for $o' \in \mathcal{O} \backslash \{o\}$, so that $S_{od}(\boldsymbol{t}_{\cdot d})$ is weakly increasing in $t_{o'd}$ for all $o \in \mathcal{O}$, and all $o' \in \mathcal{O} \backslash \{o\}$.

As for the artificial country, we see from (1.67) that demand is weakly increasing in $t_{o'd}$, for $o' \in \mathcal{O}$, if and only if the equilibrium choke price, $p_d^{\max}(\boldsymbol{t}_{\cdot d})$, is increasing in $t_{o'd}$. □

C2. Connected substitution* The Jacobian matrix $[\partial S_{od}(\boldsymbol{t}_{\cdot d}) / \partial t_{o'd}]_{o,o'}$, whose entries are given by (1.66), is invertible on T_d^*.

Proof. By assumption, all countries export to d on T_d^*. The determinant of the (symmetric) Jacobian matrix is

$$\det [\partial S_{od}(\boldsymbol{t}_{\cdot d}) / \partial t_{o'd}] = \left(-\frac{L_d}{2\gamma}\right)^O \frac{2\gamma}{2\gamma + \sum_o \eta N_{od}} \prod_o N_{od},$$

which is nonzero under the maintained assumption that each country exports, $N_{od} > 0$ for all $o \in \mathcal{O}$. □

Chapter 2

Testing for Cournot play in the market for shipping to Ecuador

2.1 Introduction

This chapter builds on the theoretical model in Chapter 1, applying a series of revealed–preference tests on shipping sector activity based on Ecuadorian imports. The empirical analysis contributes to a strand of work – dating as far as Afriat (1967) and nicely summarized in Varian (2006) and Chambers and Echenique (2016) – that attempts to rationalizes behaviour through revealed–preference tests. In particular, this chapter contributes to a strand of the literature performing a large number of pass–fail tests, and reporting the fraction of observations consistent with some set of behavioural assumptions.[1]

Afriat's results, which focus on price–taking consumer behaviour, have since been extended to various settings, including Walrasian equilibria (Brown and Matzkin, 1996), and choice over lotteries and across time (Nishimura, Ok, and Quah, 2017). My work builds on Carvajal, Deb, Fenske, and Quah (2013), who develop revealed–preference tests of the single-product Cournot model. Just as utility maximization generates a set of testable restrictions in consumer choice, Cournot competition delivers restrictions on firm outputs and prices. Carvajal et al. (2013) apply the test to OPEC production, and Matsukawa (2016) tests for Cournot play among Japanese retail electricity providers.

Deb and Fenske (2009) extend the analysis to multiproduct Cournot firms, with Carvajal et al. (2014) also considering Bertrand play. Building on these papers, I perform a series of revealed-preference tests on quarterly observations of shipping activity derived from Ecuadorian auto imports from 2007 to 2012. Initially developed in the context of multiproduct oligopolies, this test applies just as well to the demand system derived in the theory section once we identify transport services from a range sources to a given destination as set of distinct products. To the best of my knowledge, this is the first application of the

[1] Another class of tests, the "perturbation approach", exemplified in Varian (1985) and Adams et al. (2015), measures the "severity" of deviations from the proposed law.

multi-product result. Further, I extend Deb and Fenske (2009) by constructing upper and lower bounds for carrier marginal costs from the set of rationalizable observations. Turning to results, the test rejects Cournot behaviour in just under half of the cases tested, and the derived profit-margins suggest an unprofitable shipping sector ever since the Great Recession.

2.2 Data description

The theoretical framework in Chapter 1 features a highly stylized shipping sector that identifies shipping markets with origin-sector pairs trading differentiated varieties in a particular sector. The ideal shipping market for our purposes should therefore distinguish the direction of travel, and employ vessels specific to the goods they carry. This rules out, say, containerships, which transport a wide range of goods.

However, other vessels are often tailored to the goods they transport. For example, roll-on/roll-off (RoRo) ships carry wheeled cargo like cars and trailers. With this in mind, I focus on the 4-digit HS code 8703, "Motor cars and other motor vehicles; principally designed for the transport of persons (other than those of heading no. 8702), including station wagons and racing cars." I apply a revealed–preference test of the model of Cournot competition using data on Ecuadorian auto imports from 2007 to 2012, determining whether the prevailing freight rates and the number of vehicles transported by a given number of carriers over a given length of time are jointly rationalizable by time-varying transport demands and convex, time-invariant shipping cost functions.

The dataset contains the universe of Ecuador imports from 2007 to 2012. For each 8-digit HS product code, k, imported from origin o in month $m = \text{Jan2007}, \ldots, \text{Dec2012}$, I observe $\mathtt{S}^f_{o,k,m}$, the total quantity shipped by carrier f, and total spending on freight services across all carriers, denoted $\mathtt{freight\_spend}_{o,k,m}$. In this chapter, I drop the destination index; rather than speak of route $(o, \text{Ecuador})$ as in Chapter 1, I simply refer to route (or shipping market) o.

Measurement

1. *Quantities:* Focusing on the RoRo market greatly simplifies the relationship between the observed quantities, denoted $\mathtt{S}$, and the theoretical counterpart, S, from Chapter 1. I proceed, without loss of generality, by measuring transport services in units of final goods. I also aggregate observations to the quarterly level, indexed by $j = 2007\text{Q}1, \ldots, 2012\text{Q}4$. Dropping the product subscript, $\mathtt{S}^f_{o,j}$ denotes the number of vehicles that carrier f moves from origin o in quarter j.

2. *Freight rates:* Unlike the quantity data, there is no direct empirical counterpart to the freight rate, t. I therefore define the freight rate along route o in quarter j as the

Figure 2.1: Auto imports and average freight spending

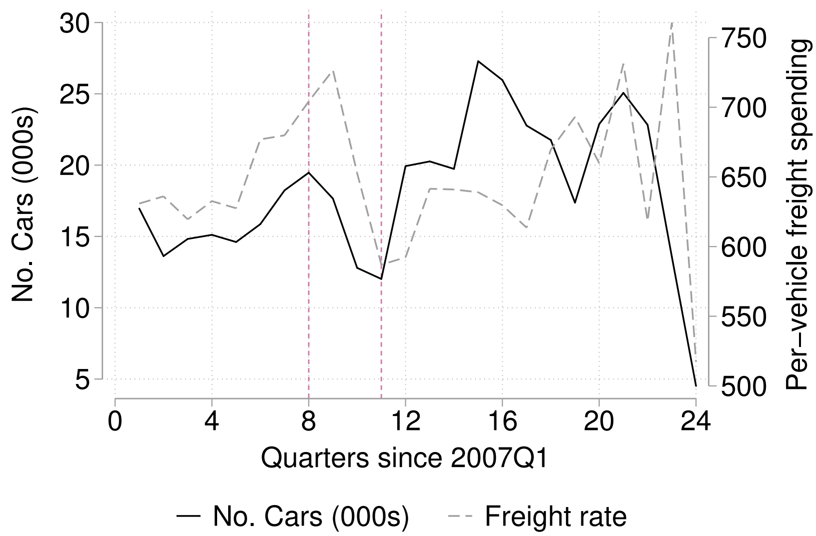

Notes: Quarterly series of auto import activity in Ecuador from 2007Q1 to 2012Q4. The solid line indicates the number of cars (in thousands), while the broken line traces the average freight rate across all source countries, defined as the ratio of freight expenditures on imports from all source countries to the number of cars imported from all countries. The two vertical lines span the Great Recession in Ecuador (Banco Central del Ecuador, 2013; Ray and Kozameh, 2012).

average freight expenditure per vehicle,

$$\mathtt{t}_{oj} \equiv \frac{\mathtt{freight\_spend}_{oj}}{\sum_f \mathtt{S}_{oj}^f}. \tag{2.1}$$

I restrict attention to vessel-operating common carriers, resulting in 18 carriers shipping automobiles from 38 source countries.

Figure 2.1 plots total auto imports, $\mathtt{S}_j \equiv \sum_{fo} \mathtt{S}_{oj}^f$ and the average freight rate across all countries, $\frac{1}{\mathtt{S}_j} \sum_o \mathtt{freight\_spend}_{oj}$ in quarter $j = 2007\text{Q1}, \ldots, 2012\text{Q4}$. The average freight rate over the sample period is \$650 per car.

The final dataset

$$\left\{\mathtt{t}_{oj}, \mathtt{S}_{oj}^f\right\}, \qquad f = 1, \ldots, 18, \quad , o = 1, \ldots, 38, \quad , j = 1, \ldots, 24, \tag{2.2}$$

is summarized in Table 2.1. Note that, unlike the symmetric configuration considered in the theoretical section, carrier presence/activity varies across source countries. Over the 24 quarters, carriers are active in anywhere from 1 to 22 countries, with the average carrier serving about 8 shipping lanes. Further, the average carrier is active in *some* market in just over half of the 24 quarters. Carriers also differ in the number of vehicles they transport, with the smallest carrier moving just one vehicle over the entire sample period, and the largest

Table 2.1: Summary statistics

	Mean	st.dev	Min	Max
Across carriers				
Number of countries served	7.7	6.5	1	22
Number of quarters active	13.7	10.1	1	24
Number of vehicles imported	24172.4	63038.4	1	235138
Across source countries				
Number of active carriers	3.7	3.3	1	16
Number of quarters active	9.3	9.0	1	24
Number of vehicles exported	11450.1	46022.9	1	271476
Across quarters				
Number of source countries	14.7	3.6	6	22
Number of active carriers	10.3	1.3	7	12
Number of vehicles exported	18129.3	5178.3	4518	27287

Notes: Within each block (in bold), the mean, standard deviation, and minimum and maximum of the given variables over the sample. For example, over the 24 quarters, the "average carrier" is active for 7.7 quarters, and transports an average of 24172.4 vehicles from 7.7 source countries.

moving more than 200,000. In particular, the three largest carriers account for just over 90 percent of the total number of vehicles imported by Ecuador from all source countries.[2]

Looking across source countries, relatively few carriers ever serve a given source country, with an average of 3.7 over the sample period. The average source country exports around 11000 automobiles in 9 of the 24 quarters, although this distribution is also skewed. South Korea, Japan, and China jointly account for 90 percent of all auto exports to Ecuador during this period.

2.3 Revealed preference test of carrier first-order conditions

Deb and Fenske (2009) and Carvajal et al. (2014) derive restrictions on data generated by Cournot competition among multiproduct firms. The test relies on three assumptions of Cournot equilibrium play among carriers. Two of these conditions restrict manufacturer demand for transport and carrier costs, while the third condition disciplines marginal costs at equilibrium output levels.

The test formulates these three conditions as a linear programme in the cross–capacity effects on manufacturer willingness–to–pay for transport, and carrier marginal costs along

[2]The distribution is highly skewed towards the largest carriers – while carriers ship an average of 24,172.4 vehicles, the median is 1265.5.

various routes, so that Cournot–rationalizability is equivalent to the feasibility of this programme. Initially developed in the context of multiproduct oligopolies, the test applies equally well in the present setting if we interpret shipping services from various origins as distinct products. A carrier shipping automobiles from China and France effectively offers two distinct products, albeit to distinct sets of consumers (here, manufacturers).

Restriction 1: Demand substitution patterns

Larger shipping capacity in any route destined for Ecuador should lower willingness–to–pay for transport among auto manufacturers in any of the 18 source countries. However, unlike the model in Section 1.2.1, I allow transport demand to vary over time. Rather than build a full-fledged dynamic model, I simply allow temporal variation in the appropriate model parameters, noting that demand for transport is derived from the interaction between consumer demand for automobiles and the characteristics of auto manufacturers. On the demand side, the size of the Ecuadorian market, or the tastes of auto consumers may vary over time, motivating quarter–specific market sizes, $L_{d,j}$, and preference parameters $(\alpha_j, \eta_j, \gamma_j)$. On the supply side, I allow variation in the number of potential o–based auto manufacturers, wages and tariffs, hence $\left(\widetilde{N}_{o,j}, w_{o,j}, \tau_{o,j}\right)$.

Regardless of the nature of the demand shifts, I assume a time-varying, differentiable inverse demand function $t_{oj}(\cdot)$, satisfying

$$\frac{\partial t_{oj}\left(S_{1j}, \ldots, S_{Oj}\right)}{\partial S_{o'j}} \equiv \partial_{o'} t_{oj} \leq 0, \tag{2.3}$$

where $\partial_{o'} t_{oj}$ is one of the unknowns of the desired linear programme. This nests the specific form (1.13) suggested by the theory, – which requires strictly negative and symmetric substitution patterns, $\partial_{o'} t_{oj} < 0$, and $\partial_{o'} t_{oj} = \partial_o t_{o'j}$ – and is therefore biased in favour of rationalizing shipping activity as a Cournot equilibrium outcome. Small magnitudes of $\partial_{o'} t_{oj}$ imply that carriers act as price takers in the relevant routes. For example, $\partial_{o'} t_{oj} \approx 0$ for $o' \neq o$ rules out the interdependent demand system posited in the theory, while $\partial_o t_{oj} \approx 0$ implies a perfectly competitive market for transport from o to Ecuador.

Restriction 2: Positive marginal costs

The second and third restrictions apply to the shipping technology, specifically, to the carrier cost functions. Departing from the stylized cost functions in the theory section, I allow supply-side cross-market interactions, only requiring that costs be increasing and convex in carrier output in each active market. This generalizes the specification in the theory section, in which carriers faced constant marginal costs along each route. Further, costs depend on the set of markets in which a carrier is active. For example, entry into a new market may divert resources from existing markets, thereby raising costs. Formally,

Definition 2.1. Carrier f transports vehicles to Ecuador from

$$\mathcal{O}_j^f \equiv \{\, o \in \mathcal{O} \,:\, \mathtt{S}_{oj}^f > 0 \,\},$$

its *set of active markets in quarter* $j \in \mathcal{J}$.

Cost functions are then parametrized by $\mathcal{O}_j^f$,

$$\Psi^f(\,\cdot|\mathcal{O}_j^f\,) : \mathbb{R}_{++}^{\left|\mathcal{O}_j^f\right|} \to \mathbb{R}_+. \tag{2.4}$$

Like the theoretical model in Section 1.2.2, shipping costs are separable across destination markets – otherwise, the analysis would Ecuadorian export data, in addition to the import data at hand. However, (2.4) allows arbitrary shipping costs interdependencies across source countries. Like the demand side restriction (2.3), this cost specification is biased in favour of Cournot rationalizability, relative to the more stringent theoretical specification.

To simplify notation, let Ψ_j^f denote the total cost to carrier f of transporting cargo to Ecuador from all its active markets, where j indexes the dependence on the set of active markets, $\mathcal{O}_j^f$. Consider the empirical counterpart to the carrier problem in (1.18),

$$(\mathtt{S}_{oj}^f)_o \in \underset{(S_o^f)_{o\in\mathcal{O}_j^f}}{\operatorname{argmax}} \sum_{o\in\mathcal{O}_j^f} t_{oj}[\,(S_o^f + \mathtt{S}_{oj}^{-f})_{o\in\mathcal{O}}\,] \times S_o^f - \Psi_j^f[\,(S_o^f)_{o\in\mathcal{O}_j^f}\,],$$

written as a function of the observable quantities, and some inverse demand function, $t_{oj}\,(\cdot)$, and carrier-specific cost function $\Psi_j^f\,(\cdot)$. Here, $\mathtt{S}_{oj}^{-f} \equiv \sum_{f'\neq f} \mathtt{S}_{oj}^{f'}$ is the (observed) time-j total shipping activity from source o among carrier f's rivals, with the understanding that $S_{oj}^f = 0$ for $o \notin \mathcal{O}_j^f$. Rationalizing the dataset (2.2) as a Cournot outcome requires that the observed freight rate be consistent with some inverse demand function evaluated at the observed aggregate output across all active carriers:

$$\mathtt{t}_{oj} = t_{oj}\,(\mathtt{S}_{1j}, \ldots, \mathtt{S}_{Oj})\,, \qquad \mathtt{S}_{oj} \equiv \sum_{f\in F} \mathtt{S}_{oj}^f$$

for some inverse demand function, $t_{oj}\,(\cdot)$. Substituting for the observed freight rate, the second test restriction follows from the empirical analogue to (1.19), carrier f's optimality condition in market $o \in \mathcal{O}_j^f$,

$$\mathtt{t}_{oj} + \sum_{o'\in\mathcal{O}_j^f} \mathtt{S}_{o'j}^f\,\partial_o t_{o'j} - \frac{\partial \Psi_j^f}{\partial S_{oj}^f} = 0, \tag{2.5}$$

which sets the marginal returns to shipping from o to zero. Specifically, I assume carriers positive marginal costs of shipping, regardless of the scale of transport services supplied; $\partial\Psi_j^f/\partial S_{oj}^f > 0$ for all $(S_o^f)_{o\in\mathcal{O}_j^f}$. In equilibrium, this implies that the right-hand side of (2.5) is strictly positive. That is,

$$\mathtt{t}_{oj} + \sum_{o' \in \mathcal{O}_j^f} \mathtt{S}_{o'j}^f \, \partial_o t_{o'j} > 0, \tag{2.6}$$

which counteracts the demand–side restriction $\partial_o t_{o'j} \leq 0$. Intuitively, given the freight rate $\mathtt{t}_{oj}$, carriers move large volumes from o' to Ecuador only if $\left|\partial_o t_{o'j}\right|$ is sufficiently small. If the negative spillovers of additional capacity from o on willingness–to–pay for transport from o' were any larger, the marginal returns to shipping from o to Ecuador would fall, suppressing carrier activity from o.

Gathering the inequalities in (2.6) across carriers (f) active in the various markets (o) over time (j), we obtain a family of restrictions on $\left(\partial_o t_{o'j}\right)_{oo'j}$. Indeed, the only unobservables in (2.6) are the cross–capacity effects, $\left(\partial_o t_{o'j}\right)_{o'}$, the unknowns in the linear programme.

Restriction 3: Convex shipping costs

Finally, I impose convex cost functions in the transport sector. Since I assume the cost functions are indexed by the set of active markets, testing for convexity requires at least two quarters in which the carrier is present in exactly the same set of markets.

Definition 2.2. Quarters j and j' are *comparable for carrier* f if carrier f is active in exactly the same set of markets in these periods. Formally,

$$j \overset{f}{\sim} j' \iff \mathcal{O}_j^f = \mathcal{O}_{j'}^f \neq \emptyset,$$

where focusing on carriers active in at least one market avoids irrelevant carriers.

Given a dataset in the form (2.2), I partition the set of observations using the $\overset{f}{\sim}$ relation. For example, consider the shipping pattern outlined in Table 2.1. Carrier X, who produces in a different set of markets each quarter, has no comparable quarters. Such carriers do not contribute any convexity-based restrictions to the test. In contrast, we may partition the set of observations according to Carrier Y's activity as quarters $\{1, 2\}$, when Y ships from Brazil and France, and quarter 3, when Y is active in just the Brazilian market. As we will see shortly, while quarter 3 places no additional restrictions on the data, Y's activity in quarters 1 and 2 introduces a pair of restrictions. To see this, expand $\Psi^f(\cdot)$ around the vector of capacities in quarter j, $(S_{1j}^f, \ldots, S_{Oj}^f)$.[3]

Definition 2.3. The cost function $\Psi^f(\cdot)$ is convex iff

$$\Psi^f(S_{1j'}^f, \ldots, S_{Oj'}^f) \geq \Psi^f(S_{1j}^f, \ldots, S_{Oj}^f) + \sum_o (S_{oj'}^f - S_{oj}^f) \frac{\partial \Psi_j^f}{\partial S_o^f}. \tag{2.7}$$

[3] I drop the dependence of the cost function on active markets since I consider comparable periods.

Using (2.5) to substitute for $\partial \Psi^f / \partial S_o^f$, and rearranging, convexity imposes

$$\Delta_{j'-j}\Psi^f \geq \sum_{o \in \mathcal{O}_j^f} [\, \mathtt{t}_{oj} + \textstyle\sum_{o' \in \mathcal{O}_j^f} \mathtt{S}_{o'j}^f \cdot \partial_o t_{o'j} \,] \times \Delta_{j'-j}\mathtt{S}_o^f, \tag{2.8}$$

where $\Delta_{j'-j}\mathtt{S}_o^f \equiv \mathtt{S}_{oj'}^f - \mathtt{S}_{oj}^f$ and $\Delta_{j'-j}\Psi^f \equiv \Psi_{j'}^f - \Psi_j^f$ are the changes in shipping activity along route o, and in total carrier costs over the period. Like the first two conditions, convexity places *linear* restrictions on the unknown cross–capacity effects $\partial_o t_{o'j}$. In contrast, it is the only condition restricting cross–capacity effects over time, and (changes in) total carrier costs.

For example, consider a carrier, f, shipping only from Brazil and France in quarters 1 and 2. Suppose f raises output along the Brazilian route ($\Delta_{2-1}\mathtt{S}_B^f \equiv \mathtt{S}_{B2}^f - \mathtt{S}_{B1}^f > 0$), but transports the same number of vehicles from France ($\Delta_{2-1}\mathtt{S}_F^f \equiv \mathtt{S}_F^f - \mathtt{S}_F^f = 0$). Combining the optimality conditions in quarters 1 and 2 bounds the change in total shipping costs over the period

$$\mathtt{t}_{B1} + \mathtt{S}_{B1}^f \cdot \partial_B t_{B1} + \mathtt{S}_F^f \cdot \partial_B t_{F1} \leq \frac{\Delta_{2-1}\Psi^f}{\Delta_{2-1}\mathtt{S}_B^f} \leq \mathtt{t}_{B2} + \mathtt{S}_{B2}^f \cdot \partial_B t_{B2} + \mathtt{S}_F^f \cdot \partial_B t_{F2}.$$

Since f raises its output over the period, convex shipping costs reject values of $\partial . t_{.,.}$ consistent with lower marginal shipping costs in quarter 2,

$$\mathtt{t}_{B1} + \mathtt{S}_{B1}^f \cdot \partial_B t_{B1} + \mathtt{S}_F^f \cdot \partial_B t_{F1} > \mathtt{t}_{B2} + \mathtt{S}_{B2}^f \cdot \partial_B t_{B2} + \mathtt{S}_F^f \cdot \partial_B t_{F2}.$$

All else equal, this is more likely if, for example, freight rates along the Brazilian route were much higher in the first quarter. Finally, note that the convexity restriction linking any two quarters is slack absent sufficient variation in carrier activity over the period in question.

Summing up, Deb and Fenske (2009) show that the dataset (2.2) satisfies the three conditions outlined if and only if the observed freight rates are consistent with Cournot equilibrium play given some time-dependent demand system and time-invariant cost functions ("Cournot rationalizable"). In other words, they show that we can, in principle, construct demand systems and cost functions given solutions to the linear programme defined by (2.3), (2.6), and (2.8).

Proposition 2.1. *(Deb and Fenske, 2009) Dataset (2.2) is Cournot rationalizable if and only if there exist*

1. *$O^2 \times J$ nonpositive cross–capacity effects $\partial_o t_{o'j}$ and*

2. *$F \times J$ positive total shipping costs Ψ_j^f*

such that for all quarters j', j'', and all carriers f,

1. *The law of demand, (2.3), holds for all source–country pairs, (o, o'), and all quarters j, with strict inequality for distinct country pairs (when $o \neq o'$).*

2. *Carriers face positive marginal costs, so that (2.6) holds in all countries $o \in \mathcal{O}_j^f$ served by any carrier, f, is active in any given quarter, j.*

3. *All carriers have convex cost functions, so that (2.8) holds for all f, and any distinct pair of comparable quarters (j, j').*

Before demonstrating the power of Proposition 2.1, note that this test is equally valid on any subset of quarters; it only imposes that carrier cost functions remain the same over the period in question. Similarly, it applies to activity among any subset of carriers. For example, if we only have data on the activity of a handful of carriers active in a given set of markets, then we simply interpret the rationalizing inverse demand, $t_{oj}(\cdot)$, as *residual* inverse demand after accounting for aggregate capacity among unobserved carriers. To summarize, we can therefore test for strategic interactions among any set of carriers over arbitrary periods.

2.3.1 Illustrating the Cournot–rationalizability test

Consider the shipping activity of carriers Y and Z in quarters 1 and 2 from Table 2.1.

Table 2.1: Consistency with Cournot equilibrium

	Freight rates		Carrier X		Carrier Y		Carrier Z	
j	Brazil	France	Brazil	France	Brazil	France	Brazil	France
1	t	t	10	0	$S_1^Y > 0$	$S_1^Y > 0$	$S_1^Z > 0$	$S_1^Z > 0$
2	$bt > t$	$bt > t$	100	100	$S_2^Y > 0$	$S_2^Y > 0$	$S_2^Z > 0$	$S_2^Z > 0$
3	40	40	0	0	100	0	50	50

Notes: Three–quarter activity of three carriers, X, Y, and Z, in routes from Brazil and France to a given destination. Between quarters 1 and 2, freight rates increase by a factor of b in both routes. In each period, carriers Y and Z set the same capacity across markets; $S_{\text{BR},j}^f = S_{\text{FR},j}^f = S_j^f$ for $i = Y, Z$ and $m = 1, 2$. However, these levels may differ over time; S_1^f is not necessarily the same as S_2^f. Taking period 1 activity as given, this pattern is consistent with Cournot equilibrium provided capacities do not differ too much in period 2, that is, $\left|S_2^X - S_2^Y\right|$ is bounded above. See Section 2.3.1.

