

Essays on International Trade: Demand Volatility, Exporting and the Adoption of Innovations

by

Michael Adetayo Olabisi

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Business Administration)
in The University of Michigan
2015

Doctoral Committee:

Professor Jan Svejnar, Co-chair
Associate Professor Jagadeesh Sivadasan, Co-chair
Associate Professor Allan N. Afuah
Assistant Professor Kyle Handley
Professor Scott E. Masten
Professor Jeffrey A. Smith

This is what you must be like. Grow wherever life puts you down.

— Ben Okri, *The Famished Road*

Il n'est pas certain que tout soit incertain.

— Blaise Pascal, *Pensées*

© Michael Adetayo Olabisi 2015

All Rights Reserved

ACKNOWLEDGEMENTS

I had the privilege of an outstanding dissertation committee. I must thank Jagadeesh Sivadasan and Jan Svejnar - my committee chairs, for their patience, encouragement and support through the journey. Jagadeesh was consistently energetic and excited about my work. It is impossible to quantify the value of leaving a meeting with your advisor feeling like a better person. His guidance was also vital to obtaining data for the first dissertation chapter. Jan remains a role model for doing good research while remaining relevant. His sound, well-refined advice improved the scope of my work and career strategy in no small way.

For his detailed comments, careful readings of various dissertation chapters and unwavering support, Jeff Smith deserves high recognition. I remain indebted for his time and insight, and for his interest in my development as a thinking person. His combination of personal warmth and attention to detail is one of the most potent testimonials to the possibility of a joyous life in academia. Kyle Handley's brilliant insights helped this dissertation. His keen questions provided food for thought that I plan to use beyond these chapters.

The third chapter of this dissertation started out as a co-authored paper with Scott Masten. His genius deserves greater recognition. Outside research, Scott's humor and wit encouraged me to consider the possibility that economics should be a fundamentally human enterprise. Allan Afuah has been more than just a mentor. His consistent advice helped me to draw on a rich well of experience that I could not reach otherwise. In many ways, he blazed the trail and made it possible for people like me to rise to new opportunities.

My first mentor in the doctoral program was Kathy Terrell. It hurts that death took her before this dissertation was completed. May academics like her never be few. May the joy and energy she shared persist into the next generation and beyond.

For further comments and encouragement I thank my colleagues from the Business Economics (GIBBERISH) Seminar, the Informal Development Seminar at the Ford School of Public Policy, the Michigan Economics International Macro Lunch and the International Policy Center's Economic Development Seminar. Notable commenters include: Achyuta Ad-

hvaryu, Manuela Angelucci, Raj Arunachalam, Tom Buchmueller, Anne Fitzpatrick, Maggie Levenstein, Johannes Norling, Christian Pröbsting, Rishi Sharma, Ajay Shenoy, Alex Persaud, Hope Thompson, Dean Yang and Xu Zhang.

Financial support for data acquisition and travel came from the Doctoral Students Office of the Ross Business School, the Department of Business Economics and Public Policy, and the Center for International Business Education and Research. The third chapter of the dissertation was developed in the course of a visit to the National University of Singapore, sponsored by the National Science Foundation.

Finally, faith, family and friends were the substance of this journey. I was consistently refreshed by the times spent with friends - friends in the doctoral program: Stephanie, Marek, Kristina, Santhosh, Jenny, Suntae, Matt, Bea, Daniel, Laura... and friends outside the doctoral program: Massy, David, Anne. Family was where I found the courage to have faith. I must recognize my co-traveler Laura for years of love and patience, as I must recognize Elijah our son. Kudos to my first teachers, Bayo and Anne Olabisi, my unwavering source of support, Lorenz and Becky Schmitt, and my band of brothers Dayo, Kayode, Bayode, Adeola and Benjamin Olabisi. Faith is the soil on which my uncertainties rested and seeded into hope, before hope sprouted into the will to persist.

This work is dedicated to the generous souls and the generous acts that made the dissertation possible.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF FIGURES	vii
LIST OF TABLES	viii
ABSTRACT	x
CHAPTER	
I. Demand Volatility and Export Entry	1
1.1 Introduction	1
1.2 Model	5
1.2.1 Adjustment Costs with Stochastic Demand	5
1.2.2 Export Entry	10
1.3 Empirics	15
1.3.1 Data and Definitions	15
1.3.2 Results	19
1.3.3 Robustness Checks	34
1.4 Discussions and Conclusion	39
Appendices	47
I.A Empirics: Data, Variable Definitions and Supplementary Tests	47
I.A.1 Exporter-Level Data	47
I.A.2 Product-Country Destination Data	49
I.A.3 Product and Country Variation in Demand Volatility	51
I.A.4 Annual Exporter Counts and Demand Volatility	53
I.A.5 Other Measures of Demand Volatility	53
I.A.6 Comparing Demand Volatility with Other Predictors	57
I.A.7 GDP per Capita and the Volatility of Imports	60
I.A.8 Demand Volatility and Prices	60
I.B Model Features	63

I.B.1	The Envelope Theorem Allows Optimal $q^* = E(q)$	63
I.B.2	Trade Volumes with Demand Volatility	65
I.B.3	Demand Volatility and Exporter Size Thresholds	67
II.	The Impact of Exporting and Foreign Investment on Product Innovation: Evidence from Chinese Manufacturers	72
2.1	Introduction	72
2.2	Related Literature	75
2.2.1	Product Innovation	75
2.2.2	Exporting and Product Innovation	76
2.2.3	FDI and Product Innovation	76
2.2.4	Exporting and FDI's Effects on Product Innovation	77
2.3	Methods, Data and Results	77
2.3.1	Data	78
2.3.2	Baseline Estimates - OLS	81
2.3.3	Tests that Controls for Selection into Exporting or FDI	84
2.3.4	Tests on Subset of Data: Foreign-Owned Firms Only	89
2.3.5	Mechanisms for Product Innovation	90
2.3.6	Learning with Corrected Biases	98
2.3.7	Other Empirical Specifications	100
2.4	Conclusions	101
Appendices	107
II.A	Covariate Balancing and Common Support	107
II.B	Product Innovation Before and After Exporting or FDI	107
II.C	More Tests with FDI Defined to Include HMT	112
III.	Bridging the Enforcement Gap in International Trade: Explaining Participation in the New York Convention on Arbitration	116
3.1	Introduction	117
3.2	Background and Related Literature	123
3.2.1	Background on the New York Convention	123
3.2.2	Related Literature	126
3.3	Model	127
3.3.1	Contagion	128
3.3.2	Treaty Membership Motivated by Trade	131
3.4	Empirics	133
3.4.1	Baseline Estimates of Government Responsiveness γ	134
3.4.2	Trade-Driven and Interaction-Driven Membership τ vs. κ	135
3.4.3	Country-Level Differences in Trade-Driven Membership	136

3.5 Conclusion	143
Appendices	149
III.A NYC Treaty Membership and Delays to Entry	149
III.A.1 Entry Date: Signature, Ratification, Accession and In-Force Dates	149
III.A.2 Entry Delay: Lag from Eligible Date to Signature Date . . .	149
III.A.3 Results for Countries Eligible in 1958	154
III.B Data	154
III.B.1 Data Sources	154
III.B.2 Explaining NYC Membership with RTA Membership	156
III.B.3 Trade with NYC members	156
III.B.4 The US case	160
III.C NYC Membership's Effect on Trade	161
III.D The NYC Membership Curve is S-Shaped	162

LIST OF FIGURES

Figure

1.1	Demand with Stochastic Shocks	7
1.2	Average Chinese Exporter Counts and Demand Volatility	22
1.3	Minimum Export Thresholds and Demand Volatility	29
1.4	Predictions for Weighted Demand Volatility	35
I.A.1	Destinations Served (and Not Served) by Chinese Exporters	50
I.A.2	The Volatility of Imports and GDP per Capita	61
I.B.1	Export Thresholds and Demand Volatility	69
2.1	Illustrative Export-Driven Innovation Pattern	91
II.A.1	Graphing Covariate Match Quality	112
3.1	NYC Members over Time: Model of Contagion vs. Data	119
3.2	NYC Members over Time: Network Good Model vs. Data	121
III.A.1	NYC Entry Delay	153
III.B.1	Number of NYC Members Over Time vs RTA Members	157
III.B.2	Number of NYC Members Over Time vs WTO Members	158
III.B.3	NYC Trade Shares of Joiners and Non-Members	159
III.B.4	Proximity and Trade as Drivers of NYC Participation over Time	160
III.C.1	NYC Membership and Trade Growth at the Intensive-Extensive Margins	162
III.D.1	S-Curve Evidence: Cubic Polynomial Fit of Relative Hazard Rate	164

LIST OF TABLES

Table

1.1	Summary of Key Variables	20
1.2	Demand Volatility and the Probability of Entry by Chinese Exporters (Dependent Variable: $\mathbf{1}[\text{Number of Exporters in Destinations} = 0]$)	21
1.3	Exporter Counts and Demand Volatility: (Dependent Variable: Log Number of Exporters in Destinations)	24
1.4	Exports and Exporter Counts vs. Demand Volatility (Dependent Variable: Log Export Measure in Destinations)	27
1.5	Export Entry with Past and Future Demand Volatility (Dependent Variable: Log Number of Exporters in Destinations)	31
1.6	Exporter Counts and Interacted Terms of Demand Volatility (Dependent Variable: Log Number of Exporters in Destinations)	33
1.7	Export Entry with Quantity-Based Demand Volatility Measure (Dependent Variable: Log Number of Exporters in Destinations)	36
1.8	Export Entry with Alternative Demand Volatility Measure (Dependent Variable: Log Number of Exporters in Destinations)	38
1.9	Exporter Counts and Demand Volatility: Poisson Estimates (Dependent Variable: Number of Exporters in Destinations)	40
I.A.1	Exporter Dynamics in China: 2000 - 2006	48
I.A.2	Additional Regression Variables	51
I.A.3	Demand Volatility: Product and Country Fixed Effects (Dependent Variable: Demand Volatility)	52
I.A.4	Annual Trade Estimates with Demand Volatility (Dependent Variable: Log Export Indicator in Destinations by Year)	54
I.A.5	Exporter Counts and Demand Volatility (China Only): (Dependent Variable: Log Number of Exporters in Destinations)	56
I.A.6	Exporter Counts and Demand Volatility (Rest of the World): (Dependent Variable: Log Number of Exporters in Destinations)	58
I.A.7	Comparing Demand Volatility and Conventional Predictors of Trade (Dependent Variable: Log Number of Exporters)	59
I.A.8	Demand Volatility and Prices, by Broad Economics Categories (Dependent Variable: Log Unit Prices)	62

I.B.1	Demand Volatility and Export Thresholds (Dependent Variable: Minimum Market Share and Destination Counts for Firms in jk)	71
2.1	Group Summaries	80
2.2	Comparing Innovation: Exporting vs. FDI	83
2.3	Predicting Exporting and FDI	86
2.4	Innovation vs. Exports and FDI: Propensity Score Matching	88
2.5	Innovation by Exporter Status: Foreign-Owned Firms Only	89
2.6	Changes at the Export and FDI Transitions	93
2.7	R&D and Asset Purchases Increase Product Innovation	95
2.8	Innovation Drivers by Stage of Export/FDI Participation	97
2.9	Matching Estimates by Stage of Export Participation	99
2.10	Comparing coefficients for FDI with HMT	101
II.A.1	Summary of Key Variables	108
II.A.2	Balancing Test for Propensity Score Matching Variables	109
II.A.3	Balancing Test for Propensity Score Variables: Transition Test I	110
II.A.4	Balancing Test for Propensity Score Variables: Transition Test II	111
II.B.1	Matching Estimates by Stage of Export Participation	113
II.C.1	Comparing coefficients for FDI with and without HMT	115
3.1	NYC Members over Time	134
3.2	Summary Statistics for Model Variables	136
3.3	NYC Membership Hazard Estimates: Contagion	138
3.4	NYC Membership Hazard Estimates: The Nature of Trade and Goods	140
3.5	NYC Membership Hazard Estimates: Country Features	142
III.A.1	Countries, NYC Membership Status and Entry Delay	151
III.A.2	Countries, NYC Membership Status and Entry Delay (continued)	152
III.A.3	States Without Treaty-Powers Annexed to the Treaty	153
III.A.4	NYC Membership Hazard Estimates: Countries Eligible in 1958	155

ABSTRACT

Essays on International Trade: Demand Volatility, Exporting and the Adoption of Innovations

by

Michael Adetayo Olabisi

Chairs: Jan Svejnar; Jagadeesh Sivadasan

This dissertation comprises three essays on two themes in international trade – volatility and innovation. I explore how volatility influences exporters’ choices of foreign destinations, how innovation is spurred by export participation and foreign investment, and how countries adopt a policy innovation – a treaty that facilitates the creation of enforceable international trade contracts.

In the first chapter, I develop a simple model of trade with heterogeneous firms facing stochastic demand. As those firms incur adjustment costs in response to market fluctuations, the model predicts that fewer firms will enter destinations with high demand volatility - those destinations are less profitable. I test the prediction using data on the universe of Chinese exports at the firm level. The data show that fewer exporters serve destinations with high demand volatility.

The second chapter of the dissertation compares the effects of foreign direct investment and exporting on product innovation. Governments in many developing economies try to stimulate innovation and private-sector growth with policies that promote exports or attract foreign investment. Using a rich firm level database of Chinese manufacturing and industrial enterprises, I show that exporting is a stronger predictor of product innovation than foreign investment. Firms that receive foreign investment tend to engage in more product innovation, but not at the same level as the firms that learn by exporting.

The third dissertation chapter discusses the adoption of innovations. The policy innovation I consider is a trade treaty that allows private parties to create enforceable international

trade contracts, even in the absence of a common legal jurisdiction. Private international arbitration between importers and exporters are enforced by the national courts of signatories to the 1958 New York Convention on the Recognition and Enforcement of Foreign Arbitral Awards (NYC). This treaty makes private international arbitration a viable legal alternative to trading on reputation or self-enforcement. Surprisingly, many countries waited decades to join the NYC despite the obvious benefits. I show in this chapter that peer effects contribute to countries' adoption of the treaty. Countries with regional trade agreement partners that are NYC members have a higher hazard of participating in the treaty.

CHAPTER I

Demand Volatility and Export Entry

Chapter Abstract

This chapter asks whether and how demand volatility affects exporters' choices of foreign destinations. All export destinations exhibit volatility, with demand from some being more volatile than others. To answer the question, I develop a simple model of trade with heterogeneous firms facing stochastic demand. Firms in the model incur adjustment costs in response to fluctuations in demand. This model predicts that fewer exporters will serve destinations with higher demand volatility, as adjustment costs decrease profits. I test this prediction using data on the universe of Chinese exports from 2000 to 2006 at the firm level. As expected, fewer exporters in the data enter destinations with high demand volatility. Additional firm level regressions show a negative and statistically significant relationship between demand volatility and aggregate trade levels.

1.1 Introduction

Demand volatility is an undeniable feature of international trade. While aggregate trade figures may hide this fact, as global trade growth remained consistently near 6% since the nineties, trade at more disaggregated levels show a different pattern. In the average year, 11% of trade destinations represented by product-country combinations experience a shock that reduces demand to zero. Generally, for imports of specific products to most countries, year-on-year growth shocks cover the full spectrum of negative and positive values.

The drivers of demand volatility include income shocks, government policy and sectoral specialization by firms in the macro-economy. Foreign destinations exhibit different levels of

demand volatility because these factors vary by product and country, and demand fluctuations reflect these factors, which are beyond individual firms' control. I characterize export destinations by the volatility of historical demand, joining a long tradition of scholarship on demand shocks that are exogenous to firm level prices (Blum et al., 2013; Rob and Vettas, 2003; Staiger and Wolak, 1992; Viner, 1922). In proposing that demand shocks may be exogenous to firms and their technical efficiency, the paper follows Foster et al. (2008).

Do exporters avoid destinations with high demand volatility? Phrased differently, do exporters choose country and product combinations with low demand volatility over similar destinations high demand volatility? The question is relevant to understanding how volatility determines the margins of international trade, and to the exporter choices that lead to aggregate trade values.

To address the question formally, this paper develops a model of trade with adjustment costs.¹ Expected profits fall with adjustment costs, e.g. the costs of hiring in times of peak demand or deactivating equipment during lulls. With lower expected profits, fewer firms will find entry into a given destination profitable. Firms in the model are heterogeneous in terms of productivity, so that the more productive are able to enter destinations across a wider range of demand volatility. The model's predictions for trade's extensive margin are unambiguous: destinations with high demand volatility will have fewer exporters and lower levels of trade.

I take the model's predictions to a unique combination of firm-level and global trade data: I observe the destination choices of exporters from the universe of Chinese export transactions between 2000 and 2006, and measure demand volatility using aggregate imports between 1995 and 2005 for each of these destinations. I define destinations as the import stream for unique product-country combinations, like the US imports of truck tires, or Kenyan imports of bicycles. The UN COMTRADE database provides this trade information for narrowly defined HS6 product categories at the country level. Using global demand to estimate demand volatility addresses possible concerns that firm-level shocks drive the measured volatility, rather than patterns of markets' demand.²

The data show a negative and statistically significant relationship between demand

¹ A robust body of work in microeconomics and macroeconomics describes the nature of labor and capital adjustment costs, and how they influence aggregate economic outcomes, e.g. (Bloom et al., 2007; Cooper and Haltiwanger, 2006; Pindyck, 1982; Lucas, 1967). Of these, Lucas (1967) raises the specific concern that adjustments change producer's per-unit costs - an idea that features notably in my model.

²Using Chinese firm-level data is informative for an empirical exercise that describes exporter behavior. China is the world's largest exporter. Furthermore, the high level of correlation between China's aggregate exports and the rest of the world suggests that its exporters behave like firms of other nationalities.

volatility and export entry. First, the likelihood of zero entry is higher for destinations with high demand volatility. It is about 20% more likely that at least one Chinese firm exports to a destination with zero demand volatility relative to an otherwise identical destination at the opposite extreme of demand volatility near 1. The average likelihood of export entry decreases by 2.8% with one deviation from the mean value of demand volatility. With the liberalization that accompanied China's WTO accession in 2001, the number of exporters nearly tripled between 2000 and 2006, therefore entry is an important measure of exporters' responses to markets in the period of trade expansion covered by my data.

Second, conditional on having at least one Chinese exporter, the number of exporters serving a destination is 5 to 10% lower for destinations with demand volatility one standard deviation above the mean, depending on the measure of volatility adopted. The negative estimated effect of demand volatility on entry holds up to alternative definitions of demand volatility, export entry and a variety of specifications that address concerns about the causal nature of this relationship. The effect of demand volatility on trade is not limited to the extensive margin – the value of exports to destinations are 2% lower with a one standard deviation increase in demand volatility in my more conservative estimates.

The findings imply that the prospect of adjustment costs deters firms from incurring the up-front costs of exporting to destinations with high demand volatility. I define adjustment costs in this paper to broadly include capital adjustment costs, as well as the costs of firing and hiring employees. From an exporter's perspective, the findings support the argument in Cuñat and Melitz (2012) that countries may derive a comparative advantage in exporting products with highly volatile demand if they have flexible labor regimes – and therefore lower labor adjustment costs. One can make broad statements about exporter behavior and demand volatility based on these findings for two reasons: Chinese exports are correlated with global exports (Amiti and Freund, 2010), and China is currently the world's largest exporter.

For an importing country, demand volatility represents a potential barrier to economic growth. Most producers in developing economies use imported inputs, usually imported capital goods (Connolly, 2003). In the data, I find that prices are slightly higher for Chinese exports of capital goods to destinations with high demand volatility. Controlling for quality differences should make these price differences starker. The literature generally finds lower prices in the developing economy destinations where high levels of demand volatility are more

common, e.g. Manova and Zhang (2012) and Harrigan et al. (2011).³ Demand volatility also leads to fewer imported varieties after controlling for other determinants of trade. Firms in destinations with high demand volatility therefore lose the potential benefits of new and more imported varieties described in Goldberg et al. (2010).

Based on the foregoing, market-specific patterns of demand volatility can contribute to our understanding of aggregate trade, its extensive margins and the presence of zeros. They complement the explanation for zeros in trade and the large body of work on the determinants of trade. Helpman et al. (2008) shows that zero bilateral trade is more likely with long distances or high marginal costs. This paper extends the idea, suggesting that high adjustment costs due to demand volatility can also increase the occurrence of zero trade. Others describe the determinants of trade in terms of firm-level productivity and geographically-driven trade costs, e.g., Anderson and Van Wincoop (2003) and Melitz (2003). This paper suggests a role for a market feature like demand volatility. In addition to providing evidence that volatility influences exporter choice, the paper contributes to the literature in three ways.

First, by focusing on the decisions of firms to serve specific combinations of products and countries, I provide an approach for explaining firm level trade choices that country-level measures like GDP, exchange rates and geographic distance may not adequately capture. My units of observation are destinations represented by product-country combinations, like the US imports of bicycles: The GDP of the US may not be relevant to a Chinese exporter of bicycles if GDP is a poor predictor of the demand for bicycles in particular. To such an exporter, historical information on imports of bicycles into each potential foreign market is more valuable.⁴ Therefore, this paper describes demand from destinations as a random walk with a constant growth trend, following Carroll et al. (2011) and Hall (2004). Exporters forecast profits from the volatility and trend observed in each destination's demand history, and enter markets on that basis. Market-specific volatility for a given product reflects income, policy or other transactional frictions not explained by equilibrium prices.

³This finding in section I.A.8 agrees with the evidence in Eaton and Kortum (2001) that capital goods have higher relative prices in developing economies (where I find the most volatility). Other papers that find lower prices for exports to developing economies do not focus on specific sub-categories like capital goods, which make up less than 11% of global exports.

⁴ Most exporters serve few foreign destinations. The median number of products and countries per exporter in the data are 5 and 4 respectively. Country-level measures like GDP mask shocks that may be important to firms that focus on a few product categories. For example, US imports of pure fructose (HS 170250) are more volatile than Rwandan imports of truck tires (HS 401120). The two countries' aggregate measures of GDP and demand volatility are at opposite ends of the spectrum.

Second, the paper extends the literature on investment under uncertainty to the context of trade. The initial costs of setting up overseas trading networks are analogous to investments made in the expectation of future returns. Future demand is not known, but rational agents characterize its expected value using historical information. This yields new testable insights on export entry. If adjustment costs are expected to be higher for destinations with high demand volatility, then fewer exporters should serve those destinations. Recent related papers show that policy uncertainty reduces exports, when trade costs are driven by policy (Handley and Limão, 2012, 2013). In that context, exporters are more likely to invest in foreign destinations with stable tariff regimes. Earlier work by Dixit (1989) shows that with uncertain prices, firms require prices above a certain threshold to expand their operations. The same relationship between uncertainty and investment holds for exchange rate uncertainty (Das et al., 2007; Frankel and Rose, 2002; Glick and Rose, 2002). The novelty in this paper is its focus on demand, as the volatility of trade costs, prices and exchange rates only explain small shares of the variation in trade.

Finally, I provide a simple and intuitive measure of volatility: i.e. the sum of squared deviations from a trend for a series. I use a linear trend, though the measure is flexible enough to admit other trend specifications. Others define volatility as the standard deviation of year-on-year growth rates, but measuring growth rates is problematic when observations include zero.⁵ The index of volatility introduced by this paper avoids such issues of measurement.

The rest of the paper is organized as follows: Section 1.2 presents a stylized model in the tradition of Melitz (2003) and Chaney (2008) to motivate the empirics. Section 1.3 follows, with the data, formal definitions for key variables, empirical specifications and results. Section 1.4 discusses the implications and concludes.

1.2 Model

1.2.1 Adjustment Costs with Stochastic Demand

In this model, exporters decide on foreign destinations using information on historical demand shocks. Exporters can form unbiased expectations of demand for each year in a forward-looking planning horizon, given the trajectory of past demand. Previous related papers model demand uncertainty in a framework that requires exporters to learn about

⁵One can correct the conventional measure of growth volatility by using a mid-point growth measure, which bounds growth between -2 and 2, but it does not help that those extreme values of growth may be outliers that skew the measure of volatility.

demand e.g. Akhmetova and Mitaritonna (2012) and Nguyen (2011). In that context, the conditional distribution of possible demand outcomes is taken as unknown. In this paper, the conditional distribution of demand outcomes depends only on historical demand realizations.

In other words, firms form expectations of volatility and future demand for each destination from its historical demand. For example, a destination that imports exactly \$1m for all years between 1995 and 2005 will lead all exporters to expect little no growth, and negligible demand volatility. (Past volatility is taken as the predictor of future volatility). This aggregate import volatility described in the paper applies to all exporters, as historical demand is common knowledge to all firms.⁶


I illustrate the section's main idea using refineries. Updating the capacity of a refinery in production is costly. For a refinery that exports, each production run incurs upfront customization costs, as gasoline and diesel blends differ by country. Therefore, before a refiner enters a foreign destination, it must consider the usual per-unit marginal costs, the upfront customization costs and the costs of capacity adjustments expected in its planning horizon for that destination. The hypothetical demand trajectories in Figure 2.1 illustrate the relevance of capacity adjustments. Judging from the plots of historical demand, the two destinations in the graph have the same expected size, but the scales of deviations from the expected trajectory of demand differ. As the producer must consider the relative costs of scaling production to match demand in each period, the destination on the right panel becomes more preferable if the historical pattern of volatility persists.

Exporters only serve destinations with non-negative expected profits – after accounting for adjustment costs. With adjustment costs, some marginally profitable destinations with high demand volatility become unprofitable.⁷ For the refinery example that motivates this section, a rational exporter may find the destination on the right panel of the figure profitable, but avoid the destination on the left. Large production scale adjustments reduce realized profits, as documented in the literature on adjustment costs (Cooper and Haltiwanger, 2006; Lucas, 1967).

⁶One could conceptually create firm-specific measures of demand volatility that use weighted composites of the demand history for each firm's target destinations. Larger shocks at the firm-level may make such a firm-specific measure higher than the currently proposed destination-specific measure – the differences may be more notable for new entrants or small firms. However, the relationship between these two conceptions of demand volatility depends on the 'diversification effect' - where demand shocks offset one another for multi-destination exporters, which account for the largest share of exports.

⁷ The model proposed here reverts to a conventional model of trade if one collapses the planning horizon to one period, and presumes perfect information about future demand. Adjustment costs disappear with these additional assumptions. In that sense, this paper extends the conventional model of trade to include multi-period sales and information.

Figure 1.1: Demand with Stochastic Shocks



The graph is purely illustrative - its two hypothetical destinations have equal average size, but different levels of demand volatility. The horizontal line represents projected demand and vertical lines represent deviations from the demand trajectory.

Destinations with high demand volatility and expected adjustment costs will attract fewer exporters – as long as expected demand volatility reflects the observed demand history. Demand volatility measures the variability in demand over time, and as expected, high demand volatility corresponds to high adjustment costs for firms serving a destination. The next steps relate demand volatility to profits.

Formally, profits for a producer i considering exports of product j to country k :

$$\Pi_{ijk} = \sum_{t=1}^H \{p_{ijk} * q_{ijkt} - \hat{c}_{ijk}(1 + \text{adjustment costs}_{ijkt})q_{ijkt}\} - S_{jk} \quad (1.1)$$

p_{ijk} = unit price

q_{ijkt} = quantity sold

S_{jk} = sunk costs of production

$\hat{c}_{ijk} = \frac{\tau_{jk}}{\phi_{ij}}$ = standard unit costs

p and q represent the prices and quantities for firm i in the planning horizon that covers periods $t \in [1, H]$. \hat{c}_{ijk} , the standard unit production cost captures τ_{jk} , the combined per-unit

costs of inputs like labor and materials, which are specific to product j , and trade factors like shipping and tariffs for destination k . It also accounts for the firm's productivity ϕ_{ij} . Firms with higher productivity ϕ will therefore have lower unit costs and higher profits per unit sold. (For parsimony, the model ignores temporal discounting and simply sums profits from period 1 to H ; a reasonable approximation if the planning horizon is short and the discount rate is small).

Adjustment costs may be overtime wage costs, equipment capacity modification or hiring and firing costs. I adopt the convex quadratic form proposed in Cooper and Haltiwanger (2006):⁸

$$\text{adjustment costs}_{ijkt} = \gamma_j \left[(q_{ijkt} - q_{ijkt}^*) / q_{ijkt}^* \right]^2 \quad (1.2)$$

The γ_j term is a product-specific scaling parameter. It enables comparisons in the cross-section of destinations. For example, the cost implications of a 20% growth shock to demand are different for an auto manufacturer, compared to a maker of tee-shirts. Each sector faces different relative costs of updating production capacity. γ_j captures these differences in scale adjustment cost differences. q_{ijkt}^* is the planned production scale or capacity for firm i in period t .⁹

In the adjustment costs function, $q - q^*$ represent deviations from the production scale, or the height of the vertical lines in Figure 2.1. Unit production costs should depend on the proximity of actual production to the expected production scale. Several recent papers show that production scale adjustments alter marginal costs (Blum et al., 2013; Soderbery, 2013; Ahn and McQuoid, 2012). (The foregoing implicitly assumes that the cost of changing production scale from q_{jkt-1}^* to q_{jkt}^* is zero, because such changes follow the planned

⁸Costs are symmetric around q^* in equation (1.2); this makes the model tractable, although cost symmetry may only be a crude approximation to the data.

⁹The additive form specified in equation (1.1) for adjustment costs ensures that unit costs will not be zero for the hypothetical destination with zero demand volatility. This structure also allows one to measure adjustment costs' effects separately from other components of unit costs, which improves on related papers that also consider changing marginal costs with market-specific shocks (Vannoorenberghe, 2012; Ahn and McQuoid, 2012; Liu, 2012).

trajectory).¹⁰

Aggregate demand is stochastic, but the growth process for demand is known, given the demand history. Therefore, exporters can characterize aggregate demand Q in each destination and estimate expected profits from equations (1.1) and (1.2):

$$Q_{jkt} = Q_{jkt}^*(1 + \nu_{jkt}) \quad (1.3)$$

Q_{jkt}^* , period t 's expected aggregate demand is estimated from the trajectory:

$$Q_{jkt}^* = Q_{jk0}^*(1 + t\hat{g}_{jk}) \quad (1.4)$$

The Q_{jk0} baseline and the expected growth trend \hat{g}_{jk} in (1.4) come from historical data.¹¹

Demand volatility σ_{jk}^2 represents the second moment of the distribution of the growth innovations ν_{jkt} in equation (1.3), given that $\nu_{jkt} \sim N(0, \sigma_{jk}^2)$. Assuming a normal distribution for ν helps to obtain a tractable form for expected profits shortly.¹² (I adopt the linear growth form for simplicity; a multiplicative growth model in (1.4) gives $Q_{jkt}^* = Q_{jk0}^*(1 + \hat{g}_{jk})^t$, which approximates linear growth for small values of \hat{g}_{jk}).

From historical data one gets \hat{g}_{jk} , which characterizes Q_{jkt}^* , as well as σ_{jk}^2 , which fully describes the expected shocks to demand, even if specific realizations of Q_{jkt} are not known. In the model, exporters estimate σ^2 for each product-country destination once, and do not update their estimates of demand volatility. This simplifying assumption helps to justify tests in the cross-section of destinations in Section 1.3. I show that this assumption is reasonable, based on tests in Section 1.3.2.3.

The next subsection derives exporters' expected demand q_{ijk}^* from the expected aggregate demand Q_{jk}^* . Firms' profits and the decision to export to jk depend on q_{ijk} and q_{ijk}^* .

¹⁰Alternatively, one could change the definition in (1.2):

$$\text{adjustment costs}_{ijk} = \gamma_j^v \left[\frac{(q_{ijk} - q_{ijk}^*)}{q_{ijk}^*} \right]^2 + \gamma_j^p \left[\frac{(q_{ijk}^* - q_{ijk}^*_{t-1})}{q_{ijk}^*_{t-1}} \right]^2$$

The inclusion of planned investment costs, scaled by γ^p suggests biased estimates if γ^p is not equal to zero. Taking γ^p as zero seems reasonable for two reasons: (1) the costs of adjusting scale upwards must be less than profits from increased scale if one is to observe more firms in larger markets in equilibrium; (2) early studies that do not consider expected demand find no statistically significant estimates for labor or capital adjustment costs (Hall, 2004).

¹¹Firms can forecast q_{ijk} , given the history of Q_{jkt} . Even with perfect information about future demand, e.g., a firm like Boeing making airplanes to order, the demand stream with greater deviations from a stable growth trajectory will incur higher adjustment costs, regular planned investments of labor and capital in each period are less costly to implement than large swings. Having information on future demand may reduce, but not eliminate those adjustment costs.

¹²Growth innovations in the data resemble a normal distribution near the mean. The $(1 + \nu)$ term correspond to growth shocks in the Euler equations proposed by Carroll et al. (2011) and Hall (2004).

1.2.2 Export Entry

The primary variable of interest is the number of exporters that find a destination profitable, and therefore export to that destination. In the empirics, this translates to gross export entry in the long run for trade destinations. This section explores the relationship between this variable and demand volatility.

Exporter i producing its unique variety of product j for market k can expect to sell q_{ijkt} in period t . Following conventional models of trade with CES demand preferences:

$$q_{ijkt} = \frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}} Q_{jkt} \quad (1.5)$$

p_{ijk} is the firm's expected price, P is the Dixit-Stiglitz aggregate price index and ε is the elasticity of substitution between varieties of j . Q_{jkt} is the aggregate demand for product j in country k . The steps that follow assume no exporter is large enough to affect the P index. As previously mentioned, Q_{jkt} is stochastic.

For each destination, $q_{ijk}^* = E(q_{ijkt})$, the optimal production scale for an exporter is the expected demand.¹³

From equations (1.2), (1.3) and (1.5):

$$\begin{aligned} \text{adjustment costs}_{ijk} &= \gamma_j \left[\frac{\frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}} (Q_{jkt} - Q_{jkt}^*)}{\frac{p_{ijk}^{-\varepsilon}}{P_{jk}^{1-\varepsilon}} Q_{jkt}^*} \right]^2 \\ &= \gamma_j (\nu_{jkt})^2 \end{aligned} \quad (1.6)$$

The expected profits over the planning horizon from equation (1.1), (with risk-neutral exporters and known sunk costs S):

$$E(\Pi_{ijk}) = E\left\{ \left[p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}} (1 + \text{adjustment costs}_{ijk}) \right] q_{ijkt} \right\} - S_{jk}$$

¹³The Envelope Theorem justifies this cost minimization, given the assumption that adjustment costs are symmetric in equation (1.2). See more on this in Appendix Section I.B.1

Substituting equation (1.6) and discarding t subscripts yields:

$$\begin{aligned}
E(\Pi_{ijk}) &= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})E(q_{ijk}) - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j E(q_{ijk}\nu_{jk}^2) - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j \{E[(q_{ijk} - q_{ijk}^*)(\nu_{jk})^2] + q_{ijk}^*E[(\nu_{ijk})^2]\} - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q_{ijk}^* \{E\left[\frac{(q_{ijk} - q_{ijk}^*)}{q_{ijk}^*}(\nu_{jk})^2\right] + E[(\nu_{ijk})^2]\} - S_{jk} \\
&= (p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}})q_{ijk}^* - \frac{\tau_{jk}}{\phi_{ij}}\gamma_j q_{ijk}^* [E(\nu_{jk}^3) + E(\nu_{jk}^2)] - S_{jk}
\end{aligned}$$

The $E(\nu_{jk}^2)$ term is σ_{jk}^2 , as defined in the notes to equation (1.3). The $E(\nu_{jk}^3)$ term is zero, being the third moment of a normal distribution.¹⁴ This gives the expected profit:

$$E(\Pi_{ijk}) = [p_{ijk} - \frac{\tau_{jk}}{\phi_{ij}}(1 + \gamma_j \sigma_{jk}^2)]q_{ijk}^* - S_{jk} \quad (1.7)$$

As in equation (1.1), standard unit costs are $\frac{\tau_{jk}}{\phi_{ij}}$: the product of variable costs τ (material, capital, labor inputs, shipping and tariffs), and the firm level productivity index $\frac{1}{\phi}$. Adjustment costs reduce expected profits.¹⁵

Firm-level prices determine expected profits, so the next steps focus on deriving prices p_{ijk} . (Note that I abstract away from period-to-period price changes – each exporter's price is assumed invariant in the planning horizon for export entry). I also assume rational risk-neutral firms that maximize expected profits.¹⁶

$$\frac{dE(\Pi_{ijk})}{dp_{ijk}} = \frac{dE(\Pi_{ijk})}{dq_{ijk}^*} \frac{dq_{ijk}^*}{dp_{ijk}} = 0 \implies \frac{dE(\Pi_{ijk})}{dq_{ijk}^*} = 0 \quad (1.8)$$

$$\begin{aligned}
0 &= p_{ijk} - \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} - q_{ijk} \left(\frac{1}{\varepsilon} \frac{p_{ijk}}{q_{ijk}^*} \right) \\
p_{ijk} &= \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} \quad (1.9)
\end{aligned}$$

¹⁴One can get (1.7) from the profit function for adjustment costs that are any real-valued function of the growth innovation ν . For a normal distribution with mean zero, σ^2 the second moment of ν can fully describe the terms of such a function i.e. higher order moments of ν .

¹⁵While the form in equation (1.7) is tractable enough to yield closed form solutions for firm level prices, the model's predictions will hold even if adjustment costs have a component that is fixed, or that does not scale linearly with expected production q_{ijk}^* . The challenge of deriving equilibrium prices with such extended forms is left for another paper.

¹⁶From equation (1.5), $\frac{dp_{ijk}}{dq_{ijk}^*} = \frac{P^{1-\varepsilon}}{-\varepsilon p_{ijk}^{1-\varepsilon}} \frac{1}{Q_{jk}^*} = \frac{-1}{\varepsilon} \frac{p_{ijk}}{q_{ijk}^*}$

Prices in (1.9) take the same form as in conventional models of trade with heterogeneous firms with one difference; unit costs include $(1 + \gamma_j \sigma_{jk}^2)$ to reflect adjustment costs. This is the expected price for firm i in destination jk . Firms with high productivity ϕ_{jk} will have lower prices, assuming no quality differences.

Expected export profits based on price p_{ijk} , from substituting p back into equation (1.7):

$$\begin{aligned}
E(\Pi_{ijk}) &= \frac{1}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} q_{ijk}^* - S_{jk} \\
&= \frac{1}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} \left[\frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} \right]^{-\varepsilon} \frac{Q_{jk}^*}{P_{jk}^{1-\varepsilon}} - S_{jk} \\
E(\Pi_{ijk}) &= \frac{Q_{jk}^*}{\varepsilon} \left[\frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} \frac{1}{P_{jk}} \right]^{1-\varepsilon} - S_{jk} \tag{1.10}
\end{aligned}$$

Only firms above a certain productivity threshold will be profitable in destination jk . Applying the zero-profit entry condition to equation (1.10) identifies those firms. One gets the productivity threshold ϕ_{jk}^* by setting the LHS to zero in (1.10):

$$\phi_{jk}^* = \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{P_{jk}} \left[\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right]^{\frac{1}{\varepsilon-1}} \tag{1.11}$$

Of the N_j firms producing j , only a fraction N_{jk} will export to destination jk . That fraction could be as low as zero if none meets the ϕ_{jk}^* threshold. Deriving the fraction N_{jk} is straightforward if one can describe the productivity of all producers of j with the distribution $G(\cdot)$. I model N_j as an exogenous variable:¹⁷

$$N_{jk} = N_j(1 - G(\phi_{jk}^*)) \tag{1.12}$$

¹⁷ In assuming an exogenous mass of exporters, I follow others – notably, Chaney (2008) and Eaton et al. (2004). Here N_j is the number of firms making product j , e.g., the number of firms that make bicycles, regardless of export status or productivity. N_{jk} represents firms whose productivity exceeds the threshold for jk , given the assumed productivity distribution. Some producers of j will not export at all, if the lowest threshold ϕ^* of all possible markets is higher than firm productivity ϕ_{ij} .

I take $G(\cdot)$ as the Pareto distribution.¹⁸

$$N_{jk} = N_j[1 - (1 - (\phi_{jk}^*)^{-\theta_j})] = N_j(\phi_{jk}^*)^{-\theta_j} \quad (1.13)$$

θ_j is the Pareto shape parameter for product j .

Equation (1.13) shows an unambiguous relationship between σ^2 and N_{jk} , (which suggests a focus on the extensive margin of trade). ϕ_{jk}^* is a function of $\tau_{jk}(1 + \gamma_j\sigma_{jk}^2)$, therefore N_{jk} is a function of σ_{jk}^2 . In contrast, Section I.B.2 in the appendix models the relationship between demand volatility and trade volumes, which is not as pointed as the relationship in (1.13).¹⁹

Substituting the threshold defined in equation (1.11) into (1.13):

$$N_{jk} = N_j \left\{ \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j\sigma_{jk}^2)}{P_{jk}} \left(\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right)^{\frac{1}{\varepsilon - 1}} \right\}^{-\theta_j}$$

Focusing on N_{jk} and σ^2 .

$$\begin{aligned} \ln(N_{jk}) &= \ln(N_j) - \theta_j \left[\ln(1 + \gamma_j\sigma_{jk}^2) + \ln\left(\frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}}{P_{jk}}\right) + \frac{1}{\varepsilon - 1} \ln\left(\frac{\varepsilon S_{jk}}{Q_{jk}^*}\right) \right] \\ \frac{d\ln(N_{jk})}{d\sigma_{jk}^2} &= \frac{-\theta_j\gamma_j}{1 + \gamma_j\sigma_{jk}^2} \end{aligned} \quad (1.14)$$

Plotting $\ln(N_{jk})$ against σ^2 should give a line with a negative slope. The elements of the RHS term in equation (1.14) are all non-negative by definition: the Pareto shape parameter, θ , the adjustment cost scaling parameter γ and the demand volatility σ^2 . Formally:

$$\frac{d\ln(N_{jk})}{d\sigma_{jk}^2} < 0 \quad (1.15)$$

Restating equation (1.15):

¹⁸ This choice follows Chaney (2008) and is consistent with the firm size distributions described in Hsieh and Ossa (2011) and Axtell (2001) Any of the general class of power law distributions should yield similar predictions, given reasonable assumptions about how the distribution is truncated.

The Pareto distribution function is $Pr(X < x) = 1 - \left(\frac{x_m}{x}\right)^\theta$ for $x \geq x_m$. The two parameters that characterize the distribution are x_m , the minimum productivity for a firm that produces j and θ , the shape parameter. For simplicity, I define the range of productivities on a scale $[1, \infty)$, this sets x_m equal to one, so $G(x) = Pr(X < x) = 1 - (x)^{-\theta}$.

¹⁹The dominance of the extensive margin is consistent with other papers that model the responses of heterogeneous firms to trade costs, e.g. Crozet and Koenig (2010) and Helpman et al. (2008). The adjustment costs associated with demand volatility increase exporters' per unit costs, just as trade costs do.

Prediction: Higher levels of destination demand volatility σ_{jk}^2 reduce the numbers of exporters in equilibrium.

To restate the hypothesis, the adjustment costs associated with demand volatility reduce profitability, such that the mass of firms that find a destination profitable decreases with increases in demand volatility. If demand volatility is zero, the model reverts to the conventional model of trade. One way to take this prediction to the data is a linear regression of N_{jk} on σ^2 ; the sign of the coefficient on demand volatility should be negative.

Lemma: Holding other factors equal, the average productivity of firms in destinations with high demand volatility is higher.

From equation (1.11), it is clear that the productivity threshold ϕ_{jk}^* for entering a destination increases with demand volatility, therefore one expects the minimum level of other proxies for productivity like exporters' share of a product's exports to increase with demand volatility, all other things being equal:

$$\phi_{jk}^* = \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{P_{jk}} \left[\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right]^{\frac{1}{\varepsilon - 1}}$$

$$\frac{d\phi_{jk}^*}{d\sigma_{jk}^2} > 0 \tag{1.16}$$

Corollary: The rate of decline in exporter numbers with demand volatility is not constant, but varies nonlinearly with gamma γ .

The cost implications of demand volatility are non-linear, so are its impacts on an exporters' profit. Entry choice by extension is nonlinear with respect to demand volatility. The non-linearity will depend on γ_j . For very small γ , such that low labor-firing costs for example make adjustments relatively costless, the slope of $\log(N_{jk})$ with respect to σ_{jk}^2 will always be approximately linear and proportional to θ_j , the productivity distribution's shape parameter. For large γ_j , the slope is increasingly non-linear, and decreases in absolute terms with σ^2 . (If the γ term was a function of σ^2 because adjustment costs are not linear proportions of ν^2 as proposed in equation (1.2), the slope could increase or decrease with σ^2).

Testing the relationship between N_{jk} and the square of the demand volatility term directly

addresses this corollary in Section 1.3.2.4.

1.3 Empirics

This section examines the relationship between exporter numbers and demand volatility. First, I describe key variables and data sources. Regression estimates follow the definitions: a baseline specification and variations that further test the model’s predictions. I address the most important alternative explanations at the end of the section.

1.3.1 Data and Definitions

The key variables come from two trade datasets: firm level export data from China to describe exporter choices, and UN COMTRADE data on global imports by product and country to describe the history of each destination in terms of size and demand volatility.

The firm level export data captures exporter numbers in each destination, derived from the universe of Chinese export transactions between years 2000 and 2006. The UN COMTRADE data include global imports of each narrowly defined HS6 product category for all countries between 1995 and 2010. Annual imports up to 2005 were collapsed to product-country-year observations – imports of bicycles (HS871200) into Kenya in the year 2000 from all countries would be one observation, for example. (Restricting historical data to 2005 and earlier is motivated by the fact that firm level export data stop at 2006). I estimate demand volatilities using this global import demand history for each destination.²⁰

The combined datasets represent the destination choices of more than 243,000 Chinese exporters, mostly in the period of export expansion that followed entry into the WTO, covering more than 390,000 of the roughly one million possible product-country combinations that define destinations. (These are imports into one of more than 200 countries in any of the 5000 narrowly defined HS6 product categories). Appendix Section I.A.1 describes these data sources further and outlines how I merge the two. Information on GDP, distance and other predictors of trade come from the CEPII gravity dataset (Head et al., 2010).

The next sub-sections describe the key variables in the empirical specification: the dependent variable, which is a multi-year count of unique exporters and the key independent variables, derived from the historical demand data described in the preceding paragraphs.

²⁰The COMTRADE database of global trade was compiled and cleaned up by Gaulier and Zignago (2010) and released to the public through the Centres d’Études Prospectives et d’Information Internationales (CEPII). I will refer to this database as COMTRADE from here. See www.cepii.fr/anglaisgraph/bdd/baci.htm.

1.3.1.1 Dependent Variable: Multi-Year Unique Exporter Counts

The paper uses two definitions of exporter numbers N_{jk} . The primary measure represents the number of Chinese firms that exported good j to a destination jk between 2000 and 2006, counting each exporter only once in the entire period. This definition reflects the equilibrium number of exporters entering a potentially profitable destination. (The measure of entry extends over multiple years because theory provides no clear guidance on how long it takes to reach equilibrium).²¹ Furthermore, exports drop to zero and rebound in the following year for many destinations. Defining entry on an annual basis may misrepresent markets that simply run on a multi-year demand cycle. I use logged values of this unique exporter count for the regressions.²²

The second measure I use is gross entry from 2001 to 2006. This is the log of the difference between [1] N_{jk} the number of unique exporters over all years and [2] the number of unique exporters observed in 2000, the first year in the data, which happens to be the year before China joined the WTO. China's WTO entry in 2001 marked a new regime of low trade costs and easier export access, evidenced by an increase in aggregate exporter numbers from around 62,000 in 2000 to more than 170,000 in 2006. By its construction, this measure of gross entry in the new trade regime controls for undocumented destination attributes that may influence export entry. For example, destinations with lower entry costs are more likely to have exporters in 2000, and to have more exporters.

This alternative measure could be more relevant, as export restrictions that may have distorted observed exporter counts were removed in the course of WTO accession. For example, before WTO accession, trading license requirements barred Chinese firms below a certain size from exporting directly (Ahn et al., 2011). In sum, exporter entry after 2000 was large, with ample variation across destinations, which helps the analysis.

²¹Later sections of the paper include tests with annual counts, which provide results comparable to typical estimates using annual data. Defining exporter counts as gross entry is a better fit to the model because the model remains silent about exits from destination after entry. Table I.A.4 relaxes this connection between the model and the empirics, using annual exporter counts as the dependent variable. These include gross entry and exit, by definition.

²²For convenience, I call N_{jk} gross entry in the tables that follow. The data show that exporters in 2000 represent only a quarter of the full set of observed unique exporters. Fortunately, China's accession to the WTO in 2001 suggests an alternative measure of gross entry in the era of liberalization, which I discuss in the next paragraph.