Both carriers are always equally active in the two markets, with carrier $f = X, Y$ transporting S_j^f in quarter j. Freight rates are identical across markets within periods, but increase by a factor $b > 1$ from quarter 1 to 2. By (2.6), and the fact that $\partial_o t_{o'j} \leq 0$,

$$\begin{aligned} j = 1: \qquad & 0 > \partial_o t_{B1} + \partial_o t_{F1} > -t/\max\left\{S_1^Y, S_1^Z\right\} \\ j = 2: \qquad & 0 > \partial_o t_{B2} + \partial_o t_{F2} > -bt/\max\left\{S_2^Y, S_2^Z\right\}. \end{aligned} \tag{2.9}$$

Applying (2.8) twice to carrier Y, first with $j=1$ and $j'=2$, and then reversing the roles of j and j' delivers bounds on f's cost difference over time,

$$\begin{array}{lcl} \Delta_{1-2}\Psi^Y & \leq & 2\left(S_1^Y - S_2^Y\right)t + \left(S_1^Y - S_2^Y\right)S_1^Y \sum_{o,o'} \partial_o t_{o'1} \\ \Delta_{1-2}\Psi^Y & \geq & 2\left(S_1^Y - S_2^Y\right)bt + \left(S_1^Y - S_2^Y\right)S_2^Y \sum_{o,o'} \partial_o t_{o'2}. \end{array} \tag{2.10}$$

Combining (2.9) and (2.10) delivers yet tighter bounds on the change in f's total costs,

$$\underbrace{2\left(1 - \frac{S_2^f}{\max\left\{S_2^Y, S_2^Z\right\}}\right)}_{\in[0,1)} b\left(S_1^Y - S_2^Y\right)t \leq \Delta_{1-2}\Psi^Y \leq 2\left(S_1^Y - S_2^Y\right)t.$$

Suppose that despite the increase in freight rates, Y is everywhere more active in the first quarter; $S_1^Y - S_2^Y > 0$. Further, (to rule out the trivial case), assume Y ships fewer vehicles than Z in quarter 2; $S_2^Y < S_2^Z$. The pattern in Table 2.1 is then Cournot rationalizable if

$$\frac{S_2^{-f} - S_2^f}{S_2^{-f}} \leq \frac{1}{b},$$

that is, if the more active carrier does not tremendously outshine its rival. Further, the gap between period-2 outputs should shrink as the jump in freight rates increases.

Finally, note that I have some leeway as to which observations to include when running the test on Table 2.1. Scanning the carriers, I exclude X because it has no comparable pairs of observations, while Y is active in the same set of markets in the first two quarters. As for Z, there is no reason, *a priori*, to run the test on its activity in the first 2 quarters. The test is equally valid in the pair of quarters $\{1,3\}$, $\{2,3\}$, or even the triple $\{1,2,3\}$; of course, rejecting Cournot behaviour, on say, $\{1,3\}$ renders the $\{1,2,3\}$ test pointless.

2.3.2 Implementing the test

Since the test is equally valid on any subset of carriers and/or quarters, I run this test on several small subsets of the data, reporting the share of "successful" datasets. This approach is common to revealed preference tests of consumer behaviour based on Afriat (1967). For example, Carvajal et al. (2013) and Matsukawa (2016) follow the same procedure when testing for Cournot behaviour on OPEC production, and output in Japanese retail electricity markets respectively.

In particular, starting with the original dataset (2.2) of 18 carriers over 24 quarters, I generate subsamples consisting of a given number of carriers over a given number of consecutive quarters. Define $\mathcal{F}_F$ as the collection of the $18!/\left(F!\left(18-F\right)!\right)$ sets containing any F of the 18 carriers. Likewise, let $\mathcal{J}_J$ be the set of $24-J+1$ sets of J consecutive quarters from the 24 available periods. Then

$$\texttt{SetObs}_{F,J} \in \mathcal{F}_F \times \mathcal{J}_J \tag{2.11}$$

is a set of observations of any F carriers over any J consecutive quarters, with a typical element denoted by $\texttt{obs}$. At one extreme, $\texttt{SetObs}_{18,24}$ contains just one set of observations, the grand dataset in (2.2). At the other extreme, $\texttt{SetObs}_{2,2}$ contains the shipping activity of all possible pairs of carriers over any two consecutive quarters.[4] Although $\texttt{SetObs}_{F,J}$ may contain an unwieldy number of datasets for moderately large values of F and J, the test disregards datasets lacking comparable periods for at least one carrier. For example, the observations in Table 2.1 correspond to some $\texttt{obs} \in \texttt{SetObs}_{3,3}$. However, since X is never active in the same set of markets, the test would only apply to datasets in $\texttt{SetObs}_{2,3}$.

Following Carvajal et al. (2013), I report the fraction of datasets in $\texttt{SetObs}_{F,J}$ that are Cournot–rationalizable. I consider narrow time periods and a small number of carriers for two reasons. Apart from easing the computational burden, using a *small* number of *consecutive* quarters, J, reduces the likelihood of significant changes to carrier costs, an assumption built-in to the Cournot test. Using a small number of carriers, F, allows me to pick up Cournot play that would otherwise be overlooked among a larger set of carriers. For example, we may expect strategic play among the three largest carriers, Y, Z, and X, who ship over 95% of the automobiles entering Ecuador over the sample, than among the 8 remaining carriers.

Table 2.2 presents the fraction of rationalizable datasets in $\texttt{SetObs}_{F,J}$, for $F = 2, 3$ and $J = 2, 3, 4$. Panel (a) considers all carriers, while Panel (b) considers the three largest carriers. As expected, the highest success rate within each panel features the smallest number of carriers and the shortest window, with 53 percent of datasets of all carriers, and 55 percent of datasets involving the "Big–3", being consistent with Cournot play. Given any cell in either panel, the success rate falls as we consider more carriers or longer horizons. At the other extreme, fewer than one quarter of 4-quarter-long datasets with 3 carriers are Cournot rationalizable. The test is powerful enough to detect deviations from Cournot play despite the flexible stance on manufacturer demand for shipping (allowing arbitrary cross-market elasticities) and carrier cost functions (non-constant marginal cost). Finally, comparing any analogous cell pairs across panels, it seems like restricting observations to the largest carriers raises the likelihood of a successful test.

[4] We need at least two carriers over two periods to run the test.

Table 2.2: Cournot rationalizability success rates

a. All carriers

Window length	Number of carriers: 2	3
2	0.53	0.45
3	0.31	0.15
4	0.26	0.12

b. Big-3 carriers

Window length	Number of carriers: 2	3
2	0.55	0.51
3	0.44	0.36
4	0.36	0.26

Notes: Given all possible subsets of the data with "column" carriers within a "row"–quarter window during the entire sample, each cell displays the fraction that are Cournot rationalizable. For example, half of the possible subsets of three carriers active within a 2–quarter interval from January 2007 to December 2012 are consistent with Cournot behaviour among the three carriers, as defined in Proposition 2.1.

2.3.3 Bounding structural parameters

Returning to the general problem, we can, at least in principle, recover carrier marginal and total costs whenever the linear program is feasible. First, given the observed freight rates and derived cross-market elasticities, we obtain a viable measure of the carrier's marginal cost in a given market from the familiar relationship between prices, markups, and marginal costs, (2.5). As usual, we expect $\partial \Psi_j^f / \partial S_{oj}^f \leq \mathtt{t}_{oj}$ for all carriers f and all markets o, with strict inequality if $\partial_o t_{o'j} < 0$ for some o'. In the extreme case that we find $\partial_o t_{\cdot j} \equiv 0$, then the restrictions predict marginal cost pricing.

However, the Deb and Fenske (2009) test does not deliver a unique value for the marginal cost because rationalizability only requires a non-empty feasible region and offers no criterion to select particular points whenever this region is nonempty.Varian (2006) terms this the "recoverability" problem. Paraphrasing his discussion of consumer choice, is there a way to describe the entire set of cost functions consistent with the data?

Rather than attempt to fully characterize the entire cost *function*, I supplement the Deb and Fenske (2009) test with a criterion that bounds each carrier's marginal cost at the observed output in an origin–quarter, selecting those feasible points that either minimize or maximize $\partial \Psi_j^f / \partial S_{oj}^f$. Specifically, carrier f's marginal cost in market o in quarter j, $\partial \Psi_j^f / \partial S_{oj}^f$, satisfies

$$\underline{\text{m.c.}}_{oj}^f \leq \partial \Psi_j^f / \partial S_{oj}^f \leq \overline{\text{m.c.}}_{oj}^f,$$

where

$$\begin{aligned} \underline{\text{m.c.}}_{oj}^f &\equiv \mathtt{t}_{oj} + \min_{\partial_o t_{\cdot j}} \textstyle\sum_{o' \in \mathcal{O}_j^f} \mathtt{S}_{o'j}^f \partial_o t_{o'j} \\ \overline{\text{m.c.}}_{oj}^f &\equiv \mathtt{t}_{oj} + \max_{\partial_o t_{\cdot j}} \textstyle\sum_{o' \in \mathcal{O}_j^f} \mathtt{S}_{o'j}^f \partial_o t_{o'j}, \end{aligned} \tag{2.12}$$

subject to $\{\partial_o t_{o'j}\}_{o'}$ satisfying (2.3), (2.6), and (2.8). Given the freight rate $\mathtt{t}_{oj}$, the lowest possible marginal cost is obtained by maximizing the markup term, $-\sum_{o'\in\mathcal{O}_j^f} \mathtt{S}_{o'j}^f \partial_o t_{o'j}$, while the highest value occurs when the markup is at its lowest.

Since I run this test once for every carrier-origin-quarter, the marginal-cost bounds for one such triple may derive from a different feasible point than another's. For example, within the (o, j) origin-quarter pair, we obtain different marginal-cost bounds across carriers to the extent that these carriers provide different levels of transport services across markets o'. In the extreme event that carriers have equal market shares in all markets, so that $\mathtt{S}_{o'j}^f \equiv \mathtt{S}_{o'j}$ for all o', we obtain identical bounds across carriers.

Deriving "tighter" bounds, relative to $0 \le \partial\Psi_j^f/\partial S_{oj}^f \le \mathtt{t}_{oj}$, requires overlapping carrier activity in the various markets. Specifically, the constraints (2.6) and (2.8) slacken as more and more carriers become inactive along the route from country o. In the extreme event that f is the sole carrier in market o in quarter j, then $\underline{\text{m.c.}}_{oj}^f$ is attained by setting $\partial_o t_{o'j}$ as small as possible subject to the sole binding restriction that the carrier's marginal cost is positive,

$$\mathtt{t}_{oj} + \sum_{o'\in\mathcal{O}_j^f} \mathtt{S}_{o'j}^f \cdot \partial_o t_{o'j} > 0,$$

which would imply $\underline{\text{m.c.}}_{oj}^f \approx 0$. Similarly, $\overline{\text{m.c.}}_{oj}^f$ is attained by setting $\partial_o t_{o'j}$ equal to zero, so that $\overline{\text{m.c.}}_{oj}^f \equiv \mathtt{t}_{oj}$. To summarize, tightening tighten the marginal cost bounds requires strictly negative cross-capacity effects, which, in turn, require persistent carrier activity over the testing horizon.

Given a carrier-origin-quarter triple, (f, o, j), I first divide the original dataset into smaller datasets. Unlike the earlier test, I solve the augmented linear programme (2.12) on each dataset in

$$\mathtt{SetObs}_J \in \mathcal{F}_{18} \times \mathcal{J}_J,$$

the set of observations of all carriers over any J consecutive quarters. In other words, unlike the previous test, I do not restrict the set of carriers. Let

$$\mathtt{SetObs}_J^+ \equiv \{\mathtt{obs} \in \mathtt{SetObs}_J \ : \ \mathtt{obs} \text{ is rationalizable}\}$$

be the set of observations that passes the augmented Cournot test, and let $\partial_\cdot t_{\cdot,\cdot}^{\mathtt{obs}}$ denote the optimizer corresponding to $\mathtt{obs} \in \mathtt{SetObs}_J^+$. The lower bound on f's quarter-j marginal cost of serving the route from o to Ecuador is the lowest marginal cost across rationalizable datasets.

$$\underline{\text{m.c.}}_{oj}^f = \mathtt{t}_{oj} + \min_{\cup_J \mathtt{SetObs}_J^+} \sum_{o'\in\mathcal{O}_j^f} \mathtt{S}_{o'j}^f \partial_o t_{o'j}^{\mathtt{obs}}.$$

Similarly, I obtain the upper bound by maximizing the marginal cost over all rationalizable datasets. Finally, to ease comparisons across carriers and markets and over time, I translate the marginal-cost bounds into profit-margin bounds,

$$\overline{\texttt{Lerner}}^f_{oj} = \frac{\texttt{t}_{oj} - \underline{\text{m.c.}}^f_{oj}}{\texttt{t}_{oj}}, \qquad \underline{\texttt{Lerner}}^f_{oj} = \frac{\texttt{t}_{oj} - \overline{\text{m.c.}}^f_{oj}}{\texttt{t}_{oj}}. \tag{2.13}$$

I then translate these bounds into bounds on the perceived capacity effects $\lambda_{o',o,j}$. Rewriting the individual first–order conditions (2.5) as

$$-\sum_{o'\in\mathcal{O}^f_j} \frac{\texttt{S}^f_{o'j}}{\texttt{t}_{oj}}\,\partial_o t_{o'j} = \frac{\texttt{t}_{oj} - \partial\Psi^f_j/\partial S^f_{oj}}{\texttt{t}_{oj}},$$

so that the bounds in (2.13) imply

$$-\,\overline{\texttt{Lerner}}^f_{oj} \leq \sum_{o'\in\mathcal{O}^f_j} \frac{\texttt{S}^f_{o'j}}{\texttt{t}_{oj}}\,\partial_o t_{o'j} \leq -\,\underline{\texttt{Lerner}}^f_{oj}. \tag{2.14}$$

Figure 2.1 illustrates these bounds in the two–country case. Conditional on the observed $\left(\texttt{S}^f_{o'j}/\texttt{t}_{1j}\right)_{o'=1,2}$, points below the outermost line correspond to $(\partial_1 t_{1j}, \partial_1 t_{2j})$ pairs consistent with the upper bound on carrier f's profit margins, while those above the innermost line correspond to pairs consistent with the lower bound. The shaded region is the set of rationalizing own– and cross–capacity effects based on carrier f's activity. This area shrinks as f's Lerner–index bounds narrow, delivering tighter bounds on $(\partial_1 t_{1j}, \partial_1 t_{2j})$.[5] Panel (b) superimposes the valid regions obtained from the activity of two carriers, shading their intersection.

[5]For reference, the demand system (1.13) imposes the additional restriction that own–capacity effects dominate, $\partial_o t_{oj} > \partial_o t_{o'j}$, restricting the set of valid points to those above the 45–degree line.

Figure 2.1: Bounding perceived own– and cross–effects of capacity in market 1

a. Individual carrier **b.** Combining information across carriers

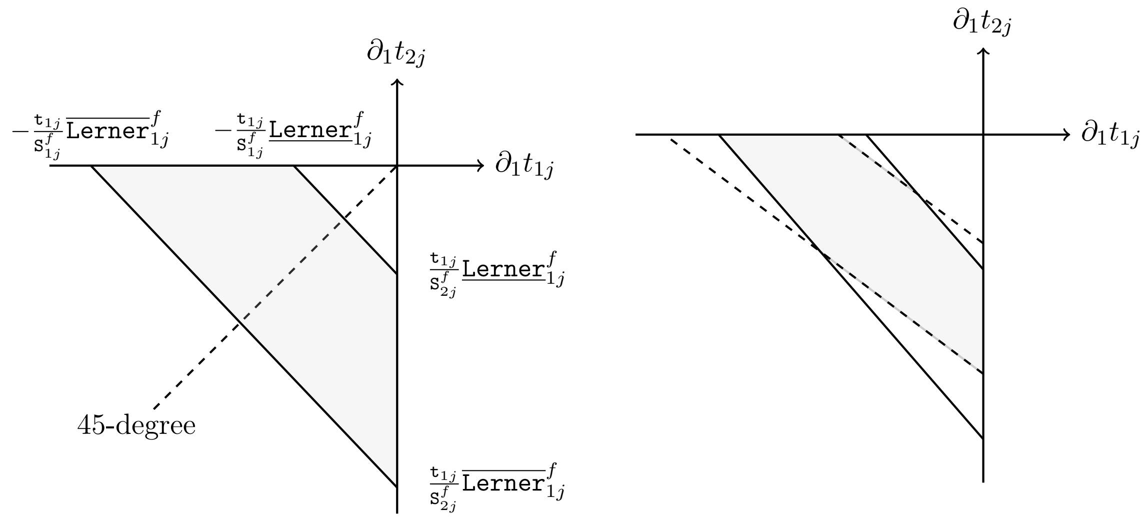

Notes: From (2.3), $\partial_1 t_{oj} \equiv \partial t_{oj}/\partial S_{1j}$ is the marginal effect of capacity in market 1 on manufacturer willingness–to–pay for transport from o to Ecuador in quarter j. **Panel (a)** Points in the shaded area are consistent with a single carrier's profit margins being in the region defined in (2.13) (the theory in Section 1.2 requires points to lie above the 45–degree line). **Panel (b)** The intersection of such areas across carriers (one solid, the other broken) yields the set of $(\partial_1 t_{1j}, \partial_1 t_{2j})$ pairs consistent with both carriers' actions. Projecting this area on the horizontal axis yields the safest bound on $\partial_1 t_{1j}$.

Figure 2.2 presents the profit margin bounds for the three largest carriers, Y, Z, and X. The broken lines contain the great recession period. Owing to carrier inactivity in certain markets, I am unable to run the test for some (f, o, j) triples. For example, Y never serves the Japanese market, while Z is active in the Korean market in only one of the 24 quarters. Similarly, X is only intermittently active in the Chinese market before the second half of 2011, and thus appears in very few consecutive quarters.

Setting these data issues aside, two broad patterns emerge. First, the data does not reject strictly positive profit margins before the recession. Admittedly, the bounds on X in Japan, and Z in China and Japan are hardly informative during this period, predicting a range of fates from breaking even at one end, to raking in twice the marginal cost at the other. In the pre-recession period, the most informative bounds concern Y and X in South Korea – as a rough measure, the mid-points of Y's intervals range from 30% to 35%, while X's lie in the 43-50% range.

Second, the profit–margin bounds collapse on zero during the Great Recession. In light of (2.14) and Figure 2.1, this suggests that carriers act as price takers in the various shipping markets, if the theory in Chapter 1 is to be believed. If, in contrast, the carriers had some market power, then marginal–cost pricing would only occur near zero levels of output.

Figure 2.2: Bounding carrier-origin-specific profit margins over time

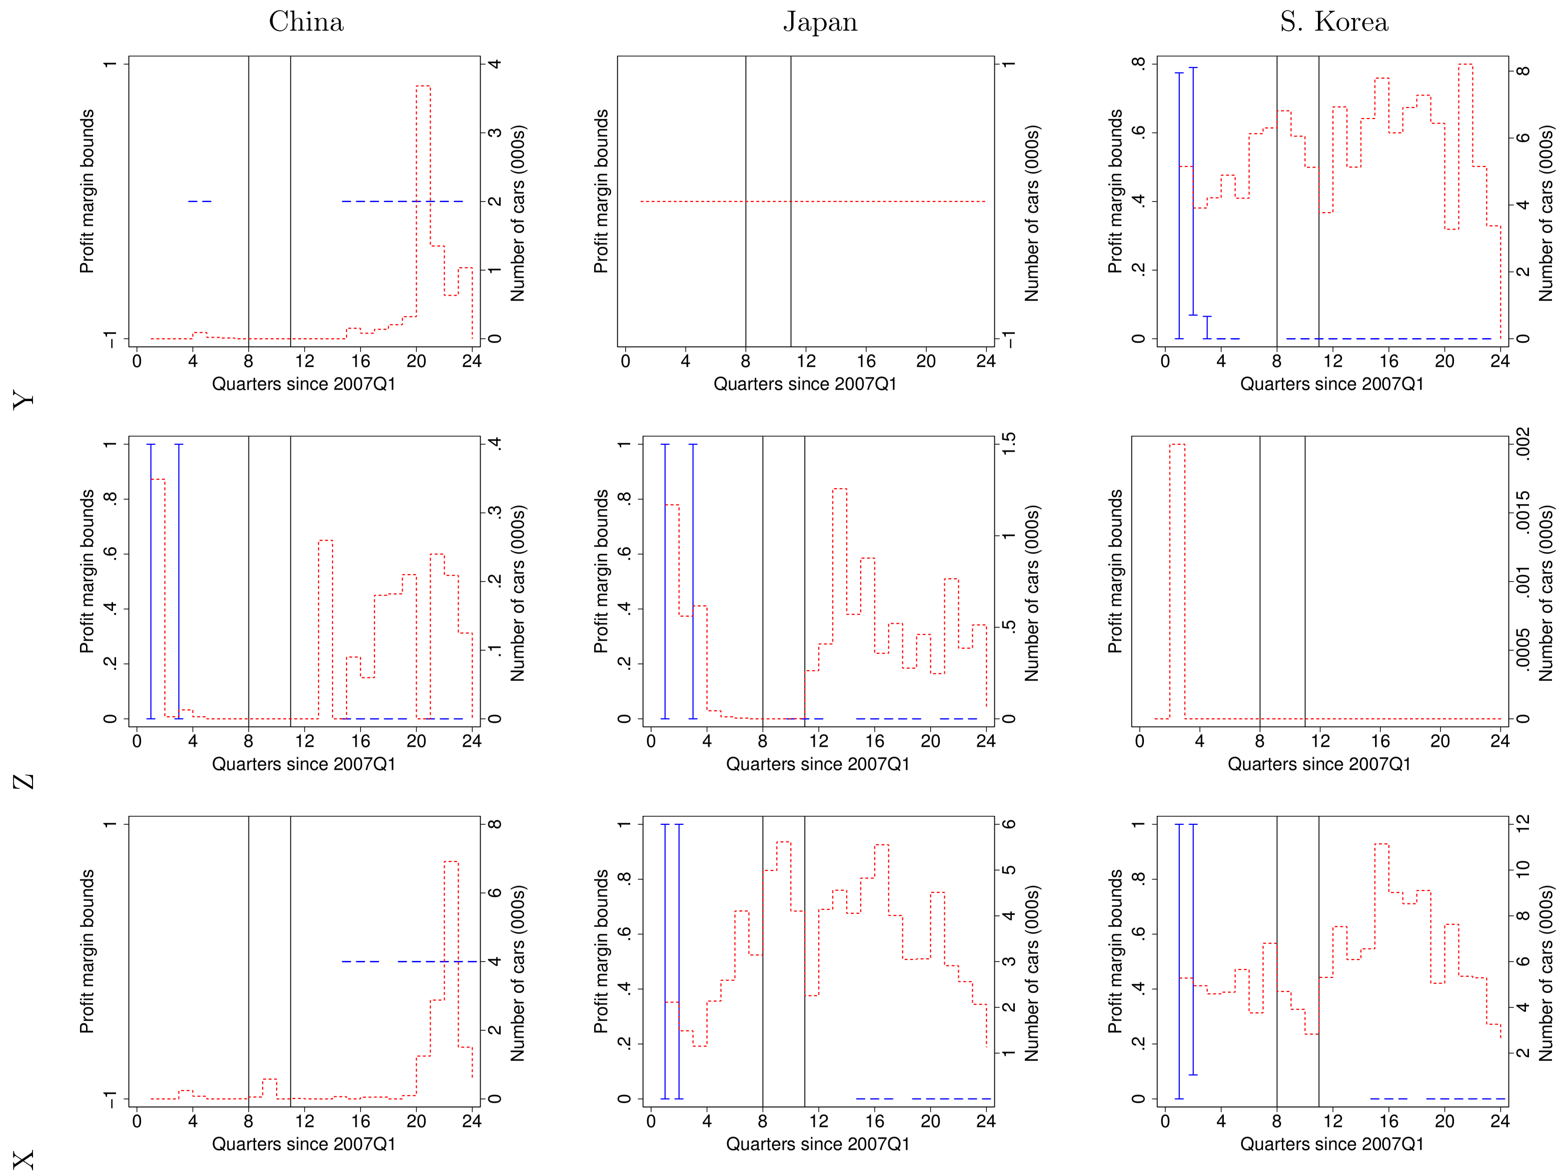

Notes: Red broken lines trace the number of automobiles moved by the row carrier from the column country. Blue bands represent bounds on the each carrier's profit margin, $\left(\mathtt{t}_{oj} - \mathrm{m.c.}_{oj}^{f}\right)/\mathtt{t}_{oj}$, where $\mathtt{t}_{oj}$ is the freight rate in market o in quarter j, and $\mathrm{m.c.}_{oj}^{f}$ is the carrier's corresponding marginal cost. The great recession persists between the broken lines.

However, it is clear that the level of carrier activity is positive in the period in question. Such low profit margins are unlikely to indicate, say, poor management relative to other carriers – after all, Y, Z, and X are the largest carriers in the market. I consider two alternative explanations for the negligible profit margins.

First, low profit margins may reflect a trough in the shipping cycle, aggravated by the Great Recession. Stopford (2009) describes typical peaks in the shipping sector as periods with freight rates as much as three times operating costs; Figure 2.1 shows an upward trend in freight rates before the Great Recession. Buoyed by the boom, shipping companies tend to order new vessels to keep up with demand for their services. Indeed, Diesenreiter and Tromborg (2009) and Samaras and Papadopoulou (2010) report unusually high demand for ocean transport, with Diesenreiter and Tromborg (2009) attributing the boom to low–cost manufacturing in China from 2005 to 2008. This may explain the large profit margins in the first few quarters in Figure 2.2.