1.3.1.2 Key Variable: Demand Volatility

I measure demand volatility as the sum of the squared deviations of demand from a linear trend over the years 1995 to 2005. The trend is calculated for the dollar value of each destination's annual demand – total imports from all countries in a given HS6 category. As Chinese exports represent less than 13% of the global total in this period, this measure mitigates concerns about reverse causality that may have resulted from defining volatility with only Chinese export data. The growth trend estimation uses a linear regression, with each destination scaled by total demand over all years to yield a scale-free measure for cross-sectional comparisons. To get σ_{jk}^2 , the volatility term, I estimate the trend and intercept for each destination jk :²³

For this exercise, the definition of destinations or destinations in the model is consistent with the empirics. Each Q_{jk} in the model corresponds to a specific product j and a country k , just as destinations in the data are defined as the combination of a narrow HS6 product and a country.

Formally, I run the following regression:²⁴

$$\frac{Q_{jkt}}{\sum_t Q_{jkt}} = \zeta_{jk}t + \alpha_{jk} + \epsilon_{jkt} \quad (1.17)$$

Using the residuals, I calculate σ^2 , (for convenience, these equations show estimated terms without the \hat{x} notation):

$$\sigma_{jk}^2 = \sum_t (\epsilon_{jkt})^2 \quad (1.18)$$

The incidental assumption in this setup is that exporters form expectations of future volatility based on observed aggregate volatility. As described in the previous section, each

²³ To avoid the bias that may result from imposing the same non-linear form on all demand trajectories, I opt for a linear regression of historical trade levels on time. This measure directly interprets the model's definition of demand volatility. Some measurement error is expected, given that I use only 11 years of demand history. However, using longer demand histories comes with the risk of including irrelevant information. On a related note, almost all destinations served by Chinese exporters had positive imports in all years, reducing concerns that observations clustered at zero would inflate the volatility measure.

²⁴ I use dollar values instead of quantities for Q in the primary estimates of volatility. This was in part because quantity data was missing for a large share of the old version of the COMTRADE data. It is also in part because of concerns that quantities are mis-measured - the incentive to measure values accurately is stronger for the customs authorities. On obtaining a more recent version of the COMTRADE data with nearly complete estimates for quantities, I repeat the baseline regressions using a quantity-based volatility measure and report the results in the robustness checks section 1.3.3.2.

exporter's expected revenue from a destination is proportional to the aggregate demand from that destination, therefore volatility at the product-country level is a reasonable measure for evaluating the profitability of each destination at the firm level.

After the main tests in the next subsection, I use the standard deviation of year-on-year growth as an alternative definition. This alternate definition is also broadly consistent with the model, and with other papers in the literature.

Nonetheless, the measure of demand volatility has the advantage of addressing the two main challenges to measuring volatility for time series: (1) making the measure independent of the size of each series and (2) separating baseline growth from volatility. I control for size by scaling all series by the total value over all periods, and control for growth by introducing the linear trend that best fits the data. Controlling for size ensures that the volatility measure is comparable across series with different initial levels. For example, using the standard deviation of historical values to compare the volatility of US aggregate imports with Rwandan imports would lead to the flawed conclusion that US imports are more volatile, simply because the absolute values are larger. (One could try to fix this by using the coefficient of variation - i.e. dividing by the mean, but that still leaves concerns about growth trends unaddressed).

Furthermore, defining demand volatility as deviations around a trend avoids mis-measurement when the data include instances of zero demand. The common measure of volatility as the standard deviation suffers from the problem of measuring growth from or to zero. If one uses the mid-point growth measure of Davis and Haltiwanger (1992), growth at these instances of zero will fall at the extreme values of -2 and 2. While those values are usable, they may represent outliers that bias the volatility measure, especially if most growth observations are clustered near zero. Measuring demand volatility as deviations from a trend simply avoids this concern by used scaled levels, rather than transforming those levels into a growth index before calculating the volatility of a series.

1.3.1.3 Key Variable: Destination Size

In estimating the effects of demand volatility, I use aggregate historical demand as an indicator of product-country destination size. Conventional estimates of international trade measure size as GDP. This would be appropriate for a model of trade where the countries define market boundaries and firms' narrow product specializations were not relevant to competition. Using aggregate historical demand for each destination as a measure of its size fits the structure provided by my model, one that emphasizes competition within narrow

product categories and exporters' choices of markets on that basis.

I define the terms as the logged sum of aggregate demand in each destination between 1995 and 2005. (I use the first 11 years of aggregate data available for the same reasons that I use those years to measure demand volatility). In principle, this logged sum represents the projected future demand for a destination. One only needs to assume that destination growth g_{jk} is consistent to use this history as a proxy for the projected aggregate demand over the exporter's planning horizon, (Q^* in the model). Formally, the projected size of a destination with average historical growth rate g_{jk} is $\log(\sum_t Q_{jkt}) \simeq \log(Q_{jk0}) + \log[\sum_t (1 + g_{jk})^t]$. As long as past growth rates are a reasonable proxy for expected growth rates, this measure allows me to control for destination size in terms of its present value and growth.

This measure offers a finer level of control for testing export entry decisions than a country-level measure like GDP. The regressions in this section will show that it explains more of the variation in export entry than conventional variables like GDP and distance. (This is in part because; historical demand is explained by GDP and distance, so that the inclusion of current GDP in an estimation exercise that includes historical demand provides little additional information). I run versions of the regressions that follow this section without this market size variable and obtain results that are similar in sign, but with larger coefficients. This is unsurprising, given the correlation between market size and volatility (see Figure I.A.1). To avoid the implied omitted variable bias, the next section reports only regressions that include this destination size variable.

1.3.2 Results

1.3.2.1 Demand Volatility and Exporter Counts

Table 1.1 summarizes the key variables.

About 40 unique exporters served the average destination between 2000 and 2006; with 35 of these being the firms that entered the destination after the year 2000. This number is highly skewed; both variables have a median value of 5. The variation in exporter counts is large; products like buttons naturally had many producers, while airplanes had few. Country-specific variations also existed; large ones like the US had more exporters. However, countries and products alone leave much of the variation in the data unexplained. Unreported regressions of export entry at the level of product-country destinations on product and country fixed effects alone yield R^2 values of 0.20 and 0.01 respectively.

Given that more than half of the possible destinations had zero Chinese exporters, the

Table 1.1: Summary of Key Variables

Variable	Mean	Std. Dev.	Min.	Max.	N
Gross Entry	39.95	190.77	1	15643	397547
Gross Entry post-2000	35.38	168.38	0	13497	397547
Log(Gross Entry)	1.96	1.66	0	9.66	397547
Log(Entry post-2000)	1.91	1.63	0	9.51	387916
Destination Size	8.65	2.63	0	20.6	397547
Demand Volatility	0.053	0.086	0	0.91	380372

Chinese exporters sent goods to 397,547 destinations between 2000 and 2006. Only 380,372 had the two or more non-zero observations required to compute demand volatility. 9,170 had no export entry after 2000. Destination size is the log of total historical demand in the COMTRADE data. Number of countries (205); products (4,902)

first order of inquiry is whether demand volatility is higher on average for those destinations that had no Chinese exporters. Figure I.A.1 in the appendix suggests that the destinations served by Chinese exporters tend to be larger and have lower demand volatility, but the visual comparisons of distributions in that figure is only suggestive at best.

Table 1.2 presents the conditional expectation. As described in section I.A.2, destinations with zero Chinese exporters have higher demand volatility on average. The specification adopted is a linear probability model with a dependent variable that is 1 if no Chinese firm exported to the destination between 2000 and 2006; it is 0 otherwise. Product fixed effects address the fact that some items are more likely to be exported than others for time-invariant reasons outside the model, and country fixed effects or variables like GDP and distance address the fact that country-level factors also determine the prevalence of zeros in trade.

The difference in the likelihood of having at least one Chinese exporter is about 20% on average for two otherwise identical destinations with levels of demand volatility at the minimum and maximum, (i.e. $0.277*(0.8 - 0.0)$). Increasing demand volatility by one standard deviation corresponds to a 2.8% decrease in the likelihood that a destination is served by Chinese exporters, after controlling for destination size and country features. The standard deviation of demand volatility is 0.11 for this set of destinations. I control for market size in columns 2 and 4 to address the concern that larger destinations will generally have more exporters, as market size is correlated with demand volatility.²⁵

Figure 1.2 shows that more Chinese exporters serve destinations with low demand volatil-

²⁵I calculated demand volatility for 708,802 destinations. The balance of destinations were not usable because the demand volatility variable is not defined for destinations with only one or two years of imports.

Table 1.2: Demand Volatility and the Probability of Entry by Chinese Exporters
(Dependent Variable: $\mathbf{1}[\text{Number of Exporters in Destinations} = 0]$)

VARIABLES	(1)	(2)	(3)	(4)
Demand Volatility	0.391*** (0.007)	0.277*** (0.007)	0.335*** (0.006)	0.259*** (0.006)
Destination Size		-0.045*** (0.001)		-0.040*** (0.001)
Log(GDP)	-0.086*** (0.001)	-0.054*** (0.001)		
Log(GDP per capita)	0.018*** (0.001)	0.022*** (0.001)		
Log(Distance)	0.103*** (0.002)	0.098*** (0.002)		
Constant	-0.390*** (0.023)	-0.561*** (0.023)		
Observations	573,997	573,997	708,802	708,802
R-squared	0.465	0.480	0.515	0.525
Country FE			Y	Y
Product FE	Y	Y	Y	Y

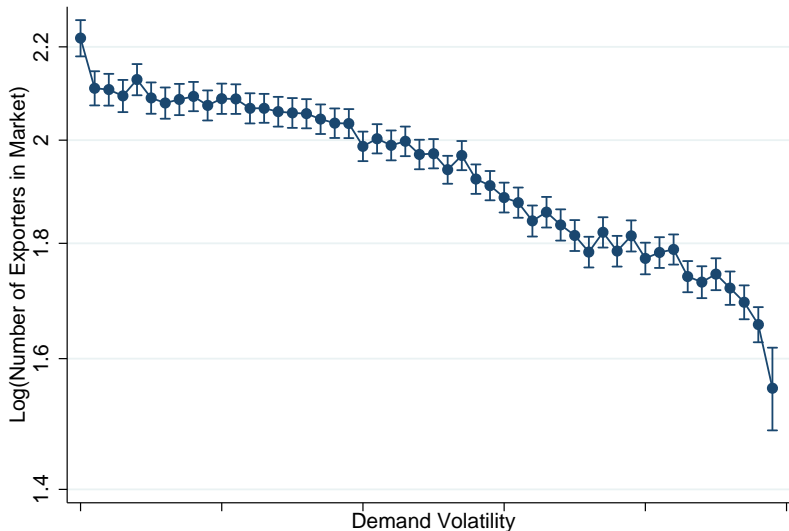
Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The dependent variable is a dummy equal to 1 for destinations with no Chinese exporter. The units of observation are destinations: unique HS6-product and country combinations.

ity, on average. The plot sets the logged number of exporters that served destinations between 2000 and 2006 against demand volatility. The predicted averages in the plot control for market size, country features and product-fixed effects. Each average is calculated separately for 50 equal-frequency bins of demand volatility.²⁶

Figure 1.2: Average Chinese Exporter Counts and Demand Volatility



Predicted unique exporter counts over 2000 - 2006. Estimated means and 95% confidence intervals after controlling for market size, country factors and HS6 product fixed effects. Observations are destinations grouped into 50 Quantiles, least to largest by demand volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

I estimate the following baseline specification:

$$\log(N_{jk}) = \beta_0 \sigma_{jk}^2 + \beta_1 X_{jk} + \alpha_j + \alpha_k + \epsilon_{jk} \quad (1.19)$$

X_{jk} = a vector of gravity model variables e.g., GDP, distance

α_j = product fixed effects

α_k = country fixed effects

To identify the effects of demand volatility on export entry, this reduced form specification plays on differences between destinations and the fact that certain products or countries tend to have higher volatility. It is necessary to control for product-specific factors; the number of potential entrants and the sunk costs of entry vary significantly along this dimension.

²⁶To interpret the graph, note that the average destination's logged export entry is 1.96 (with a standard deviation of 1.66). The curvature of the graph is addressed later in this section.

For example, between narrowly defined HS6 product categories, the number of exporting firms ranges from 1 for nuclear reactor fuel cartridges (HS 840130) to more than 40,000 for miscellaneous plastic articles (HS 392690). Estimates of the Pareto Distribution parameter also ranged from less than 5 to greater than 15, with varying degrees of fit for these product categories. Applying product fixed effects in the cross-section helps to address these differences.

Differences in export entry by country are expected, given factors like GDP, distance, language and currency. Furthermore, differences in volatility by country exist, as shown in Figure I.A.2. (Papers like di Giovanni and Levchenko (2012a); Koren and Tenreyro (2007) provide similar evidence). To ensure differences in export entry due to these factors are not conflated with demand volatility at the product-country level, I introduce either country fixed effects or direct measures of these variables (for the year 2006).

The specification with product fixed effects can simply be described as a comparison for a product like bicycles, using countries like Portugal and Greece that have similar GDP, GDP per Capita and distance from China, if bicycle imports into these countries differ in terms of demand volatility. The way the data is set up makes it possible to identify which country has the higher level of demand volatility for bicycles, knowing that the similar comparisons for other products are not guaranteed to be identical. Table 1.3 presents the results.

Fewer exporters enter destinations with high demand volatility, after controlling for the common predictors of exporter numbers. The observed number of exporters is about 5.1% lower on average for destinations one standard deviation above the mean. (The estimated effect in Table 1.7 is 11% when demand volatility is measured with quantities not values). Columns 1 and 2 show the number of unique exporters observed between 2000 and 2006 as the dependent variable; the other columns use gross entry after 2000, the first year in the data.

The results in Table 1.3 translate to about 2 fewer exporters in the average destination if demand volatility increases by one standard deviation. (The response is calculated as $\{1 - \exp[(-0.612) * 0.086]\}$). The next paragraphs describe the findings further. They also show how I identify the effects of demand volatility. Columns 3 and 4 use gross entry as the dependent variable, i.e. the number of new exporters in a destination after 2000. The columns have fewer observations because the log transform excludes destinations with no new exporters after 2000.²⁷ As previously discussed, gross entry in columns 3 and 4 measures

²⁷Columns 1 and 3 also have fewer observations due to missing GDP, distance or other control variables from the CEPII gravity dataset.

Table 1.3: Exporter Counts and Demand Volatility:
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.612*** (0.035)	-0.613*** (0.032)	-0.584*** (0.035)	-0.582*** (0.032)
Destination Size	0.421*** (0.003)	0.421*** (0.003)	0.412*** (0.003)	0.412*** (0.003)
Log(GDP)	0.001 (0.001)		0.001 (0.001)	
Log(GDP per capita)	0.001 (0.002)		0.001 (0.002)	
Log(Distance)	0.002 (0.006)		0.001 (0.006)	
Constant	-1.645*** (0.079)		-1.619*** (0.079)	
Observations	272,926	371,531	266,915	363,381
R-squared	0.547	0.545	0.546	0.544
Country FE		Y		Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures the net export entry that accompanied China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

how many exporters found a destination potentially profitable in the less restrictive trade regime that started in 2001, the year of China's WTO accession.

The last two columns of Table 1.3 agree in sign and significance with the first two, although the estimated coefficient on demand volatility decreases. Destination size takes away most of the statistical significance associated with gravity variables like GDP and distance. The variable, which I measure as the logged sum of imports between 1995 and 2005, represents both observed and projected demand growth, like Q_{jk}^* in the model. By its definition, it also addresses concerns that historical average growth rates affect exporter entry.²⁸ GDP may not adequately represent the size of specific destinations, even if it represents market size well in country-level gravity models. I do not show colonial relationships, WTO membership and other gravity variables in the table to conserve space. The gravity model variables all come from the year 2006.

The demand volatility measure varies in the cross-section across destinations, but not over time. (This also matches the model and baseline empirical specification, which described each destination destination with one demand volatility parameter and one measure of export entry). To justify the implicit assumption that exporter estimates are not updated over time, I provide two tests. First, in Section 1.3.2.3 I repeat the main regression using demand volatilities defined over various periods and show that the predicted outcomes remain largely unchanged. Second, Section 1.3.3.1 shows that changing the weights allocated to the deviation from trend in calculating demand volatility does not substantively alter the findings of the paper. These tests suggest that demand volatility is itself a stable feature of destinations, or that exporters do not substantially update their perceptions of destinations.

In identifying the effect of demand volatility, note that the demand data is a global aggregate, which mitigates concerns about reverse causation, as described in the variable definition. As mentioned earlier, product fixed effects capture differences in the γ adjustment parameter, the mass and distribution of exporters N_j and θ_j , as well as the setup costs and fixed costs associated with specific products. Country fixed effects also control for factors that include exchange rates, exchange rate volatility, country size, trade costs and policies like tariffs and trade agreements. Testing in the cross-section helps to avoid concerns about

²⁸ A possible challenge to the definition of this variable is that total absorption each destination includes imports and the domestic production. That poses no real problem for this paper, the fact that imports and domestic production are generally close substitutes within the narrow product categories suggests that imports can be used as a proxy for aggregate demand. Regressing destination demand on annual lagged demand gives an R^2 of 0.93, indicating that imports patterns are persistent or autocorrelated, even in the absence of data on domestic production.

other time-varying factors, as long as the variables I use are stable over the period under review. Column 1 allows the GDP, GDP per capita and other gravity variables to explain country-specific determinants of trade costs. The gravity model variables are either constant, like distance, or highly auto-correlated.²⁹ Columns 2 and 4 apply both country and product fixed effects simultaneously.³⁰ The country fixed effects absorb the gravity model variables, as expected.

The foregoing shows that Chinese exporters entered destinations with lower demand volatility in greater numbers. This is after accounting for product and country characteristics that account for the potential costs and profits from exporting. The results hold whether the dependent variable is a count of unique exporters or gross entry – the increase in unique exporter counts between 2001 and 2006. As I do not consider how entry into multiple markets may mitigate firm-level volatility, the estimated effects are on the conservative side.³¹

In describing whether trade levels fall with increasing demand volatility, one must consider how much of its predicted effect lies on the extensive margin – the number of exporters as shown above, or the intensive margin - exports per exporter. The model predicts that more of demand volatility’s effects are observed in exporter entry – the extensive margin. The higher expected prices associated with demand volatility imply lower expected demand and profits, given non-zero sunk costs. In a world with heterogeneous firms, those with lower productivity will generally self-select out of destinations with high demand volatility.

Table 1.4 presents the results. Higher demand volatility is associated with fewer exporters, (exporter numbers shows the largest coefficients and the highest level of explained variation in the table). Total exports from China summed across all years is lower for destinations with high demand volatility (columns 1 and 2); with fewer exporters as predicted, exports per exporter increase (columns 5 and 6). The coefficients in columns 3 and 5 sum to column 1, as the regressions are linear in logs. Destination size explains much of the

²⁹ Many small economies were missing GDP and GDP per capita, hence the differences in the number of observations between the even and odd-numbered columns

³⁰For computational efficiency, I follow the algorithm proposed by Guimarães and Portugal (2009) for multiple high-dimensional fixed effects. This algorithm iteratively estimates the coefficients, unlike conventional OLS estimation that directly calculates the matrix inverse. The coefficients represent a vector for the selected fixed effects and independent variables that yield the least squared error, within a 1e-6 tolerance. The effects are not fully interacted, as that would eliminate all degrees of freedom in the data.

³¹Most exporters serve more than one destination – and destinations are not perfectly correlated – one must consider that entering two destinations simultaneously generally yields a portfolio volatility that is lower than the demand volatility of either. (The upper bound of the portfolio volatility being the higher of the two markets). This indicates that this paper’s predicted effects are muted. When $\beta_0 = [\log(N_{jk}) - (\beta_1 X_{jk} + \alpha_j + \alpha_k)]/\sigma_{jk}$ in equation (1.19); if the true volatility perceived by exporters $\sigma_{jk}^* \leq \sigma_{jk}$, then the true coefficient $|\beta_0^*| \geq |\beta_0|$

variation in exporter counts in this table, just as in Table 1.3. Section I.B.2 develops the model to show a relationship between equilibrium trade levels and demand volatility; the relationships established in the model for both variables are found in this table.

Table 1.4: Exports and Exporter Counts vs. Demand Volatility
(Dependent Variable: Log Export Measure in Destinations)

VARIABLES	(1) Log(Exports)	(2)	(3) Log(Exporters)	(4)	(5) Log(Exports per Exporter)	(6)
Demand Volatility	-0.219*** (0.075)	-0.189*** (0.067)	-0.612*** (0.035)	-0.613*** (0.032)	0.392*** (0.057)	0.424*** (0.050)
Destination Size	0.834*** (0.005)	0.836*** (0.005)	0.421*** (0.003)	0.421*** (0.003)	0.413*** (0.003)	0.415*** (0.003)
Log(GDP)	0.001 (0.003)		0.001 (0.001)		-0.001 (0.002)	
Log(GDP per capita)	0.003 (0.004)		0.001 (0.002)		0.002 (0.003)	
Log(Distance)	0.008 (0.011)		0.002 (0.006)		0.006 (0.008)	
Constant	4.176*** (0.159)		-1.645*** (0.079)		5.821*** (0.113)	
Observations	272,926	371,531	272,926	371,531	272,926	371,531
R-squared	0.490	0.488	0.547	0.545	0.390	0.387
Country FE		Y		Y		Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exports in columns 1-2 are for all Chinese exporters. Exporters in columns 3-4 is the log of unique firms with recorded exports to a destination between 2000 and 2006. It is the same number used for columns 1-2 of Table 1.3. Destination size is the log of total demand from a destination between 1995 and 2005. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

In sum, trade levels are lower for destinations with high demand volatility, and most of the effect comes from the extensive margin, represented by columns 3 and 4. Columns 1 and 2 of the table are consistent with the prediction in equation (1.29). The imperfect matching of countries between the trade and CEPII gravity data set means that GDP and

GDP per capita are missing for many observations. (Section I.A.1 describes the matching). Unreported regressions on the even-numbered columns give nearly identical coefficients on a sample restricted to those with no missing variables in the odd-numbered columns.

To assuage concerns that counts of unique exporters over multiple years may not be comparable to conventional estimates of gravity models with annual export measures, Appendix Section I.A.4 presents regressions that include estimates with annual exporter counts, annual control variables and country-year fixed effects. These show that the reported estimates are not due to periodic shocks, or exchange rate volatility – other potential drivers of trade in the literature. The first two columns of Table 1.4 represent firm level gravity model regressions, (as do the first two of Table I.A.4). The two tables show that fewer exporters enter destinations with high demand volatility.

Finding a consistent pattern of lower exporter counts with demand volatility suggests that prices will be higher in those destinations. However, reliable tests of demand volatility’s effect on prices are difficult with no data on quality, given the well-documented link between prices and quality (Hallak and Schott, 2011; Hallak and Sivadasan, 2009). Regressions of price on demand volatility yield statistically insignificant coefficients for the largest product categories. This was after including various sets of controls that included gravity model variables, product-year fixed effects, country year fixed effects and firm fixed effects. (See appendix section I.A.8.)

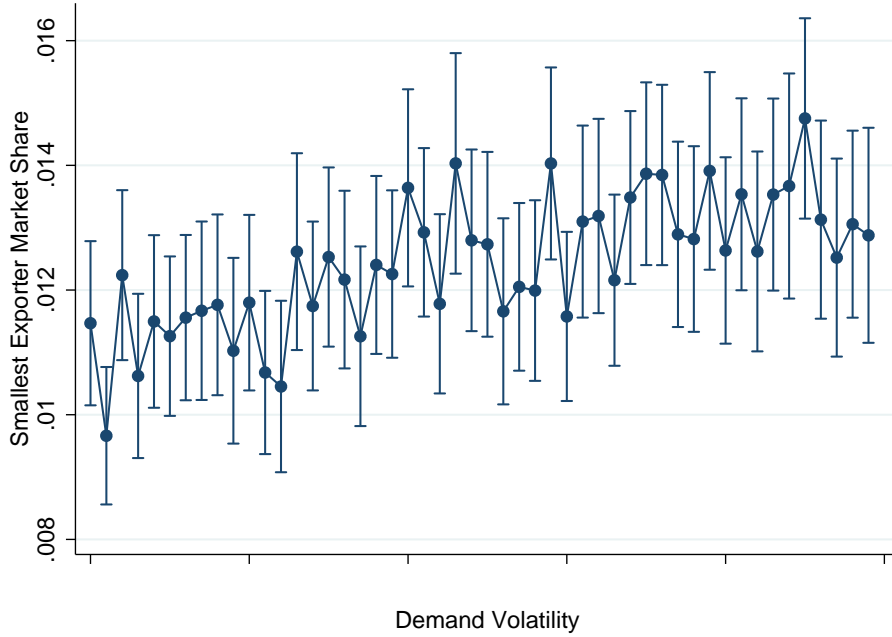
Demand volatility compares favorably with conventional variables like GDP and geographic distance in predicting trade. I run separate regressions (reported in Table I.A.7) of exporter counts on demand volatility, GDP, geographic distance and destination size, in the absence of additional controls. The omitted-variable regressions involve only the dependent variable, the selected variable and product fixed effects. The respective adjusted R^2 values are 0.29, 0.22, 0.22 and 0.54. Only destination size explains more exporter count variation on this crude test.