Typically, the cycle wanes as shipbuilders deliver the vessels ordered during the boom, creating excess supply. Hoffmann (2009) and Kalgora and Christian (2016) report an unusually large number of orders between 2005 and 2007, claiming that there would have been a significant downturn in freight rates, even without Great Recession. Finally, Wong (2017) also documents persistent over–capacity in container shipping, with 30 percent more space on ships than cargo. According to BRS (2008), the RoRo sector was not immune to this enthusiasm. Finally, the trough of the shipping cycle features excess shipping capacity, unusually low freight rates, and even the sale of vessels at prices below their book value. This seems to describe market conditions during the Great Recession, with Kalgora and Christian (2016) reporting a collapse in demand and freight rates, and ports filled up with fleets of empty freighters. Indeed, Figure 2.1 features a sharp decline in freight rates during the Great Recession, followed by a steady recovery to pre–crash periods.

Setting aside cycles in the shipping–sector, the second set of explanations represents a significant departure from the Cournot mode of competition presented in Chapter 1. In particular, I consider the implications of Bertrand competition and contestable markets, oligopolistic markets that mimic perfectly competitive outcomes like marginal cost pricing under certain conditions. In both models, the erosion of market power relies on the assumption that shipping services are homogenous, that carriers have identical shipping technologies, and that carriers can capture (and then serve) the entire market by undercutting rivals. RoRo shipping is fairly standardized. It is well known that Bertrand competition under these conditions delivers marginal cost pricing.

The theory of contestable markets, described in (Baumol et al., 1982), further restricts shipping technologies, and imposes behavioural restrictions on carriers. In particular, it assumes free entry and exit along any route. Second, it assumes that potential entrants evaluate the profitability of entry at existing prices. Since all investments are instantly reversible, carriers are always ready to enter profitable markets. Since carriers are identical and have the same technologies, potential entrants opt to stay out of the market because it is unprofitable to undercut incumbents. Further, marginal cost pricing prevails in routes served by at least two carriers, as is the case in the three shipping markets I consider. Crucially, sunk costs must be completely absent. This model has been used to explain quasi-competitive market performance in other transport sectors.[6] Bailey (1981) and Bailey and Panzar (1981) argue that airline carriers can easily move aircraft across routes, thus recovering most capital costs. Closer to home, Davies (1986) suggests that the containership sector has all the ingredients of a contestable market.

2.4 Conclusion

This chapter tests for Cournot behaviour among carriers shipping automobiles to Ecuador. It fails to reject such equilibrium behaviour among small sets of carriers and over short time horizons. However, the test conclusively rejects Cournot equilibrium once we depart from either of these settings. When part of the data is Cournot rationalizable, it provides potential candidates for demand elasticities, and bounds on carrier costs. As expected, it is easier to rationalize shipping activity among smaller groups of carriers over shorter horizons as Cournot equilibrium outcomes. I then bound carrier marginal costs using the set of rationalizable observations, and find evidence of dwindling profit margins since the beginning of the Great Recession, consistent with reports on the state of the shipping sector during that time.

[6] See Tye (1985) on barges.

Chapter 3

Property rights and hold-up in international shipping

Introduction

The international trade literature has delved into the mechanisms behind the effects of transportation costs, policy and institutional barriers, and information costs on international trade, surveyed in Anderson and van Wincoop (2004). These trade costs have tangible effects on global value chains, the set of tasks involved in bringing a good or service from its conception to its end use (Global Value Chain Initiative, 2017). Modern supply chains (manufacturing and distribution processes) are more susceptible to such barriers, given the surge in offshoring and just-in-time production, which require several interrelated shipments to accomplish previously straightforward tasks. The viability of supply chain technology improvements therefore depends on the relative magnitudes of the cost savings from "unbundling" production across borders and the sum of trade costs incurred at each interface.

Trade policy intervention has unmistakably alleviated some of these costs. For example, many countries offer duty drawbacks, refunding import duties on intermediate productions upon the exportation of the resulting goods.[1] However, it is unrealistic to expect silver-bullet policies that address all possible impediments to supply chains. This chapter studies the distribution component of supply chain management, which has long taken a back seat in the minds of policy makers, who often focus on the organization of the manufacturing phase. However, distribution physically links one manufacturing phase in the supply chain to the next, and successful distribution relies on executing various costly and often unpredictable logistical and administrative tasks. It is incumbent upon buyers and sellers, interacting

[1] However, despite plaudits for historically low tariffs, Bown and Crowley (2016) conclude that non-tariff barriers like quantitative restrictions, antidumping regulations, temporary trade barriers, and "behind-the-border" policies (national subsidies and taxes, labour and environmental standards, and antitrust regulations) still present significant policy-based barriers.

across national borders, to coordinate these tasks when sending goods from the seller's location to their intended final destination.

I focus on trade costs arising from the *organization of distribution in global supply chains,* defined as the allocation of delivery-related tasks between buyers and sellers party to international transactions. This allocation matters whenever trading partners differ in their ability to execute the various tasks, and cannot directly compensate each other for their efforts towards smooth logistical operations. Such scenarios abound outside the present setting, with parties often resorting to indirect mechanisms to encourage valuable effort. For example, they may link payment to observable outcomes affiliated with the underlying productive effort. Applying this insight to distribution, buyers and sellers may condition payments on the state of the shipment at the destination. However, it may be difficult to verify shipment quality in all but the extreme cases when goods are damaged beyond repair, or worse, lost in transit. To compound the problem, sellers may attribute goods that arrive in poor condition to the unpredictability of long-distance shipping.

This chapter studies a potential workaround to such two-sided moral hazard problems in an extreme contracting environment where parties can only contract on the volume of shipment, the allocation of delivery tasks, and an ex-ante payment. Despite their limitations, such contracts encourage the desirable but otherwise unverifiable behaviour because responsibility for a given distribution-related task often confers valuable rights over the shipment for the duration of the task. The allocation of tasks and the associated power therefore offers buyers and sellers an alternative means to encourage productive efforts in the absence of quality–contingent contracts.

Does the organization of distribution have observable implications for trade flows? If so, what determines the allocation of tasks? I use Colombian transaction data to demonstrate that this allocation is a relevant margin of trade, explaining around 2 percent of the variation in firm-level trade, even after controlling for buyer, seller, and product characteristics. I then build an incomplete-contracting model of production and delivery, where contracts between exporters and importers specify the shipment volume, and designate one of the parties as consignor. The party in charge of delivery then signs a freight contract that affords them the right to modify delivery following unforeseen events, ultimately determining the shipment's fate. Once this auxiliary contract is in place, the exporter incorporates an unverifiable level of quality into the agreed-upon volume of goods. The buyer and seller then bargain over the value added during manufacturing, proceeding to the delivery phase only if they come to a mutually beneficial agreement.

I assume that the exporter possesses all the bargaining power in the first phase of production, thus sidestepping the decision to integrate production and sales into a single firm (see, for example, Antràs, 2003). Instead, I focus on the optimal allocation of control over delivery-related activities. This assignment directly affects buyer and seller behaviour during distribution, and, as we will see, also affects the forward-looking seller's manufacturing decisions. After manufacturing, delivery requires some "maintenance" activity by at least one of the parties. Again, such efforts are unverifiable and thus prone to hold-up. The parties therefore bargain over the value added by their joint efforts. Armed with the rights to dictate the ultimate fate of the goods, the party in charge of distribution may threaten to take possession of the shipment and put it to some alternative use. Foreseeing these control-dependent bargaining externalities, the exporter and importer allocate these scarce rights to minimize overall distortions from their first-best levels.

This chapter contributes to the literature at the intersection of International Trade and Organizational Economics. While existing work covers areas as diverse as ownership/integration and sourcing (Antràs (2003), Antràs and Helpman (2004); Grossman and Helpman (2002); McLaren (2000); and Schwarz and Suedekum (2014) for theoretical contributions; Feenstra and Hanson (2005) for empirical tests) and the internal organization of firms (Marin and Verdier, 2003), I am the first to study the organization of *distribution* in international trade logistics.

I also contribute to the emerging literature on contractual frictions in sequential production processes. Fally and Hillberry (Forthcoming) consider the tradeoffs between arms-length transaction costs and in-house coordination costs in determining the complexity of global supply chains. I present a model similar in its property-rights foundations to work by Antràs and Chor (2013) and Alfaro, Antràs, Chor, and Conconi (Forthcoming). These papers focus on independent agents acting at each stage, and describe the effects of investment in upstream stages on subsequent investment decisions. In contrast, and motivated by the observation that international distribution requires joint efforts, I allow multiple parties to undertake productive actions at a given stage. This departure introduces strategic interactions *within* a given stage.

In a broader sense, my work is related to research on the determinants of vertical integration, surveyed in Lafontaine and Slade (2007) and Klein (2008). The Industrial Organization literature traditionally stresses economies of scale and scope, foreclosure, and double marginalization as the main reasons to integrate activities. Such motives are bound to play some role here. For example, as Malfliet (2011) notes, it is reasonable to assign greater responsibility to the larger or more experienced of the two trading parties, with the hope of leveraging its buyer power to earn quantity discounts from the carrier. While such direct costs play a role, this interpretation downplays the indirect costs of allocating control among parties, as highlighted by the incomplete contracting literature.

3.1 Institutional background and motivation

Although there are several ways to allocate roles to each party, most international transactions fall under one of the International Chamber of Commerce's *International Commerce Terms (INCOTERMS)*. Widely adopted in international transactions, these *delivery terms* reduce shipping-related confusion by outlining each party's rights and obligations during the delivery process. Conveniently, any two terms can be ranked in terms of the exporter's responsibility. At one extreme, the exporter is a passive observer, and assumes an additional role under each subsequent step. These are (in order): arranging for carriage to the port ("pre-carriage"), customs clearance at the origin, loading the shipment onto the vessel, international freight and insurance, unloading at the destination port, customs clearance at the destination, and carriage to the importer's premises ("on-carriage"). There terms are usually classified into four groups, *E*, *F*, *C*, and *D*, ranked in increasing order of exporter burden. To avoid confusion, I recode the terms so that *E, F, C,* and *D* correspond to the 1st, 2nd, 3rd, and 4th broad groups.

The first group consists of the *Exworks (EXW)* term, where the exporter simply packs the goods and makes them available at his factory's gate. In some cases, he may help with loading (at the importer's risk), and/or obtaining customs clearance. Given the bureaucratic hurdles sometimes associated with customs clearance, the importer may be at the exporter's mercy despite trading under this term. The importer can certainly lower this dependence by hiring agents or obtaining some other presence at the origin, but often finds it easier to rely on the exporter's effort in ensuring compliance with local regulations.

The second group consists of two terms, *Free Carrier (FCA)*, and *Free on Board (FOB)*. Under FCA, the exporter loads the goods at his premises, arranges inland freight to the port, and clears the goods through customs.[2] Under FOB, the exporter assumes the additional role of loading the goods on board the vessel. The importer arranges the remaining portions of the trip (international carriage and insurance, and customs clearance and on-carriage). It is crucial that the two parties coordinate the handover if they are to avoid any costs associated with delays, such as extra storage until any errors are corrected. Such coordination requires some effort, especially in countries with poor external support from dedicated freight-forwarders or "door-to-door" services.

The third group also consists of two terms: *Carriage Paid To (some port of destination) (CPT)*, and *Cost, Insurance and Freight (CIF)*. Under these terms, the exporter assumes the extra responsibility of arranging international carriage to the destination port. The importer obtains import clearance and arranges inland freight at the destination. The parties' efforts in finding reliable and affordable carriers can make all the difference when deciding between the second and third delivery term groups. First, as Malbon and Bishop (2014) explain, shippers have little bargaining power when negotiating carriage contracts with carriers, unless they ship exceptionally large quantities and have built a history with a particular carrier. Second, entrusting shipments to reliable carriers goes some way to preventing future hassle, given the uncertain nature of international freight.

Finally, the fourth class consists of *Delivered at Terminal (DAT)*, and *Delivered Duty Paid (DDP)*. Relative to the third group of terms, DAT extends the exporter's responsibility

[2] In some cases, the exporter delivers the shipment at some point before the port.

to unloading the goods at the destination port, leaving import clearance and inland freight to the importer. The DDP term places the greatest burden on the exporter, requiring that he also clear the goods through customs at the destination, effectively rendering the importer a spectator in the delivery process.

By delineating the various delivery-related tasks, INCOTERMS indicate that at least one of the parties must exert some costly effort to ensure successful delivery. With this background in hand, the remainder of this section uses the universe of Colombian firm-level transaction data from 2009 to 2013 to establish that variation in delivery terms constitutes a relevant margin of trade. I observe the date that each shipment was cleared through Colombian customs, the associated delivery term, the contents of the shipment (quantities and FOB values of each 10-digit HS product code), a unique tax identifier tied to the Colombian exporter, and, in some cases, the name of the foreign importer. For most of the analysis, I aggregate trade flows to the exporter-product-year level. I also include importer identities for the subset of transactions destined for Spain.

Table 3.1 shows the popularity of the various arrangements among Colombian exporters using exports at the transaction level. The last row shows that a nontrivial fraction of shipments involve customized delivery terms. According to Ramberg (2011), buyers and sellers sometimes make minor modifications to the standard delivery terms either because of standard practice in the industry, or to accommodate one party's exceptional needs. Among the standard delivery terms, the second group, where the exporter's responsibility ends at the origin port, is by far the most popular, accounting for no less than three quarters of annual export values.

Table 3.1: Delivery term popularity

Term group: Final exporter task	2009	2010	2011	2012	2013
1: None	0.8	0.5	0.4	0.4	0.5
2: Origin port	74.2	78.4	82.2	82.1	80.8
3: Destination port	11.6	12.6	7.2	7.7	7.3
4: Inland at destination	4.3	3.8	5.6	6.0	6.5
N/A	9.1	4.8	4.5	3.8	5.0

Notes: Aggregate shares of annual export values under each of the delivery term groups. The first column indicates the group of terms (1E, 2F, 3C, 4D) and the final exporter task (also the additional exporter task relative to the preceding group). For example, the exporter assumes responsibility for getting the goods to the destination port in moving from group 2 to group 3. Transactions in the "N/A" row involve custom arrangements and do not fall under any of the four traditional commerce terms.

Table 3.1 masks variation in delivery-term choice across exporters. Consider a Colombian firm, x, exporting $\texttt{exports}_{x,\mathcal{O}}$ worth of goods to the rest of the world under a delivery term in group $\mathcal{O} = 1, 2, 3, 4, \text{N/A}$ in 2013 (this is representative of other years). Let

$$\texttt{share}_{x,\mathcal{O}} \equiv \texttt{exports}_{x,\mathcal{O}} \Big/ \sum_{\mathcal{O}'=1,2,3,4,\text{N/A}} \texttt{exports}_{x,\mathcal{O}'}, \tag{3.1}$$

denote the firm-level share of 2013 export values under term $\mathcal{O}$.

Figure 3.1: Distribution of delivery-term popularity, $\texttt{share}_{x,\mathcal{O}}$, at the exporter-level

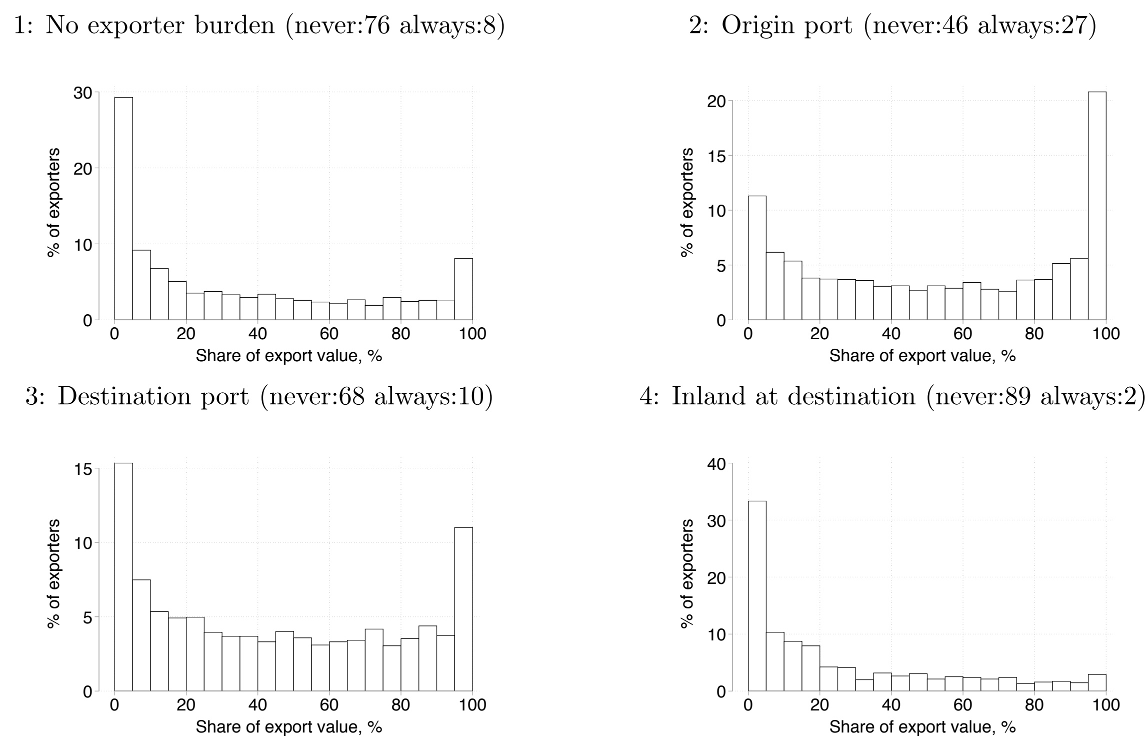

Notes: Panel headers indicate the standard delivery term labels, the corresponding final exporter responsibility, the percentage of exporters that never used a term in a given group, and the percentage of exporters that exclusively used a given term in 2013. For example, 76% of exporters performed *some* delivery-related task in all their transactions in 2013, while 8% did not help with delivery in any of their transactions. The histogram in Panel $\mathcal{O}$ is the distribution of $\texttt{share}_{x,\mathcal{O}} \equiv \texttt{exports}_{x,\mathcal{O}} / \sum_{\mathcal{O}'=E,F,C,D,N/A} \texttt{exports}_{x,\mathcal{O}'}$ among exporters with $0 < \texttt{share}_{x,\mathcal{O}} < 1$ for the $\mathcal{O}$ in question. When computing the shares, the base includes exports under unknown delivery terms (the "N/A" column in Table 3.1); results are similar if I exclude unclassified transactions.

Panel $\mathcal{O} = 1, 2, 3, 4$ of Figure 3.1 presents the intensive-margin distribution of $\texttt{share}_{x,\mathcal{O}}$ for exporters that use at least one other term in 2013. The panel headers show the share of exporters that entirely avoid a given term, and those that trade exclusively under the term. At one extreme of exporter burden, 76 percent of exporters were involved in some aspect of distribution in each of their transactions, while 8 percent did not take part in any such tasks at any point in 2013. At the other extreme, only 2 percent of exporters undertook all delivery-related tasks in all their transactions, while 89 percent never ventured beyond the destination port at any point during the year.

Figure 3.2: (Exporter-level) total exports, $\texttt{exports}_x$

a. Average annual exports by term popularity

b. $\texttt{exports}_x$ conditional on number of delivery terms

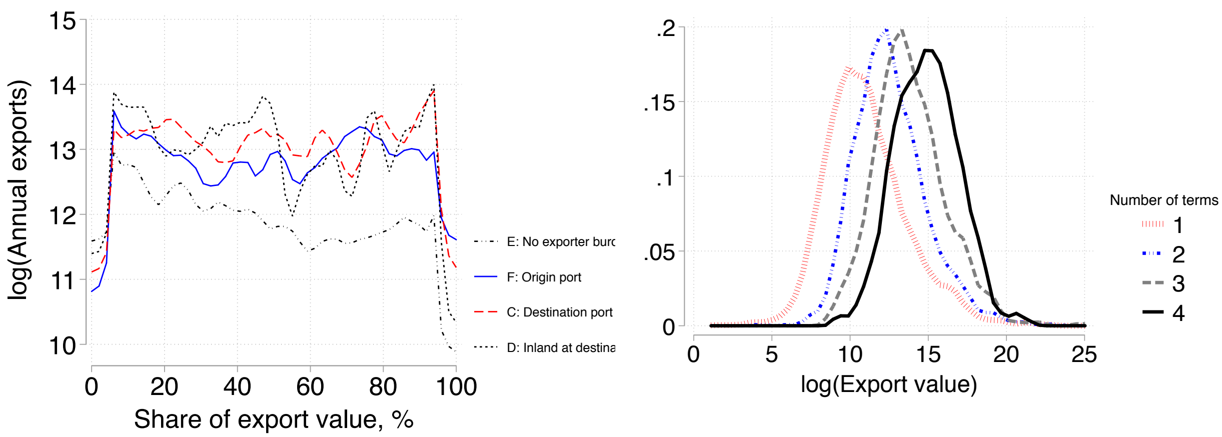

Notes: Panel (a) plots a local polynomial approximation of the conditional mean of total exports conditional on each term's share of exports, $\mathbb{E}\left[\log\left(\texttt{exports}_x\right) | \texttt{share}_{x\mathcal{O}}\right]$. Each exporter appears in each of the four regressions. Unlike Figure 3.1, the conditional means includes all exporters, not just those with intermediate shares ($0 < \texttt{share}_{x,\mathcal{O}} < 1$). Letting $\texttt{numterms}_x$ denote the number of terms used by x in 2013 (i.e., those $\mathcal{O}$ with $\texttt{exports}_{x\mathcal{O}} > 0$), Panel (b) plots the density $f\left(\log\left(\texttt{exports}_x\right) | \texttt{numterms}_x\right)$ of (log) exports conditional on the number of terms. Each exporter appears in exactly one distribution. Figures are representative of other years in the 2009-2013 window, and similar patterns emerges if I include the unclassified term as a fifth option.

Turning to the intensive margin within each panel, the share distributions for all terms are bimodal, with most exporters employing a given term either very rarely, or very often. However, Figure 3.2 shows that such firms account for a small fraction of annual exports. The vertical axis in Panel (a) measures the (log of) annual firm exports, while the horizontal axis measures the fraction of the value these flows that were traded under a particular delivery term. The four curves trace sample means of firm-level annual exports, conditional on the fraction of annual exports under a given delivery term. Regardless of delivery term, exporters with less diverse delivery term portfolios (those at either end of the horizontal axis) have below-average annual exports. Setting aside the particular deliver term, Panel (b) shows the distribution of annual exports, conditional on the *number* of delivery terms used in a given year. Distributions associated with more "diverse" exporters dominate those using fewer delivery terms.

These figures offer a cursory glance at the data, and the remaining part of the introduction offers a more formal analysis, decomposing the variation in (i) annual trade, and (ii) the popularity of predominantly exporter-controlled trade into various effects. To ease comparison with work predating Melitz (2003), I begin by decomposing the variation in annual trade flows into product, destination, and delivery term effects. I then introduce exporter effects, before using matched exporter-importer data on sales from Colombia to Spain to explore the explanatory power of importer effects. Finally, I repeat the analysis, this time decomposing the variation in the share of trade under the two terms with the greatest exporter burden.

Table 3.2 decomposes the variation in aggregate exports into the variance attributable to 10-digit HS product categories, destinations, and delivery terms. In particular, consider an arbitrary transaction characteristic, g, which may denote a single attribute like exporter identity, or a composite like an exporter-product pair. Given g, consider the following models for the value of exports of product p to destination d under delivery term $\mathcal{O}$:

$$\begin{array}{llcl} (1)\ \text{Raw} & : \texttt{exports}_{\mathcal{O}pd} = \alpha_g^1 & + u_{\mathcal{O}pd}^1 & \left(R^2_{\text{only}\,g}\right) \\ (2)\ \text{Exclude}_g & : \texttt{exports}_{\mathcal{O}pd} = \alpha_{\mathcal{O}pd-g}^2 & + u_{\mathcal{O}pd}^2 & \left(R^2_{\text{except}\,g}\right) \\ (3)\ \text{Joint} & : \texttt{exports}_{\mathcal{O}pd} = \alpha_g^3 + \alpha_{\mathcal{O}pd-g}^4 & + u_{\mathcal{O}pd}^3 & \left(R^2_{\text{full}\,g}\right). \end{array} \quad (3.2)$$

The first model projects annual exports on a set of g fixed effects, the second explains this variation using fixed effects for the remaining observation characteristics, while the last model includes both pairs of fixed effects. The *semi-partial R-squared for characteristic g* is the difference $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$ in the explained variation between the model that includes both sets of fixed effects, and that including the remaining characteristics. This statistic offers a rough measure of the explanatory power of transaction characteristics included in g.

Table 3.2: Explaining variation in aggregate exports

g	Raw effect	$R^2_{\text{full}\,g}$	$R^2_{\text{except}\,g}$	Isolated effect
Product	0.37	0.46	0.11	0.34
Destination	0.05	0.57	0.50	0.07
Delivery terms	*0.03*	*0.62*	*0.60*	*0.02*

Notes: The raw effect is $R^2_{\text{only}\,g}$ in (3.2), the R-squared from the regression of $\texttt{exports}_{\mathcal{O}pd}$ on a set of g-fixed effects. The semi-partial R-squared, $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$, is the difference between the R^2's of regressions of $\texttt{exports}_{\mathcal{O}pd}$ (i) on g fixed effects and fixed effects for the remaining characteristics; and (ii) on just the remaining characteristics. See equations (3.2).

Of the three individual effects, product classifications have the greatest explanatory power, with delivery terms accounting for a small fraction of the variation in annual exports. Table 3.3, which adds exporter effects, confirms that some of the variation initially attributed to delivery terms in Table 3.2 is actually due to exporter-level variation.

Table 3.3: Explaining variation in exporter-level exports

g	Raw effect	$R^2_{\text{full}\,g}$	$R^2_{\text{except}\,g}$	Isolated effect
Exporter	0.40	0.70	0.55	0.15
Product	0.36	0.68	0.56	0.12
Destination	0.04	0.74	0.71	0.03
Delivery terms	*0.02*	*0.77*	*0.76*	*0.01*

Notes: $R^2_{\text{only}\,g}$ is the R-squared from the regression of $\texttt{exports}_{\mathcal{O}xpd}$ on a set of g-fixed effects. The semi-partial R-squared, $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$, is the difference between the R^2's of regressions of $\texttt{exports}_{\mathcal{O}pd}$ (i) on g fixed effects and fixed effects for the remaining characteristics; and (ii) on just the remaining characteristics.

This implies that delivery terms constitute a small but significant margin of trade at the aggregate, and exporter levels. However, it is entirely plausible that importer-level heterogeneity also explains the volume of trade. Indeed, (Bernard, Moxnes, and Ulltveit-Moe, 2018) document exactly such a phenomenon. With this in mind, Table 3.4 summarizes the role of differences across importers. I interpret these results with caution, since we lose the destination dimension by focusing on Colombian exports to a single destination, Spain. Nonetheless, importer-level differences explain some of the variation in trade, confirming results in Bernard et al. (2018). More importantly for our purposes, delivery terms retain their explanatory power.

Table 3.4: Explaining variation in exporter-importer trade

g	Raw effect	$R^2_{\text{full}\,g}$	$R^2_{\text{except}\,g}$	Isolated effect
Exporter	0.61	0.76	0.60	0.16
Importer	0.56	0.78	0.73	0.05
Product	0.61	0.80	0.58	0.22
Term	*0.09*	*0.89*	*0.87*	*0.02*
Exporter-importer	0.61	0.82	0.62	0.19
Exporter-product	0.72	0.78	0.56	0.22
Exporter-term	*0.62*	*0.78*	*0.63*	*0.15*

Notes: $R^2_{\text{only}\,g}$ is the R-squared from the regression of $\texttt{exports}_{\mathcal{O}xmp}$ on a set of g-fixed effects. The semi-partial R-squared, $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$, is the difference between the R^2's of regressions of $\texttt{exports}_{\mathcal{O}xmp}$ on (i) on g fixed effects and fixed effects for the remaining characteristics; and (ii) on just the remaining characteristics.

Margins of delivery-term choice

In this section, I decompose the variation in $\texttt{share}_{i,\mathcal{O}}$, where i is the level of observation. Given the paucity of trade flows under terms at either extreme of the exporter-burden spectrum, I group terms so that

$$\texttt{share\_expcontrol}_i \equiv \sum_{\mathcal{O}=3,4} \texttt{share}_{i,\mathcal{O}} \tag{3.3}$$

is the share of predominantly exporter-controlled i-transactions (where the exporter controls port-to-port distribution). I begin by decomposing variation in $\texttt{share\_expcontrol}_i$ into exporter, product, destination, and year effects, before considering the importer dimension (again, at the expense of the destination effects).

Table 3.5: Explaining variation in share of exporter-controlled transactions

g	Raw effect	$R^2_{\text{full}\,g}$	$R^2_{\text{except}\,g}$	Isolated effect
Exporter	0.48	0.65	0.38	0.27
Product	0.17	0.88	0.87	0.01
Destination	0.07	0.63	0.62	0.01
Time	0.00	0.83	0.83	0.00
Exporter-product	0.58	0.61	0.09	0.52
Exporter-destination	0.73	0.76	0.21	0.55
Exporter-year	0.57	0.69	0.36	0.34

Notes: $R^2_{\text{only}\,g}$ is the R-squared from the regression of $\texttt{share\_expcontrol}_{xpdt}$ on a set of g-fixed effects. The semi-partial R-squared, $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$, is the difference between the R^2's of regressions of $\texttt{share\_expcontrol}_{xpdt}$ (i) on g fixed effects and fixed effects for the remaining characteristics; and (ii) on just the remaining characteristics.

Table 3.5 shows the resulting R-squared statistics using unmatched exporter-importer data, where i is an exporter-destination-product-year. Restricting attention to individual effects, exporters-level heterogeneity best explains the variation in the share of exporter-controlled trade. This suggests that any model explaining the choice of delivery terms should, at the very least, allow for differences across exporters along some dimension. Turning to joint effects, we see that augmenting either a product or destination dimension substantially improves the model fit. Finally, allowing exporter-product-destination level heterogeneity results in a semi-partial R-squared statistic of 0.83.[3]

Table 3.6 shows the analogous results from matched Spanish data, where i is an exporter-importer-product-year. As with the unmatched dataset, differences across exporter are the best single predictors of variation in the share of predominantly exporter-controlled shipments. However, the matched data demonstrates that importer-level heterogeneity also accounts for a substantial fraction of the variation in $\texttt{share\_expcontrol}$. Turning to joint effects, differences across exporter-importer pairs accounts for more variation than either set of individual effects, suggesting that interactions between exporter and importer charac-

[3] Since we observe trade at the exporter-destination-product-year level, the results from the individual time effects imply $R^2_{\text{full}\,xpd} = R^2_{\text{full}\,t} = 0.83$, and $R^2_{\text{except}\,xpd} = R^2_{\text{only}\,t} = 0.00$.

Table 3.6: Explaining variation in share of exporter-controlled transactions (exporter-importer trade)

g	Raw effect	$R^2_{\text{full}\,g}$	$R^2_{\text{except}\,g}$	Isolated effect
Exporter	0.69	0.89	0.66	0.23
Importer	0.73	0.94	0.79	0.15
Product	0.37	0.97	0.96	0.01
Time	0.00	0.87	0.87	0.00
Exporter-importer	0.86	0.91	0.43	0.47
Exporter-product	0.74	0.94	0.83	0.11
Exporter-year	0.82	0.95	0.77	0.18

Notes: $R^2_{\text{only}\,g}$ is the R-squared from the regression of $\texttt{share\_expcontrol}_{xmpt}$ on a set of g-fixed effects. The semi-partial R-squared, $R^2_{\text{full}\,g} - R^2_{\text{except}\,g}$, is the difference between the R^2's of regressions of $\texttt{share\_expcontrol}_{xmpt}$ on (i) on g fixed effects and fixed effects for the remaining characteristics; and (ii) on just the remaining characteristics.

teristics are important in explaining the choice of delivery term. Lastly, allowing exporter-importer-product level heterogeneity results in a semi-partial R-squared statistic of 0.87.[4]

To summarize, the delivery-term margin accounts for a significant share of the variation in both aggregate and firm-level trade, even in the presence of previously studied margins. Further, the popularity of various delivery terms depends, at the very least, on buyer, seller, and product characteristics. These results motivate the upcoming model, which studies buyer-seller pairs that self-select into delivery terms based on their distribution capabilities and the nature of the product being traded.

3.2 Model

This section describes consumer demand for final goods, the supply chain technologies, and the buyer-seller contracting problem.

Demand for final goods

Each market consists of L consumers, who spend their income across various industries, with each industry consisting of a variety of differentiated products. Following Antoniades (2015), the representative consumer derives sub-utility

$$\begin{aligned} U_d \;=\;& q_0^c + \int \alpha \left(q^c(\omega) + z(\omega)\right) - \frac{\gamma}{2}\left(q^c(\omega)^2 + z(\omega)^2\right) + \gamma z(\omega)\, q^c(\omega)\, \mathrm{d}\omega \\ & - \frac{\chi}{2}\left(\int q^c(\omega) - \frac{1}{2} z(\omega)\, \mathrm{d}\omega\right)^2 \end{aligned} \tag{3.4}$$

[4]Since observations are at the $xmpt$ level, the results from the time effects imply $R^2_{\text{full}\,xmp} = R^2_{\text{full}\,t} = 0.87$, and $R^2_{\text{except}\,xmp} = R^2_{\text{only}\,t} = 0.00$.

from consuming q_0^c units of a numéraire good and $q^c(\omega)$ units of quality $z(\omega)$ of variety ω in a given sector. The parameters $\alpha, \chi > 0$ reflect preferences for the differentiated varieties relative to the numéraire, while $\gamma > 0$ measures love-of-variety within a sector. Let y_d^c denote individual consumer income from inelastically supplying a unit of labour. In addition to these labour returns, workers have an exogenous endowment, $\overline{q}_0^c > 0$, of the numéraire. I assume that this endowment is large enough to guarantee positive demand for the numéraire, thereby eliminating any income effects in demand for the differentiated good. Consumers maximize utility U_d subject to the budget constraint

$$q_0^c + \int p(\omega) q^c(\omega) \, \mathrm{d}\omega \leq y_d^c + \overline{q}_0^c.$$

Conditional on quality, these preferences deliver linear inverse-demand and quadratic revenue functions,

$$p(q,z) = A + \gamma z - \frac{\gamma}{L} q, \qquad r(q,z) = \left(A + \gamma z - \frac{\gamma}{L} q\right) q, \tag{3.5}$$

where $A \equiv (\gamma + \chi N)^{-1} (\alpha\chi + \chi N \overline{p} - \gamma\chi N \overline{z}/2) > 0$, which depends on the destination-wide average price, $\overline{p}$, and quality, $\overline{z}$, is an exogenous demand shifter from the perspective of the seller of any given variety. Note that γ, which measures love-of-variety, also determines the marginal effect of quality on sales revenue, $\partial r(q,z)/\partial z = \gamma q$. This observation will drive many of the subsequent results.

Supply-chain technology

Having met exogenously, a potential exporter and importer, indexed by X and M, work together to serve the L consumers described above. The importer has direct access to the final-goods market, while the exporter owns a manufacturing plant with independent physical-unit-production and quality-creation techniques, in the sense that the marginal product of any given input into quality creation is independent of the scale of production, and vice versa.

The exporter's factory capabilities are summarized by the pair (c, ψ_0), where $c > 0$ is the marginal cost of producing physical units, and $\psi_0 > 0$ shifts the marginal cost of quality innovation. Specifically, the total cost of producing q units with initial/factory-set quality z is

$$C(q,z) = cwq + \frac{1}{2}\psi_0 z^2, \tag{3.6}$$

where w is the prevailing wage in the source country.

I index shipments by their volume, quality, and location, so that the pair (q_i, z_i) represents q_i units of quality z_i at location $i \in \{0,1\}$, where $i = 0$ corresponds to the manufacturing plant in the source country, and $i = 1$ is the destination market. Although quality is a vertical characteristic according to consumer preferences in (3.4), I assume that it is suitably tailored to some subset of the population linked to the initial importer M. The exporter is therefore subject to hold-up if he produces a bundle with any given importer in mind. Once the exporter produces (q_0, z_0), one of the parties takes possession of the bundle and oversees distribution to the destination.

In addition to transporting the goods across space, delivery may alter their physical characteristics. Throughout, I will assume that the volume of the shipment is fixed at its factory level q_0 (let q denote this fixed level), while its quality may change during transit. In particular, $j \in \{X, M\}$ may exert e_j units of unverifiable effort at a cost $\psi_j e_j^2/2$ to improve shipment quality. Individual efforts then combine via the aggregator

$$E(e_X, e_M) = \left(\eta\, e_X^{\rho} + (1-\eta)\, e_M^{\rho}\right)^{\frac{1}{\rho}}, \qquad \rho \in (0,1), \tag{3.7}$$

where $\eta \in (0,1)$ measures the relative importance of exporter effort in quality-maintenance, while ρ is related to the substitutability of individual efforts. For example, large values of η may indicate origin-specific regulations that explicitly require exporter participation. Given the partial specificity of quality to M's intended consumers, the marginal product of such maintenance effort depends on the pair's relationship surviving past the distribution phase. In particular, quality at the destination is proportional to a Cobb-Douglas composite of factory-set quality and the aggregate maintenance effort, and is equal to

$$z_1(E|\, z_0) = z_0^{1-\beta} E^{\beta}, \qquad \beta \in (0,1), \tag{3.8}$$

if the relationship survives transit-stage bargaining, and $\delta\, z_1(E|\, z_0)$ if the relationship breaks down, where $\delta \in (0,1)$ measures the "salvage value" of $(X, M)-$specific quality. The parameter β measures the importance of quality maintenance efforts relative to the initial quality z_0, in determining final quality.

3.2.1 First-best contracts

If initial quality and maintenance efforts are verifiable to third parties, the importer proposes a contract (q, z_0, e_X, e_M, s) that specifies the desired physical output, initial quality, each party's maintenance efforts, and a payment $s \in \mathbb{R}$ to the exporter

$$\begin{aligned} \max_{q, z_0, e_X, e_M, s} \; & r\left(q, z_1\left(E(e_X, e_M)|\, z_0\right)\right) - \frac{1}{2}\psi_M e_M^2 - s \\ \text{s.t. } \; & s - \left(cwq + \frac{1}{2}\psi_0 z_0^2 + \frac{1}{2}\psi_X e_X^2\right) \geq 0. \end{aligned} \tag{3.9}$$

The importer chooses the transfer s that just secures exporter participation, which implies that the importer maximizes sales revenues net of the joint (across parties) production and distribution costs.

Conditional on the first-best shipping volume and initial quality, $(q_{FB}, z_{0,FB})$, j's optimal maintenance effort is

$$e_{j,FB}\left(q_{FB},z_{0,FB}\right)=\left(\gamma q_{FB}\,\beta z_{0,FB}^{1-\beta}\,\Phi_{FB}^{\beta-\rho}\right)^{\frac{1}{2-\beta}}\;\phi_{j,FB},\tag{3.10}$$

where

$$\Phi_{FB}\equiv\left(\eta\,\phi_{X,FB}^{\rho}+\left(1-\eta\right)\,\phi_{M,FB}^{\rho}\right)^{\frac{1}{\rho}},\qquad\qquad\phi_{j,FB}\equiv\left(\frac{\eta_j}{\psi_j}\right)^{\frac{1}{2-\rho}}.\tag{3.11}$$

The term Φ_{FB} is a share-weighted index of individual distribution capabilities, $\phi_{j,FB}$, itself a share-weighted measure of j's marginal efficiency. Large values of $\phi_{j,FB}$ indicate that j's effort is particularly important, and/or cheaper on the margin. Returning to (3.10), first-best individual efforts are increasing in shipment volume, q_{FB}, consumer love-of-variety, γ, and, if $\beta>\rho$, in the exporter-importer pair's joint capabilities, Φ_{FB}. The first two effects follow from the fact, alluded to when discussing consumer preferences, that γq measures the marginal returns to quality (in terms of higher sales revenue).

Given the CES effort aggregator, j's effort, relative to $-j$'s, and to the aggregate, are

$$\frac{e_{j,FB}}{e_{-j,FB}}=\frac{\phi_{j,FB}}{\phi_{-j,FB}}=\left(\frac{\eta_j}{1-\eta_j}\frac{\psi_{-j}}{\psi_j}\right)^{\frac{1}{2-\rho}},\qquad\qquad\frac{e_{j,FB}}{E_{FB}}=\frac{\phi_{j,FB}}{\Phi_{FB}}.\tag{3.12}$$

All else equal, j contributes relatively more if their effort is more important ($\eta_j>1/2$), or they are more efficient at the margin ($\psi_j/\psi_{_{-j}}<1$).

While individual efforts are of interest in their own right, we are ultimately interested in aggregate effort, which combines with factory-set quality, z_0, to determine final quality. Substituting (3.10) into (3.7), the first-best aggregate effort,

$$E_{FB}\left(q_{FB},z_{0,FB}\right)=\left(\gamma q_{FB}\,\beta z_{0,FB}^{1-\beta}\,\Phi_{FB}^{2-\rho}\right)^{\frac{1}{2-\beta}},\tag{3.13}$$

is increasing in the shipping volume q_{FB}, initial factory quality, $z_{0,FB}$, and aggregate productivity, Φ_{FB}. This follows from two simple observations. First, the marginal return to aggregate effort ultimately derives from the resulting increase in sales revenue due to higher quality goods at the destination, γq. Second, aggregate effort and initial quality are complements in the production of final quality, so that $z_{0,FB}$ increases the marginal returns to aggregate effort.

Having established the optimal maintenance efforts conditional on $(q_{FB},z_{0,FB})$, we now turn to initial quality. Simplifying,

$$z_{0,FB}\left(q\right)=\Theta_{z0,FB}\,\gamma q,\qquad\qquad\Theta_{z0,FB}\equiv\beta^{\frac{\beta}{2}}\left(1-\beta\right)^{\frac{2-\beta}{2}}\left(\frac{1}{\psi_0}\right)^{\frac{2-\beta}{2}}\Phi_{FB}^{\frac{\beta(2-\rho)}{2}}>0.\tag{3.14}$$

The term $\Theta_{z0,FB}$ summarizes the role of the supply chain technology. Holding the shipment volume constant, higher quality goods leave the factory whenever quality creation is particularly cheap (low ψ_0), or when the parties are adept at distribution (high Φ_{FB}). Conditional on distribution capabilities, the exporter creates higher quality goods when producing large volumes. Again, this follows from the complementarity between final quality and volume in revenue generation.

Combining the first-best factory quality (3.14) and aggregate effort during transit,

$$E_{FB}(q) = \Theta_{E,FB}\,\gamma q, \qquad \Theta_{E,FB} \equiv \left(\beta\,\Theta_{z0,FB}^{1-\beta}\,\Phi_{FB}^{2-\rho}\right)^{\frac{1}{2-\beta}}, \tag{3.15}$$

according to the Cobb-Douglas technology (3.8), yields quality-at-destination

$$z_{1,FB}(q) = \Theta_{z1,FB}\,\gamma q, \qquad \Theta_{z1,FB} \equiv \Theta_{z0,FB}^{1-\beta}\Theta_{E,FB}^{\beta} = \beta^{\beta}(1-\beta)^{1-\beta}\left(\frac{1}{\psi_0}\right)^{1-\beta}\Phi_{FB}^{\beta(2-\rho)} > 0, \tag{3.16}$$

which, like initial quality and aggregate maintenance effort, is linear in the shipment volume, and increasing in the (first-best) joint distribution capabilities Φ_{FB}.

Finally, the first-best shipping volume is the unique q that equates the marginal revenue and marginal cost of output. From (3.5), an increase in shipment volumes changes revenues by $A-(2\gamma/L)\,q$, and, operating through the final quality given in (3.16), further raises revenue by $(z_{1,FB} + q\,\partial z_{1,FB}/\partial q)\,\gamma$.[5] Similarly, the marginal cost of output consists of the marginal cost of producing the physical units, cw, and the induced marginal costs of factory-quality, $(\partial/\partial q)\left\{\psi_0\, z_{0,FB}^2/2\right\}$, and distribution efforts, $(\partial/\partial q)\left\{\sum_j \psi_j e_{j,FB}^2/2\right\}$. Again, (3.14) and (3.13) imply that the last two terms are both positive because larger volumes result in higher levels of factory- and transit-efforts.

At this stage, it is worth comparing the current setup to Antoniades (2015), which assumes that quality is fixed at the factory level (equivalent to letting $\beta \to 0$). Subsequently, the marginal return to output in that paper is simply $A-(2\gamma/L)\,q-cw$. Whether shipment volumes differ from this benchmark depend on the sign of the quality-mediated effect, which is summarized in Lemma 3.1.

Lemma 3.1. *Setting aside the standard net return to volume, $A-(2\gamma/L)\,q-cw$, the* net quality-mediated marginal gain to shipping volumes *simplifies to*

$$\begin{aligned}\left(z_{1,FB} + q\frac{\partial z_{1,FB}}{\partial q}\right)\gamma - \frac{\partial}{\partial q}\left\{\frac{\psi_0}{2}z_{0,FB}^2 + \sum_j \frac{\psi_j}{2}e_{j,FB}\right\} \\ = \underbrace{\beta^{\beta}(1-\beta)^{1-\beta}(1/\psi_0)^{1-\beta}\,\Phi_{FB}^{\beta(2-\rho)}}_{\equiv 2\Theta_{q,FB}}\gamma^2 q,\end{aligned} \tag{3.17}$$

which is

1. *always positive;*

[5] It is clear from (3.16) that this additional term is positive.

2. *increasing in γ, which measures the marginal effect of quality on sales revenue;*

3. *decreasing in the marginal cost of quality-creation, ψ_0;*

4. *increasing in the trading pair's joint distribution capability Φ_{FB}.*

Proof. See Section 3.A. □

The first part of Lemma 3.1 implies that first-best quantities in this paper exceed those in Antoniades (2015). Specifically, the first-best shipment volume solves

$$\underbrace{A - \frac{2\gamma}{L}q - cw}_{\text{Antoniades (2015)}} + \underbrace{2\gamma^2\Theta_{q,FB}q}_{\text{through quality choice}} = 0 \quad \Longleftrightarrow \quad q_{FB} = \frac{L}{2\gamma}\frac{A - cw}{1 - L\gamma\Theta_{q,FB}} \tag{3.18}$$

The (X, M) pair trade whenever

$$\begin{aligned} &\text{Positive margin :} && c < A/w \\ &\text{Limited quality-scope :} && L\,\gamma\,\Theta_{q,FB} < 1 \end{aligned} \tag{3.19}$$

As in Melitz and Ottaviano (2008), the first condition determines trade participation – given demand conditions A and labour costs w, the exporter/manufacturer must be sufficiently productive for exporting to be profitable. The second condition ensures declining marginal revenue, thus curbing the forces that generate sufficiently steep "quality ladders" in Antoniades (2015), where firms are more likely to innovate in quality if they face low innovation costs (low ψ_0), markets are large (L large), varieties are sufficiently differentiated (γ large). In contrast, the (3.14) and (3.16) guarantee quality innovation regardless of market size, exporter capabilities, or consumer preferences. The upper bound on $\Theta_{q,FB}$ in (3.19) ensures a strictly decreasing marginal revenue, guaranteeing a unique joint–welfare–maximizing level of output; otherwise, the parties wish to trade as large a volume as possible.[6]

3.2.2 Holdup and the role ownership

In this section, I assume that maintenance costs and the value of the goods are unverifiable to outside parties, so that the exporter and importer cannot sign quality-contingent contracts. Further, suppose the parties cannot commit to a revenue-sharing scheme. Instead, the contract between X and M simply specifies the desired level of physical output, q, the consignor, $\mathcal{O}$, and some initial payment, s, from the importer to the exporter.

Figure 3.1 illustrates the order of play. First, the importer proposes a contract $(q, \mathcal{O}, s)$.[7] The exporter accepts the contract if his expected payoff from the ensuing production and

[6] I show in Section 3.A that (3.18) delivers the Antoniades (2015) equilibrium, which assumes that quality is fixed at the factory level ($\beta \to 0$).

[7] The Principal's identity is irrelevant if we assume that both parties have quasilinear preferences and unlimited wealth.

distribution stages exceeds his reservation utility, which is normalized to zero. Substituting the binding participation constraint, and letting $z_{0,\mathcal{O}}$ and $e_{j,\mathcal{O}}$ denote the equilibrium (volume-contingent) levels of initial quality and maintenance by party j, the second-best contract solves

$$\max_{q,\mathcal{O}} \quad r\left(q, z_1\left(E\left(e_{X,\mathcal{O}}, e_{M,\mathcal{O}}\right) \middle| z_{0,\mathcal{O}}\right)\right) - \left(cwq + \frac{1}{2}\psi_0 z_{0,\mathcal{O}}^2 + \sum \frac{1}{2}\psi_j e_{j,\mathcal{O}}^2\right). \tag{3.20}$$

Unlike the first-best (3.9), the importer cannot decree that the parties take particular unverifiable actions. Instead, she must induce the exporter (and herself) to choose the desired levels of these unverifiable inputs in a manner consistent with their selfish interests.

Combined with the diminished value of the existing bundle to alternative buyers, this contractual incompleteness implies that the parties potentially bargain twice over any potential surplus from maintaining their relationship. They first bargain *after* the exporter has hired the $l = cq$ workers consistent with the desired output and sunk the initial effort z_0, but *before* z_0 has been incorporated into the q units. For example, the importer, well aware that outside parties value only a fraction δ of the initial quality z_0, may want to renegotiate the terms of trade after the exporter has already exerted some effort towards z_0. If the parties arrive at a mutually beneficial arrangement, they initiate the delivery stage with the bundle (q, z_0). However, if they disagree on the terms of trade, the relationship is terminated, and the exporter proceeds independently with the bundle $(q, \delta z_0)$.[8]

3.2.2.1 Distribution phase

If they maintain their relationship beyond the factory, the exporter and importer exert some maintenance effort towards (q, z_0). However, just as with the exporter's quality *creating* effort, the parties bargain over some unforeseen contingency after exerting maintenance effort but before incorporating these efforts into the factory-set bundle (q, z_0). If the relationship survives this second round of bargaining, the parties produce $(q, z_1(E|\, z_0))$, which the importer sells for

$$r^{IN}\left(E|\, q, z_0\right) \equiv r\left(q, z_1\left(E|\, z_0\right)\right) = \left(A + \gamma z_1\left(E|\, z_0\right) - \frac{\gamma}{L} q\right) q. \tag{3.21}$$

Unlike the factory-bargain, their disagreement payoffs depend on consignor's identity. Let $v^1_{j,\mathcal{O}}\left(E|z_0, q\right)$ denote j's disagreement payoff when $\mathcal{O}$ controls delivery, taking the aggregate maintenance effort, E, and initial quality and volume, (z_0, q), as given. Control over distribution determines disagreement outcomes because most contracts of carriage grant the party in charge the right to decide how to proceed following unforeseen events during shipment. For example, Marcet and de Ochoa Martínez (2006) note that such *residual rights of control* stem from the carriers' obligation – when reasonable – to await the consignor's in-

[8] I do not consider factory integration; see the vast literature on vertical integration in the face of holdup.