1.3.2.2 Demand Volatility and Exporter Productivity Thresholds

Figure 1.3 supports the model’s claim that demand volatility filters out producers with low productivity. The high costs of adjustment mean that in general, fewer firms will have the profit margins required to succeed in destinations with high volatility. (As the data allows no direct measures of productivity, I use producers’ market shares within product categories as a proxy).

To facilitate comparisons in the cross-section, I compute each exporter’s share of Chinese

Figure 1.3: Minimum Export Thresholds and Demand Volatility



The minimum market share of the exporters in each destination jk . I use firms' share of exports within a narrow HS6 product category as a proxy for productivity, so that destinations served by firms with the least market share get the lowest value of this proxy for ϕ_{jk}^* . The y-axis averages these minimum market share measure for all destinations in each quantile. Destinations Grouped Into 50 Quantiles, Least to Largest by Demand Volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

exports in 2006 within its HS6 category: $Share_{ij} = \frac{q_{ij}}{\sum_i q_{ij}} = \frac{\sum_k q_{ijk}}{\sum_i \sum_k q_{ijk}}$. I represent the productivity threshold ϕ_{jk}^* by the smallest market share recorded by any firm in destination jk . That is, for each destination, ϕ_{jk}^* is represented by $\min(Share_{ij})$. Therefore, destinations served by only the largest exporter in the product category will report a higher threshold than the destination served by both the largest and smallest exporter, (if more than one firm exports the product from China).

I use simple regression methods to check whether this measure of productivity thresholds increases with demand volatility. The regressions mimic equation (1.19), replacing the demand volatility measure with a dummy for each of 50 equal-frequency bins for destinations, ranked by demand volatility. Product fixed effects control for differences in the distributions of market shares by product. Section I.B.3 in the Appendix includes additional plots that use other proxies for firm-level productivity. That section also includes results of regression exercises that support the finding of higher productivity thresholds with increasing demand volatility.

The plot shows an increasing trend in the predicted size-rank of exporters with increasing demand volatility; the standard errors suggest that the trend is statistically significant. In other words, the destinations with the highest demand volatility have on average, exporters that are among the largest producers for the related product. This is after controlling for destination size, and variables like GDP and Distance. This supports the mechanism proposed in the model that expectations of adjustment costs drive more low-productivity firms away from the most volatile destinations.

1.3.2.3 Temporal Variations in Demand Volatility

The previous test steps assume that expected demand volatility does not change over time from the exporter's perspective. As this assumption is important to how I simplify the estimation, I test it using the specification in (1.20):

$$\log(N_{jk}) = \beta_0 \sigma_{jk}^{t_1-t_2} + \beta_1 X_{jk} + \alpha_j + \alpha_k + \epsilon_{jk} \quad (1.20)$$

$$\sigma_{jk}^{t_1-t_2} = \text{Demand volatility for years } t_1 \text{ to } t_2$$

the other terms mirror the definitions in equation (1.19)

The variants of $\sigma_{jk}^{t_1-t_2}$ use (1995,2000), (2000,2005) and (2003,2008) as (t_1, t_2) pairs. These shorter demand histories yield less precise estimates. The measure for (1995,2000) gives more weight to historical demand before exporters in the data made their entry choices. The two additional 6-year history samples draw on years that put more weight on in-sample and after-the-sample data, i.e. 2000-2005 and 2003-2008. (I stop at 2008 because the severe drop in trade across many products for 2009 is exceptional).

Table 1.5 shows that the assumption of stable expected demand volatility is not far-fetched. The coefficient of the 1995-2000 measure is only one standard deviation away from the 2003-2008 measure. The standard errors and explained variations are similar across all three measures of demand volatility, whether the measure is weighted toward the past, or toward the future. While the size of demand volatility's coefficient is higher for the mid-data sample, the sign and significance remain unchanged.

Compared to the predicted 5.1% change in exporter numbers from Table 1.4, the predicted declines in exporter numbers using only 1995-2000 to measure demand volatility drops to 1.6% if only years 1995-2000 were used to form expectations. The respective predictions for the other sample periods are 4.1 and 1.9%. Correlations between the measures of demand volatility are 0.37 (95-00, 00-05), 0.44 (95-00,03-08) and 0.59 (00-05,03-08).

Table 1.5: Export Entry with Past and Future Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3)	(4)	(5)	(6) Log(Gross Entry Post-2000)
Dem. Volat.(95-00)	-0.182*** (0.029)			-0.136*** (0.029)		
Dem. Volat.(00-05)		-0.493*** (0.030)			-0.483*** (0.031)	
Dem. Volat.(03-08)			-0.220*** (0.033)			-0.235*** (0.033)
Destination Size	0.435*** (0.003)	0.424*** (0.003)	0.430*** (0.003)	0.426*** (0.003)	0.415*** (0.003)	0.420*** (0.003)
Observations	356,899	369,240	368,012	349,331	361,372	360,395
R-squared	0.545	0.544	0.544	0.545	0.543	0.543
Country-Year FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. Gross entry post-2000 captures the difference between all firms that served a product market and firms that served the destination in 2000.

In sum, the paper’s main qualitative findings are robust to the period used for measuring demand volatility. Had the model and estimates included updating of expectations by exporters, the findings are expected to remain consistent. (Tables 1.3 and 1.5 broadly agree in terms of the sign and significance of the estimated coefficients of demand volatility). Section 1.3.3.1 provides further evidence – using the demand history from 1995-2005, but applying different weights to each year.

1.3.2.4 Non-Linear Responses to Demand Volatility

The non-linear form of σ^2 in equation (1.14) suggests a test of the dependent variable on higher orders of demand volatility, as well as its interactions with market size, among other variables.

Table 1.6 includes the square of demand volatility as a regressor, as well as interactions with destination size. (If the γ term is constant, the model implies a positive coefficient on the squared volatility term. However, negative coefficients are possible if the adjustment costs term is not a simple linear function of deviations from expectation).

The results confirm expectations of non-linearity. The coefficients of the squared demand volatility term are statistically significant in all specifications. However, unlike what a naive interpretation of the model would suggest, Table 1.6 indicates a steeper decline in exporter counts at the higher levels of demand volatility. This suggests a more complex structure to adjustment costs than I outline in the model. Concave relationships between demand volatility and the log of exporter numbers are possible with adjustment costs that depend on aggregate demand, or adjustment costs that are non-linear functions of squared deviations from projected trend. (If demand volatility enters exporters’ considerations of profit strictly as described in equation (1.14) and γ is a constant, the curve should be convex). The more remarkable finding is that the higher order terms are non-zero and statistically significant.

Interactions with destination size also show consistency with the main predictions. While the coefficient of the linear demand volatility term switches signs, the corresponding change associated with the market size interaction is larger. The squared volatility measure remains statistically significant and negative. Columns 4 to 6 repeat the same pattern for gross entry from 2001.

This section’s results corroborate the predictions of the model for the effects of demand volatility on the number of exporters. The estimated effects of demand volatility on exporter choice are non-linear, and more exporters enter destinations with low demand volatility, holding other factors equal. The next steps check the robustness of the main findings.

Table 1.6: Exporter Counts and Interacted Terms of Demand Volatility
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Log(Gross Export Entry)			Log(Gross Entry Post-2000)		
Demand Volatility	-0.613*** (0.032)	-0.301*** (0.073)	3.294*** (0.142)	-0.582*** (0.032)	-0.231*** (0.074)	3.321*** (0.144)
(Demand Volatility) ²		-0.687*** (0.131)	-0.529*** (0.135)		-0.777*** (0.132)	-0.637*** (0.135)
Destination Size *Demand Volatility			-0.503*** (0.016)			-0.495*** (0.017)
Destination Size	0.421*** (0.003)	0.424*** (0.003)	0.442*** (0.003)	0.412*** (0.003)	0.415*** (0.003)	0.433*** (0.003)
Observations	371,531	371,531	371,531	363,381	363,381	363,381
R-squared	0.545	0.545	0.548	0.544	0.544	0.547
Country FE	Y	Y	Y	Y	Y	Y
Product FE	Y	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a product market between 2000 and 2006. Destination size is the log of total demand from a destination between 1995 and 2005. The coefficients suggest that γ in equation (1.14) is a non-linear functions of squared deviations from projected trend. If γ was a constant, as described in the model, the coefficient of the squared σ^2 term should be positive.

1.3.3 Robustness Checks

This section addresses potential concerns about how exporter counts and demand volatility are measured. The appendix sections I.A.4 to I.A.6 include additional tests.

1.3.3.1 Demand Volatility Weighted by Recency

In estimating volatility, firms may ascribe greater weight to recent information (Bloom et al., 2007). Therefore, recent shocks may carry a disproportionate share of exporter’s demand volatility estimates, (or less in times of high uncertainty). Figure 1.4 tests the idea by plotting the coefficient and R^2 values obtained for definitions of demand volatility with different weight indices η . The weights w_t , indexed from 1 to 10 put more emphasis on recent information with higher values of η ; setting η to 1 reverts to the default scheme of equal weights. By design, the weighting scheme does not affect destinations with uniform deviations from the demand trajectory in all periods. (η specifies the relative size of the first and last terms of an arithmetic series that sums to H . H is the most recent period).

$$\sigma_{jk}^2 = \sum_{t=1}^H w_t (\epsilon)^2, \quad \epsilon \text{ is the residual defined in equation (1.17):} \quad (1.21)$$


$$w_t = \frac{2}{\eta + 1} + \frac{2(t - 1)(\eta - 1)}{(H - 1)(\eta + 1)}, \quad \text{for } t \in [H, 1], \text{ where } \eta \in [1, 10] \text{ and } \sum w_t = H = 11 \quad (1.22)$$

To create the figure, I repeat the baseline regression in Table 1.3 for each weighted variant of σ^2 and collate the estimated coefficients and R^2 s.

The left panel of Figure 1.4 indicates that increasing the weight of recent information in the estimation of demand volatility does not increase the explained variation in the number of exporters serving a destination. (The R^2 remain between 0.545 and 0.546). The right panel remains consistent. The estimated coefficient of demand volatility on exporter counts remains statistically significant, but decreases slightly in scale from the default estimate of 0.06 to 0.04 for index 10, the volatility estimate that ascribes 10 times the weight of the first year (1995) to the most recent year of demand history (2005).

These results suggest that the particular period used to estimate demand volatility does not affect the paper’s main qualitative predictions.

Figure 1.4: Predictions for Weighted Demand Volatility



(a) R^2 values

Estimates after controlling for Market Size, Country and HS6 Product fixed effects. The demand volatility estimated with weight index 1 assigns equivalence to all periods, (as do all tables in the paper). The demand volatility estimated with weight 10 assigns deviations in 2005 10 times the weight of deviations in 1995. The weights are a linear density function. Data Sources: China GAC Export Data (2000-2006), UN COMTRADE

(b) Coefficient β_0 on demand volatility

1.3.3.2 Alternative Demand Volatility Measures

Table 1.7 replicates Table 1.3, measuring demand volatility as the sum of squared deviations from trend for quantities, not values in each destination. There are advantages and disadvantages to this approach: using quantities helps the econometrician avoid conflating exchange rates and price volatility with demand volatility, and quantities are a more faithful representation of demand volatility from the model in Section 1.2. However, quantity data is not always as reliable or available as data on trade values. The customs authorities that collect trade data have a stronger incentive to collect accurate information on values, and units of measurement for quantities are not always reported consistently by the original trade data sources (Gaulier and Zignago, 2010).

The two measures of demand volatility are nevertheless similar, with the quantity measure having a mean and standard deviation of 0.085 and 0.130, and the value-measure having 0.053 and 0.086. The quantity-based volatility measure is larger, as expected, given the inverse correlation between prices and quantities. That inverse correlation ensures that the volatility of prices and exchange rates, if any, reduces the observed volatility of export values relative to the volatility of quantities.

The coefficients in Table 1.7 are consistent with the numbers in Table 1.3. The estimated

Table 1.7: Export Entry with Quantity-Based Demand Volatility Measure
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.882*** (0.022)	-0.889*** (0.021)	-0.866*** (0.022)	-0.871*** (0.021)
Destination Size	0.391*** (0.003)	0.392*** (0.003)	0.383*** (0.003)	0.384*** (0.003)
Log(GDP)	0.002 (0.001)		0.001 (0.001)	
Log(GDP per capita)	0.000 (0.002)		0.000 (0.002)	
Log(Distance)	0.002 (0.006)		0.001 (0.006)	
Constant	-0.666*** (0.077)		-0.663*** (0.077)	
Observations	280,060	381,196	273,567	372,373
R-squared	0.539	0.538	0.538	0.536
Country-Year FE		Y		Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 captures the difference between all firms that served a destination and firms that served the destination in 2000. Demand volatility is constructed as in Table 1.3, but with quantities not dollar values of demand from 1995-2005 in each destination. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

effect sizes are also larger: increasing this measure of demand volatility by one standard deviation corresponds to a 10.7% decline in exporter numbers, (calculated as $1 - \exp(-0.88 * 0.130)$). The quantity-based definition yields estimates with greater economic and statistical significance than the value-based definition of demand volatility. This supports previous statements in the paper that the estimated effects in Table 1.3 tend to be conservative.

Table 1.8 replicates Table 1.3, measuring demand volatility as the standard deviation of annual growth rates of destinations. Section 1.2 summarizes the intuition behind this new measure of demand volatility. For a destination with a constant growth rate, measured demand volatility and adjustment costs would be zero, as the constant growth rate would match exporters' projections. The standard deviation of period-to-period growth is a common proxy for volatility, from finance and microeconomics (Guiso and Parigi, 1999; Sharpe, 1966) to macroeconomics and international trade (di Giovanni and Levchenko, 2009; Koren and Tenreyro, 2007). It may not fit the model in this paper exactly, but it is fairly simple to derive and explain. di Giovanni and Levchenko (2012b) use this measure to construct export-weighted average measures of risk for countries, in a study that links trade volatility to aggregate macroeconomic volatility. Koren and Tenreyro (2007) apply the same definition to GDP volatility in explaining differences between rich and poor countries. I include tests using this measure for consistency with related works on trade and volatility.

Table 1.8 shows that the estimated relationship between exporter numbers and demand volatility is robust to using this popular alternative measure. For each destination, I measure demand volatility as the standard deviation of demand growth between 1995 and 2005. The mean and standard deviation of this new measure are 0.78 and 0.55 respectively. (Differences between these numbers and Table 1.1 are indicative of how each measure is constructed).

The coefficients in Table 1.8 are consistent with the numbers in Table 1.3, both in sign and significance. The predicted effects sizes are also similar: a 4.8% decline in exporter numbers is expected for the average destination for a standard deviation increase in this measure of demand volatility, compared with 5.1% in the first table. In sum, this new definition yields results that broadly agree with the former definition of demand volatility.

I replicate Tables 1.4 to 1.9 using this measure and obtain the same pattern: slightly smaller coefficients with the same sign and significance as the original measure.

Section I.A.5 in the Appendix presents results that show the main findings are robust to the selection of markets used to define demand volatility. That section replicates Table 1.3 using two sets of demand histories to define demand volatility: product-country volatility for Chinese goods, and destination volatility for the rest of the world. The first definition of

Table 1.8: Export Entry with Alternative Demand Volatility Measure
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.095*** (0.011)	-0.152*** (0.012)	-0.082*** (0.010)	-0.134*** (0.012)
Destination Size	0.283*** (0.004)	0.240*** (0.003)	0.278*** (0.004)	0.236*** (0.003)
Log(GDP)	0.224*** (0.003)		0.224*** (0.003)	
Log(Distance)	-0.540*** (0.005)		-0.514*** (0.005)	
Constant	6.410*** (0.092)		6.030*** (0.091)	
Observations	311,337	362,490	304,545	354,337
R-squared	0.710	0.748	0.708	0.745
Product FE	Y	Y	Y	Y
Country FE		Y		Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Exporter counts represent the log of the number of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 captures the difference between all firms that served a destination and firms that served the destination in 2000. Demand volatility is constructed as the standard deviation of destinations' annual growth rates from 1995-2005. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership.

volatility corresponds to a model where Chinese exporters' expectations of demand volatility is linked to only the history of goods imported from China for any given destination. The definition of demand volatility that excludes imports from China follows an alternative scenario in which expectations neither correspond to firms' own historic sales to a destination, nor to past supply shocks from China.

The two sets of results in section I.A.5 yield the same finding as the previous tables in the paper, we see that destinations with high demand volatility are served by fewer exporters. The coefficients for the main variables in the tables retain their sign and statistical significance, though there are small differences in the R^2 values, as with the number of observations. This is not unexpected, as the number of destinations with years of usable history changes.

1.3.3.3 Poisson Regressions to Include Zeros

Tables 1.3 to 1.8 ignored destinations with zero exporters – the estimations used logarithms of a count variable. Table 1.9 addresses concerns of possible bias from ignoring these destinations with a Poisson regression. The inclusion of observations with zero entries after 2000 should accentuate the claims made in the previous section because demand volatility is higher on average for destinations with zero entry, as shown in Figure I.A.1. Those destinations are part of this new regression.

The estimates using Poisson regressions with fixed effects are consistent with those from Tables 1.3, and larger, as expected. Interpreting the coefficients in Table 1.9 suggests that a 19% decline in the number of exporters should be associated with a standard deviation increase in demand volatility from the mean, holding other factors constant. (The response is calculated as $\{1 - \exp[(-2.5) * 0.086]\}$, using negative terms to represent the negative predicted effect). The estimates are higher than the comparable numbers in Table 1.4. The destinations with no Chinese exporters had higher levels of demand volatility and the estimates are consistent with expectations in terms of size, sign and significance.

1.4 Discussions and Conclusion

This paper examines how the margins of trade are shaped by observed volatility – the patterns of exogenous fluctuations in aggregate demand from potential export destinations. The nature of these exogenous fluctuations is well documented (Carroll et al., 2011; Hall, 2004). Demand volatility has country and product-specific patterns, such that imports for

Table 1.9: Exporter Counts and Demand Volatility: Poisson Estimates
(Dependent Variable: Number of Exporters in Destinations)

VARIABLES	(1) Gross Export Entry	(2) Gross Entry Post-2000
Demand Volatility	-2.518*** (0.092)	-2.408*** (0.093)
Destination Size	0.288*** (0.006)	0.276*** (0.006)
Observations	688,005	686,997

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women’s cotton overcoats (HS 610220). All 708,000 destinations with measurable demand volatility were included, many of which had zero Chinese exporters. Product and country fixed effects limited the usable observations to about 690,000. Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The gross entry post-2000 variable captures the unique number of entrants from the year China joined the WTO.

some raw materials like steel are particularly volatile for most countries, and countries with low GDP per capita generally have high import volatility, as documented by papers that link national output volatility to GDP per capita.³² Nevertheless, neither product categories nor countries as isolated variables explain most of the variation in the demand volatility of destinations – unique combinations of products and countries.

I provide evidence that higher demand volatility in a destination leads to fewer Chinese exporters and lower trade volumes. Further, I show that a notable share of the variation in trade is explained by observed volatility. The R^2 figures explaining the variable’s relevance at the beginning of the paper come from Table I.A.7 in the appendix, which shows that demand volatility’s explanatory power compares reasonably with those of established variables like GDP and distance.

In showing a statistically significant first-order relationship between the extensive margin of trade and demand volatility, this paper suggests a new factor that explains observed trade

³²Papers in this vein hypothesize that aggregate volatility is due to the correlation between national policy shocks and sectoral shocks, the capacity of economic institutions, as well as specialization in sectors that are unusually volatile (di Giovanni and Levchenko, 2012b; di Giovanni and Levchenko, 2009; Krishna and Levchenko, 2009; Koren and Tenreyro, 2007).

volumes, prices and quantities. In explaining trade, demand volatility complements existing explanations that rely on exporter productivity, geography and economic size (Anderson and Van Wincoop, 2003; Melitz, 2003). As I show in Section I.A.7, import volatility is higher in developing economies, consistent with the findings of di Giovanni and Levchenko (2012b) and Koren and Tenreyro (2007). The paper’s main finding suggests that these economies may be hurt by the absence of imported varieties.

The estimates of exporter responses to demand volatility use Chinese firm-level export data. This choice was necessitated in part by the absence of a dataset that reports exporter numbers for all countries. Furthermore, China is the world’s largest exporter, and its exports are highly correlated with global exports. (The correlation is 0.995 for China’s aggregate exports and the rest of the world’s between 2000 and 2006). Therefore, the empirical exercises in the chapter should be informative about the global pattern of exporter behavior, which is what the model describes.

As imports can be a source of productive efficiency in developing economies (Goldberg et al., 2010; Connolly, 2003), the foregoing implies that low demand volatility can be an advantage for producers. Firms in locations with low demand volatility can enjoy more variety in inputs and possibly lower prices. This matters because imported inputs are a common feature of production processes – for developing and developed economies alike. One could extend this idea to assessing the impacts of trade diversification policies designed to reduce aggregate economic volatility, as discussed in Cadot et al. (2011).

Rational economic agents respond to volatility, and firms engaged in international trade are no exception. This paper extends the conventional model of trade to include adjustment costs in the presence of stochastic demand. (Adjustment costs reflect the fact that demand shocks affect unit costs: firms must pay overtime wage premiums and equipment deactivation costs when demand either peaks or crashes). These costs motivate exporters’ choices, with clear implications for the number of firms that export, and the level of trade between countries.

Firm level data on the universe of Chinese export transactions between 2000 and 2006 support this hypothesis: Fewer exporters serve destinations with high demand volatility. The results are robust to how exporter numbers are counted – as the unique number of exporters in all years, as the unique number of new exporters after 2000 when China adopted liberalization policies in preparation for WTO membership, or as annual counts of exporters in each of the years. The results are also robust to how demand volatility is defined, and to several empirical specifications. Furthermore, the data confirm the model’s predictions

about how the relationship between exporter counts and demand volatility is non-linear.

The estimated effects of demand volatility are statistically and substantially significant. Increasing demand volatility by one standard deviation for the average destination in my conservative specification predicts a 5% drop in the number of exporters that serve the destination. The explained variation in exporter counts and predicted effects of this new variable are comparable to those obtained from conventional predictors of trade like GDP and distance.

These findings suggest further work to evaluate how volatility affects the development process, given how instrumental trade has been to growth in the last half-century. The paper also motivates an inquiry into how differences in labor and capital adjustment costs explain the heterogeneity in producers' responses to demand volatility.

Bibliography

- Ahn, J., Khandelwal, A. K., Wei, S.-J., 2011. The Role of Intermediaries in Facilitating Trade. *Journal of International Economics* 84 (1), 73–85.
- Ahn, J., McQuoid, A. F., 2012. Capacity Constrained Exporters: Micro Evidence and Macro Implications. Working Paper.
- Akhmetova, Z., Mitaritonna, C., 2012. A Model of Firm Experimentation under Demand Uncertainty with an Application to Multi-Destination Exporters.
- Amiti, M., Freund, C., 2010. The anatomy of china's export growth. In: *China's Growing Role in World Trade*. University of Chicago Press, pp. 35–56.
- Anderson, J. E., Van Wincoop, E., 2003. Gravity with Gravitas: A Solution to the Border Puzzle. *The American Economic Review* 93 (1), 170–192.
- Axtell, R. L., 2001. Zipf Distribution of US Firm Sizes. *Science* 293 (5536), 1818–1820.
- Bas, M., 2009. Trade, Foreign Inputs and Firms Decisions: Theory and Evidence. Working Paper.
- Bloom, N., Bond, S., Van Reenen, J., 2007. Uncertainty and Investment Dynamics. *Review of Economic Studies* 74 (2), 391–415.
- Blum, B. S., Claro, S., Horstmann, I. J., 2013. Occasional and Perennial Exporters. *Journal of International Economics* 90 (1), 65–74.
- Cadot, O., Carrère, C., Strauss-Kahn, V., 2011. Export Diversification: What's Behind the Hump? *Review of Economics & Statistics* 93, 590–605.
- Carroll, C. D., Slacalek, J., Sommer, M., 2011. International Evidence on Sticky Consumption Growth. *Review of Economics And Statistics* 93 (4), 1135–1145.
- Chaney, T., 2008. Distorted Gravity: the Intensive and Extensive Margins of International Trade. *American Economic Review* 98 (4), 1707–1721.
- Chevassus-Lozza, E., Gaigné, C., Le Mener, L., 2013. Does Input Trade Liberalization Boost Downstream Firms' Exports? Theory and Firm-Level Evidence. *Journal of International Economics* 90 (2), 391–402.
- Connolly, M., 2003. The Dual Nature of Trade: Measuring its Impact on Imitation and Growth. *Journal of Development Economics* 72 (1), 31–55.
- Cooper, R. W., Haltiwanger, J. C., 2006. On the Nature of Capital Adjustment Costs. *Review of Economic Studies* 73 (3), 611–633.