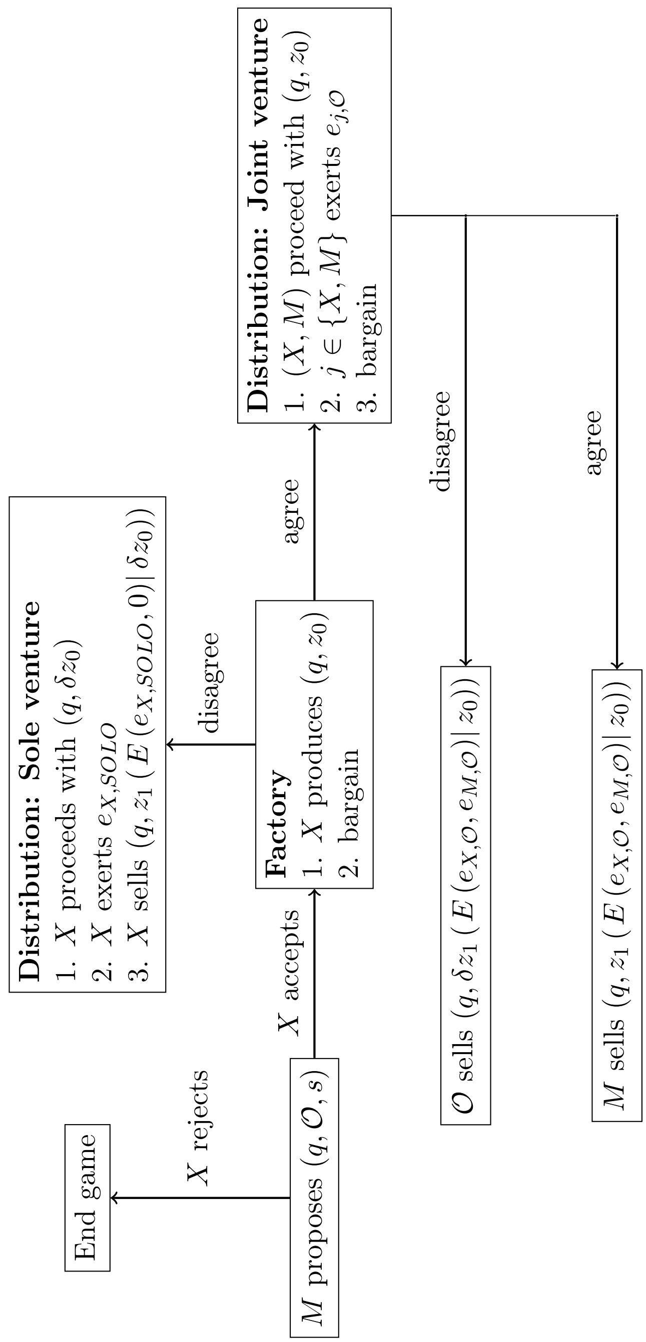

Figure 3.1: Order of play

structions "when transportation cannot be carried out, or impediments to the delivery arise." These freight contracts also effectively confer ownership over the shipment, as carriers must obey the consignor's wishes as to the intended recipient. Freight contracts therefore grant the consignor many of the rights typically associated with ownership. With this in mind, I refer to the party controlling delivery as the consignor or owner.

Because some of the effort is lost if the parties fail to reach an agreement, j's disagreement payoff, gross of the sunk maintenance cost, is

$$v^1_{j,\mathcal{O}}\left(E|q,z_0\right) = \mathbb{1}_{\mathcal{O}=j}\, r^{OUT}\left(E|\, q,z_0\right) \equiv \mathbb{1}_{\mathcal{O}=j}\left(A+\gamma\,\delta z_1\left(E|\, z_0\right)-\frac{\gamma}{L}q\right)q, \tag{3.22}$$

where $\mathbb{1}_{\mathcal{O}=j}$ indicates that j controls distribution. That is, $\mathcal{O}$ earns the sales revenue from a bundle embodying a fraction δ of the aggregate maintenance effort, while the non-controlling party is left empty-handed. Consumer demand (3.5) implies that higher initial quality renders the consignor's outside option more valuable in proportion to the shipment volume. However, larger volumes do not necessarily imply a more valuable outside option – the final salvageable quality must be sufficiently large, exceeding the threshold $z_1^{MIN} = 2q/L - A/\gamma$, an exceedingly difficult task when shipping in more differentiated sectors or to smaller markets.

The property rights literature in the tradition of Grossman and Hart (1986) and Hart and Moore (1990) stresses the distortionary effects of control rights in environments where parties make ex-ante non-contractible investments, as I assume here. Several variations of these models focus on environments where parties undertake too little of some productive activity (relative to the first-best) because they anticipate earning but a fraction of the marginal value of their investments. As we will see, the exporter and importer may *over*-invest, owing to the exporter's ability to influence future outcomes through his factory-based choices.

Definition 3.1. The *transit-stage renegotiation surplus under* $\mathcal{O}-$*control* is the difference between the value of reaching an agreement during transit-stage bargaining and the joint disagreement payoffs

$$\begin{aligned} R^1\left(E|q,z_0\right) &= r^{IN}\left(E|\, q,z_0\right) - r^{OUT}\left(E|\, q,z_0\right) \\ &= \left(1-\delta\right)\cdot\gamma q\cdot z_1\left(E|\, z_0\right). \end{aligned} \tag{3.23}$$

This surplus is the difference between the value added through maintenance efforts during transit within and outside the relationship. Inspecting the revenue function (3.5), the marginal return to quality is proportional to γq. The parties are thus more eager to reach an agreement when shipping large volumes, and this effect is magnified when the goods in question are highly differentiated.

Further, the transit surplus is increasing in $1-\delta$, the specificity of effort to the particular exporter-importer pair. The surplus approaches the entire valued added in transit as these efforts become increasingly specialized to M's consumer base. At the other extreme, there is nothing at stake during bargaining if quality is just as valuable outside the relationship ($\delta = 1$).

Lastly, recall that $z_1(E|\, z_0) \equiv z_0^{1-\beta} E^{\beta}$ measures destination quality, given factory quality z_0 and aggregate transit efforts E. As a result, bargaining is pointless if initial quality, z_0, is zero, or if neither party performs maintenance ($E = 0$). In contrast, the renegotiation surplus is large whenever high-quality bundles leave the factory, and/or the parties exert a great deal of effort before bargaining. Gathering these observations yields the following sufficient condition for "successful" transit-stage bargaining.

Proposition 3.1. *The parties reach an agreement during transit-stage bargaining whenever (i) maintenance effort is partially-specific to the relationship; (ii) the shipment embodies positive quality levels upon leaving the factory; and (iii) at least one party exerts effort towards transit-stage quality maintenance.*[9]

Under simple Nash bargaining over the transit pie, j earns their disagreement payoff, plus half of the renegotiation surplus. Taking the other party's choice as given, j anticipates earning

$$u^1_{j,\mathcal{O}}(e_X, e_M|q, z_0) = \mathbb{1}_{\mathcal{O}=j} \cdot r^{OUT}(E|\, q, z_0) + \frac{1}{2} R^1(E|q, z_0) - \frac{1}{2}\psi_j e_j^2, \tag{3.24}$$

from choosing choosing e_j.

Since the marginal returns to quality on both the owner's outside option and the surplus is γq, the exporter's best-response solves

$$\mu_{j,\mathcal{O}}\, \gamma q\, \beta z_0^{1-\beta} E(e_X, e_M)^{\beta-\rho}\, \eta_j e_j^{\rho-1} = \psi_j e_j, \qquad \mu_{j,\mathcal{O}} \equiv \mathbb{1}_{\mathcal{O}=j}\, \delta + \frac{1}{2}(1-\delta) \tag{3.25}$$

In equilibrium, j equates the marginal cost of maintenance to their share of value added during transit, adjusting for ownership rights by $\mu_{j,\mathcal{O}}$. This adjustment factor ranges from a high of $\frac{1}{2}(1+\delta)$ when j controls delivery, to a low of $\frac{1}{2}(1-\delta)$ when the other party is in charge. Delivery rights encourage owner effort at the expense of the other party's efforts. Finally, if quality maintenance is entirely relationship-specific ($\delta = 0$), then $\mu_{j,\mathcal{O}} = \frac{1}{2}$ does not vary across parties or ownership structures.

It is worth highlighting the differences between the current setup and Antràs and Chor (2013), who also model sequential production. In their model, each stage – analogous to our "factory" and "delivery" phases – is operated by a distinct agent. In their baseline model, each agent only considers the effect of its investment on final sales revenue, so that effort at a given stage depends only on effort in preceding stages.[10]

---

[9] While parts (ii) and (iii) seem to rely on the particular functional form for $z_1(E|\, z_0)$, the result follows given our demand function provided $\delta < 1$ and $z_1(\cdot|\, z_0)$ is increasing.

[10] They do consider an extension to agents who internalize the effect of their choice on downstream production, but drive this forward-looking behaviour to zero by considering a continuum of stages.

In contrast, this paper recognizes the fact that both the exporter and importer may enhance the shipment's quality. This observation introduces *within*-stage strategic interactions. Specifically, individual maintenance efforts interact via the CES effort aggregator $E(e_X, e_M)$. If $\beta = \rho$, then the parties have dominant strategies, so that the analysis follows à-la-Grossman and Hart (1986). Setting aside this knife-edge case, best responses are upward sloping whenever $\beta > \rho$, and downward sloping otherwise.

The equilibrium *aggregate* effort is

$$E_{\mathcal{O}}(q, z_0) = \left(\gamma q \cdot \beta z_0^{1-\beta}\, \Phi_{\mathcal{O}}^{2-\rho}\right)^{\frac{1}{2-\beta}}, \tag{3.26}$$

where

$$\Phi_{\mathcal{O}} \equiv E(\phi_{X,\mathcal{O}}, \phi_{M,\mathcal{O}}) = \left(\eta\, \phi_{X,\mathcal{O}}^{\rho} + (1-\eta)\, \phi_{M,\mathcal{O}}^{\rho}\right)^{\frac{1}{\rho}}, \qquad \phi_{j,\mathcal{O}} \equiv \left(\mu_{j,\mathcal{O}} \frac{\eta_j}{\psi_j}\right)^{\frac{1}{2-\rho}} \tag{3.27}$$

are the *control-adjusted efficiency index,* and the *individual control-adjusted efficiency.* Like its first-best counterpart (3.11), $\Phi_{\mathcal{O}}$ is a share-weighted average of individual capabilities, with weights corresponding to the importance of a party's maintenance effort. These indices differ in the $\mu_{j,\mathcal{O}}$ terms, which summarize the effects of relationship-specific efforts and ownership on aggregate productivity. Specifically, relative to the first best, $\Phi_{\mathcal{O}}$ scales down individual productivities by the effective contribution to value-added, $0 < \mu_{j,\mathcal{O}} < 1$. Effort non-contractibility is therefore equivalent to a reduction in individual distribution capabilities that disproportionately targets the non-controlling party.

Aggregate efficiency differs across ownership structures depending on the relative importance of exporter effort, the relative exporter marginal cost, and the substitutability of individual efforts. Figure 3.2 plots the first- and second-best joint capabilities as functions of the exporter's share in aggregate effort. The panels differ in the identity of the relatively more efficient trading partner, with Panel (a) corresponding to a more efficient importer ($\psi_M < \psi_X$).

The first-best aggregate distribution capability exceeds the second-best under either party's control, regardless of the relative importance of exporter effort, η. Further, the joint capability under exporter-control eventually surpasses that under importer control as exporter effort becomes more important (as $\eta \to 1$). Lemma 3.2 shows that the critical value of η depends on the relative marginal cost of effort and the substitutability of individual efforts.[11]

Lemma 3.2. *The joint distribution capability, Φ, is greater under exporter-control if and only if exporter effort is sufficiently important. Specifically, the exporter's share of aggregate effort, η, must exceed a threshold, $\eta_\Phi^* = \eta_\Phi^*(\psi_X/\psi_M, \rho)$, which*

1. *increases in the exporter's relative marginal cost of effort, ψ_X/ψ_M, and*

[11] See 3.B for the proof and explicit formula for the cutoff

Figure 3.2: Aggregate productivity and the contracting environment

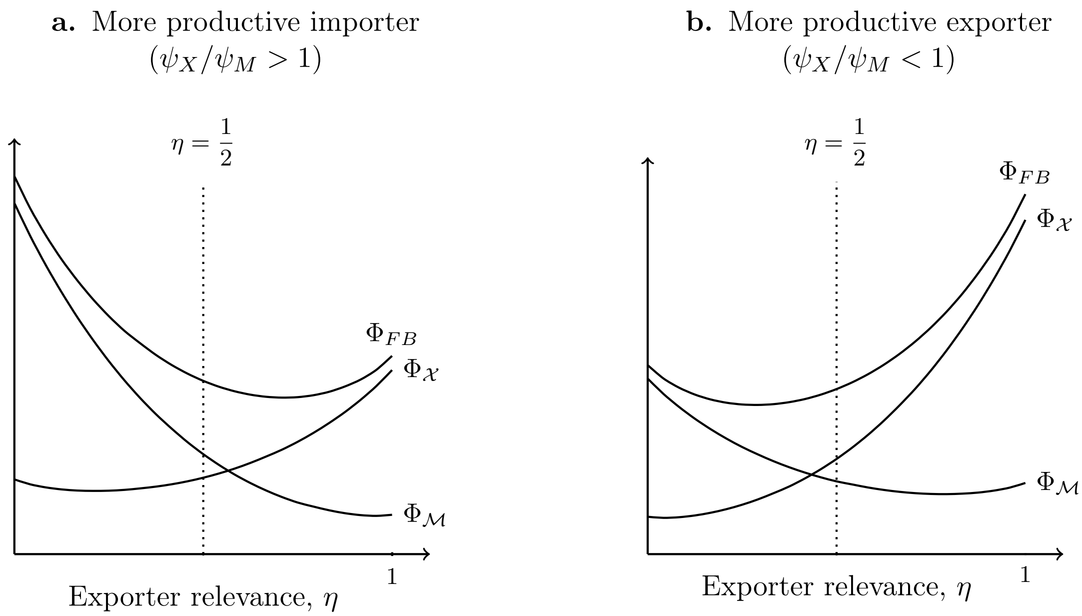

Notes: Aggregate productivity in the first-best, and under exporter ($\Phi_{\mathcal{X}}$) and importer ($\Phi_{\mathcal{M}}$) control. The importer is relatively more productive at the margin in Panel a, while the exporter is more productive in Panel b. First-best aggregate productivity exceeds the second-best under any ownership arrangement, regardless of the value of η. Second-best aggregate productivity is higher under exporter control when his contribution to aggregate effort, η, exceeds the threshold defined by the intersection of $\Phi_{\mathcal{X}}$ and $\Phi_{\mathcal{M}}$.

2. *decreases in the elasticity of substitution between individual efforts if the exporter is more efficient* ($\psi_X/\psi_M < 1$)*, and increases in the elasticity of substitution if the importer is more efficient* ($\psi_X/\psi_M > 1$)*.*

In particular, $\eta^*_\Phi(1,\rho) = 1/2$*; if the exporter and importer are equally productive, the pair is better at distribution under exporter control if and only if the exporter makes the more important investment* ($\eta > 1/2$)*.*

All else equal, the gains from transferring ownership to the exporter are increasing in his relative productivity, ψ_X/ψ_M. If the exporter is less productive than the importer, then joint capability falls whenever he assumes control, with a more pronounced decline as individual efforts become increasingly substitutable. Intuitively, when ρ is large, transferring ownership to the exporter discourages the importer from exerting that is just as valuable on the margin. Further, $\Phi_{\mathcal{O}}$ is increasing in δ, the fraction of effort valuable outside the relationship, whenever $\mathcal{O}$'s effort is relatively more important. This follows from the complementarity between δ and $\Phi_{\mathcal{O}}$ in the salvageable destination quality, $\delta z_1 \left(E_{\mathcal{O}} | \, z_{0,\mathcal{O}} \right)$. In spite of this, $\Phi_{\mathcal{X}}/\Phi_{\mathcal{M}}$ is independent of δ, so that δ does not single-handedly determine the ranking of aggregate capabilities across contractual forms. In the extreme case where alternative buyers do not value the pair's particular quality improvements ($\delta = 0$), then

$\mu_{j,\mathcal{O}} = \frac{1}{2}$, which renders $\Phi_{\mathcal{O}}$ independent of the contractual form $\mathcal{O}$.

Applying Lemma 3.2 to the expression for $E_{\mathcal{O}}$ in (3.26) provides a ranking of aggregate effort across ownership structures:

Proposition 3.2. *Aggregate maintenance effort is greater under exporter control if and only if the exporter has a sufficiently large share of aggregate effort.*

Returning to the investment game, individual effort is

$$e_{j,\mathcal{O}}(q, z_0) = \left(\gamma q \cdot \beta z_0^{1-\beta}\, \Phi_{\mathcal{O}}^{\beta-\rho}\right)^{\frac{1}{2-\beta}} \phi_{j,\mathcal{O}} \qquad \left(= \frac{\phi_{j,\mathcal{O}}}{\Phi_{\mathcal{O}}} E_{\mathcal{O}}(q, z_0)\right). \tag{3.28}$$

As with aggregate effort, ownership effectively changes the marginal productivity of effort. Further, in what will become a recurring theme, individual effort is increasing in initial quantity and quality. Finally, note that relative efforts,

$$\frac{e_{j,\mathcal{O}}}{e_{-j,\mathcal{O}}} = \frac{\phi_{j,\mathcal{O}}}{\phi_{-j,\mathcal{O}}} = \left(\frac{\mu_{j,\mathcal{O}}}{\mu_{-j,\mathcal{O}}}\right)^{\frac{1}{2-\rho}} \frac{\phi_{j,FB}}{\phi_{-j,FB}} = \left(\frac{\mu_{j,\mathcal{O}}}{\mu_{-j,\mathcal{O}}}\right)^{\frac{1}{2-\rho}} \frac{e_{j,FB}}{e_{-j,FB}}, \tag{3.29}$$

are pinned down by the relative control-adjusted efficiencies, as shown in Panel (a) of Figure 3.3. Like the first-best in (3.12), relative equilibrium efforts under $\mathcal{O}$-control lie along a ray through the origin. However, the slopes of the rays under the three regimes (first best, $\mathcal{O} = \mathcal{M}$, and $\mathcal{O} = \mathcal{X}$) differ due to the scarcity of control rights. Assigning ownership to one party necessarily deprives the other of control, so that each party's relative effort is higher when it controls distribution:

$$\frac{\mu_{M,\mathcal{M}}}{\mu_{X,\mathcal{M}}} > 1 > \frac{\mu_{M,\mathcal{X}}}{\mu_{X,\mathcal{X}}} \implies \frac{e_{M,\mathcal{M}}}{e_{X,\mathcal{M}}} > \frac{e_{M,FB}}{e_{X,FB}} > \frac{e_{M,\mathcal{X}}}{e_{X,\mathcal{X}}}. \tag{3.30}$$

Panel (b) of Figure 3.3 illustrates the first-best, and second best equilibrium efforts under the two control structures, assuming that $\beta > \rho$ (maintenance efforts are strategic complements), $\psi_X = \psi_M$ (the parties are equally productive at the margin), and $\eta < \eta^*_\Phi < 1/2$ (exporter effort is not important enough to increase aggregate productivity). The solid and broken lines indicate best responses under exporter- and importer-controlled shipments respectively.

Comparing the first-best efforts and the intersections either the broken ($\mathcal{O} = \mathcal{M}$) or solid ($\mathcal{O} = \mathcal{X}$) lines, the first-best individual efforts exceed their second-best counterparts under either ownership arrangement. Further, comparing the second-best equilibria, exporter control leads to lower importer effort but greater exporter effort. When the importer makes the more important investment, which is the case in the figure, the fall in importer effort outweighs the increase in exporter effort, as shown by the new equilibrium lying below the downward sloping iso-aggregate-effort curve through the initial equilibrium point. In this scenario, there is a clear ranking of *aggregate* effort across the various arrangements : $E_{FB} > E_{\mathcal{M}} > E_{\mathcal{X}}$.

Figure 3.3: Best-response curves when efforts are strategic complements ($\beta > \rho$)

a. Changes in γ, q, z_0 **b.** Changes in ownership

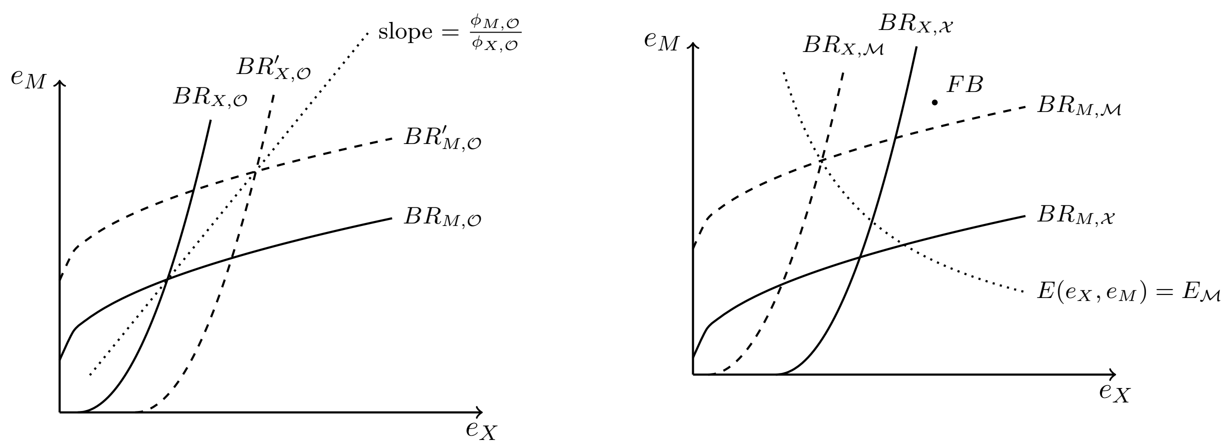

Notes: $BR_{j,\mathcal{O}}$ is j's best response under a $\mathcal{O}-$controlled shipment. Individual efforts are strategic complements when $\beta > \rho$, hence the upward sloping best-response functions. Panel (a) traces equilibrium efforts as the delivery-stage state variables, q and z_0, change, holding ownership fixed. Panel (b) illustrates the role of ownership. The downward sloping line traces combinations of exporter and importer efforts that result in the equilibrium level of aggregate maintenance effort under importer control. Aggregate effort is lower under X-control because M's effort is more important ($\eta < 1/2$) and M and X are equally productive at the margin ($\psi_X = \psi_M$).

3.2.2.2 Manufacturing phase

This section characterizes the equilibrium factory-set quality, taking shipment volume and the Nash equilibrium in subsequent transit efforts as given. The analysis delivers a control-specific policy rule $z_{0,\mathcal{O}}(q)$, and derives comparative statics with respect to shipment volume q and consignment rights $\mathcal{O}$.

Having accepted the importer's contract, the exporter chooses initial quality, aware that the parties will bargain soon thereafter, and, if successful, proceed to the delivery stage and play the strategies $e_{j,\mathcal{O}}(q, z_0)$ derived in the previous section. The exporter thus chooses z_0 to maximize his payoffs across both bargaining stages. Let

$$U^1_{j,\mathcal{O}}(q, z_0) \equiv u^1_{j,\mathcal{O}}\left(e_{X,\mathcal{O}}(q, z_0), e_{M,\mathcal{O}}(q, z_0) \middle| q, z_0\right) \tag{3.31}$$

denote the corresponding equilibrium payoffs from bargaining in transit, where, recall, $u^1_{j,\mathcal{O}}$ in (3.24) is j's objective in the distribution investment game.