- Crozet, M., Koenig, P., 2010. Structural Gravity Equations with Intensive and Extensive Margins. *Canadian Journal of Economics* 43 (1), 41–62.
- Cuñat, A., Melitz, M. J., 2012. Volatility, Labor Market Flexibility, and The Pattern Of Comparative Advantage. *Journal of the European Economic Association* 10 (2), 225–254.
- Das, S., Roberts, M. J., Tybout, J. R., 2007. Market Entry Costs, Producer Heterogeneity, and Export Dynamics. *Econometrica* 75 (3), 837–873.
- Davis, S. J., Haltiwanger, J., 1992. Gross Job Creation, Gross Job Destruction, and Employment Reallocation. *Quarterly Journal of Economics* 107 (3), 819–863.
- di Giovanni, J., Levchenko, A. A., 2009. Trade Openness and Volatility. *Review of Economics and Statistics* 91 (3), 558–585.
- di Giovanni, J., Levchenko, A. A., 2012a. Country Size, International Trade, and Aggregate Fluctuations in Granular Economies. *Journal of Political Economy* 120 (6), 1083–1132.
- di Giovanni, J., Levchenko, A. A., 2012b. The Risk Content of Exports: A Portfolio View of International Trade. In: *NBER International Seminar on Macroeconomics*. Vol. 8. University of Chicago Press, pp. 97–151.
- Dixit, A., 1989. Entry and Exit Decisions Under Uncertainty. *Journal of Political Economy*, 620–638.
- Eaton, J., Kortum, S., 2001. Trade in Capital Goods. *European Economic Review* 45 (7), 1195 – 1235, international Seminar On Macroeconomics.
- Eaton, J., Kortum, S., Kramarz, F., 2004. Dissecting Trade: Firms, Industries, and Export Destinations. *American Economic Review* 94 (2), 150–154.
- Foster, L., Haltiwanger, J., Syverson, C., 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98 (1), 394–425.
- Frankel, J., Rose, A., 2002. An Estimate of the Effect of Common Currencies on Trade and Income. *Quarterly Journal of Economics*, 437–466.
- Gaulier, G., Zignago, S., 2010. BACI: International Trade Database at the Product-Level. The 1994-2007 Version. CEPII Working Papers (Working Papers 2010-23).
- Glick, R., Rose, A. K., 2002. Does a Currency Union Affect Trade? The Time-Series Evidence. *European Economic Review* 46 (6), 1125–1151.
- Goldberg, P. K., Khandelwal, A. K., Pavcnik, N., Topalova, P., 2010. Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. *Quarterly Journal of Economics* 125 (4), 1727–1767.

- Guimarães, P., Portugal, P., 2009. A Simple Feasible Alternative Procedure to Estimate Models with High-Dimensional Fixed Effects. SSRN 1329573.
- Guiso, L., Parigi, G., 1999. Investment and Demand Uncertainty. *Quarterly Journal of Economics* 114 (1), 185–227.
- Hall, R. E., 2004. Measuring Factor Adjustment Costs. *The Quarterly Journal of Economics*, 899–927.
- Hallak, J. C., Schott, P. K., 2011. Estimating Cross-Country Differences in Product Quality. *Quarterly Journal of Economics* 126, 417–474.
- Hallak, J. C., Sivadasan, J., 2009. Firms’ Exporting Behavior Under Quality Constraints. NBER Working Paper 14928.
- Handley, K., Limão, N., 2012. Trade and Investment Under Policy Uncertainty: Theory and Firm Evidence. NBER Working Paper 17790.
- Handley, K., Limão, N., 2013. Policy Uncertainty, Trade and Welfare: Theory and Evidence for China and the US. NBER Working Paper 19376.
- Harrigan, J., Ma, X., Shlychkov, V., 2011. Export Prices of US Firms. NBER Working Paper.
- Head, K., Mayer, T., Ries, J., 2010. The Erosion of Colonial Trade Linkages after Independence. *Journal of International Economics* 81 (1), 1–14.
- Helpman, E., Melitz, M., Rubinstein, Y., 2008. Estimating Trade Flows: Trading Partners and Trading Volumes. *The Quarterly Journal of Economics* 123 (2), 441–487.
- Helpman, E., Melitz, M. J., Yeaple, S. R., 2004. Export versus FDI with Heterogeneous Firms. *The American Economic Review* 94 (1), 300–316.
- Hsieh, C.-T., Ossa, R., 2011. A Global View of Productivity Growth in China. NBER Working Paper 16778.
- Hummels, D., Klenow, P., 2005. The Quality and Variety of a Nations Trade. *American Economics Review* 95 (3), 704–723.
- Jones, C. I., 2011. Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics*, 1–28.
- Koren, M., Tenreyro, S., 2007. Volatility and Development. *Quarterly Journal of Economics* 122 (1), 243–287.
- Krishna, P., Levchenko, A. A., 2009. Comparative Advantage, Complexity and Volatility. NBER Working Paper 14965 (w14965).

- Lee, J.-W., 1995. Capital Goods Imports and Long-run Growth. *Journal of Development Economics* 48 (1), 91–110.
- Liu, Y., 2012. Capital Adjustment Costs: Implications for Domestic and Export Sales Dynamics. Working Paper.
- Lucas, R. E., 1967. Adjustment Costs and the Theory of Supply. *Journal of Political Economy*, 321–334.
- Manova, K., Yu, Z., 2012. Firms and Credit Constraints along the Global Value Chain: Processing Trade in China. NBER Working Paper 18561.
- Manova, K., Zhang, Z., 2012. Export Prices Across Firms and Destinations. *Quarterly Journal of Economics* 127 (1), 379–436.
- Melitz, M. J., 2003. The Impact Of Trade On Intra-Industry Reallocations And Aggregate Industry Productivity. *Econometrica* 71 (6), 1695–1725.
- Milgrom, P., Segal, I., 2002. Envelope Theorems for Arbitrary Choice Sets. *Econometrica* 70 (2), 583–601.
- Nguyen, D. X., 2011. Demand Uncertainty: Exporting Delays and Exporting Failures. *Journal of International Economics* 86 (2), 336–344.
- Pindyck, R. S., 1982. Adjustment Costs, Uncertainty, and The Behavior Of The Firm. *American Economic Review*, 415–427.
- Rob, R., Vettas, N., 2003. Foreign Direct Investment and Exports with Growing Demand. *Review of Economic Studies* 70 (3), 629–648.
- Sharpe, W. F., 1966. Mutual Fund Performance. *Journal of Business*, 119–138.
- Soderbery, A., 2013. Market Size, Structure, and Access: Trade with Capacity Constraints. Working Paper.
- Staiger, R. W., Wolak, F. A., 1992. The Effect of Domestic Antidumping Law in the Presence of Foreign Monopoly. *Journal of International Economics* 32 (3), 265–287.
- Vannoorenberghe, G., 2012. Firm-level Volatility and Exports. *Journal of International Economics* 86 (1), 57–67.
- Viner, J., 1922. The Prevalence of Dumping In International Trade. *Journal of Political Economy* 30 (5), 655–680.

Appendix

I.A Empirics: Data, Variable Definitions and Supplementary Tests

I.A.1 Exporter-Level Data

Firm level export choices are taken from the universe of Chinese export transactions between 2000 and 2006, collapsed to annual values for firms in each product-country. This dataset identifies the year of trade, firms by unique IDs, the countries to which they export, and the products sent to each destination. The full data set exceeds 24 million observations. The raw data report the f.o.b. value of exports in nominal U.S. dollars in an unbalanced panel over 7 years of more than 240,000 firms serving 390,000 destinations comprising 200 economies and about 4,100 HS6 product categories. This rich dataset provides no identifiers for buyers in overseas markets, unfortunately. Thus, each destination conceptually stands for one representative consumer.

To link this data to the destinations identified in the COMTRADE data, I match the two sources on countries and product categories. The product categories are originally reported as eight-digit HS categories in the firm-level data, of which the leading 6 digits correspond to standardized categories.³³ To ensure that the product category definitions remain consistent over time, I convert all years to the 1992 HS standard using the concordances provided by the UN at <http://unstats.un.org/unsd/trade/conversions/HSCorrelationandConversiontables.htm>. The destination data that I describe next are reported using the 1992 HS standard. I complete the matching between the two data sources by mapping the country-codes in the export data to the standardized ISO categories used by COMTRADE.

For Chinese exporters, 2000 to 2006 was a period of remarkable growth and entry into

³³The 6-digit harmonized system (HS6) is a global standard used for reporting trade between most countries; revisions to its roughly 5,000 product categories occurred in 1996, 2002 and 2007. Each country may have more detailed HS8 or HS10 categories that further refine the HS6 product categories

foreign destinations. It spans China’s entry into the WTO at the end of 2001, which lowered trade costs, reduced internally mandated barriers to trade and created export opportunities for Chinese firms. The nominal dollar value of Chinese goods exports nearly quadrupled to \$968bn in 2006 from \$250bn 2000. A large share of this growth was at the extensive margin —the number of exporters went from 62,600 to more than 170,000. (Table I.A.1 decomposes this growth in exporter numbers into its intensive and extensive margins).³⁴ Among firms that remained exporters, entry into new destinations was pervasive. Four out of five exporters entered new product-country destinations in the average year.³⁵

Table I.A.1: Exporter Dynamics in China: 2000 - 2006

Year	Exporter Count $A = B + C$	Incumbents $B = A_{t-1} - D$	Entrants C	Leavers D
2000	62,603			
2001	68,347	52,201	16,146	10,402
2002	78,567	57,263	21,304	11,084
2003	95,627	68,506	27,121	10,061
2004	120,363	82,858	37,505	12,769
2005	143,583	103,724	39,859	16,639
2006	170,642	124,419	46,223	19,164

	Exporter-ProdMkt Count $A = B + C$	Incumbents $B = A_{t-1} - D$	Entrants C	Exits D
2000	1,782,803			
2001	2,011,808	696,379	1,315,429	1,086,424
2002	2,464,544	828,853	1,635,691	1,182,955
2003	3,076,358	1,059,347	2,017,011	1,405,197
2004	3,827,074	1,307,810	2,519,264	1,768,548
2005	4,846,699	1,593,626	3,253,073	2,233,448
2006	5,895,393	1,907,010	3,988,383	2,939,689

Table I.A.1 also helps to justify the paper’s focus on export entry in Tables 1.3 to 1.6. The first panel shows exporter numbers nearly tripled from 2000 to 2006. The last three columns in each panel break down the annual changes into entry (column C), exits (column D) and

³⁴See Ahn et al. (2011); Manova and Yu (2012) for fuller descriptions of how WTO accession reduced trade costs for Chinese exporters.

³⁵The average exporter in 2000 served 28 destinations - 13 countries and 7 HS6 products, while the corresponding number for 2006 was 34 (16 countries and 8 products). Destinations had a skewed distribution of Chinese exporter participation —40 firms on average served each destination. The median exporter count was 5. The reported central moments exclude destinations with zero Chinese exporters.

holdovers (column B). The increase in the number of exporters fits the pattern expected for trade liberalization with China's WTO accession.

Entry is also the dominant dynamic at the finer level of firm destination combinations (in the second panel). Here, turnover rates are higher. The fact that many exporters exit a destination in one year only to return in a later year suggests the need to redefine entry over periods longer than a year.

I.A.2 Product-Country Destination Data

Global trade in nominal US dollar terms grew at an average annual rate of 7% to reach \$12tn in 2006, covering more than 220 countries and 5,000 HS6 products. Approximately 990,000 unique destinations registered imports in the COMTRADE database, though many had zero demand in several years. The COMTRADE database indicates a moderate expansion of destinations from 695,000 to 738,000 between 2000 and 2006. (In the same period, destinations served by Chinese exporters increased by 64% – from 179,000 to 292,000).

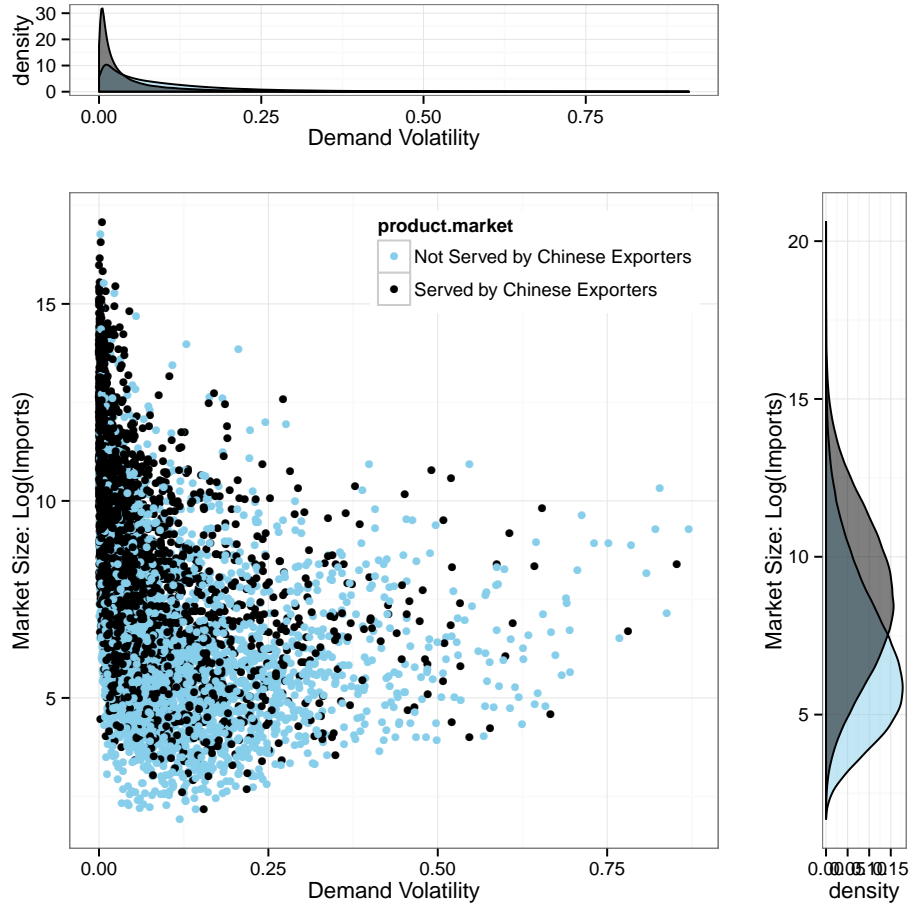
The original data set reports more than 63 million observations of trade at the HS6 product level for importing and exporting country-pairs in years from 1995 to 2005. (The full dataset goes to 2012, but only the first 11 years are usable as history because the firm level data stops in 2006). I collapsed this data to importing country-HS6 combinations for each of the years, noting that HS6 categories remain consistent over time. The UN COMTRADE data is reported in the 1992 version of the HS6 system for all years (Gaulier and Zignago, 2010). This collapsed form represents the history of imported demand into each destination in the analysis. I expand the data to a balanced panel of 11 years for all 990,000 destinations. This identifies instances of zero-demand. (Missing data is coded as zero). Market sizes and demand volatility come from this expanded data.

About 390,000 unique destinations were served at least once by Chinese exporters between 2000 and 2006. About 200,000 of these had no Chinese exporters in 2000 and only show up later in the exporter-level data. It is important to evaluate the differences between the destinations with Chinese exporters and the other destinations, of the 990,000 possibilities. Some of the variables of interest in this comparison of destinations that drew Chinese exporters include size and demand volatility.

Figure I.A.1 graphs the distribution of sizes and demand volatility simultaneously for all destinations in the data, comparing those served by Chinese exporters with the remainder.

The density plots in Figure I.A.1 indicate that the regressions in the main body of the paper provide broad coverage of destinations in terms of size and demand volatility. The

Figure I.A.1: Destinations Served (and Not Served) by Chinese Exporters



The scatter plot only shows a random 1% sample of the more than 900,000 destinations. The density plots for market size and demand volatility use the full data set. Demand Volatility represents deviations around the import trend as defined in Section 1.3.1.2, and market size is the log of the total USD value of imports into the destination between 1995 and 2005.

Data Sources: China GAC Export Data (2000-2006), UN COMTRADE

ranges of demand volatility and destination sizes covered by the two categories appear similar in the graph. Nevertheless, from the distribution at the top of the graph, one sees that a larger share of Chinese exporters – more than the global average – were concentrated in destinations with low demand volatility. The average demand volatility for destinations served by Chinese exporters is .053, while markets not served by Chinese exporters have an average of .110. Table 1.2 makes the same point with linear regressions. The density on the right panel shows that the destinations Chinese exporters serve are also larger on average.

To facilitate replication, here are country-level variables used but not reported in Tables 1.3, 1.4 and 1.8 to test the relationship between demand volatility and the number of Chinese exporters that enter a destination. They are from Head et al. (2010).

Table I.A.2: Additional Regression Variables

VARIABLES	N	Mean	Std. Dev.	Min.	Max.
GDP	289801	10.06095	2.438002	4.258556	16.39586
GDP per Capita	288048	8.018416	1.572457	4.634483	11.11192
Log(Distance)	346217	8.989019	.5693675	6.925665	9.857974
Log(Remoteness)	346217	-9.105768	.4982548	-10.53251	-8.298696
Contiguity	346217	.0830606	.2759742	0	1
Language	346217	.0277225	.1641769	0	1
Legal Origin	346217	.1771115	.3817636	0	1
GATT/WTO	346217	.7430744	.4369386	0	1

GDP: GDP of the destination economy

GDP per Capita: of the destination economy

Distance: from China to the destination economy

Remoteness: Geographic remoteness, i.e. Destination’s GDP-weighted distance from all countries

Contiguity: Dummy indicating whether country has shared borders with China

Language: Dummy indicating whether country shares ethnic or official languages with China

Legal Origin: Dummy indicating whether country shares legal origin with China

GATT/WTO: Whether destination economy was a member of WTO

I.A.3 Product and Country Variation in Demand Volatility

Country-specific factors explain as much of the variation in demand volatility as product-specific factors. This is consistent with the patterns observed for output volatility in Koren and Tenreyro (2007).

Table I.A.3 presents linear regressions of demand volatility on destination size. The first panel in the table is limited to destinations with at least one Chinese exporter, while the second panel extends the regressions to include all destinations with aggregate demand

volatility data. Although destination size explains a notable share of the variation in demand volatility in the first panel, it is interesting to note that country-specific factors also explain a larger share than product specific factors. This comes from a comparison of columns 1 and 2. When combined with destination size, the two sets of fixed effects explain comparable incremental shares of the variation in the dependent variable, if one compares the R^2 values in columns 4 and 5.

Table I.A.3: Demand Volatility: Product and Country Fixed Effects
(Dependent Variable: Demand Volatility)

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Destination Size			-0.02*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
Constant	0.08*** (0.00)	0.08*** (0.00)	0.21*** (0.00)	0.23*** (0.00)	0.19*** (0.00)	
Observations	259,963	259,963	259,963	259,963	259,963	259,963
R-squared	0.10	0.15	0.15	0.25	0.20	0.31
Product FE	Y			Y		Y
Country FE		Y			Y	Y
For All Destinations						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Destination Size			-0.02*** (0.00)	-0.02*** (0.00)	-0.01*** (0.00)	-0.01*** (0.00)
Constant	0.08*** (0.00)	0.08*** (0.00)	0.21*** (0.00)	0.22*** (0.00)	0.19*** (0.00)	
Observations	708,802	708,802	708,802	708,802	708,802	708,802
R-squared	0.09	0.14	0.15	0.24	0.20	0.29
Product FE	Y			Y		Y
Country FE		Y			Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Demand volatility is the sum of the squared deviations of demand from a linear trend over the years 1995 to 2005. Destination size is the sum of aggregate demand in each destination from 1995 to 2005.

The second panel of Table I.A.3 follows the same pattern as the first panel; smaller

destinations tend to have higher demand volatility. (This called for the inclusion of market size as a control in the regressions that I report in the main body of the paper). Furthermore, product-fixed effects generally explain less of the variation in demand volatility if one does not control for destination size. As expected, destination size is driven in part at the country-level. Larger economies import more of most product categories.

I.A.4 Annual Exporter Counts and Demand Volatility

To exploit the annual variations in export flows and exporter counts, Table I.A.4 provides results of regressions structured after conventional gravity model estimations. The regressions include estimates with country-year fixed effects in the even-numbered columns.

The results indicate that demand volatility decreases export volumes and exporter counts, just as in Table 1.4. Unlike Table 1.4, the dependent variables here measure annual export volumes and exporter numbers like most gravity model estimations.³⁶ In this table, exports per exporter decrease with increased volatility, another point of difference with the cumulative version in the main body of the paper. However, this table does not include full-fixed effects for China, only its GDP is included as an additional control for changes over time.

To address concerns about time-varying factors like exchange rate volatility, country-specific shocks or trade deals, the even-numbered columns include country-year fixed effects. These improve the estimated coefficients, usable observations and explained variation. A caveat is necessary: The demand volatility term used in this table does not change over time, nor do the country-level measures like GDP change within products for each country. In other words, the estimated variables are not well matched. (Variants of this table that include the destination size term, or a measure of logged annual imports from all destinations also predict smaller exporter numbers with higher demand volatility, even if they do not consistently predict lower volumes).

I.A.5 Other Measures of Demand Volatility

To avoid confronting criticism that that the measure of demand volatility does not reflect the expectations of Chinese exporters, or is biased, this section repeats the estimates in Table 1.3 using two variations on the measure of demand volatility.

³⁶ One could try to reconcile this to the model with claims that each year represents its own equilibrium; a claim that requires justification, but that may change how demand volatility should be defined.

Table I.A.4: Annual Trade Estimates with Demand Volatility
(Dependent Variable: Log Export Indicator in Destinations by Year)

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Log(Exports)		Log(Exporter Counts)		Log(Exports per Exporter)	
Demand Volatility	-2.619*** (0.084)	-2.398*** (0.083)	-1.201*** (0.040)	-0.921*** (0.041)	-1.418*** (0.056)	-1.477*** (0.054)
Log(GDP)	0.687*** (0.006)		0.343*** (0.003)		0.345*** (0.003)	
Log(GDP per Capita)	-0.095*** (0.006)		0.016*** (0.003)		-0.111*** (0.004)	
Log(GDP China)	2.983*** (0.051)		1.848*** (0.027)		1.134*** (0.034)	
Log(Distance)	-0.573*** (0.011)		-0.372*** (0.005)		-0.200*** (0.007)	
Constant	19.190*** (0.275)		8.133*** (0.127)		11.056*** (0.186)	
Observations	1,485,716	1,590,273	1,485,716	1,590,273	1,485,716	1,590,273
R-squared	0.435	0.442	0.570	0.575	0.329	0.340
Country FE		Y		Y		Y
Product FE	Y	Y	Y	Y	Y	Y

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destination - years: unique year, HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220) in 2006. The dependent variable is log (the number of firms with recorded exports to a destination) in each of the years between 2000 and 2006. Destination size is the value of all imports into each destination, for the respective year.

I.A.5.1 Chinese Aggregate Demand Volatility

First, I use only imports from China to estimate the volatility measure for each destination. This goes with the idea that Chinese goods are poor substitutes for other countries' products, *à la* the Armington national differentiation narrative, so that in each destination, there is a market for Chinese goods, and Chinese exporters expect demand fluctuations to reflect the history of demand for only Chinese goods.

Formally, this alternative measure of demand volatility:

$$\sigma_{jk}^2 = \sum_t (\epsilon_{jkt})^2$$

where ϵ_{jkt} is derived from:

$$\frac{Q_{jkt}}{\sum_t Q_{jkt}} = \zeta_{jkt} + \alpha_{jk} + \epsilon_{jkt}$$

Q_{jkt} is the sum of imports into destination jk from only Chinese sources in year t . I use the Chinese component of the UN COMTRADE data for this purpose.

Table I.A.5 shows that destinations with high demand volatility have fewer exporters, just as in Table 1.3. Destination size, GDP, and Distance keep the same signs as the comparable results from the body of the chapter, with more of the coefficients being statistically significant in this specification.

There are other differences, all of which are reasonable for this specification. More of the variation is explained by the right-hand side variables. (R^2 values in Table I.A.5 are higher by about 0.24 on average). There are fewer usable observations, as there are destinations with multiple years of available demand history globally, but for which demand history for Chinese exporters alone is not sufficient to construct a volatility measure.

One possible weakness of Table I.A.5 is that the results may be biased for destinations served by only a few Chinese exporters. For those destinations, the measure of volatility largely reflects shocks to the Chinese exporters; firm-level shocks not related to the destination's attributes could make the measured volatility higher than its implicit value.

Nevertheless, the estimates support the main finding that exporters on average, enter destinations with low demand volatility in greater numbers.

Table I.A.5: Exporter Counts and Demand Volatility (China Only):
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-1.140*** (0.022)	-1.113*** (0.020)	-1.081*** (0.022)	-1.062*** (0.020)
Destination Size	0.366*** (0.002)	0.333*** (0.002)	0.352*** (0.002)	0.319*** (0.002)
Log(GDP)	0.200*** (0.002)		0.205*** (0.002)	
Log(GDP per capita)	0.002 (0.002)		-0.001 (0.002)	
Log(Distance)	-0.321*** (0.004)		-0.308*** (0.004)	
Constant	3.515*** (0.063)		3.289*** (0.064)	
Observations	244,139	279,780	240,832	275,951
R-squared	0.784	0.814	0.777	0.805
Country-Year FE	N	Y	N	Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures the net export entry that accompanied China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

I.A.5.2 Aggregate Demand Volatility: Rest of the World

Second, I use imports from the rest of the world, excluding China, to measure demand volatility for each destination. Compared with Section I.A.5.1, this takes a different perspective in assuming that Chinese goods are substitutes for items in the same narrow product category from other countries. By that reasoning, aggregate imports for each product-country combination form a reasonable basis for setting expectations of demand fluctuations for Chinese exporters.

The definition for demand volatility for this table is analogous to that for Table 1.3, except that the Q_{jk} used to construct the measure represents each destination's imports from all countries, excluding China. The exclusion implies that expectations of volatility exclude exporters' own sales and are tied to shocks exogenous to the exporting firm.

Table I.A.6 presents results that are largely similar to Tables I.A.5 and 1.3. As in the previous table, the main finding stands: exporters serve destinations with low demand volatility in higher numbers. Destination size is also a statistically significant predictor of export entry between 2000 and 2006, as in the two previous tables.

The differences between Tables I.A.6 and 1.3 follow the pattern seen in Table I.A.5 for the most part. A greater share of the variation in the dependent variable is explained by the right-hand side variables, and more of these have statistically significant coefficients, while retaining the same sign as in Table 1.3.

In sum, the results are robust to definitions of demand volatility that exclude the history of imports from China, or demand from the rest of the world. High demand volatility is consistently linked to lower levels of export entry into foreign destinations by Chinese exporters.

I.A.6 Comparing Demand Volatility with Other Predictors

Table I.A.7 shows simple OLS regressions of exporter counts on demand volatility, GDP, Distance, and Destination Size. This regression with omitted variables provides only correlations to facilitate comparisons. The correlations suggest that demand volatility is informative for analyses of trade and exporter counts.

The R^2 values reported in the table alone indicate that demand volatility compares favorably with conventional predictors in explaining the variation in exporter counts. Comparing each variable's column and column 5, which includes all predictors, provides further evidence. The sign and statistical significance of demand volatility remains consistent between

Table I.A.6: Exporter Counts and Demand Volatility (Rest of the World):
(Dependent Variable: Log Number of Exporters in Destinations)

VARIABLES	(1) Log(Gross Export Entry)	(2)	(3) Log(Gross Entry Post-2000)	(4)
Demand Volatility	-0.967*** (0.033)	-0.900*** (0.029)	-0.930*** (0.033)	-0.869*** (0.030)
Destination Size	0.251*** (0.003)	0.216*** (0.003)	0.245*** (0.003)	0.211*** (0.003)
Log(GDP)	0.262*** (0.003)		0.260*** (0.003)	
Log(GDP per capita)	-0.029*** (0.003)		-0.031*** (0.003)	
Log(Distance)	-0.562*** (0.005)		-0.538*** (0.005)	
Constant	6.351*** (0.087)		6.028*** (0.086)	
Observations	331,153	385,885	324,042	377,196
R-squared	0.695	0.736	0.693	0.733
Country-Year FE	N	Y	N	Y
Product FE	Y	Y	Y	Y

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). Gross entry is the log of unique firms with recorded exports to a destination between 2000 and 2006. The change in exporter count captures the difference between all firms that served a destination and firms that served the destination in 2000. This difference measures the net export entry that accompanied China's trade liberalization from 2001 onwards. Control variables used but not shown in the table include geographic remoteness and dummies for shared borders, common languages and WTO membership. Missing observations in Columns 1 and 3 are because GDP data is not always available. However, estimates on the largest common sample are almost identical to columns 2 and 4.

Table I.A.7: Comparing Demand Volatility and Conventional Predictors of Trade
(Dependent Variable: Log Number of Exporters)

VARIABLES	(1)	(2)	(3)	(4)	(5)
Demand Volatility	-5.156*** (0.054)				-0.607*** (0.034)
Log(GDP)		-0.001 (0.001)			0.001 (0.001)
Log(Distance)			0.002 (0.005)		0.001 (0.004)
Destination Size				0.431*** (0.003)	0.421*** (0.003)
Constant	2.286*** (0.003)	1.938*** (0.011)	1.917*** (0.041)	-1.822*** (0.026)	-1.717*** (0.047)
Observations	371,531	289,801	346,217	371,531	276,459
R-squared	0.287	0.217	0.215	0.544	0.546
Product FE	Y	Y	Y	Y	Y

Robust standard errors in parentheses. Errors clustered by HS6 products.

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are destinations: unique HS6-product and country combinations, e.g., Ethiopian imports of women's cotton overcoats (HS 610220). The dependent variable is the log of the count of unique firms with recorded exports to a product market between 2000 and 2006. Destination size is the value of all imports into each destination in 2000-2006. The other variables follow conventional definitions.

columns (1) and (5). Of the other variables, only destination size explains as much of the variation in exporter counts.

I.A.7 GDP per Capita and the Volatility of Imports

How does demand volatility as observed by exporters relate to economic development?

Figure I.A.2 suggests that in the context of trade, import demand volatility declines with increasing GDP per capita. The trend in the graph is downwards, indicating a negative correlation between the volatility of imports and GDP per capita. Not reported in this paper is a graph that shows a strikingly high correlation between the volatility of imports and exports. The small island nations on the right of the graph like Bermuda (BMU), Vanuatu (VUT) and Palau (PLW) also follow the trend, even if it appears shifted. The pattern is comparable to plots of output volatility against GDP per capita in related papers (di Giovanni and Levchenko, 2012b; Koren and Tenreyro, 2007). The vertical axis shows the standard deviation of each country's year-on-year aggregate import growth, while the horizontal axis represents the GDP per capita in 2006.


Economic output is generally more volatile in low-income economies (Koren and Tenreyro, 2007; di Giovanni and Levchenko, 2009; Krishna and Levchenko, 2009). Some pointers on the direction of causation come from di Giovanni and Levchenko (2009), which shows that openness to trade generally increases aggregate output volatility; trade is usually a higher share of output for smaller economies. Krishna and Levchenko (2009) and Koren and Tenreyro (2007) both provide evidence that low-GDP economies tend to specialize in volatile sectors, with the former attributing this tendency to the quality of institutions.

From the perspective of firms importing inputs for industrial processes, the reduction in exporter counts associated with demand volatility matters. Where fewer exporters are available to supply the inputs required for industrial production, the resulting higher prices for those inputs discourage potential growth in the related productive sectors. Several papers show that imports are critical to economic growth and the expansion of exports in developing economies (Goldberg et al., 2010; Bas, 2009; Chevassus-Lozza et al., 2013).

I.A.8 Demand Volatility and Prices

Table I.A.8 tests for a relationship between prices and demand volatility separately for each of the UN broad economic categories or BEC goods classifications. The classifications are important for this paper, given how the contributions of each of these imported categories

Figure I.A.2: The Volatility of Imports and GDP per Capita



The unit of observation is a country, represented by its 3-letter ISO code. The vertical axis shows the volatility of imports over 1995-2005 and the horizontal axis shows 2006 GDP per capita. Here, the volatility of imports is the sum of the deviations of imports from trend at the country-level, as in equation (1.18).
 Data Sources: UN COMTRADE, CEPII (Head et al., 2010)

to economic development differ (Jones, 2011). Imports of capital goods have been cited often as a source of productivity growth in developing economies (Eaton and Kortum, 2001; Lee, 1995).

Table I.A.8: Demand Volatility and Prices, by Broad Economics Categories
(Dependent Variable: Log Unit Prices)

VARIABLES	(1) Capital Goods	(2)	(3) Consumer Goods	(4)	(5) Intermediates	(6)
Demand Volatility	0.099** (0.040)	0.079* (0.042)	-0.080*** (0.030)	-0.077** (0.032)	-0.014 (0.022)	0.023 (0.022)
Prod. Mkt. Size	-0.021*** (0.002)	-0.012*** (0.003)	-0.022*** (0.002)	0.000 (0.002)	-0.013*** (0.001)	-0.005*** (0.001)
Log (GDP)	0.009*** (0.002)	0.176* (0.100)	0.027*** (0.001)	-0.022 (0.087)	0.019*** (0.001)	-0.009 (0.054)
Log(GDP per cap)	0.022*** (0.002)	-0.128 (0.096)	0.027*** (0.002)	0.096 (0.080)	0.023*** (0.001)	0.043 (0.052)
Log(Distance)	0.009** (0.004)	-0.604 (41.277)	0.019*** (0.002)	-3.412 (215.224)	0.015*** (0.002)	0.106 (35.871)
Constant	2.450*** (0.054)		0.033 (0.039)		0.101*** (0.027)	
Observations	2,123,105	2,123,105	10,752,290	10,752,290	8,214,698	8,214,698
R-squared	0.964	0.964	0.920	0.920	0.919	0.919
Firm-Product-Year FE	Y	Y	Y	Y	Y	Y
Country FE		Y		Y		Y

Robust standard errors in parentheses. Errors clustered by product-year.

*** p<0.01, ** p<0.05, * p<0.1

The units of observation are unique combinations of firms, HS8 products, destination countries and years, e.g., exports of women's cotton overcoats (HS 61022000) to Ethiopia in 2006 by firm #311996528A. The dependent variable is the log of value divided by quantity for each observation. I used the finer HS8 rather than HS6 product categorization was used because the units for measuring quantity are consistent within HS8 but not HS6 categories. Repeating the regressions for products with only one type of quantity unit within HS6 categories gives substantially similar estimates. The classification into capital, consumer and intermediate goods follows the UN's correspondence between HS6 products and broad economic categories (BEC).

The results indicate that prices are higher for capital goods in destinations with high demand volatility, if only slightly so. This result holds even in the absence of explicit controls for product quality, which could matter for a category with quality differentiation potential like capital goods. Prices are lower on average with high demand volatility for consumer

goods, which also tend to have quality differentiation potential but represent a larger share of exports. The estimated effect of volatility on prices is not as clear for intermediate goods and other products that do not fit into any of these three categories (like automobiles). While the findings for these other categories do not conform to the model, they could be explained by related works that imply destinations with low GDP per capita receive lower import prices (Manova and Zhang, 2012; Hummels and Klenow, 2005).

The coefficients of other variables in Table I.A.8 are consistent with related work on prices in international trade; prices increase with GDP per capita for the specifications that are statistically significant, increase with distance and show mixed effects with GDP.

The firm-product-year fixed effects address several potential concerns with the estimate, including possible changes in HS8 categories from one year to the next (Amiti and Freund, 2010), quality differentiation by firm (Manova and Zhang, 2012) and product differences. That prices vary by product is rudimentary, but when some HS8 products are reclassified to other categories from one year to the next, it is important to include fully-interacted HS8-year effects as a control, even if the reclassifications affect only a small share of product categories. This pair interacted with firms to create fixed effects that reflect consistent differences due to firm productivity, and investments in quality. The country fixed effects address country-specific factors that may consistently affect prices like exchange rates, but are not captured by GDP, distance or GDP per capita.

I.B Model Features

I.B.1 The Envelope Theorem Allows Optimal $q^* = E(q)$

Exporters maximize expected profits in the model by working with two parameters: [1] they set prices p , which is equivalent to setting quantities q in monopolistic competition and [2] they set production scale q^* , given that deviations of q from q^* are costly, as defined in equation (1.2).

The Envelope Theorem justifies my approach of treating this optimization as a one-parameter choice, with the second parameter, in this case q^* fixed at the optimal level. Formally, if profits are a function of both the production scale q^* and prices (which predicts actual quantities sold), the set of optimal profits with respect to prices should be at values of q^* that maximize profits.

Formally, one may define profits as the objective, prices as the parameter that determines profits and the production scale q^* as the maximizer. In optimizing, i.e. setting the derivative

equal to zero, the derivative of the profit objective with respect to the production scale equals the partial derivative of profits with respect to prices or quantities, holding the maximizer fixed at its optimal level. Expected profits $E(\Pi)$ is a function of both p and q^* :

$$\max_p E(\Pi) = \max E(\Pi) \implies \frac{\delta E(\Pi)}{\delta p} = 0 \Big|_{\frac{\delta E(\Pi)}{\delta q^*} = 0} \quad (1.23)$$

Prices should map one-to-one to quantities, given equation (1.5), thus one can maximize the preceding equation with respect to q .

In the main body of the paper, I assume production scale will always be set to $E(q)$, and justify the claim with the Envelope Theorem. Here I provide the formal derivation. The optimization exercise fixes quantities and prices for trade in monopolistic competition, with the optimal exporters' production scale q^* :

$$\begin{aligned} \frac{dE(\Pi)}{dq^*} &= 0 \\ &= \frac{dE(\Pi)}{dE(\text{adjustment costs})} \frac{dE(\text{adjustment costs})}{dq^*} \end{aligned} \quad (1.24)$$

From equation (1.1), $\frac{dE(\Pi)}{dE(\text{adjustment costs})} \neq 0$, therefore:

$$0 = \frac{dE(\text{adjustment costs})}{dq^*} \quad (1.25)$$

$$\begin{aligned} \frac{dE(\Pi)}{dq^*} &= \frac{d[\gamma \left(\frac{E(q)-q^*}{q^*}\right)^2]}{dq^*} \\ &\implies q^* = E(q) \end{aligned} \quad (1.26)$$

This leaves us with a one-parameter optimization, as long as the production scale is fixed at expected quantities. (One can generalize equation (1.26) to a set of production scales, one for each period of the exporter's planning horizon).

The Envelope Theorem has also been applied to the analysis of incentive constraints in contract theory and non-convex production problems (Milgrom and Segal, 2002).

I.B.2 Trade Volumes with Demand Volatility

In equilibrium, trade volumes are the integral of firm level sales over the distribution of productivities that meet the threshold ϕ_{jk}^* :

$$X_{ijk} = p_{ijk}q_{ijk}^* = p_{ijk}^{1-\varepsilon} \frac{Q_{jk}^*}{P_{jk}^{1-\varepsilon}}$$

Summing across all firm varieties for destination jk , where $G(\cdot) = 1 - \phi_{ij}^{-\theta_j}$.

$$\begin{aligned} X_{jk} &= \frac{Q_{jk}^*}{P_{jk}^{1-\varepsilon}} \int_{\phi_{jk}^*}^{\infty} p_{ijk}^{1-\varepsilon} dG(\phi_{ij}) \\ &= \frac{Q_{jk}^*}{P_{jk}^{1-\varepsilon}} \int_{\phi_{jk}^*}^{\infty} \left[\frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\phi_{ij}} \right]^{1-\varepsilon} (\theta_j \phi_{ij}^{-\theta_j - 1}) d\phi_{ij} \\ &= \frac{\theta_j Q_{jk}^*}{P_{jk}^{1-\varepsilon}} \left[\frac{\varepsilon \tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\varepsilon - 1} \right]^{1-\varepsilon} \int_{\phi_{jk}^*}^{\infty} (\phi_{ij}^{-\theta_j - 1 + (\varepsilon - 1)}) d\phi_{ij} \\ &= \frac{\theta_j Q_{jk}^*}{P_{jk}^{1-\varepsilon}} \left[\frac{\varepsilon \tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\varepsilon - 1} \right]^{1-\varepsilon} \frac{-\phi_{jk}^{*(\varepsilon - \theta_j - 1)}}{\varepsilon - \theta_j - 1} \end{aligned} \quad (1.27)$$

substituting ϕ_{jk}^* from equation (1.11):

$$X_{jk} = \frac{-\theta_j \varepsilon S_{jk} \left(\frac{Q_{jk}^*}{\varepsilon S_{jk}} \right)^{\frac{\theta_j}{\varepsilon - 1}}}{P_{jk}^{-\theta_j} (\varepsilon - \theta_j - 1)} \left[\frac{\varepsilon \tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\varepsilon - 1} \right]^{-\theta_j} \quad (1.28)$$

Taking logs, (while noting that $\varepsilon - 1 < \theta$, for sales to be finite):

$$\ln(X_{jk}) = -\theta_j \ln \left[\frac{\varepsilon \tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{\varepsilon - 1} \right] + \frac{\theta_j}{\varepsilon - 1} \ln \left(\frac{Q_{jk}^*}{\varepsilon S_{jk}} \right) + \ln \left[\frac{-\theta_j \varepsilon S_{jk}}{P_{jk}^{-\theta_j} (\varepsilon - \theta_j - 1)} \right] \quad (1.29)$$

Equation (1.29) implies that trade should decrease with increasing demand volatility. The slope of X with respect to σ^2 in levels and logs should be negative. However, the ε , Q and P , τ and S terms may lead to estimated effects that are smaller than the extensive margin. This pattern is consistent with the assertion in Chaney (2008) that if trade levels change due to changes in costs, the extensive margin dominates.

The intensive margin on the other hand, depends on the productivity distribution. In

practice, it should depend on the combination of the productivity slope parameter θ_j and the elasticity of demand ε . For example, exports per exporter $\frac{X_{jk}}{N_{jk}}$ may rise if high demand volatility leads to higher prices that reduce demand, but the slope of the productivity distribution is high enough that the changing export productivity threshold leaves few exporters to meet demand, leading to higher exports per exporter. In principle, it is independent of demand volatility.

Formally;

$$\ln(X_{jk}) - \ln(N_{jk}) = \ln \left[\frac{-\theta_j \varepsilon S_{jk}}{(\varepsilon - \theta_j - 1) N_j} \right]$$

So that:

$$\frac{d \ln \left(\frac{X_{jk}}{N_{jk}} \right)}{d \sigma_{jk}^2} = 0 \tag{1.30}$$

The relationship in (1.30) is a distinctive feature of the Pareto distribution of productivity (or the class of power laws in general). For these productivity distributions, any change in the exports per exporter that would have resulted from the changes in prices due to demand volatility is perfectly offset by the change in the number of exporters. This requires the usual assumption of large numbers of atomistic exporters. In the firm level data, the average product is associated with 1300 exporters; the median has 400. In sum, deviations from a Pareto productivity distribution and a small pool of potential exporters may skew the findings away from the prediction in (1.30). However, it is still expected that the extensive margin dominates the intensive margin, regardless of the exact nature of the productivity or the size of N_j . Chaney (2008) derived a similar relationship between trade levels and trade costs.

Total exported value to destination jk is the integral over the distribution of productiv-

ities of firm-level exports, as given in equation (1.27). The elasticity:

$$\begin{aligned} \frac{d \ln X_{jk}}{d \ln \sigma_{jk}^2} &= \frac{d X_{jk}}{d \sigma_{jk}^2} \frac{\sigma_{jk}^2}{X_{jk}} \\ &= \frac{\int_{\phi_{jk}^*}^{\infty} \frac{d p_{ijk}^{1-\varepsilon}}{d \sigma_{jk}^2} \sigma_{jk}^2 dG(\phi_{ij})}{\int_{\phi_{jk}^*}^{\infty} p_{ijk}^{1-\varepsilon} dG(\phi_{ij})} - \frac{p_{ijk}^{1-\varepsilon}(\phi_{jk}^*) (1 + \gamma_j \sigma_{jk}^2) \frac{d \phi_{jk}^*}{d(1 + \gamma_j \sigma_{jk}^2)} dG(\phi_{ij}^*)}{\int_{\phi_{jk}^*}^{\infty} p_{ijk}^{1-\varepsilon} dG(\phi_{ij})} \end{aligned} \quad (1.31)$$

intensive margin extensive margin

$$= \frac{\gamma_j (1 - \varepsilon)}{(1 + \gamma_j \sigma_{jk}^2)} - \frac{\gamma_j (\theta_j - \varepsilon + 1)}{(1 + \gamma_j \sigma_{jk}^2)} \quad (1.32)$$

Because,

$$\frac{\int_{\phi_{jk}^*}^{\infty} \frac{d p_{ijk}^{1-\varepsilon}}{d \sigma_{jk}^2} \sigma_{jk}^2 dG(\phi_{ij})}{\int_{\phi_{jk}^*}^{\infty} p_{ijk}^{1-\varepsilon} dG(\phi_{ij})} = \frac{\gamma_j (1 - \varepsilon)}{1 + \gamma_j \sigma_{jk}^2}$$

and,

$$\frac{p_{ijk}^{1-\varepsilon}(\phi_{jk}^*) (1 + \gamma_j \sigma_{jk}^2) \frac{d \phi_{jk}^*}{d(1 + \gamma_j \sigma_{jk}^2)} dG(\phi_{ij}^*)}{\int_{\phi_{jk}^*}^{\infty} p_{ijk}^{1-\varepsilon} dG(\phi_{ij})} = \frac{\gamma_j (\theta_j - \varepsilon + 1)}{(1 + \gamma_j \sigma_{jk}^2)}$$

given that

$$\frac{d \phi_{jk}^*}{d(1 + \sigma_{jk}^2)} = \frac{\phi_{jk}^*}{(1 + \gamma_j \sigma_{jk}^2)}$$

and,

$$\frac{d p_{ijk}^{1-\varepsilon}}{d \sigma_{jk}^2} = \frac{\gamma_j (1 - \varepsilon)}{(1 + \gamma_j \sigma_{jk}^2)} p_{ijk}^{1-\varepsilon}$$

Equation (1.32) shows why the extensive margin is prominent.

I.B.3 Demand Volatility and Exporter Size Thresholds

One implication of the model is that destinations with higher demand volatility also have higher exporter productivity thresholds.

From equation (1.11)

$$\phi_{jk}^* = \frac{\varepsilon}{\varepsilon - 1} \frac{\tau_{jk}(1 + \gamma_j \sigma_{jk}^2)}{P_{jk}} \left[\frac{\varepsilon S_{jk}}{Q_{jk}^*} \right]^{\frac{1}{\varepsilon-1}}$$

As demand volatility σ^2 increases, so does ϕ^* . This has two implications: first, that more volatile destinations have fewer exporters, holding other factors constant, and that those exporters on average will be the most productive within their product categories. This complements a similar argument for productivity thresholds and sunk costs in related papers (Helpman et al., 2004; Melitz, 2003).

Figure 1.3 in the main body of the paper illustrates the idea, derived from this empirical specification:

$$\phi_{jk}^* = \beta_0^\phi \hat{\sigma}_{jk}^2 + \beta_1^\phi X_{jk} + \alpha_j^\phi + \alpha_k^\phi + \epsilon_{jk}^\phi \quad (1.33)$$

where ϕ_{jk} represents the lowest index of productivity for firms in destination jk ,

the index could be market share or the number of destinations in 2006

X_{jk} = variable(s) representing market size

α = product or country fixed effects


$\hat{\sigma}_{jk}^2$ = demand volatility or dummy indicating one of 50 quantiles for demand volatility

Figure I.B.1 presents a pattern that is similar to the trend observed in Table 1.3 in the main body of the paper. In the first panel of the figure, we see that the minimum number of destinations served by firms selling goods to a location jk is higher for destinations with high demand volatility. Taking the number of destinations served by a firm as a proxy for productivity follows the finding in ? that the most productive firms serve more export destinations. (Note that the number of destinations in this first panel is limited to 200, as only destinations for the specific product j are considered). For example, say firm C sells chocolate to 140 countries, coffee to 15 countries and cream to 60 countries, while firm D sells chocolate to 20 countries, and coffee to 100. Both firms sell chocolate to the US in this example. Firm C is taken to be the more productive chocolate exporter and the 'least productive' firm in the US market for imported chocolate is firm D with 20 destinations, if firms C and D are the only sellers to the US. If only firm C sells chocolate to Australia, then that destination registers an exporter threshold proxy of 140 - for the destinations getting chocolate from C. The graph shows that destinations with high demand volatility tend to

be limited to the largest exporters that serve the most destinations.