Since I rule out factory integration, disagreement at this stage leaves the importer empty-handed. That is, her factory-disagreement payoff is $v^0_M(q, z_0) = 0$. In contrast, the exporter can appropriate a fraction δ of the relationship-specific factory-set quality.[12] Thus, if the

[12] At the cost of extra notation, I could allow the salvageable components of factory and transit effort to differ. All subsequent results are robust to the simplifying assumption in the main text.

parties disagree, the exporter independently initiates the delivery phase with the bundle $(q, \delta z_0)$, eventually offloading the final bundle on a less enthusiastic buyer.[13] In the absence of the initial importer, aggregate transit effort is $E(e_X, 0) = \eta^{1/\rho} e_X$, which implies destination quality of $z_1\left(\eta^{1/\rho} e_X \middle| \delta z_0\right)$. The exporter's disagreement payoff in the factory-based Nash bargaining game is

$$v_X^0(q, z_0) = \max_e \left\{ r\left(q, z_1\left(E(e,0)\middle|\, \delta z_0\right)\right) - \frac{1}{2}\psi_X e^2 \right\}, \tag{3.32}$$

the maximized profit from selling the bundle to some alternative buyer. In this branch of play, exporter effort in the transit phase is characterized by a single-agent first-order condition rather than a pair of best-response functions as in (3.25). The optimal "breakaway" transit effort is

$$e_{X,SOLO}(q, z_0) = \left(\gamma q \cdot \beta z_0^{1-\beta} \Phi_{SOLO}^{\beta-\rho}\right)^{\frac{1}{2-\beta}} \phi_{X,SOLO}, \tag{3.33}$$

where

$$\Phi_{SOLO} \equiv \left(\eta \phi_{X,SOLO}^{\rho} + (1-\eta)\, \phi_{M,SOLO}^{\rho}\right)^{\frac{1}{\rho}}, \qquad \phi_{X,SOLO} \equiv \left(\frac{\eta \delta^{1-\beta}}{\psi_X}\right)^{\frac{1}{2-\beta}}, \quad \phi_{M,SOLO} = 0$$

assume the roles of the ownership-adjusted efficiencies $\Phi_{\mathcal{O}}$ and $\phi_{j,\mathcal{O}}$, with $\left(\delta^{1-\beta}, 0\right)$ assuming the role of $(\mu_{X,\mathcal{O}}, \mu_{M,\mathcal{O}})$ in the cooperative outcome. The restriction $\phi_{M,SOLO} = 0$ reflects the importer's inactivity in the exporter's sole venture and the exporter's sole access to the shipment. Like the cooperative maintenance efforts, this threat-point-maximizing effort is increasing in the shipment volume and the degree of product differentiation. Note that aggregate effort is

$$E_{SOLO}(q, z_0) = \left(\gamma q \cdot \beta z_0^{1-\beta} \Phi_{SOLO}^{2-\rho}\right)^{\frac{1}{2-\beta}}. \tag{3.34}$$

Having derived the exporter's disagreement payoff, we now turn to the value of cooperation during factory-phase bargaining. If the parties reach an amicable settlement, they proceed to the delivery phase with the higher quality bundle (q, z_0), and then play their equilibrium strategies (3.28), earning $U_{j,\mathcal{O}}^1(q, z_0)$. Letting $e_{j,\mathcal{O}}$ and $E_{\mathcal{O}}$ denote the equilibrium individual and aggregate levels in (3.28) and (3.26), this branch of play earns the parties

$$\begin{aligned} \sum_j U_{j,\mathcal{O}}^1(q, z_0) &= r^{OUT}\left(E_{\mathcal{O}}\middle|\, q, z_0\right) + R^1\left(E_{\mathcal{O}}\middle| q, z_0\right) - \sum_j \frac{1}{2}\psi_j e_{j,\mathcal{O}}^2 \\ &= r^{IN}\left(E_{\mathcal{O}}\middle|\, q, z_0\right) - \sum_j \frac{1}{2}\psi_j e_{j,\mathcal{O}}^2 \end{aligned} \tag{3.35}$$

where I use the fact that transit-phase bargaining is a constant-sum game, in which the exporter and importer divide the sales revenue from maintaining their relationship through

[13] For simplicity, I assume that the exporter uses some exogenous delivery system to get the shipment to its alternative buyers, rendering the existing carriage contract worthless.

delivery, net of the total effort cost.

The threat point (3.32) and the value of cooperation (3.35) define the factory-based bargaining game, where the parties reach a mutually beneficial agreement as long as their joint future payoffs exceed the exporter's immediate outside option.

Definition 3.2. The *factory-stage renegotiation surplus under $\mathcal{O}-$control,* is the difference between the value of reaching an agreement during factory-based bargaining and exporter's immediate outside option:

$$R^0_{\mathcal{O}}(q, z_0) \equiv \sum_j U^1_{j,\mathcal{O}}(q, z_0) - v^0_X(q, z_0) = r^{IN}(E_{\mathcal{O}} | q, z_0) - \sum_j \frac{1}{2}\psi_j e^2_{j,\mathcal{O}} - v^0_X(q, z_0). \quad (3.36)$$

Note that, unlike the distribution-phase surplus (3.23), which is analogous to the surplus in standard single-stage production models, the factory surplus accounts for the joint payoffs from the subsequent delivery stage.

As the unique actor in the factory phase, the exporter has considerable leeway in influencing future play to suit his needs. He considers the effect of his choice of initial quality on his immediate outside option, $v^0_X(q, z_0)$, his future bargaining payoff, $U^1_{X,\mathcal{O}}$, and – with symmetric Nash bargaining – half the value of allowing production to advance to the delivery stage, $\frac{1}{2}\sum_j U^1_{j,\mathcal{O}}(q, z_0)$.

The exporter's payoff from factory-based bargaining, net of the cost of effort, is

$$\begin{aligned} u^0_{X,\mathcal{O}}(q, z_0) &= v^0_X(q, z_0) + \frac{1}{2}R^0_{\mathcal{O}}(q, z_0) - \frac{1}{2}\psi_0 z_0^2 \\ &= \frac{1}{2}v^0_X(q, z_0) + \frac{1}{2}\sum_j U^1_{j,\mathcal{O}}(q, z_0) - \frac{1}{2}\psi_0 z_0^2, \end{aligned} \quad (3.37)$$

where the second line follows from substituting (3.36). The exporter places some weight on the off-the-equilibrium-path event that factory-based bargaining breaks down. Unlike the familiar one-shot production/trade models, the exporter chooses initial quality, z_0, to maximize his joint bargaining payoffs across the production and delivery phases. Assuming no discounting, the forward-looking exporter maximizes $u^0_{X,\mathcal{O}}(q, z_0) + U^1_{X,\mathcal{O}}(q, z_0)$, so that his factory-stage objective is a weighted sum of three income streams: (i) his income from a solo venture, $v^0_X(q, z_0)$; (ii) his own payoff in the delivery phase, $U^1_{X,\mathcal{O}}(q, z_0)$; and (iii) the importer's delivery-phase payoff, $U^1_{M,\mathcal{O}}(q, z_0)$. Specifically, using the second line of (3.37),

$$u^0_{X,\mathcal{O}}(q, z_0) + U^1_{X,\mathcal{O}}(q, z_0) = \underbrace{\frac{1}{2}v^0_X(q, z_0)}_{\text{Sole-venture incentive}} + \underbrace{\frac{3}{2}U^1_{X,\mathcal{O}}(q, z_0) + \frac{1}{2}U^1_{M,\mathcal{O}}(q, z_0)}_{\text{Joint-venture incentive}} - \frac{1}{2}\psi_0 z_0^2.$$

The first term on the right-hand side captures the gains accruing from the exporter's immediate outside option, which involves proceeding to the shipping phase alone. The second term measures the gains from sustaining the existing relationship. Adopting the perspective of the exporter (who chooses z_0), I refer to $\frac{3}{2}U^1_{X,\mathcal{O}}(q, z_0)$ as the *own-payoff incentive,* and $\frac{1}{2}U^1_{M,\mathcal{O}}(q,z_0)$ as the *rival-payoff incentive.* Ignoring the positive weights unless absolutely necessary, I first sign the own- and rival-payoff incentives, $\mathrm{d}U^1_{X,\mathcal{O}}(z_0,q)/\mathrm{d}z_0$ and $\mathrm{d}U^1_{M,\mathcal{O}}(z_0,q)/\mathrm{d}z_0$ under an arbitrary contractual form, and then describe changes in these incentives as control transfers from the importer to the exporter, $\mathrm{d}\left[U^1_{j,\mathcal{X}}(q,z_0)-U^1_{j,\mathcal{M}}(q,z_0)\right]/\mathrm{d}z_0$.

Sole-venture effect. In this branch of play, the exporter opts to terminate the relationship before delivery begins. The marginal return to quality on the exporter's threat point is

$$\frac{\partial v^0_X(q,z_0)}{\partial z_0}=\gamma q\,\frac{\partial z_1\left(E_{SOLO}(q,z_0)|\,\delta z_0\right)}{\partial z_0}=\gamma q\;(1-\beta)\left(\gamma q\cdot\beta z_0^{1-\beta}\Phi_{SOLO}^{2-\rho}\right)^{\frac{1-\beta}{2-\beta}}>0,$$

where the envelope theorem allows us to disregard the effects of z_0 on v^0_X through the optimally chosen effort, $e_{X,SOLO}(q,z_0)$.

The factory-stage threat point elicits greater initial quality in transactions involving large volumes in differentiated sectors. Note that ownership rights, which only matter if the relationship survives beyond the manufacturing phase, are irrelevant for the sole-venture channel.

Joint-venture effect. The incentives to alter future play depend on the effects of factory-set quality on the transit-stage best responses (cross-stage strategic interactions), and on strategic interactions between individual efforts within the transit stage. Recall that this income stream arises from the exporter appropriating j's transit-stage payoff, where $j=X,M$. The total effect of a change in z_0 on j's transit-stage payoff is

$$\frac{\mathrm{d}U^1_{j,\mathcal{O}}(z_0,q)}{\mathrm{d}z_0}=\mu_{j,\mathcal{O}}\cdot\gamma q\left(\frac{\partial z_1\left(E_{\mathcal{O}}|\,z_0\right)}{\partial z_0}+\frac{\partial z_1\left(E_{\mathcal{O}}|\,z_0\right)}{\partial e_{-j}}\frac{\partial e_{-j,\mathcal{O}}}{\partial z_0}\right), \tag{3.38}$$

where the envelope theorem eliminates the effect of z_0 on $U^1_{j,\mathcal{O}}$ through j's own choice $e_{j,\mathcal{O}}$. The total effect is the sum of the *direct effect* of initial quality on $U^1_{j,\mathcal{O}}$, and the *strategic effect*, mediated by $-j$'s response, $e_{-j,\mathcal{O}}$.

The direct effect combines j's share of value added in transit, the marginal returns to quality, and the marginal returns to effort in quality creation,

$$\mu_{j,\mathcal{O}}\cdot\gamma q\frac{\partial z_1\left(E_{\mathcal{O}}|\,z_0\right)}{\partial z_0}=\gamma q\cdot(1-\beta)\left(\gamma q\cdot\beta z_0^{-1}\Phi_{\mathcal{O}}^{2-\rho}\right)^{\frac{\beta}{2-\beta}}\mu_{j,\mathcal{O}}. \tag{3.39}$$

This effect is unambiguously positive, and stronger when shipping large volumes of differentiated goods. After all, factory effort increases initial quality, which increases sales revenue disproportionately in sectors where consumers enjoy variety. Thus, holding ownership rights fixed, the direct effect encourages the exporter to create high quality goods.

Changes in contractual form affect the direct channel through (i) the pair's joint ca-

Figure 3.4: Joint-venture: Changes in quality-creation incentives induced by transfer of ownership

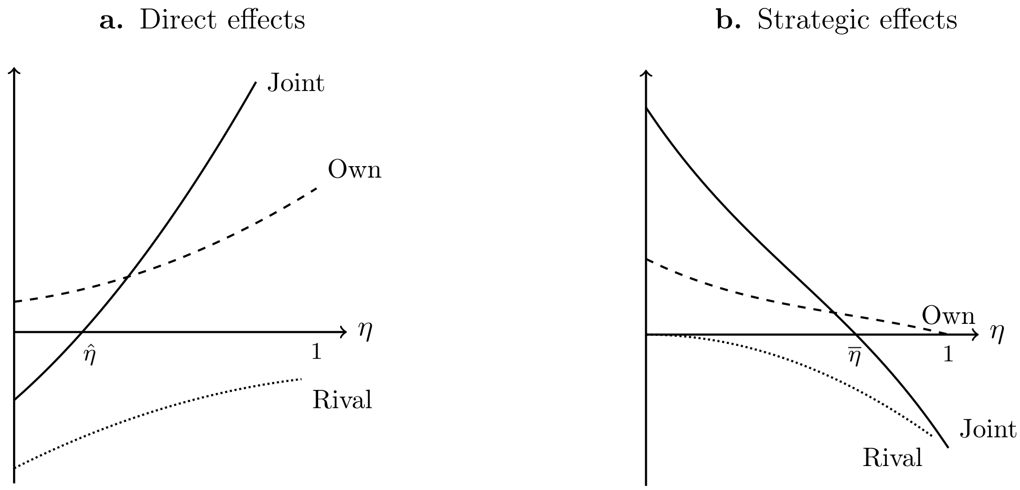

Notes: Effects of transferring control to the exporter on (a) the direct (3.39); and (b) the strategic (3.40) channels for initial quality choice. I ignore $\mathcal{O}$-independent terms in these expressions, focusing on changes in $\Phi_{\mathcal{O}}^{\beta(2-\rho)/(2-\beta)}\,\mu_{j,\mathcal{O}}$ in Panel (a), and changes in $\Phi_{\mathcal{O}}^{2(\beta-\rho)/(2-\beta)}\,\mu_{j,\mathcal{O}}\,\eta_{-j}\phi_{-j,\mathcal{O}}^{\rho}$ in Panel (b). The *own-payoff incentive* ($j = X$) is the effect through the exporter's transit-stage payoff, while the *rival-payoff incentive* ($j = M$) operates through the exporter's share of the importer's future payoff. The *joint venture incentive* effect is three times the own incentive, plus the rival incentive.

pability, $\Phi_{\mathcal{O}}$, which, according to Lemma 3.2, is greater in exporter-controlled shipments whenever exporter effort is particularly useful; and (ii) j's share of value added, $\mu_{j,\mathcal{O}}$, which is greater when j controls delivery. In principle, these effects may oppose each other, with different implications for the own and rival incentives.

Panel (a) of Figure 3.4 summarizes the effects of transferring consignment rights to the exporter on the direct effect. The own-payoff incentive (dashed line) is unambiguously positive if exporter effort is sufficiently important, that is, if $\eta > \eta_{\Phi}^{*}$. Transferring control to the exporter raises his incentives to invest, and raises aggregate productivity. Perhaps surprisingly, it remains positive even if exporter effort is not important enough to guarantee that $\Phi_{\mathcal{X}} > \Phi_{\mathcal{M}}$. In other words, the exporter's desire to extract a larger share of the transit-stage surplus, operating through $\mu_{X,\mathcal{O}}$, outweighs potential efficiency concerns. In contrast, the rival-payoff incentive (dotted line) is negative. The importer earns a smaller share of value-added under exporter control ($\mu_{M,\mathcal{X}} < \mu_{M,\mathcal{M}}$), which compounds the efficiency loss when $\eta < \eta_{\Phi}^{*}$, and outweighs any efficiency gains when $\eta > \eta_{\Phi}^{*}$ (again, the $\mu_{M,\mathcal{O}}$ term, which captures the battle over transit-surplus, dominates). If exporter effort is not too important (less than $\hat{\eta}$ in Figure 3.4), the direct effect results in lower quality goods leaving the factory under exporter control despite the exporter attaching three times as much weight to the own-payoff incentive than to the rival-payoff incentive.[14]

[14]The cutoff $\hat{\eta}$ is increasing in β and ρ.

The strategic effect combines j's share of value added in transit, the marginal returns to quality, and $-j$'s optimal effort,

$$\mu_{j,\mathcal{O}} \cdot \gamma q \frac{\partial z_1\left(E_{\mathcal{O}}\middle| z_0\right)}{\partial e_M} \frac{\partial e_{-j,\mathcal{O}}}{\partial z_0} = \frac{\beta\left(1-\beta\right)}{2-\beta}\left(\left(\gamma q \cdot \beta z_0^{-1}\right)^{\beta} \Phi_{\mathcal{O}}^{2(\beta-\rho)}\right)^{\frac{1}{2-\beta}} \gamma q\, \mu_{j,\mathcal{O}}\, \eta_{-j} \phi_{-j,\mathcal{O}}^{\rho}. \tag{3.40}$$

The *own-payoff incentive* requires that the exporter manipulate the importer's future behaviour, while the *rival-payoff incentive* requires that the exporter alter his own future behaviour. The apparent switch in perspective follows from the envelope theorem. When appropriating some of his rivals future payoff $U^1_{M,\mathcal{O}}$, the exporter ignores changes in $e_{M,\mathcal{O}}$, leaving only his own action. In any case, the strategic effect is also positive because higher levels of initial quality shift the maintenance best responses outwards ($\partial e_{j,\mathcal{O}}/\partial z_0 > 0$), resulting in higher quality destination goods ($\partial z_1\left(E\left(e_X, e_M\right)\middle| z_0\right)/\partial e_j > 0$). Intuitively, although I hold the allocation of delivery rights fixed, so that the exporter earns the same share of the surplus, he is better off because the parties now share a larger pie. Yet again, this effect is magnified when shipping large volumes of differentiated goods. Therefore, given an allocation of property rights, the exporter creates higher quality goods than he would without such strategic considerations.

In our discussion of the effects of ownership on the direct channel, we established that the desire to earn a larger share of the transit-stage surplus, operating through $\mu_{j,\mathcal{O}}$, outweighs efficiency concerns. The same reasoning applies in the strategic effect. For example, looking at the own-payoff incentive, we see that transferring ownership to the exporter raises his share of value added, $\mu_{X,\mathcal{X}} > \mu_{X,\mathcal{M}}$, while lowering the importer's efficiency, $\phi_{M,\mathcal{X}} < \phi_{M,\mathcal{M}}$. The key departure from that discussion concerns $-j$'s share-weighted efficiency, $\eta_{-j}\phi_{-j,\mathcal{O}}$, and the effects of changes in overall efficiency.

The $\phi_{-j,\mathcal{O}}$ term, which is greater whenever $-j$ controls delivery, appears because only $-j$'s choice has a first-order effect on $U^1_{j,\mathcal{O}}$. This seemingly presents an additional force, proportional to the importance of $-j$'s effort (η_{-j}), against the $\mu_{j,\mathcal{O}}$-driven battle over the distribution-phase surplus. Turning to overall efficiency concerns, giving the exporter control of shipping (even when his effort is sufficiently important; $\eta > \eta^*_{\Phi}$) strengthens the strategic channel if and only if efforts are strategic complements.

Panel (b) of Figure 3.4 summarizes the effects of transferring control to the exporter for various levels of η. Note that the rival effect vanishes as η approaches zero – if exporter effort has little effect on the aggregate, then the exporter has little incentive to restrict his future behaviour as this will have a negligible effect on the importer's actions. Similarly, the own effect vanishes as $1-\eta$ approaches zero. If exporter effort is important enough (greater than $\overline{\eta}$ in Figure 3.4), the strategic effect results in lower quality goods leaving the factory under exporter control. These qualitative properties hold for all values of β and ρ; that is, regardless of whether maintenance efforts are strategic complements or substitutes as determined by the sign of $\beta-\rho$. Instead, β and ρ affect the cutoff $\overline{\eta}$ beyond which the strategic effect leads to lower quality goods under exporter control. In particular, $\overline{\eta}$ is increasing in β and decreasing in ρ.

To summarize, transferring control to the exporter affects initial quality through direct and strategic channels. The direct channel encourages higher quality goods under exporter control when exporter effort is sufficiently important, while the strategic channel discourages higher quality goods if exporter control is too important. The remainder of this section combines these two counteracting forces, deriving a sufficient statistic for the effect of contractual form on the choice of initial quality. We will see that initial quality is higher under exporter-control regardless of the importance of exporter effort, η, or the nature of strategic interactions within the delivery stage.

Relegating the details to Section 3.C, the first-order condition for initial quality, which equates the marginal cost and benefit of z_0, delivers

$$z_{0,\mathcal{O}}\left(q\right) = \Theta_{z0,\mathcal{O}}\,\gamma q, \tag{3.41}$$

where

$$\Theta_{z0,\mathcal{O}} \equiv \beta^{\frac{\beta}{2}}\left(1-\beta\right)^{\frac{2-\beta}{2}}\left(\frac{1}{\psi_0}\right)^{\frac{2-\beta}{2}}\left(\omega_{SOLO}\Phi_{SOLO}^{\frac{\beta(2-\rho)}{2-\beta}} + \sum_j \omega_{j,\mathcal{O}}\Phi_{\mathcal{O}}^{\frac{\beta(2-\rho)}{2-\beta}}\right)^{\frac{2-\beta}{2}} > 0 \tag{3.42}$$

summarizes the effects of production and distribution technologies. Setting aside $\Theta_{z0,\mathcal{O}}$ for the moment, the exporter responds to higher shipping volumes by raising the initial quantity. Further, initial quality, like all other effort levels we have considered thus far, is higher in differentiated sectors.

Returning to $\Theta_{z0,\mathcal{O}}$, the first three terms, which also appear in the first-best decision rule $\Theta_{z0,FB}$ in (3.14), show that initial quality is decreasing in the marginal cost of quality creation, ψ_0. It is comforting to know that this purely technologically-driven conclusion is independent of the contracting environment. The term $\Theta_{z0,\mathcal{O}}$ differs from the first-best $\Theta_{z0,FB}$ through a weighted power mean of the exporter's individual capability, $\Phi_{SOLO}^{\beta(2-\rho)/2}$, and the (ownership-dependent) joint capability, $\Phi_{\mathcal{O}}^{\beta(2-\rho)/2}$. The weights are given by

$$\omega_{SOLO} \equiv \frac{1}{2}\delta^{1-\beta}, \qquad \omega_{j,\mathcal{O}} \equiv \lambda_j \mu_{j,\mathcal{O}}\left(1 + \frac{\beta}{2-\beta}\eta_{-j}\left(\frac{\phi_{-j,\mathcal{O}}}{\Phi_{\mathcal{O}}}\right)^{\rho}\right), \qquad \left(\lambda_X, \lambda_M\right) = \left(\tfrac{3}{2}, \tfrac{1}{2}\right). \tag{3.43}$$

Here, the term ω_{SOLO} summarizes exporter incentives due to solo venture effect. It is increasing in the fraction of initial quality useful outside the existing relationship, which, for simplicity, is identical to the fraction of aggregate maintenance effort salvageable in case the relationship breaks down during the transit phase.[15] Holding the value of future cooperation fixed, the exporter creates higher quality goods if he expects to fetch more for the goods in

[15] Recall that none of the subsequent results hinge on this simplifying assumption.

the event of an early break in the relationship.

The value of maintaining the relationship beyond the manufacturing phase affects the choice of initial quality through the ownership-dependent term $\Omega_{\mathcal{O}} \equiv \sum_j \omega_{j,\mathcal{O}}$, where $\omega_{j,\mathcal{O}}$ measures exporter incentives through his share of $U^1_{j,\mathcal{O}}$. The term $\omega_{j,\mathcal{O}}$ comprises the direct effect, which is proportional to $\lambda_j \mu_{j,\mathcal{O}}$, and the strategic effect, which is proportional to $\lambda_j \mu_{j,\mathcal{O}} \frac{\beta}{2-\beta} \eta_{-j} \left(\frac{\phi_{-j,\mathcal{O}}}{\Phi_{\mathcal{O}}} \right)^{\rho}$.

Lemma 3.3. *Ownership affects the choice of initial quality through $\Omega_{\mathcal{O}} \equiv \sum_j \omega_{j,\mathcal{O}}$, where $\omega_{j,\mathcal{O}}$ is given in (3.43). Regardless of the importance of exporter effort, η, or the nature of transit-phase strategic interactions,* $\operatorname{sign}\{\beta - \rho\}$, *the term $\Omega_{\mathcal{O}}$ is*

1. *positive for all ownership arrangements; and*

2. *greater under exporter control.*

Part (1) of this result implies that the prospect of proceeding to the distribution phase with his current partner encourages the exporter to create higher quality goods. This holds despite the observation (Figure 3.4) that the rival strategic channel may discourage initial quality when exporter effort is particularly important. Part (2) shows that this incentive is greater whenever he controls delivery. Finally, applying Lemma 3.3 to the second-best initial quality (3.41) provides comparative statics of $z_{0,\mathcal{O}}(q)$ with respect to the contractual form $\mathcal{O}$.

Proposition 3.3. *Conditional on shipment volume, q, the exporter creates higher quality goods when controlling delivery if and only if his contribution to aggregate effort exceeds some threshold, $\eta^*_z = \eta^*_z(\psi_X/\psi_M, \rho, \beta, \delta) \in [0, 1)$. The critical value η^*_z is*

1. *increasing in β, the sensitivity of final quality to maintenance efforts relative to factory-set quality levels*

2. *zero if final quality is primarily determined by factory-set quality rather than by maintenance efforts during delivery (β is sufficiently low)*

3. *increasing in δ, the fraction of aggregate effort useful outside the existing relationship.*

3.2.2.3 Optimal second-best contract

Having analyzed the the optimal strategies in the manufacturing and distribution subgames, I now return to the optimal contract (3.20), which maximizes sales revenue less total manufacturing and distribution costs. To clarify the analysis, I first describe the dependence between each component of joint welfare and the agreed-upon volume, q.

First, sales revenue depends on the contract $(q, \mathcal{O})$ directly through the quantity sold, and indirectly through the optimal final quality. The shipment's final quality, in turn, depends on the induced aggregate effort, now written as a a function of shipping volume by substituting the optimal initial quality (3.41) into (3.26):

$$E_{\mathcal{O}}\left(q\right) \equiv E_{\mathcal{O}}\left(z_{0,\mathcal{O}}\left(q\right),q\right) = \Theta_{E,\mathcal{O}}\,\gamma q, \qquad \Theta_{E,\mathcal{O}} \equiv \left(\beta\Theta_{z0,\mathcal{O}}^{1-\beta}\Phi_{\mathcal{O}}^{2-\rho}\right)^{\frac{1}{2-\beta}} \tag{3.26'}$$

Substituting $z_{0,\mathcal{O}}\left(q\right)$ and $E_{\mathcal{O}}\left(q\right)$ into the final-quality production function (3.8) delivers final quality as a function of the agreed shipment volume

$$z_{1,\mathcal{O}}\left(q\right) \equiv z_1\left(E_{\mathcal{O}}\left(q\right)\middle|\, z_{0,\mathcal{O}}\left(q\right)\right) = \Theta_{z1,\mathcal{O}}\,\gamma q, \qquad \Theta_{z1,\mathcal{O}} \equiv \Theta_{z0,\mathcal{O}}^{1-\beta}\Theta_{E,\mathcal{O}}^{\beta} = \left(\beta^{\beta}\Theta_{z0,\mathcal{O}}^{2(1-\beta)}\Phi_{\mathcal{O}}^{\beta(2-\rho)}\right)^{\frac{1}{2-\beta}}. \tag{3.44}$$

Second, with constant returns to production, manufacturing costs, cwq, depend on the exporter-specific marginal cost and the shipment volume. Lastly, quality related costs are

$$\frac{1}{2}\psi_0\left(z_{0,\mathcal{O}}\left(q\right)\right)^2 + \sum_j \frac{1}{2}\psi_j\left(e_{j,\mathcal{O}}\left(q\right)\right)^2 = \frac{1}{2}\psi_0\left(\Theta_{z0,\mathcal{O}}\,\gamma q\right)^2 + \frac{1}{2}\Psi_{\mathcal{O}} \times \left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}}\gamma q\right)^2, \tag{3.45}$$

where the CES effort aggregator allows us to write total maintenance costs using the index $\Psi_{\mathcal{O}} \equiv \sum_j \psi_j\phi_{j,\mathcal{O}}^2$. This cost index inherits many properties from $\Phi_{\mathcal{O}}$, including

$$\Psi_{\mathcal{X}} \geq \Psi_{\mathcal{M}} \iff \Phi_{\mathcal{X}} \geq \Phi_{\mathcal{M}} \iff \eta \geq \eta^*\left(\psi_X/\psi_M,\rho\right). \tag{3.46}$$

Putting it all together, the optimal contract solves

$$\max_{q,\mathcal{O}} \left\{\left(A + \gamma z_{1,\mathcal{O}}\left(q\right) - \frac{\gamma}{L}q\right)q - \left(cwq + \frac{1}{2}\psi_0 z_{0,\mathcal{O}}\left(q\right)^2 + \sum\frac{1}{2}\psi_j e_{j,\mathcal{O}}\left(q\right)^2\right)\right\} \tag{3.20'}$$

I first characterize $q_{\mathcal{O}}$, the optimal shipping volume conditional on ownership rights, and then describe the optimal allocation of consignment rights, $\mathcal{O}^*$. Relegating the details – which follow the same steps as the first-best solution – to Section 3.D, the $\mathcal{O}-$conditional optimal shipping volume is

$$q_{\mathcal{O}} = \frac{L}{2\gamma}\frac{A - cw}{1 - L\gamma\Theta_{q,\mathcal{O}}}, \tag{3.47}$$

where

$$\Theta_{q,\mathcal{O}} \equiv \underbrace{\Theta_{z0,\mathcal{O}}^{1-\beta}\Theta_{E,\mathcal{O}}^{\beta}}_{\text{Quality–induced increase in revenue}} - \left(\underbrace{\frac{1}{2}\psi_0\Theta_{z0,\mathcal{O}}^2}_{\text{Quality creation costs}} + \underbrace{\frac{1}{2}\Psi_{\mathcal{O}} \times \left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}}\right)^2}_{\text{Quality enhancement costs}}\right) \tag{3.48}$$

is the second-best analogue to the first-best quality-effect of shipping volume defined in (3.17). It summarizes the effect of shipment volume on joint welfare through quality, and has the following properties, which allow us to compare the trade volumes under the first-best and the two ownership arrangements.

Figure 3.5: Optimal volumes, and Joint payoffs under different rights allocations

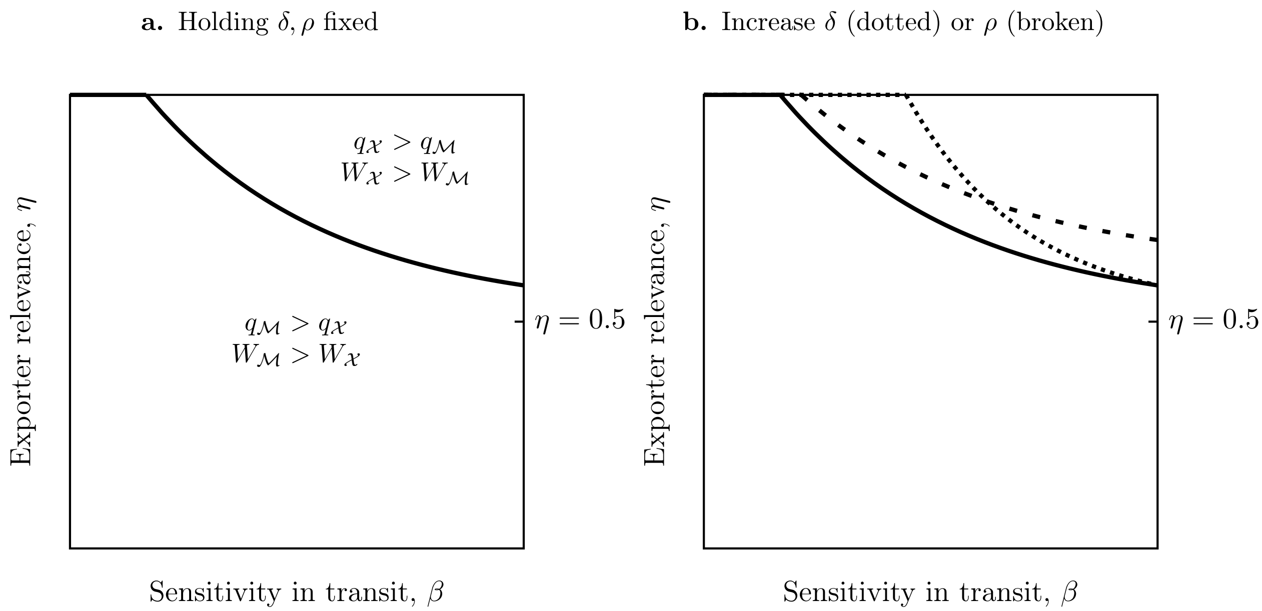

Notes: Threshold in **Panel (a)** divides (β, η)–space according to rankings of shipment volumes, $q_{\mathcal{O}}$ (see (3.47)), and joint payoffs, $W_{\mathcal{O}}$ (see (3.49)), where $\mathcal{O} = \mathcal{X}, \mathcal{M}$ indicates the party in charge of arranging delivery. Remaining parameters: (i) the importer is more efficient at maintenance; $\psi_M < \psi_X$; (ii) $\rho = 0.4$ governs substitutability among exporter and importer efforts; and (iii) $\delta = 0.3$ so that the shipment loses 70% of the relationship–specific quality in secondary markets. All else equal, exporter control is optimal whenever exporter effort is important for in–transit quality improvements. The dotted line in **Panel (b)** traces threshold for higher values of δ, the fraction of quality useful outside the relationship. The broken line plots the threshold at higher values of ρ, the substitutability of exporter and importer efforts.

Lemma 3.4. *The second-best quality-mediated effect of shipping volume on joint welfare,* $\Theta_{q,\mathcal{O}}$*, is*

1. *always positive, regardless of the importance of exporter effort,* η*;*
2. *independent of the marginal cost of physical output.*

Further, $\Theta_{q,\mathcal{X}} > \Theta_{q,\mathcal{M}}$*, if and only if exporter relevance,* η*, exceeds some threshold* η_q^**.*

Panel (a) of Figure 3.5 displays the threshold rule as a function of η, the exporter's importance for aggregate effort, and β, the relative significance of transit effort for final quality. Here, the importer is more efficient ($\psi_M < \psi_X$), quality–enhancing efforts are sufficiently substitutable ($\rho = 0.4$), and only 30 percent of the relationship–specific quality is valuable in secondary markets ($\delta = 0.3$).

Consider the limiting case as final quality becomes insensitive to enhancements during transit (as $\beta \to 0$). In this case, the parties do not bother exerting effort during delivery (3.26'), and the marginal effect of additional volume simplifies to the increase in sales revenue due to higher factory–set quality, net of the cost the exporter incurs when creating said goods

$$\Theta_{q,\mathcal{O}} \underset{\beta \to 0}{\longrightarrow} \underbrace{\Theta_{z0,\mathcal{O}}}_{\text{Quality–induced increase in revenue}} - \underbrace{\frac{1}{2}\psi_0 \Theta_{z0,\mathcal{O}}^2}_{\text{Quality creation costs}} .$$

Substituting for $\Theta_{z0,\mathcal{O}}$ from (3.42) and taking limits, $\Theta_{q,\mathcal{X}} \geq \Theta_{q,\mathcal{M}}$ whenever $\delta \leq 0$, which is impossible.[16] Intuitively, despite quality remaining fixed at its initial level, the exporter is all too eager to create high–quality goods when in control, provided he does not discount his future payoff too heavily. Unfortunately for the trading pair, the increase in sales revenues does not justify the costs of creating such high–quality goods. Subsequently, exporter–control is never optimal when perceived quality is insensitive to enhancements made during delivery.[17]

At the other extreme, if quality at the destination is wholly determined during transit ($\beta \to 1$), the shipment leaves the exporter's factory devoid of quality, and the marginal effect of additional volume consists of higher sales revenue due to higher efforts during delivery, net of the joint cost of effort during delivery

$$\Theta_{q,\mathcal{O}} \underset{\beta \to 1}{\longrightarrow} \underbrace{\Phi_{\mathcal{O}}^{2-\rho}}_{\text{Quality–induced increase in revenue}} - \underbrace{\frac{1}{2}\Psi_{\mathcal{O}} \Phi_{\mathcal{O}}^{2(1-\rho)}}_{\text{Quality enhancement costs}} .$$

[16] If the exporter discounts future payoffs by a factor $0 \leq \kappa < 1$, then

$$\Theta_{q,\mathcal{X}} \geq \Theta_{q,\mathcal{M}} \iff \delta \leq 1 - \kappa,$$

so that exporter control is optimal whenever quality is sufficiently relationship specific.

[17] By continuity, importer–control is always optimal for low values of β.

Exporter control is then optimal if transferring control raises enough revenue to offset any rise in delivery–related costs, which simplifies to

$$\Phi^{\rho}_{\mathcal{X}} - \Phi^{\rho}_{\mathcal{M}} \geq \frac{1}{2}\left(\Psi_{\mathcal{X}} - \Psi_{\mathcal{M}}\right),$$

which is more likely when exporter effort is important (η sufficiently large).[18]

Panel (b) of Figure 3.5 illustrates changes in the threshold rule in response to changes in δ, the fraction of quality valued in secondary markets, and in ρ, which indexes substitutability between exporter and importer effort. Compared to the baseline in Panel (a), exporter effort must be even more important for quality enhancement if the importer's effort is otherwise just as good (when ρ is large). Similarly, changes in the salvage value, indexed by δ, raise the η cutoff as long as quality is at least partially determined in the factory ($\beta < 1$). Finally, applying Lemma 3.4 to the expression for the optimal shipment volume (3.47) yields a familiar result from the vast literature on productivity–based sorting into various activities in international trade.

Proposition 3.4. *Conditional on being productive enough to produce for export, more productive (low c) exporters trade larger volumes under any control assignment.*

As Mrázová and Neary (Forthcoming) point out, this result follows from the simple observation that c affects the marginal returns to q solely through the marginal cost of manufacturing the physical units.

Armed with these observations, I know characterize the optimal $\mathcal{O}-$conditional volume of trade, and the optimal contract. Joint welfare under $\mathcal{O}$'s ownership, $W_{\mathcal{O}}$, consists sales revenues, less production and distribution costs:

$$\begin{aligned} W_{\mathcal{O}} &\equiv r\left(q_{\mathcal{O}}, z_{1,\mathcal{O}}\left(q_{\mathcal{O}}\right)\right) - \left[cwq_{\mathcal{O}} + \frac{1}{2}\psi_0 z_{0,\mathcal{O}}\left(q_{\mathcal{O}}\right)^2 + \sum \frac{1}{2}\psi_j e_{j,\mathcal{O}}\left(q_{\mathcal{O}}\right)^2\right] \\ &= \frac{L}{4\gamma}\frac{\left(A - cw\right)^2}{1 - L\gamma\Theta_{q,\mathcal{O}}}. \end{aligned} \tag{3.49}$$

where the equality follows from substituting the optimal shipment volume (3.47) into (3.44) to determine sales revenue, and into (3.45) to determine total costs (see 3.E). As a result, the sign of $W_{\mathcal{X}} - W_{\mathcal{M}}$ – and hence the party in charge of delivery – is determined by the same threshold behind the ranking of shipment volumes across contractual forms (see Figure 3.5). After all, with linear demand, constant marginal costs, and quadratic quality–creation and quality–upgrading costs, profits from sale of final goods – which coincide with joint welfare – are linear in output.

Proposition 3.5. *The exporter controls delivery if and only if his effort is sufficiently important, exceeding the same threshold η^*_q that determines whether the optimal volume of trade is greater under exporter control.*

[18] See Section 3.E for details.

In line with the standard result from the property-rights literature, the exporter should assume control of delivery if his effort is sufficiently more important than the importer's.

3.3 Conclusion

This paper studies the organization of international shipping when agents exert unverifiable effort in a sequential production process. It characterizes the optimal contract, and derives optimal production and quality maintenance decisions by self-interested parties subject to hold-up.

Individuals exert greater effort in the second stage when trading large volumes of highly differentiated goods. The exporter, who acts as a Stackelberg leader in the first stage, magnifies this sensitivity to volume and love-of-variety by trying to influence subsequent play in his favour. Both parties then exert too much costly effort, which often outweighs any potential offsetting sales revenue. As a result the exporter is more likely to assume control when effort is especially important in getting the goods to the destination in good condition.

Symbol	Meaning (First appearance)
Demand	
$A > 0$	Demand shifter; (Eq. (3.5)).
$\gamma > 0$	Consumer love of variety / proportional to willingness-to-pay for quality (Eq. 3.4).
Technology	
$C(q,z) = cwq + \frac{\psi_0}{2} z$	Factory costs (Eq. (3.6)).
$\frac{\psi_j}{2} e_j$	Maintenance costs.
$E(e_X, e_M) = (\eta\, e_X^{\rho} + (1-\eta)\, e_M^{\rho})^{\frac{1}{\rho}}, \quad \rho \in (0,1)$	Maintenance effort aggregator (Eq. (3.7)).
$z_1(E\|\, z_0) = z_0^{1-\beta} E^{\beta}, \qquad \beta \in (0,1)$	Final-quality production function (Eq. 3.8).
$\phi_{j,FB} \equiv (\eta_j/\psi_j)^{1/(2-\rho)}$	Share-weighted first-best individual efficiency. (Eq. (3.11)).
$\Phi_{FB} \equiv \left(\eta\, \phi_{X,FB}^{\rho} + (1-\eta)\, \phi_{M,FB}^{\rho}\right)^{\frac{1}{\rho}}$	First-best aggregate efficiency (Eq. (3.11)).
$\phi_{j,\mathcal{O}} \equiv (\mu_{j,\mathcal{O}}\eta_j/\psi_j)^{1/(2-\rho)}$	Control–adjusted, share-weighted first-best individual efficiency. (Eq. (3.27)).
$\Phi_{\mathcal{O}} \equiv \left(\eta\, \phi_{X,\mathcal{O}}^{\rho} + (1-\eta)\, \phi_{M,\mathcal{O}}^{\rho}\right)^{\frac{1}{\rho}}$	Second-best aggregate efficiency (Eq. (3.27)).
Bargaining	
$r^{IN}(E\|z_0, q) = \left(A + \gamma \cdot g(E\|z_0) - \frac{\gamma}{L} q\right) q$	Value of shipment within relationship (Eq. (3.21)).
$r^{OUT}(E\|z_0, q) = \left(A + \gamma \cdot \delta\, g(E\|z_0) - \frac{\gamma}{L} q\right) q$	Value of shipment outside relationship (Eq. (3.22)).
$\delta \in (0,1)$	Fraction of effort valuable outside relationship.
$\mu_{j,\mathcal{O}} \equiv \mathbb{1}_{\mathcal{O}=j}\, \delta + \frac{1}{2}(1-\delta)$	Control–adjustment factor (Eq. (3.25)).
$u_{j,\mathcal{O}}^{1}(e_X, e_M\|z_0, q),\; U_{j,\mathcal{O}}^{1}$	Ex-ante transit-phase payoff; corresponding value function (Eq. (3.24)).
$v_X^{0}(z_0, q)$	Exporter's factory-stage outside option (Eq. (3.32)).
W	Joint welfare (Eq. (3.49)).

Commonly used symbols

Appendix

3.A First–best shipping volume

Recall the first-best destination quality (3.16) is

$$z_{1,FB}\left(q\right)=\Theta_{z1,FB}\,\gamma q,$$

and individual efforts

$$e_{j,FB}\left(q_{FB}\right)\equiv e_{j,FB}\left(q_{FB},z_{0,FB}\left(q_{FB}\right)\right)=\left(\beta\,\Theta_{z0,FB}^{1-\beta}\,\Phi_{FB}^{\beta-\rho}\right)^{\frac{1}{2-\beta}}\;\gamma q\,\phi_{j,FB}.$$

Thus, aggregate maintenance costs (conditional on shipping volume q_{FB}) are

$$\sum_j\frac{1}{2}\psi_j e_{j,FB}\left(q_{FB}\right)^2=\frac{1}{2}\left(\left(\beta\Theta_{z0,FB}^{1-\beta}\,\Phi_{FB}^{\beta-\rho}\right)^{\frac{1}{2-\beta}}\cdot\gamma q_{FB}\right)^2\sum_j\psi_j\phi_{j,FB}^2.$$

The net marginal benefit of shipping volume on joint first-best welfare equals marginal revenue less production and distribution costs,

$$A-\frac{2\gamma}{L}q-cw+\left[\left(z_{1,FB}+q\frac{\partial z_{1,FB}}{\partial q}\right)\gamma-\frac{\partial}{\partial q}\left\{\frac{\psi_0}{2}z_{0,FB}^2+\sum_j\frac{\psi_j}{2}e_{j,FB}\right\}\right],\qquad(3.50)$$

where

$$\left(z_{1,FB}+q\frac{\partial z_{1,FB}}{\partial q}\right)\gamma=\left(\Theta_{z1,FB}\,\gamma q+q\frac{\partial\Theta_{z1,FB}\,\gamma q}{\partial q}\right)\gamma=2\Theta_{z1,FB}\,\gamma^2 q,\qquad(3.51)$$

and

$$\frac{\partial}{\partial q}\left\{\frac{\psi_0}{2}z_{0,FB}^2+\sum_j\frac{\psi_j}{2}e_{j,FB}\right\}=\left[\psi_0\left(\Theta_{z0,FB}\gamma\right)^2+\left(\gamma\beta\Phi_{FB}^{\beta-\rho}\right)^{\frac{2}{2-\beta}}\left(\Theta_{z0,FB}\gamma\right)^{\frac{2(1-\beta)}{2-\beta}}\sum_j\psi_j\phi_{j,FB}^2\right]q.\qquad(3.52)$$

Adding (3.51) and (3.52), the term in brackets in (3.50) evaluates to

$$\left(z_{1,FB} + q\frac{\partial z_{1,FB}}{\partial q}\right)\gamma - \frac{\partial}{\partial q}\left\{\frac{\psi_0}{2}z_{0,FB}^2 + \sum_j \frac{\psi_j}{2}e_{j,FB}\right\}$$
$$= \beta^\beta (1-\beta)^{1-\beta}\left(\frac{1}{\psi_0}\right)^{1-\beta}\left(1+\beta-\beta\Phi_{FB}^{-\rho}\sum_j \psi_j\phi_{j,FB}^2\right)\Phi_{FB}^{\beta(2-\rho)}\gamma^2 q. \tag{3.53}$$

Applying the definition of Φ_{FB} in (3.11),

$$\Phi_{FB}^{-\rho}\sum_j \psi_j\phi_{j,FB}^2 = \frac{\sum_j \psi_j\phi_{j,FB}^2}{\sum_j \eta_j\phi_{j,FB}^\rho} = \frac{\sum_j \psi_j\left(\eta_j/\psi_j\right)^{\frac{2}{2-\rho}}}{\sum_j \eta_j\left(\eta_j/\psi_j\right)^{\frac{\rho}{2-\rho}}} = 1,$$

so that (3.53) simplifies to $\beta^\beta (1-\beta)^{1-\beta}\left(\frac{1}{\psi_0}\right)^{1-\beta}\Phi_{FB}^{\beta(2-\rho)}\gamma^2 q$, which, in turn delivers the expression in (3.17) once we define

$$\Theta_{q,FB} \equiv \frac{1}{2}\beta^\beta (1-\beta)^{1-\beta}\left(\frac{1}{\psi_0}\right)^{1-\beta}\Phi_{FB}^{\beta(2-\rho)}. \tag{3.54}$$

Connection with Antoniades (2015)

Antoniades (2015) ignores quality changes during transit, so that quality is fixed at the factory level ($\beta \to 0$). In this scenario,

$$\Theta_{z0,FB} \to \frac{1}{\psi_0} \equiv \Theta_{z0,\text{Antoniades}}, \qquad \Theta_{q,\text{Antoniades}} \equiv \frac{1}{2\psi_0} \tag{3.55}$$

Then

$$q = \frac{L}{2\gamma}\frac{2\psi_0\left(A-cw\right)}{2\psi_0-\gamma L}, \tag{3.56}$$

where A is the marginal cost threshold between the firms that produce and those that exit. Antoniades (2015) assumes $2\psi_0 > \gamma L$ to ensure positive qualities and quantities, which is equivalent to the assumption $L\gamma\Theta_{q,FB} < 1$ in the main text.

3.B Comparing aggregate productivity across ownership structures

Recall that transferring control to a party raises that party's effective productivity. Aggregate productivity is thus greater under X-control if the resulting (weighted) gains in exporter productivity exceed the loss in importer productivity.

$$\Phi_{\mathcal{X}} > \Phi_{\mathcal{M}} \iff \eta\phi_{X,FB}\left(\mu_{X,\mathcal{X}}^{\frac{\rho}{2-\rho}} - \mu_{X,\mathcal{M}}^{\frac{\rho}{2-\rho}}\right) > (1-\eta)\,\phi_{M,FB}\left(\mu_{M,\mathcal{M}}^{\frac{\rho}{2-\rho}} - \mu_{M,\mathcal{X}}^{\frac{\rho}{2-\rho}}\right).$$

If exporter and importer marginal costs differ ($\psi_X \neq \psi_M$), the above condition delivers a quadratic equation in η, whose solution

$$\eta^*_\Phi \left(\psi_X/\psi_M, \rho\right) \equiv \frac{\sqrt{(\psi_X/\psi_M)^\rho} - (\psi_X/\psi_M)^\rho}{1-(\psi_X/\psi_M)^\rho}, \qquad \psi_X/\psi_M \neq 1,$$

is increasing in ψ_X/ψ_M and increasing in ρ. If $\psi_X = \psi_M$, then define $\eta^*_\Phi(1,\rho) \equiv 1/2$, so that aggregate productivity is higher under X-control if and only if X effort is more important.

Continuum of ownership arrangements

This section abstracts from the all-or-nothing ownership structure in the main text, instead allowing a continuum of possible ownership arrangements, indexed by the exporter's share of the renegotiation surplus, $\tau \in [0,1]$. Let

$$\mu_{j,\tau} \equiv \begin{cases} \tau\delta + \frac{1}{2}(1-\delta) & j = X \\ (1-\tau)\,\delta + \frac{1}{2}(1-\delta) & j = M \end{cases}$$

be j's effective share under arrangement τ, and let

$$\Phi_\tau \equiv \left(\eta\, \phi^\rho_{X,\tau} + (1-\eta)\; \phi^\rho_{M,\tau}\right)^{\frac{1}{\rho}}, \qquad \phi_{j,\tau} \equiv \left(\mu_{j,\tau}\frac{\eta_j}{\psi_j}\right)^{\frac{1}{2-\rho}}$$

be aggregate and individual ownership-adjusted efficiencies. A simple threshold rule determines whether increasing the exporter's share increases control-adjusted efficiency.