The second panel of Figure I.B.1 presents a stronger pattern for the relationship between the minimum number of destinations served by firms and demand volatility. The measure of productivity in this panel of the graph uses all destinations served by firms, rather than the number of destinations within a given product. This alternative measure recognizes that many exporters are multi-product firms, so that their productivity within a narrow product category may not fully represent the ability to compete. For the example in the last paragraph, firm D is still the least productive exporter in the US chocolate destination, but the measure of productivity is 120 (100 coffee+ 20 chocolate). In the same vein, if both exporters sold coffee to the US, Firm D would still be least productive for the US coffee destination, as its 120 destinations is fewer than firm C's 215. For the measure in the left panel, Firm C would have been the least productive firm and the measure of minimum productivity would be 15. For this measure of productivity, the same pattern of higher productivity thresholds for destinations with high demand volatility is observed.

Figure I.B.1: Export Thresholds and Demand Volatility



The minimum number of destinations served by exporters in each destination jk . The number of destinations served by an exporter is a proxy for its productivity, so that a destination served by only the most productive firm gets the highest value of this measure. The y-axis averages these threshold counts for all destinations in each quantile. Destinations Grouped Into 50 Quantiles, Least to Largest by Demand Volatility. Data Sources: China GAC Export Data (2000-2006), COMTRADE

The increase in thresholds associated with demand volatility in the right panel of Figure I.B.1 appears mostly in the upper half of the 50 volatility quantiles used for the plot. The first half of the panel may not indicate a strong relationship because, as the model indicates, the effect of demand volatility on profits is non-linear. Nevertheless, the first order implications

of both panels in the graph consistently agree with the paper's main finding.

Table I.B.1 presents the results of the exercises described in equation (1.33). As expected, columns 2 and 3 of the table show that the threshold for firm productivity, measured by the number of destinations firms serve, is higher in destinations with high demand volatility. I control for destination specific drivers of the exporter threshold with destination size, as used for other estimates in the paper. I also include country and product fixed effects to represent product or country-related factors broadly determine the exporters that can access a foreign destination. While demand volatility does not yield statistically significant predictions if the firm share of output for a product is the assumed measure of productivity, in column 2 the fewest destinations served by firms in a destination increases by a number that is statistically significant at the 90% level. A difference in demand volatility of 0.086, i.e. one standard deviation, should correspond to a 0.05 increase from the mean of 3.8 in the threshold of markets served within each narrow product category. The table also predicts that larger markets are expected to have lower minimum export thresholds, in line with the literature.³⁷

Column 3 of Table I.B.1 completes this exercise, showing that the minimum thresholds for destinations with high demand volatility are higher. Given the mean value of 262 for destination counts and standard deviation of 0.086 for demand volatility, a difference in demand volatility of one standard deviation should correspond to an increase of 30 in the threshold of destinations served.

The increase in ϕ^* associated with demand volatility in (1.11) implies that the average rank should be higher for destinations with high demand volatility. The higher threshold implies a marginal increase in the average productivity, and size of observed exporters.

³⁷Repeating the regression with firm shares multiplied by 1000 to scale up the variable, given its mean of 0.0017, yields 0.493 as the coefficient on σ_{jk}^2 - positive but not statistically significant.

Table I.B.1: Demand Volatility and Export Thresholds
(Dependent Variable: Minimum Market Share and Destination Counts for Firms in jk)

VARIABLES	(1)	(2)	(3)
Demand Volatility	0.000 (0.002)	0.616** (0.270)	347.170*** (41.265)
Market Size	0.000 (0.000)	-0.414*** (0.018)	-55.225*** (2.421)
Observations	285,675	285,661	285,661
R-squared	0.377	0.235	0.086
Country-Year FE	Y	Y	Y
Product FE	Y	Y	Y

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

The units of observation are destinations: unique HS6-product and country combinations. The dependent variable in column 1 is the minimum share of exports within a product category provided by any firm in the destination represented by an observation. Firm share of exports within a product category is a proxy for productivity in this column as in Figure 1.3. (The average market share is 0.0017, so the zero coefficient is unsurprising). In column 2, it is the minimum number of destinations within the product category served by any exporter with sales to a given destination. In column 3 it is the minimum number of destinations served by any exporter with sales to a given destination. The count of export destinations for the firms are not limited to any products for this column.

CHAPTER II

The Impact of Exporting and Foreign Investment on Product Innovation: Evidence from Chinese Manufacturers

Chapter Abstract

To understand the drivers of product innovation at the firm level, I compare the effects of foreign direct investment and exporting on product innovation using a rich firm level database of manufacturing and industrial enterprises. The chapter focuses on product innovation, as it is vital to economic development. Estimates from linear regressions and propensity score matching tests show that learning-by-exporting is a stronger predictor of product innovation. Firms that receive foreign investment also tend to engage in more product innovation, but not at the same level as the firms that export. Additional tests confirm that as they start and stop exporting, firms change their patterns of investment in the drivers of product innovation - fixed capital and research.

2.1 Introduction

Emerging countries are no longer content to be sources of cheap hands and low-cost brains. Instead they too are becoming hotbeds of innovation . . . They are redesigning products. . . They are redesigning entire business processes to do things better and faster than their rivals in the West. Forget about flat – the world of business is turning upside down.

The Economist Magazine - (Masters of Innovation: 2010)

In developing economies, exporters and foreign owned firms contribute disproportionately to product innovation.¹ The mechanism behind this pattern is not clear, nor is it clear that technology transfer through foreign ownership translates to more product innovation at the firm level compared to homegrown efforts. It is clear however, that product innovation is vital to development. Economies that consistently create more varieties have better growth outcomes (Hummels and Klenow, 2005). This may explain the interest of policymakers in developing economies charged with promoting innovation-driven private-sector led development in export promotion and foreign investment (FDI). I compare the relative efficacy of these two well-known approaches.²

China is an excellent case for this study: it has grown to be the world's largest exporter, is the number one FDI destination among developing economies, all while expanding the scope of its industrial output (Amiti and Freund, 2010; Hu and Jefferson, 2002). Chinese exporters featured in 85% of US imported manufactured goods categories in 2005, (up from 9% in 1972) (Schott, 2008). Firm level evidence buttresses the point. In the Chinese annual survey of manufacturing firms between 2005 and 2007, 13% of firms reported creating new product varieties and 10% by value of aggregate output in the data was from the product varieties that were new to the firms. Thus, one cannot ignore product innovation in the narrative of China's growth experience.

To understand the firm level drivers of product innovation in China, this paper compares two firm categories – exporters and foreign-owned firms. The literature on product innovation motivated this comparison. Gorodnichenko et al. (2010) and Damijan et al. (2010) indicate that exporters tend to do more product innovation, while others attribute product innovation to foreign ownership, e.g. (Guadalupe et al., 2012). There are good reasons for both arguments, and the reverse could be true. Firms that start exporting may learn the methods required for product innovation, as may firms that receive foreign capital. Likewise, large, productive firms may be more likely to introduce new products, export and find foreign owners.

I measure the effects of exporting and FDI using propensity score matching estimators. This approach uses selection on observed characteristics of firms to address concerns about endogeneity in the relationship between product innovation and exporting or foreign

¹Several papers provide evidence that exporters undertake more product innovation, most notably Gorodnichenko et al. (2010). Others show that foreign ownership leads to product innovation (Guadalupe et al., 2012). These papers suggest that product innovation from either exporting or foreign ownership is due to technology transfer through international relationships.

²Section 2.2 discusses the relationship between product innovation and economic growth briefly.

ownership (or FDI). Effectively, I limit comparisons of product innovation by exporters or foreign-owned firms to firms with very similar observed characteristics. I used a set of control variables that was large and relevant enough that one could plausibly argue that any difference between exporters and non-exporters with the same set of characteristics was close to random. For example, in comparing only firms in the same industry and with nearly the same size, the approach addresses concerns that larger, more innovative firms in a particular sector are more likely to experience the exporting ‘treatment’ (Bernard and Jensen, 2004; Aitken et al., 1997) or the foreign ownership ‘treatment’ (Guadalupe et al., 2012). I use Chinese firm level data from the NBS annual survey of industrial enterprises between 2005 and 2007.

I find that export participation leads to a higher likelihood of product innovation. The matching estimates show that new products are a greater share of output for exporters: 20% for exporters, versus 14% for non-exporters with matched propensities. New products are 12.9% of the output of majority foreign-owned firms, compared to 19.0% for Chinese-owned peers chosen to control for selection into FDI status (i.e. foreign ownership).³ This provides an interesting contrast to papers that find higher levels of product innovation for foreign-owned firms in other contexts like Eastern or Western Europe, e.g. (Commander and Svejnar, 2011; Guadalupe et al., 2012). The differences suggest that context may determine the level of product innovation that foreign owners undertake.

I emphasize two causal mechanisms for product innovation – research and development (R&D) and investments in fixed capital. This builds on earlier papers that provide evidence of a positive correlation between exporting and R&D, e.g., (Lileeva and Trefler, 2010; Aw et al., 2008, 2000). I use difference-in-differences estimates to show that on average, both of these inputs to product innovation increase as firms start exporting, and decrease for the firms that stop exporting. The same pattern does not register for foreign ownership.

I organize the rest of the paper as follows: Section 2.2 discusses the related literature, while the subsequent section covers the methods, data and results. The paper concludes in Section 2.4 after several robustness checks in Section 2.3.6.

³Tables 2.1 and 2.4 delve further into these comparisons. In the main results, I show that these differences between categories do not depend on whether I measure the intensity or the incidence of product innovation.

2.2 Related Literature

This paper focuses on the direct impacts of trade and foreign direct investment (FDI) on firms that exported goods or received foreign capital respectively. (I will not discuss spillovers from FDI and exporting; if these exist, they are likely to bias my estimates toward zero and leave the main findings unchanged. It is reasonable to expect that the direct impacts of FDI and export participation vastly exceed the spillover effects).

2.2.1 Product Innovation

Product innovation is vital to economic development. It is no accident that larger economies produce and consume greater numbers of product varieties, as documented by (Hummels and Klenow, 2005). This follows the Schumpeterian view of development (Schumpeter, 1942); economies grow because firms successfully create new varieties as the old ones disappear. Madsen (2008) finds support for a Schumpeterian growth hypothesis that links R&D and the creation of new product varieties to economic growth. That paper used international data from OECD economies. The argument in that paper builds on earlier work like Segerstrom (1991), which motivates an unambiguous positive relationship between innovation and economic growth. Benhabib et al. (2014) also provides a model of firm-level growth that is driven by innovation in a related paper. Other papers provide formal models and evidence that link product innovation to welfare through consumers' love of variety e.g., (Broda and Weinstein, 2006; Krugman, 1980).

In the Chinese case, product innovation helped increase the scope, volume and sophistication of aggregate exports (Amiti and Freund, 2010; Schott, 2008). For firms, the creation of new varieties adds new profit streams and increases the utilization of human and physical capital (Bernard et al., 2011; Eckel and Neary, 2010). They can also help to diversify a firm's portfolio against potential adverse product-specific shocks. Given the importance of product innovation to growth, especially for China, this paper tries to understand the factors driving the creation of new varieties, starting from its well-documented drivers – FDI and exporting. To the best of my knowledge, this is the first paper to compare these two drivers of product innovation in the Chinese context.

In considering exporting as a potential driver of product innovation, this paper comes close to the learning-by-exporting literature, which I describe next.

2.2.2 Exporting and Product Innovation

Much of the work on learning-by-exporting focuses on revenue productivity, notable examples include (De Loecker, 2013, 2007; Clerides et al., 1998). Most of these papers argue that in equilibrium exporters are more productive because firms learn to be more productive as they export, contrary to the view that exporters are larger because the most productive firms self-select into exporting.

Few papers have tested learning-by-exporting with respect to product innovation. One is Damijan et al. (2010), which examines whether the higher level of product innovation by exporters is due to selection, or learning-by-exporting. (That paper found evidence in support of learning-by-exporting, using Slovenian data).

Gorodnichenko et al. (2010) provides evidence that exporters engage more in product innovation, identifying the causal mechanism as information exchange through vertical linkages to foreign firms. Their tests use 2002 and 2005 data from a set of World Bank's firm level surveys in 27 transition economies from Eastern and Central Europe. Others have reported similar results for Italy (Bratti and Felice, 2012; Castellani and Zanfei, 2006) and Slovenia (Damijan et al., 2010).⁴ This paper extends the research objective of Gorodnichenko et al. (2010) to Chinese industrial enterprises, in combination with the question of foreign investment's impact on product innovation.

2.2.3 FDI and Product Innovation

Guadalupe et al. (2012) uses selection on observed variables to test for the effects of foreign investment on product innovation, but does not include a comparison with exporting, like this paper. Furthermore, their paper does not test for a causal mechanism that drives product innovation in foreign-owned firms. Furthermore, we define product innovation differently: I define product innovation as a continuous measure of output share, while the Guadalupe et al. (2012) paper uses a dummy that indicates whether a firm introduced new products. Despite these differences, their conclusions are similar to what I find.

Several earlier works suggest that FDI or foreign ownership should lead to more product innovation (Girma et al., 2012; Iacovone et al., 2009; Girma et al., 2008; Lai, 1998). The reasons offered by this literature include: [1] Foreign owners support subsidiaries' R&D efforts, [2] FDI enables access to needed credit or finance for innovation, [3] foreign multinationals

⁴All these papers support the learning-by-exporting hypothesis. That said, one must emphasize the distinction between the product innovation and productivity dimensions of learning-by-exporting. Keller (2004) reviews the debate on learning-by-exporting for productivity.

transfer their innovations to subsidiaries to facilitate low cost production. As a parallel to the learning-by-exporting literature, papers that link FDI to productivity have a history that goes back to Iacovone et al. (2009); Javorcik (2004); Djankov and Hoekman (2000); Aitken and Harrison (1999).

2.2.4 Exporting and FDI's Effects on Product Innovation

This paper contributes to the literature by directly comparing the direct impacts of exporting and foreign direct investment (FDI) on product innovation. The papers cited above generally examine the role of trade in product innovation, without exploring the effect of foreign ownership. The following papers argue that foreign investment promotes product innovation, also without providing a comparison to exporting (Guadalupe et al., 2012; Girma et al., 2008). Note that I use the term 'foreign ownership' to describe FDI in most of the paper; the term seems more relevant to firm level descriptions.

Commander and Svejnar (2011) compare the effects of foreign ownership and exporting like this paper, but for the ratio of sales to inputs. In their analysis, both exports and foreign ownership are associated with higher efficiencies or throughput ratios. However, the foreign ownership variable takes away the statistical significance of the export variable in a regression model with both variables. Furthermore, this paper's context is China, while Commander and Svejnar (2011) looks at Eastern Europe.

The tests that follow recognize that FDI and exporting are not orthogonal features of firm level data. The prevalence of export-platform FDI implies that in many cases, exports happen because of FDI. Conversely, one can make the case for foreign investment that follows a successful exporting relationship. In the main tests for the effects of exporting, I introduce control variables for foreign ownership. Similarly, I control for firms' export orientation in testing for the effects of foreign ownership on product innovation. Examples of the first scenario include Kneller and Pisu (2007) which uses aggregate data for Europe and Sun (2009), which uses Chinese firm level data to show that FDI increases exports as a share of total output.

2.3 Methods, Data and Results

This section reports three sets of results: (1) OLS regressions that test the effects of FDI and exporting on product innovation, (2) Propensity Score Matching tests that show the same idea more robustly and (3) tests that show drivers of product innovation before and

after export entry.

The baseline OLS exercise helps to establish that FDI and exporting as drivers of innovation are relevant to the Chinese context, as documented in the literature. It is a simple comparison of foreign-owned and exporting firms with all other firms in the data. Correlation between these categories and product innovation does not imply causation, so I use propensity score matching (PSM) to mitigate bias that may result if the firms most likely to introduce product innovations based on their observed traits, are also likely to be foreign owned or exporters.

In principle, being an exporter or foreign-owned leads to product innovation because these categories of firms *do things differently* – using new methods, equipment, or processes. Therefore, in Section 2.3.5 I further support the claim of a causal relationship between exporting and product innovation by testing whether firms that start exporting also change their pattern of spending on innovation drivers. The innovation drivers I use for this paper are R&D and asset investment outlays. (I show before these tests that the selected drivers are strong predictors of product innovation).

2.3.1 Data

The data comprises all annual surveys of Chinese industrial firms from 2005 to 2007. China’s National Bureau of Statistics compiled this firm level data. The sample approximates a census of all firms with revenues greater than 5 million Yuan (about \$600,000), supplemented with a stratified random sample of firms below this threshold. The entire dataset is an unbalanced panel of 763,036 firm-year observations, covering over 329,000 unique firms. 55% of the firms are present in all three years, while another 20% show up in at least two.⁵

I identify exporters from the reported sales and exports values for each firm-year. Foreign ownership is determined from the reported components of ownership capital. The data cover a period of strong export participation and foreign investment for Chinese firms: this was after China’s WTO accession in December 2001. To illustrate the significance of the timing, the number of firms in the data increased from 249,028 to 311,186 between 2005 and 2007,

⁵Before these assessments, I dropped 12,293 observations with one or more of these issues: negative sales, negative ownership capital, foreign capital that exceeded total ownership capital, and exports that exceeded sales. (These observations accounted for 1% of the output observed in the data). This was after I excluded observations for industries outside manufacturing, to avoid comparability issues. The relevant Chinese two-digit industry codes are between Food Manufacturing(14) and Instruments and Office Equipment Manufacturing(41).

and the share of those numbers that were exporters in 2007 was 25%. Firms with majority foreign ownership were 8% of the sample in 2007.⁶

Only a minority of firms undertake product innovation - 90% of firm-year observations registered zero new products. The nearly 76,000 observations with positive values of new products belong to 45,340 firms that count for 115,315 of the total firm-year observations. (The firms that undertook product innovation between 2005 and 2007 did so in only 2 of 3 years on average).

To preview whether product innovation co-occurs more with foreign ownership or exporting, one could sort the data into four groups that combine the two sets of categories: from Chinese-owned non-exporters to Chinese-owned exporters and from foreign-owned non-exporters to foreign-owned exporters. A non-parametric comparison of average innovation intensities for these groups provides the first hint of what to expect in the results.

Table 2.1 summarizes the differences in levels of product innovation for the four exclusive subgroups created by the two categories of interest. New products as a share of total output value vary between these groups, with the exporting sub-groups having higher averages. Foreign owned firms do not appear to undertake more product innovation than the average firm according to the table, although they are larger and more likely to export. The pattern for size and propensity to export is in line with the literature, e.g. (Guadalupe et al., 2012; Commander and Svejnar, 2011; Gorodnichenko et al., 2010).

The numbers in Table 2.1 imply that the two sets of categories are meaningfully distinct, i.e. foreign ownership is *not nearly* a perfect predictor of export participation and *vice versa*. The distinction is necessary for any meaningful comparison of the nature proposed by this paper.

Table 2.1 also provides the first hint of a reasonable overlap between exporters and non-exporters, as well as firms with and without foreign ownership. (The overlap is necessary for the tests that match on observed characteristics in subsequent sections of the paper). 24% of exporters have foreign capital, more than a third of foreign-invested firms do not

⁶The dataset reports firms' ownership capital in each of six source categories - individual, collective, national, other corporations/legal persons, non-Chinese foreign and Chinese-foreign i.e. Hong Kong, Macau and Taiwan. The first four categories correspond to private and state-owned sources of funds from mainland China. I define foreign-owned firms as those with majority stakes from non-Chinese sources, i.e. outside mainland China, Hong-Kong, Macau and Taiwan. I do not consider capital from Hong Kong, Taiwan and Macau as foreign. The strong historical ties and similar business cultures suggest that these locations should be considered Chinese. An additional rationale for defining foreign capital as I do is round tripping. Xiao (2004) suggests that, to avoid regulation, some persons invest funds from mainland China through entities in these locations, so that ownership is only nominally from outside mainland China. Sections 2.3.7 and appendix section II.C report estimates with alternative definitions of foreign ownership.

Table 2.1: Group Summaries

Group	Attribute	2005	2007
Chinese-owned Non-Exporter	Product Innovation	.058	.056
	Group Share of Total Output	.410	.448
	Number of Firms	151,975	205,033
	Group of Share of Total Number	.677	.719
Chinese-owned Exporter	Product Innovation	.173	.207
	Group Share of Total Output	.421	.380
	Number of Firms	54,134	57,156
	Group of Share of Total Number	.241	.201
Foreign-Owned Non-Exporter	Product Innovation	.054	.043
	Group Share of Total Output	.031	.033
	Number of Firms	5,966	7,911
	Group of Share of Total Number	.027	.028
Foreign-Owned Exporter	Product Innovation	.130	.135
	Group Share of Total Output	.138	.140
	Number of Firms	12,264	14,966
	Group of Share of Total Number	.055	.053

export and more than a quarter of wholly Chinese-owned firms participate in the export market. As foreign-owned firms and exporters are larger than average, there are reasonable odds of finding large non-exporters as a comparison group for exporters of any ownership category - several large foreign-owned firms should help to populate the counterfactual category. Similarly, large foreign-owned firms would have no small measure of comparably large Chinese-owned firms as a comparison group. (To illustrate output per firm comparisons; exporters being 29% of firms, accounted for 55% of output in 2005 and the 8% of firms that were foreign-owned in the same year accounted for 17% of output).⁷

Other variables of interest that are not summarized above include R&D, and investments in production assets. (I use these in Section 2.3.5). First, one must note that R&D is fairly uncommon. Barely more than 10% of observations register any R&D expenditure.⁸ Like R&D, investments in production assets is another measure for which the majority of firms

⁷ From the group estimates, one may deduce that 4% of total output in all years was new to the producing firms. Here are other summary statistics not in the table: 27.4% of firm-years involved exporting, 8% involve foreign ownership, and the hypothetical average firm employed 193 persons to produce 103 million Yuan of output per year.

⁸ 61 observations were dropped for reporting negative R&D expenditure. Note that the full dataset has more than 763000 observations.

report zero. (I take the measure from firms' reported cash outlays for asset investments - this data set is restricted to 2005-2007 in part because the investments variable is not available for years before 2005). These two variables can be considered inputs into the process of product innovation. Exporters have higher average values for these innovation inputs, as section 2.3.5 will show.

Some firms switched exporting or foreign-ownership status between 2005 and 2007. These 'transition firms' help with the estimation procedures that follow the OLS regressions and propensity score estimation in the next two subsections.

2.3.2 Baseline Estimates - OLS

The simple OLS approach below provides the first formal test of the paper's main question. It is easy to interpret. It reports the conditional mean share of output due to new products, or the likelihood of undertaking product innovation with exporting and foreign ownership as competing explanatory factors.

Formally,

$$Product\ Innovation_{it} = \alpha + \beta Exporting_{it} + \gamma FDI_{it} + S_{pst} + \varepsilon_{it} \quad (2.1)$$

Product Innovation measures the share of output represented by products each firm produced only for the first time that year. It could also be a dummy to show firm-years with positive product innovation.⁹ Defining product innovation as the share of output that is new to a firm follows other papers in the literature like Girma et al. (2012). Using a dummy variable to represent product innovation also follows other papers in the literature (Gorodnichenko and Schnitzer, 2013; Guadalupe et al., 2012; Gorodnichenko et al., 2010). While these measures suffer the same weakness as any firm-level index that is survey-based, their strength lies in how they correspond closely to the outputs of firms innovative processes. (This represents a contrast to previous work that focus on inputs into innovation processes like R&D). Furthermore, the measures directly addresses the objective raised in the introduction to the chapter of creating new varieties for the benefit of consumers.¹⁰

⁹Being tax-irrelevant, this measure comes with fewer concerns about misreporting. Nevertheless, the definition is firm specific - one firm's new product may be another firm's staple. The official guidance advises firms to report only substantially new products under this heading. If the measure is misreported in a particular manner for certain firms or industries, the firm-fixed effects in the section 2.3.6 estimates address the concern.

¹⁰ The firm fixed effects used in the baseline analysis and subsequent tests address concerns about reporting bias in this survey data, even if it differs across firms, so long as the bias is fixed over time.