Lemma 3.5. *Φ_τ is increasing in τ for $\tau < \tau^*$, and decreasing in τ for $\tau \geq \tau^*$, where the unique threshold exporter share, $\tau^* = \tau^*(\psi_X/\psi_M, \eta)$, is*

1. *increasing in exporter relevance, η, and*
2. *decreasing in exporter's relative marginal cost of effort, ψ_X/ψ_M.*

Proof. Ownership-adjusted productivity, Φ_τ, is differentiable in τ, so that

$$\frac{\partial \Phi_\tau}{\partial \tau} = \frac{\delta}{2-\rho}\Phi_\tau^{1-\rho}\left[\eta\frac{\phi^\rho_{X,\tau}}{\mu_{X,\tau}} - (1-\eta)\frac{\phi^\rho_{M,\tau}}{\mu_{M,\tau}}\right],$$

which is positive if and only if

$$\left(\frac{\eta}{1-\eta}\right)^2 \left(\frac{\psi_M}{\psi_X}\right)^\rho > \left(\frac{\tau\delta + \frac{1}{2}(1-\delta)}{(1-\tau)\,\delta + \frac{1}{2}(1-\delta)}\right)^{2-2\rho}.$$

Figure 3.B.1 plots these terms as functions of τ. With the exception of ρ (effort substitutability), which appears on both sides of the inequality, this expression compares (i) on

Figure 3.B.1: Ownership-adjusted productivity

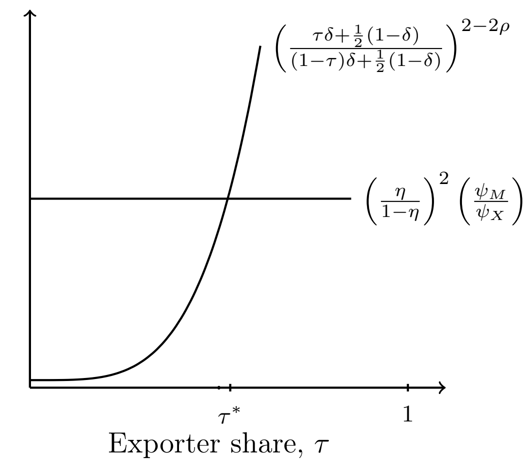

Notes:

the left-hand side, the effects of shipping technology (exporter relevance, η, and individual marginal costs ψ_j); and (ii) on the right-hand side, the effects of contractual incompleteness (relationship specificity, $1-\delta$, and the exporter's share of renegotiation surplus, τ). Since LHS is increasing in the exporter's relative importance, and the importer's relative *in*efficiency, it is possible that $LHS < RHS$ for sufficiently low η or large ψ_X/ψ_M. Intuitively, the gains from transfer additional control to X decrease as his effort becomes inconsequential, or if the importer is more efficient at the margin. In either of these cases, define $\tau^* = 0$. □

3.C Second–best initial quality

Solving for optimal factory quality

The marginal cost of initial quality is $\psi_0 z_0$, while the marginal benefit is

$$\frac{(1-\beta)\gamma q}{z_0^{\beta}}\left(\Omega_{\mathcal{O}}\, E_{\mathcal{O}}(q,z_0)^{\beta}+\omega_{SOLO}\, E_{SOLO}(q,z_0)^{\beta}\right), \tag{3.57}$$

where $E_{\mathcal{O}}(q,z_0)$ and $E_{SOLO}(q,z_0)$, the aggregate maintenance efforts under a joint transit phase and X's solo venture, are given in (3.34) and (3.26), and $\Omega_{\mathcal{O}} \equiv \sum_j \omega_{j,\mathcal{O}}$.

The factors

$$\omega_{SOLO} \equiv \frac{1}{2}\delta^{1-\beta}, \qquad \omega_{j,\mathcal{O}} \equiv \lambda_j \mu_{j,\mathcal{O}}\left(1+\frac{\beta}{2-\beta}\eta_{-j}\left(\frac{\phi_{-j,\mathcal{O}}}{\Phi_{\mathcal{O}}}\right)^{\rho}\right).$$

summarize the effects of the immediate outside option and future cooperation. Substituting

$E_{\mathcal{O}}(q, z_0)$ and $E_{SOLO}(q, z_0)$,

$$\psi_0 z_0 = (1-\beta)\left((\beta)^\beta (\gamma q)^2 z_0^{-\beta}\right)^{\frac{1}{2-\beta}} \left(\omega_{SOLO} \Phi_{SOLO}^{\frac{\beta(2-\rho)}{2-\beta}} + \Omega_{\mathcal{O}} \Phi_{\mathcal{O}}^{\frac{\beta(2-\rho)}{2-\beta}}\right),$$

and then solving for z_0 yields (3.41) in the main text.

Proof of Proposition 3.3 (comparing initial quality across ownership structures)

Conditional on shipment volume, q, the first-best, and $\mathcal{O}-$controlled factory qualities are

$$z_{0,FB}(q) = \Theta_{z0,FB}\,\gamma q, \qquad z_{0,\mathcal{O}}(q) = \Theta_{z0,\mathcal{O}}\,\gamma q, \quad \mathcal{O} \in \mathcal{X}, \mathcal{M}.$$

Comparing interior solutions,

$$z_{0,FB} > z_{0,\mathcal{O}} \iff \omega_{SOLO}\left(\frac{\Phi_{SOLO}}{\Phi_{FB}}\right)^{\frac{\beta(2-\rho)}{2-\beta}} + \Omega_{\mathcal{O}}\left(\frac{\Phi_{\mathcal{O}}}{\Phi_{FB}}\right)^{\frac{\beta(2-\rho)}{2-\beta}} < 1,$$

and

$$z_{0,\mathcal{X}} > z_{0,\mathcal{M}} \iff \left(\frac{\Phi_{\mathcal{X}}}{\Phi_{\mathcal{M}}}\right)^{\frac{\beta(2-\rho)}{2-\beta}} > \frac{\Omega_{\mathcal{M}}}{\Omega_{\mathcal{X}}}.$$

Conditional on shipment volume, differences in factory quality are independent of the marginal cost of initial quality, ψ_0, and the solo venture effect (which is independent of ownership during delivery). Therefore,

$$\begin{aligned} \operatorname{sign}\{z_{0,\mathcal{X}}(q) - z_{0,\mathcal{M}}(q)\} &= \operatorname{sign}\left\{\Phi_{\mathcal{X}}^{\frac{\beta(2-\rho)}{2-\beta}}\Omega_{\mathcal{X}} - \Phi_{\mathcal{M}}^{\frac{\beta(2-\rho)}{2-\beta}}\Omega_{\mathcal{M}}\right\} \\ &= \operatorname{sign}\left\{\Omega_{\mathcal{M}}\left(\Phi_{\mathcal{X}}^{\frac{\beta(2-\rho)}{2-\beta}} - \Phi_{\mathcal{M}}^{\frac{\beta(2-\rho)}{2-\beta}}\right) + \Phi_{\mathcal{X}}^{\frac{\beta(2-\rho)}{2-\beta}}(\Omega_{\mathcal{X}} - \Omega_{\mathcal{M}})\right\}. \end{aligned}$$

The first term is positive whenever the exporter's share of aggregate effort exceeds some threshold $\eta_\Phi^*(\psi_X/\psi_M, \rho)$, and increasing in η, as demonstrated in Section 3.B, while the second term is always positive. Figure 3.5 depicts the threshold rule in (β, η)–space.

Lemma 3.6. *Conditional on shipping volume, q, the exporter creates higher quality goods when in charge of shipping if aggregate productivity is higher under his control (i.e., if $\eta > \eta_\Phi^*(\psi_X/\psi_M, \rho)$).*

This is a sufficient, but not necessary condition for higher quality goods under exporter-control. Suppose $\eta < \eta_\Phi^*(\psi_X/\psi_M, \rho)$, so that $\Phi_{\mathcal{X}}^{\beta(2-\rho)/(2-\beta)} - \Phi_{\mathcal{M}}^{\beta(2-\rho)/(2-\beta)} < 1$. Then there exists $\eta_z^* \equiv \eta_z^*(\psi_X/\psi_M, \rho, \beta, \delta) < \eta_\Phi^*(\psi_X/\psi_M, \rho)$ such that

$$\eta \in [\eta_z^*, \eta_\Phi^*] \implies \Omega_{\mathcal{M}}\left(\Phi_{\mathcal{M}}^{\frac{\beta(2-\rho)}{2-\beta}} - \Phi_{\mathcal{X}}^{\frac{\beta(2-\rho)}{2-\beta}}\right) \leq \Phi_{\mathcal{X}}^{\frac{\beta(2-\rho)}{2-\beta}}(\Omega_{\mathcal{X}} - \Omega_{\mathcal{M}}),$$

with equality at $\eta = \eta_z^*$.

3.D Second–best shipment volume

Assuming an interior solution, $q_{\mathcal{O}}$ solves

$$p\left(q, z_{1,\mathcal{O}}\left(q\right)\right) + q\frac{\partial p\left(q, z_{1,\mathcal{O}}\left(q\right)\right)}{\partial q} = \frac{\partial}{\partial q}\left[cwq + \frac{1}{2}\psi_0 z_{0,\mathcal{O}}\left(q\right)^2 + \sum \frac{1}{2}\psi_j e_{j,\mathcal{O}}\left(q\right)^2\right], \quad (3.58)$$

where

$$\begin{array}{llcl} (1) & z_{1,\mathcal{O}}\left(q\right) & = & \left(\left(\beta\right)^{\beta} \Theta_{z0,\mathcal{O}}^{2(1-\beta)} \Phi_{\mathcal{O}}^{\beta(2-\rho)}\right)^{\frac{1}{2-\beta}} \gamma q \\ (2) & p\left(q, z_{1,\mathcal{O}}\left(q\right)\right) & \equiv & A + \gamma z_{1,\mathcal{O}}\left(q\right) - \frac{\gamma}{L} q \\ (3) & \sum \frac{1}{2}\psi_j e_{j,\mathcal{O}}\left(q\right)^2 & = & \frac{\Psi_{\mathcal{O}}}{2}\left(\frac{E_{\mathcal{O}}\left(q\right)}{\Phi_{\mathcal{O}}}\right)^2 = \frac{\Psi_{\mathcal{O}}}{2}\left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}}\right)^2 \left(\gamma q\right)^2 \\ (4) & \Psi_{\mathcal{O}} & \equiv & \sum_j \psi_j \phi_{j,\mathcal{O}}^2. \end{array}$$

The left-hand side of (3.58) is marginal revenue, evaluated at the optimal final quality $z_{1,\mathcal{O}}$. The right hand-side is the marginal cost of shipping volume, which, in addition to the direct effect on factory costs, cwq, also incorporates the effects on downstream efforts.

To facilitate comparison with the first-best outcome, gather terms so that we can perform the standard marginal benefit vs (direct) marginal cost

$$A - \frac{2\gamma}{L}\left(1 - L\gamma\Theta_{q,\mathcal{O}}\right) q = cw, \quad (3.59)$$

where

$$\Theta_{q,\mathcal{O}} \equiv \Theta_{z0,\mathcal{O}}^{1-\beta}\Theta_{E,\mathcal{O}}^{\beta} - \frac{1}{2}\left(\psi_0\Theta_{z0,\mathcal{O}}^2 + \Psi_{\mathcal{O}} \times \left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}}\right)^2\right). \quad (3.60)$$

Solving for for q,

$$q_{\mathcal{O}} = \frac{L}{2\gamma}\frac{A - cw}{1 - L\gamma\Theta_{q,\mathcal{O}}}.$$

3.E Optimal ownership

This section expresses joint welfare solely as a function of ownership, $\mathcal{O}$, by first expressing welfare as a function of the volume shipped. From (3.44)

$$z_{1,\mathcal{O}}\left(q\right) = \Theta_{z1,\mathcal{O}}\gamma q, \qquad \Theta_{z1,\mathcal{O}} \equiv \Theta_{z0,\mathcal{O}}^{1-\beta}\Theta_{E,\mathcal{O}}^{\beta} = \left(\beta^{\beta}\Theta_{z0,\mathcal{O}}^{2(1-\beta)}\Phi_{\mathcal{O}}^{\beta(2-\rho)}\right)^{\frac{1}{2-\beta}}.$$

Then final sales revenue is

$$r_{\mathcal{O}} \equiv r\left(q_{\mathcal{O}}, z_1\left(E_{\mathcal{O}}\left(q_{\mathcal{O}}\right) \middle| z_{0,\mathcal{O}}\left(q_{\mathcal{O}}\right)\right)\right) = \left(A + \gamma\Theta_{z1,\mathcal{O}}\gamma q_{\mathcal{O}} - \frac{\gamma}{L}q_{\mathcal{O}}\right) q_{\mathcal{O}}$$

so that joint welfare is given by

$$\begin{aligned} W_{\mathcal{O}} &\equiv r_{\mathcal{O}} - \left(cwq_{\mathcal{O}} + \frac{1}{2}\psi_0 \left[z_{0,\mathcal{O}} \left(q_{\mathcal{O}} \right) \right]^2 + \sum \frac{1}{2}\psi_j \left[e_{j,\mathcal{O}} \left(q_{\mathcal{O}} \right) \right]^2 \right) \\ &= \left(A + \gamma \Theta_{z1,\mathcal{O}} \gamma q_{\mathcal{O}} - \frac{\gamma}{L} q_{\mathcal{O}} \right) q_{\mathcal{O}} \\ &\quad - \left(cwq_{\mathcal{O}} + \frac{1}{2} \left(\psi_0 \left(\Theta_{z0,\mathcal{O}} \gamma q_{\mathcal{O}} \right)^2 + \Psi_{\mathcal{O}} \times \left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}} \gamma q_{\mathcal{O}} \right)^2 \right) \right) \\ &= \left[A - cw - \frac{\gamma}{L} \left(1 - L\gamma \Theta_{q,\mathcal{O}} \right) q_{\mathcal{O}} \right] q_{\mathcal{O}}, \end{aligned}$$

where

$$\Theta_{q,\mathcal{O}} \equiv \Theta_{z1,\mathcal{O}} - \frac{1}{2} \left(\psi_0 \Theta_{z0,\mathcal{O}}^2 + \Psi_{\mathcal{O}} \times \left(\frac{\Theta_{E,\mathcal{O}}}{\Phi_{\mathcal{O}}} \right)^2 \right).$$

Assuming interior solutions under both ownership structures, $\min \{q_{\mathcal{X}}, q_{\mathcal{M}}\} > 0$, and substituting the optimal shipping volume into the term in brackets,

$$\begin{aligned} W_{\mathcal{O}} &\equiv \left[A - cw - \frac{\gamma}{L} \left(1 - L\gamma \Theta_{q,\mathcal{O}} \right) \times \frac{L}{2\gamma} \frac{A - cw}{1 - L\gamma \Theta_{q,\mathcal{O}}} \right] q_{\mathcal{O}} \\ &= \frac{1}{2} \left(A - cw \right) q_{\mathcal{O}} \\ &= \frac{1}{2} \left(A - cw \right) \frac{L}{2\gamma} \frac{A - cw}{1 - L\gamma \Theta_{q,\mathcal{O}}} \\ &= \frac{L}{4\gamma} \frac{\left(A - cw \right)^2}{1 - L\gamma \Theta_{q,\mathcal{O}}}. \end{aligned}$$

To compare welfare across ownership structures, factor $L \left(A - cw \right)^2 / 4\gamma \geq 0$, so that

$$W_{\mathcal{X}} > W_{\mathcal{M}} \iff \Theta_{q,\mathcal{X}} > \Theta_{q,\mathcal{M}},$$

that is, the parties are jointly better off under exporter control if shipping volumes are greater under exporter control. This tight relationship between joint payoffs and shipment volumes follows because joint welfare equals profits, which, with linear demand, are linear in output.

Limiting behaviour as $\beta \to 1$ If quality at the destination is wholly determined during transit ($\beta \to 1$), exporter control is then optimal if

$$\Phi_{\mathcal{X}}^{\rho} - \Phi_{\mathcal{M}}^{\rho} \geq \frac{1}{2} \left(\Psi_{\mathcal{X}} - \Psi_{\mathcal{M}} \right).$$

Substituting for $\Phi_{\mathcal{O}}$ and $\Psi_{\mathcal{O}}$, this condition is equivalent to

$$\sum_j \eta_j \left(\frac{\eta_j}{\psi_j} \right)^{\frac{\rho}{2-\rho}} \left(\mu_{j,\mathcal{X}}^{\frac{\rho}{2-\rho}} - \mu_{j,\mathcal{M}}^{\frac{\rho}{2-\rho}} \right) > \sum_j \frac{\psi_j}{2} \left(\frac{\eta_j}{\psi_j} \right)^{\frac{2}{2-\rho}} \left(\mu_{j,\mathcal{X}}^{\frac{2}{2-\rho}} - \mu_{j,\mathcal{M}}^{\frac{2}{2-\rho}} \right).$$

All else equal, this is more likely when exporter effort is important (η large).

Contract-specific fixed costs

According to the optimal-$\mathcal{O}$ rule "$W_{\mathcal{X}} > W_{\mathcal{M}} \iff \Theta_{q,\mathcal{X}} > \Theta_{q,\mathcal{M}}$", the degree of product differentiation, γ, and destination market conditions, (L, A), do not affect the choice of ownership, conditional on all other model parameters.

Consider instead

$$W_{\mathcal{O}} = \frac{L}{4\gamma} \frac{(A - cw)^2}{1 - L\gamma\Theta_{q,\mathcal{O}}} - f_{\mathcal{O}},$$

where $f_{\mathcal{O}} > 0$ is an ownership-specific fixed cost that may vary across buyer-seller pairs. Then

$$W_X > W_M \iff \frac{\Theta_{q,\mathcal{X}} - \Theta_{q,\mathcal{M}}}{(1 - L\gamma\Theta_{q,\mathcal{X}})(1 - L\gamma\Theta_{q,\mathcal{M}})} > \frac{4(f_{\mathcal{X}} - f_{\mathcal{M}})}{[L(A - cw)]^2}.$$

Setting $f_{\mathcal{X}} = f_{\mathcal{M}}$ delivers the previous result.

Bibliography

Adams, A., R. Blundell, M. Browning, and I. Crawford (2015): "Prices versus Preferences: Taste Change and Revealed Preference," Working paper, University College London.

Afriat, S. N. (1967): "The construction of utility functions from expenditure data," *International Economic Review*, 8, 67–77.

Alfaro, L., P. Antràs, D. Chor, and P. Conconi (Forthcoming): "Internalizing Global Value Chains: A Firm-Level Analysis," *Journal of Political Economy*.

Anderson, J. E. and E. van Wincoop (2004): "Trade Costs," *Journal of Economic Literature*, 42, 691–751.

Antoniades, A. (2015): "Heterogeneous firms, quality, and trade," *Journal of International Economics*, 95, 263–273.

Antràs, P. (2003): "Firms, Contracts, And Trade Structure," *The Quarterly Journal of Economics*, 118, 1375–1418.

Antràs, P. and D. Chor (2013): "Organizing the global value chain," *Econometrica*, 81, 2127–2204.

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

Bagnoli, M. and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic theory*, 26, 445–469.

Bailey, E. E. (1981): "Contestability and the design of regulatory and antitrust policy," *The American Economic Review*, 71, 178–183.

Bailey, E. E. and J. C. Panzar (1981): "The Contestability of Airline Markets During the Transition to Deregulation," *Law and Contemporary Problems*, 44, 125–146.

Banco Central del Ecuador (2013): "Boletín Anuario," Retrieved 11/10/2018 from https://www.bce.fin.ec/index.php/component/k2/item/327-ver-boletín-anuario-por-años.

Baumol, W., J. Panzar, and R. Willig (1982): "Contestable markets and the theory of market structure," *Nueva York, Harcourt Brace Javanovich, Inc.*

Behrens, K., C. Gaigné, and J.-F. Thisse (2009): "Industry location and welfare when transport costs are endogenous," *Journal of Urban Economics*, 65, 195–208.

Behrens, K. and P. M. Picard (2011): "Transportation, freight rates, and economic geography," *Journal of International Economics*, 85, 280–291.

Bernard, A. B., A. Moxnes, and K. H. Ulltveit-Moe (2018): "Two-sided heterogeneity and trade," *Review of Economics and Statistics*, 100, 424–439.

Berry, S., A. Gandhi, and P. Haile (2013): "Connected substitutes and invertibility of demand," *Econometrica*, 81, 2087–2111.

Bown, C. P. and M. A. Crowley (2016): "The empirical landscape of trade policy," *Handbook of Commercial Policy*, 1, 3–108.

Brancaccio, G., M. Kalouptsidi, and T. Papageorgiou (2017): "Geography, search frictions and endogenous trade costs," Tech. rep., National Bureau of Economic Research.

Brown, D. J. and R. L. Matzkin (1996): "Testable restrictions on the equilibrium manifold," *Econometrica*, 1249–1262.

BRS (2008): "Shipping and Shipbuilding Markets in 2007," Retrieved 02/12/2017 from http://www.brsbrokers.com/review_archives.php.

Bulow, J. I., J. D. Geanakoplos, and P. D. Klemperer (1985): "Multimarket oligopoly: Strategic substitutes and complements," *Journal of Political economy*, 93, 488–511.

Carvajal, A., R. Deb, J. Fenske, and J. K.-H. Quah (2013): "Revealed preference tests of the Cournot model," *Econometrica*, 81, 2351–2379.

——— (2014): "A nonparametric analysis of multi-product oligopolies," *Economic Theory*, 57, 253–277.

Chambers, C. P. and F. Echenique (2016): *Revealed preference theory*, vol. 56, Cambridge University Press.

Cheng, L. (1985): "Inverting systems of demand functions," *Journal of Economic Theory*, 37, 202–210.

Cosar, A. K. and B. Demir (2018): "Shipping inside the box: Containerization and trade," *Journal of International Economics*, 114, 331–345.

Davies, J. E. (1986): "Competition, contestability and the liner shipping industry," *Journal of Transport Economics and Policy*, 299–312.

Deb, R. and J. Fenske (2009): "A Nonparametric Test of Strategic Behavior in the Cournot Model," Working paper.

Diesenreiter, F. and E. Tromborg (2009): "World Biofuel Maritime Shipping Study,"
.

Eckel, C. and J. P. Neary (2010): "Multi-product firms and flexible manufacturing in the global economy," *The Review of Economic Studies*, 77, 188–217.

Fally, T. and R. Hillberry (Forthcoming): "A Coasian model of international production chains," *Journal of International Economics*.

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

Friedt, F. L. and W. W. Wilson (2017): "Trade, Transportation and Trade Imbalances: An Empirical Examination of International Markets and Backhauls," Working paper, University of Oregon.

Global Value Chain Initiative (2017): Retrieved 12/1/2017 from https://globalvaluechains.org/concept-tools.

Grossman, G. M. and E. Helpman (2002): "Integration versus Outsourcing in Industry Equilibrium," *The Quarterly Journal of Economics*, 117, 85–120.

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

Hart, O. and J. Moore (1990): "Property Rights and the Nature of the Firm," *Journal of Political Economy*, 98, 1119–58.

Head, K. and T. Mayer (2014): "Gravity equations: Workhorse, toolkit, and cookbook," in *Handbook of International Economics*, Elsevier, vol. 4, 130–195.

Hoffmann, J. (2009): "Shipping out of the economic crisis," *Brown J. World Aff.*, 16, 121.

Hummels, D., V. Lugovskyy, and A. Skiba (2009): "The trade reducing effects of market power in international shipping," *Journal of Development Economics*, 89, 84–97.

Ishikawa, J. and N. Tarui (2018): "Backfiring with backhaul problems: Trade and industrial policies with endogenous transport costs," *Journal of International Economics*, 111, 81–98.

Kalgora, B. and T. M. Christian (2016): "The financial and economic crisis, its impacts on the shipping industry, lessons to learn: the container-ships market analysis," *Open Journal of Social Sciences*, 4, 38.

Klein, P. G. (2008): "The make-or-buy decisions: Lessons from empirical studies," in *Handbook of new institutional economics*, Springer, 435–464.

Lafontaine, F. and M. Slade (2007): "Vertical Integration and Firm Boundaries: The Evidence," *Journal of Economic Literature*, 45, 629–685.

Malbon, J. and B. Bishop (2014): *Australian Export*, Cambridge University Press.

Malfliet, J. (2011): "Incoterms 2010 and the mode of transport: how to choose the right term," in *Management Challenges in the 21st Century: Transport and Logistics: Opportunity for Slovakia in the Era of Knowledge Economy*, City University of Seattle Bratislava, 163–179.

Marcet, J. J. and A. de Ochoa Martínez (2006): *The Handbook of Logistics Contracts: A Practical Guide to a Growing Field*, Palgrave Macmillan.

Marin, D. and T. Verdier (2003): "Globalization and the new enterprise," *Journal of the European Economic Association*, 1, 337–344.

Matsukawa, I. (2016): "An Application of a Revealed Preference Test of the Cournot Model to a Retail Electricity Market: Evidence from Japan," .

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

Melitz, M. J. (2003): "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 71, 1695–1725.

Melitz, M. J. and G. I. Ottaviano (2008): "Market size, trade, and productivity," *The Review of Economic Studies*, 75, 295–316.

Mesquita Moreira, M., C. Volpe Martincus, and J. S. Blyde (2008): *Unclogging the arteries: the impact of transport costs on Latin American and Caribbean trade*, Inter-American Development Bank.

Mrázová, M. and J. P. Neary (Forthcoming): "Selection effects with heterogeneous firms," *Journal of the European Economic Association*.

Nishimura, H., E. A. Ok, and J. K.-H. Quah (2017): "A comprehensive approach to revealed preference theory," *American Economic Review*, 107, 1239–63.

Okuguchi, K. (1987): "Equilibrium prices in the Bertrand and Cournot oligopolies," *Journal of Economic Theory*, 42, 128–139.

Ramberg, J. (2011): *ICC Guide to Incoterms 2010: Understanding and Practical Use*, International Chamber of Commerce.

Ray, R. and S. Kozameh (2012): "Ecuador's Economy since 2007," .

Samaras, I. and E. Papadopoulou (2010): "The global financial crisis–The effects on the liner shipping industry and the newly adopted leading practices," in *1st Olympus International Conference on Supply Chains*, 1–2.

Schwarz, C. and J. Suedekum (2014): "Global sourcing of complex production processes," *Journal of International Economics*, 93, 123–139.

Stopford, M. (2009): *Maritime Economics*, Routledge.

Tye, W. B. (1985): "The Applicability of the Theory of Contestable Markets to Rail/Water Carrier Mergers," *Logistics and Transportation Review*, 21.

UNCTAD (2017): *Review of maritime transport*, United Nations Publications.

Varian, H. R. (1985): "Non-parametric analysis of optimizing behavior with measurement error," *Journal of Econometrics*, 30, 445–458.

——— (2006): "Revealed preference," *Samuelsonian economics and the twenty-first century*, 99–115.

Wong, W. F. (2017): "The Round Trip Effect: Endogenous Transport Costs and International Trade," Working paper, University of Oregon.