Exporting is a dummy variable equal to one for firm-years with non-zero exports. By comparison, *FDI* indicates whether the share of a firm's capital owned by entities outside China, Hong Kong, Taiwan and Macau exceeds 50%. The introduction to this section motivated the definition of foreign capital sources. Desai et al. (2004) motivated the choice of majority-ownership as the threshold for indicating foreign ownership. Their paper argues that majority- or wholly owned foreign affiliates experience more technology transfer from parent companies than minority-owned affiliates. ε_{it} is the error term.

Other control variables include industry, year and province: the S_{pst} term represents fully interacted province p , industry sector s and year t fixed effects. The default level of product innovation is usually industry-specific. For example, makers of cotton yarn are not expected to introduce new product varieties at the same rate as the firms that turn the yarn into clothing. (Hering and Poncet, 2010)'s description of the persistent and large differences between Chinese provinces in terms of economic development and R&D motivated this specification, as did the possibility of year-to-year changes in the investments that support product innovation. I leave out other variables to avoid clutter in this first-stage comparison of the firm categories.

Table 2.2 reports positive relationships between product innovation and *Exporting*. *FDI* shows a similar pattern. The conclusions do not depend on whether one measures product innovation as a share of output, or with a dummy variable. Column 1 of the table suggests that new products as a share of exporters' output will be twice the average for firms in the same sector, province and year. To interpret this term, consider that product innovation's mean value in the data is 3.9%, while 28% of firms export in the average year. Column 4 reports nearly identical predictions: firms that export are 13% more likely to introduce a new product on average, compared to non-exporters. By comparison, 10% of firm-years in the data register product innovation, which implies that exporters have about twice the rate of the average firm.

Column 2 reports on the *FDI* term, yielding a lower R^2 , and a coefficient that indicates new products are 0.3% higher as a share of output for foreign-owned firms', relative to firms in the same sector, province and year. The direction and size of the coefficient agree with prior works, e.g. (Guadalupe et al., 2012; Girma et al., 2008). 8% of firm-years fall in this majority foreign-owned category. Column 5 suggests that 0.5% more of the foreign-owned firm-years report product innovation.

Columns 3 and 6 include *FDI* and *Exporting* in the same regression. The point estimates strongly suggest that exports had a much bigger impact on innovation, and the FDI variable's

Table 2.2: Comparing Innovation: Exporting vs. FDI

	(1)	(2)	(3)	(4)	(5)	(6)
	Product Innovation			Product Innovation > 0		
Exporter	0.039*** (0.000)		0.041*** (0.000)	0.126*** (0.001)		0.133*** (0.001)
FDI		0.003*** (0.001)	-0.012*** (0.001)		0.005*** (0.001)	-0.045*** (0.001)
Constant	0.029*** (0.000)	0.040*** (0.000)	0.030*** (0.000)	0.065*** (0.000)	0.099*** (0.000)	0.066*** (0.000)
Observations	760,777	760,777	760,777	762,883	762,883	762,883
R-squared	0.081	0.072	0.081	0.141	0.110	0.142
Province-Year-Sector FEs	Y	Y	Y	Y	Y	Y

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Product innovation measures new products as a share of total output. This is the dependent variable in Columns 1-3. In columns 4 -6 the dependent variable is a dummy that indicates whether product innovation is greater than zero. The Exporting and FDI variables indicate exporting and majority-ownership by foreign entities respectively. Not shown are the coefficients for fully interacted two-digit industry, province and year fixed effects.

contribution changes signs to negative. A comparison with Commander and Svejnar (2011) is interesting: In that paper, the coefficient of the export variable effectively became zero when an FDI variable was added to the regression. The reverse is observed here. Differences in the role of export-platform FDI as well as the nature of the transition to trade in Eastern Europe may be responsible for this difference - which invites a separate study to compare drivers of product innovation in China and Eastern Europe.

This exploratory step is highly informative, but comes with many caveats: province, year and sector-fixed effects are the only controls, and the observed correlation does not clearly account for the possibility that the most innovative firms may self-select into exporting or foreign ownership. The next step approaches the question using a propensity score matching estimator.

2.3.3 Tests that Controls for Selection into Exporting or FDI

To mitigate concerns about self-selection, I use a propensity score matching estimator to measure the effect of foreign ownership and exporting on product innovation. Some definitions are in order: The causal effect of exporting on product innovation is the difference between the average performance of firms given the export treatment and comparable non-exporters. (This is the average treatment effect on the treated observations). Propensity score matching relies on contrasts between exporters and non-exporters that are similar on observed measures. The approach relies on having a robust set of descriptors for firms, such that differences between matched exporters and non-exporters that are not captured by the observed traits used for matching should be essentially random, i.e. the Conditional Independence Assumption.¹¹

I use 12 variables suggested by the literature for this matching process. The first set of variables are related to size and productivity. Size is a consistently positive predictor of exporting and foreign ownership in related studies (Arnold and Javorcik, 2009; Greenaway and Kneller, 2007; Bernard and Bradford Jensen, 1999). I apply the following measures of size in the data: the value of total assets for each firm-year and a simple count of the number of employees. I also follow Arnold and Javorcik (2009) in using productivity to predict foreign ownership. I use the same measures to predict export status: total output per asset value, in logs, and output per employee, in logs. Like the same paper, I also use firm age, the square of age and state ownership as predictors of FDI at the firm level. State-owned firms and older firms generally tend to be less open to foreign investment and exporting.

¹¹Leuven and Sianesi (2012) explains methods and tools for propensity score matching further.

Other variables that I use to predict exporting and foreign ownership status describe the sector, location and nature of investments by the firm. Firms that spend more on R&D as a share of sales are expected to innovate more, while having higher propensities to foreign ownership and exporting. I also use vintage, the ratio of the book value of equipment to their original purchase values. All firms in the sample are required to use the same accounting standards, so the measure provides a relatively uniform measure of capital equipment vintage. I also consider investments in human capital, recorded as employee training expenses divided by the total wage bill. Firms with high investments in employees' skills are expected to innovate more and export more. Finally, I use categorical variables to represent each observation's province and four-digit sector. Provinces capture proximity to foreign markets, ports, as well as the Special Economic Zones and Open Coastal Cities that were established to promote Chinese exports. Sectors are represented by categorical variables for the 445 industry groups aggregated at the US-equivalent of four-digit codes. (Appendix section II.A has a summary table for these variables).

While evaluating the propensity to export, I include a variable to capture the fraction of ownership capital owned by foreign entities. Similarly, the test step for FDI includes a measure of export intensity. One could argue that exporting firms get more visibility to potential foreign investors, and that linkages to foreign investors increase the likelihood of exporting. This also follows the convention in Arnold and Javorcik (2009).

Table 2.3 show that these observable characteristics are different for foreign-owned firms, as well as exporters. (The propensity scores used for matching are estimated in columns 1 and 3). More productive in terms of output per employee tend to be exporters or foreign-owned firms, as are firms that have equipment of recent vintage and spend less on employee training as a share of wages. The two categories also tend to produce less output with each yuan of total assets. Firms with more employees tend to be exporters, while foreign owned firms tend to have less. R&D is higher on average for exporters, but it is less for foreign-owned firms. Foreign owners tend to acquire younger firms, (which agrees with the finding in Arnold and Javorcik (2009)), while exporters tend to be older. As expected, firms that are more export-intensive also tend to be foreign-owned, while firms with a greater share of foreign ownership also tend to be exporters.

The business literature suggests that foreign investors make the decision to select overseas affiliates on observable characteristics like these, which makes the selection of variables in this paper a good starting point for estimating propensity scores. Similarly, the size, productivity, location, industry and knowledge investment variables used in this paper represent

Table 2.3: Predicting Exporting and FDI

	(1)	(2)	(3)	(4)
	Exporter Dummy		Foreign Ownership Dummy	
Output per Employee	0.106*** (0.003)	0.230*** (0.011)	0.072*** (0.005)	0.124*** (0.007)
Output per Asset Value	-0.022*** (0.002)	-0.126*** (0.007)	-0.083*** (0.003)	-0.132*** (0.005)
Log(Assets)	0.070*** (0.003)	-0.131*** (0.011)	0.259*** (0.004)	0.162*** (0.007)
Log(Employee No)	0.265*** (0.003)	0.354*** (0.011)	-0.017*** (0.005)	0.009 (0.007)
Vintage	-0.353*** (0.011)	-0.359*** (0.038)	-0.788*** (0.015)	-0.941*** (0.022)
Log(R&D)	0.035*** (0.001)	0.017*** (0.003)	-0.036*** (0.001)	-0.045*** (0.002)
Training Expenses	-0.314*** (0.037)	-0.053 (0.111)	-0.445*** (0.060)	-0.389*** (0.090)
Log(Age)	0.117*** (0.003)	0.221*** (0.011)	-0.265*** (0.004)	-0.346*** (0.006)
Foreign Ownership Share	1.043*** (0.007)	0.536*** (0.043)		
Exports Share of Sales			0.969*** (0.007)	0.565*** (0.011)
State-Owned Dummy	-0.286*** (0.007)			
Constant	-3.552*** (0.143)	2.849 (89.357)	-3.293*** (0.191)	-1.509*** (0.250)
Observations	760,567	61,207	694,327	198,681
Pseudo-R2	0.268	0.148	0.232	0.175
Exporters Only				Y
Foreign-Owned Only		Y		

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The table reports probit coefficients and standard errors for an empirical model that uses an exporting dummy as the dependent variable in columns 1-2 and majority foreign ownership for columns 3-4. Columns 1 and 3 report the propensity scores for all the usable observations, while columns 2 and 4 are for targeted subsets of the data. Column 2 shows that the propensity to export is similar on most observable firm traits for firms that majority foreign-owned, while Column shows that the propensity to be foreign-owned follows a similar pattern for exporters. The output per assets and employee variables are measures of productivity, while the log assets or employee numbers are measures of firm size. Vintage represents the age of firm assets, R&D and training expenses measure investments in knowledge or human capital. To estimate the propensity to export, I include a variable that measures foreign owners' share of firm capital, and for the propensity for foreign-ownership, I include a measure of exports as a share of sales. A dummy that indicates whether the majority of a firm's capital is state-owned is also used to predict export status. Province and industry fixed-effects applied to all columns not shown.

the common predictors of exporting activity in the economics literature. In sum, the list of matching variables - including location and 4-digit sectors not shown in the table, provide some support for the conditional independence assumption.¹²

I match exporters and FDI recipients to their nearest-neighbors. Nearest neighbors are the counterfactual items whose propensity scores are most similar to the reference observation. The propensity score is the predicted value of the exporting or FDI dummy in a first-stage probit regression using the ‘treatment’ variables that I describe in the next paragraph. Table 2.4 presents the propensity score matching estimates, which show the effects of export participation and foreign ownership in columns 1 and 2 respectively.¹³

This matching estimate of the treatment effects shows that export participation predicts an additional 6.1% of outputs that are new products (19.9% for exporters versus 13.8% for comparable non-exporting firms). Firms with majority foreign ownership under-perform relative to their peers. New products account for 12.9% of their output, compared to 19.0% for Chinese-owned firms with similar propensities. Understandably, foreign-owned firms are larger and more likely to do R&D, so the innovation benchmark is set higher than for exporters.

To address the possibility that only foreign-owned exporters account for the estimated effects of exporting, I repeat the propensity score tests on the subset of the data that is foreign-owned only. (This gives 7,527 observations on the common support, much less than the 90,461 used in column 1 of Table 2.4). Among foreign-owned firms, exporters enjoy a product innovation advantage that is comparable but less than that in the full sample (5.0%); suggesting that this subset’s average cannot account for all the export treatment

¹²It is noteworthy that the signs and statistical significance of coefficients in columns 2 and 4 largely mirror those in columns 1 and 3 of Table 2.3. The implication is that using subsets of the data do not change the pattern of export or foreign ownership propensities. Among foreign-owned firms, those with high output per employee, high R&D, newer equipment and more employees are more likely to be exporters. Similarly, among exporters, the propensity to be foreign-owned follows a pattern that is similar to the population of firms.

¹³The simple nearest-neighbor match suits this paper’s purpose. The number of observations is large, with many firms in the control and treatment categories sharing similar observable attributes. Therefore, one expects counterfactuals that roughly approximate each tested firm-year. If the overlap between control and treatment items was worse or observations fewer, one could have considered kernel matching or other N-neighbor matching to average out the control observations used.

The results from N-nearest neighbor matching are largely similar. (I use N=3 and N=5, but do not report these to conserve space). Abadie and Imbens (2009) and Caliendo and Kopeinig (2008) explain the advantages of N-nearest neighbor matching over simple nearest neighbor matching.

Table 2.4: Innovation vs. Exports and FDI: Propensity Score Matching

Dependent Variable:	Product Innovation		Product Innovation > 0	
	(1)	(2)	(3)	(4)
Exporting	0.062*** (0.004)		0.186*** (0.007)	
FDI		-0.061*** (0.006)		-0.112*** (0.009)
Constant	0.138*** (0.002)	0.190*** (0.04)	0.262*** (0.002)	0.355*** (0.002)
Observations on Common Support	90,461	78,499	82,932	73,337

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Product innovation – the dependent variable measures new products as a share of total output. Columns 3 and 4 use a dummy as the outcome variable. The reported effects are the estimated average treatment effects on treated observations (ATT).

The Exporting and FDI variables indicate exporting and majority-ownership by foreign entities respectively. Section II.A of the appendix summarizes the variables used to correct for self-selection.

effect in Table 2.4. (See Section 2.3.4 for these results).¹⁴

Section II.A in the appendix supports these results by showing that the sample selected for matching is balanced in terms of the observed covariates, and graphically illustrates the common support on the propensity score for firms that received the export or foreign ownership treatments.

Comparing the results of this set of tests with the baseline OLS estimates, the 6.0% difference obtained from the matching setup in column (1) is more than the 3.8% from the OLS regression for exporters. It is reassuring to see the two tests yield coefficients with the same sign.

¹⁴ The summary statistics on firms that change exporting/foreign-ownership are also informative. For the 61,000 observations that represent majority foreign-owned firms, new products represent 4.7% of output, while the comparable share is 6.3% for the 2,300 observations representing foreign-owned firms that changed status from non-exporter to exporter. While the product innovation measure is 6.8% for the 209,000 observations representing exporters, the 3,400 observations in the subset that changed from Chinese-owned to foreign-owned exporters reported an average of 5.8%. The data in no way suggests that foreign-owned exporters are primarily responsible for the main findings.

2.3.4 Tests on Subset of Data: Foreign-Owned Firms Only

Table 2.5 reports the estimated effect of exporting on product innovation for the subset of firms that are foreign-owned. Like Table 2.4, this table uses propensity score matching estimates. The controls or counterfactuals for each observation are the most similar observations in terms of characteristics that predicted selection into the treatment.

Table 2.5: Innovation by Exporter Status: Foreign-Owned Firms Only

Dependent Variable:	Product Innovation	Product Innovation > 0
	(1)	(2)
Exporting	0.054*** (0.014)	0.150*** (0.021)
Constant	0.122*** (0.005)	0.236*** (0.008)
Observations on Common Support	12,010	12,010

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

The product innovation measure represents new products as a share of total output. Columns 3 and 4 use a dummy as the outcome variable. The reported effects are the estimated average treatment effects on treated observations (ATT). The Exporting variable is a dummy that indicates firm-years with non-zero exports. Section II.A describes the variables used to correct for self-selection.

Within the group of majority foreign-owned firms, exporters on average have an additional 5.0% of outputs that are new products. (17.4% for exporters versus 12.0% for comparable non-exporting firms). The difference is statistically significant, with a t-statistic of 3.84. Using the product innovation dummy as the outcome of interest yields a starker difference between exporters and non-exporters. The gap of 15% between these two subgroups is also statistically significant; 38% of foreign-owned exporters create new products while only 23% of comparable non-exporters do so. (The equivalent results on the subset of exporting firms yield negative coefficients on the foreign ownership dummy, which is not surprising, given the clear differences in the summary statistics from Table 2.1).

As expected, the propensities for matched estimates were well balanced. Within the category of foreign-owned firms, exporters and non-exporters were similar in terms of size, location, employee numbers and other observed traits. The 7,527 observations on the common support, is smaller than the 90,461 used in column 1 of Table 2.4 largely because foreign-owned firms are less than 8.5% of the sample. Similar gains in product innovation

for exporters are observed if the sample was chosen to be all firms with any level of foreign ownership. I do not tabulate those results to avoid clutter.

In sum, even within the group of foreign-owned firms, exporters introduce more product innovations. This remains consistent with this paper's conclusion that while foreign ownership may lead to product innovation, the effect of exporting on product innovation is larger.

2.3.5 Mechanisms for Product Innovation


Given the findings that link higher levels of product innovation to exporting, this section explores possible mechanisms that enable product innovation. Intangible factors associated with exporting or foreign ownership may drive the decision to create new products, but the act of creating new products must require measurable changes to the factors of production. Examples of those tangible changes could be investments in R&D to develop or improve products. It could also be investments in equipment to change production processes and methods.

In other words, I ask whether firms learn product innovation by investing in R&D and new production assets. The focus on these two potential mechanisms is arguably justified by the context. Aggregate R&D as a share of GDP in China was growing throughout this period. At the firm level, investing in R&D clearly indicates a commitment to learning, which could translate into product innovation. In the same vein, asset purchases could reflect technology diffusion through the acquisition of assets with embodied knowledge, as is well documented for China (Brahmbhatt and Hu, 2010; Augier et al., 2013). A large number of Chinese producers import their production equipment, which usually embodies associated production methods (Woo, 2012).¹⁵

To the extent that exporting is the causal driver for product innovation, it should also be causal to these changes in production, observed and unobserved. In other words, if firms learn to undertake actions like R&D necessary for innovation as they export, the observed measures of these mechanisms should increase when firms begin to export, grow as firms continue to export and decline for firms that stop exporting. Figure 2.1 illustrates this pattern of learning-by-exporting. (In contrast, the selection hypothesis would predict small increases on transition into the treatment, and no changes thereafter). Following the argument in De Loecker (2007), R&D and new assets could be mechanisms that firms learn

¹⁵ Other causal drivers of product innovation may exist outside the two that are central to this section of the paper. The approach to estimating the causal relationship in equation (2.2) addresses the possibility of other unobserved causes that do not change in the short run.

Figure 2.1: Illustrative Export-Driven Innovation Pattern



This graph is purely illustrative. It was not created from real data.

as they export, and in learning, become more productive.

Formally, for a set of mechanisms that lead to product innovation Z :

$$Z_{it} = \gamma_i + \lambda_t + \beta S_{it} + \bar{\varepsilon}_{ist} \quad (2.2)$$

S represents the exporting or foreign ownership treatment status; γ helps to address selection - it is the average difference between exporters and non-exporters (or foreign vs. domestically owned firms). β is the parameter of interest, it measures the extent to which firm i changes Z because its ownership or exporting status changed. Z represents the set of causal factors like R&D, and investment in fixed capital. Firms may not report all elements of Z in the data.

$\beta = E(Z_{after,treated} - Z_{before,treated} + Z_{after,untreated} - Z_{before,untreated})$ is the identifying assumption in (2.2), i.e. $E(\bar{\varepsilon}) = 0$. This is reasonable, especially if one includes firm fixed effects.

In other words, R&D spending and asset investments should experience a positive shock right about when a firm starts to export, the positive trend should continue on a reduced scale for firms that keep exporting and one should see an incomplete reversal of the increased patterns of investment for firms that stop exporting. The reversal should be incomplete if

those that stopped learned from their export experience.

Testing this idea is a regression model that extends the specification used by Bernard and Bradford Jensen (1999) and De Loecker (2007). The primary differences in this case are: (1) I test for innovation drivers, not productivity on the left hand side, and (2) I include lagged values of the dependent variable to reduce concerns about endogeneity.

Formally:

$$\ln Z_{it} = a + \ln Z_{it-1} + \beta_1 Start_{it} + \beta_2 Stay_{it} + \beta_3 Stop_{it} + \beta_4 Size_{it} + e_{it} \quad (2.3)$$

Start, *Stay* and *Stop* capture all the possible treatment status options for a firm in (2.3). For the exporting treatment, *Start* indicates firms that do not export in year $t-1$ but export in year t and *Stay* denotes firms that export in year $t-1$ and continue to export in year t . *Stop* flags firms that exported in year $t-1$, but failed to register exports in year t . X is a placeholder for the matrix of firm characteristics that include size, industry and location. To interpret the regression, one should consider that only observations in 2006-2007 are usable: of these, 4% of observations fit the *starting exports* category, 25% fall in the *Stay* category and 4% are observations corresponding to firm-years where exporting stopped. Firm-years unrelated to exporting make up the remaining 67% of observations.

Given a causal relationship between exporting or foreign ownership and R&D for example, one must still show R&D is causally linked to product innovation. Correlation would be sufficient if reverse causation were impossible. In this case, it is possible that firms undertake R&D or asset purchases after embarking on a course of product innovation for another reason.

Formally:

$$Product\ Innovation_{it} = \alpha_i + \alpha_2 Z_{it} + \hat{\epsilon}_{it} \quad (2.4)$$

Firms fixed effects α_i help to identify the relationship in (2.4) as causal. (This approach also mitigates bias due to omitted elements of Z that are firm specific). If $\alpha_2 > 0$ in (2.4) and $\beta > 0$ in (2.3), one could argue that the variables in Z_{it} are possible mechanisms through which exporters or foreign-owned firms undertake product innovation. The rest of this section focuses on estimating (2.3) and (2.4).

I rely on a Tobit empirical specification for (2.2) as relatively few firms undertake research and development. R&D expenses are greater than zero for only 83,176 of the 763,036 firm-year observations in the data. These expenses are attributable to 45,340 firms. Even these firms do not spend on R&D in every year; they account for 127,883 observations, which

Table 2.6: Changes at the Export and FDI Transitions

VARIABLES	Before Exporting	After Exporting	Before FDI	After FDI
	Group Averages			
Product Innovation	.049	.077	.044	.050
I(Product Innovation > 0)	.104	.206	.094	.106
R&D	475.93	773.75	464.47	632.10
I(R&D > 0)	.133	.165	.113	.127
Log(Asset Investments)	7.322	7.572	7.286	7.421
I(Asset Investments > 0)	.243	.262	.282	.295
N	15726		5688	

suggests that for them, R&D expense occurs on average in about 2 of 3 years. I use the cash flow into investments reported in the data to represent asset investments - the value of this measure is positive for about 20% of observations, although they tend to be higher for exporting firms.

Table 2.6 presents some non-parametric comparisons before the regression exercises. It shows differences in exporting, R&D and asset purchases for firms that changed exporting or FDI status. Only 15,700 and 5,700 firms fit each of these categories, but those numbers are large enough to be instructive in this summary table format. As the dataset is a short 3-year panel, no distinction is made between firms that started exporting in 2006 or 2007. The table provides suggestive evidence of a strong relationship between the transition to exporting and product innovation, with exporting having a stronger association than FDI. 20% undertake product innovation in a year they export, compared to 10% for the same firms before exporting. (The comparable numbers are 10.6 and 9.4% for foreign ownership). Table II.B.1 in the appendix provides further support, with propensity score matching estimates at the transition to and from exporting/FDI. The share of output due to new products also increases in Table 2.6, while an additional 3% of firms start spending on R&D in the year of exporting relative to the year before exporting. About 12.7% of firms that received foreign capital undertake R&D; in the year before receiving foreign capital the fraction is 11.3% – so the incidence of R&D increases with foreign ownership, just not as much as with exporting.

The complimentary measure relies on the idea that for industrial firms, new equipment are strongly associated with new production processes and product innovation. To capture this,

I measure asset investments: using the recorded value of cash outflows into investments.¹⁶ The question of interest here is whether positive values of asset investments are correlated with product innovation and exporting.

Table 2.7 shows that product innovation increases for firms with R&D and assets investments, status notwithstanding. To avoid discarding the nearly 90% of observations with zero R&D expenditure, I represent this firm-year attribute with two variables: One variable takes a value of zero if reported R&D is zero but is the log of R&D expense otherwise. I include a dummy variable for non-zero R&D expense in the OLS regression to address potential bias from the prevalence of zeros. This is less of an issue for asset purchases, but I use a similar dummy as a precaution. I also controlled for size - measured as the log of total asset values and employee numbers. I also kept the usual dummies for fully interacted 2-digit industry year, as well as province dummies.

Column 1 of the table indicates that conditional on having any R&D expense, firms that do more R&D also tend to do more product innovation. For example, if the average firm increases its R&D spending by one standard deviation (1.9), it would increase its output of new products by 25% and its likelihood of product innovation by 2%. Firm fixed effects help to address concerns about endogenous R&D or reverse causality. The set of firms that do not report any R&D at all tend to engage in some product innovation after controlling for firm size in the first two specifications. This supports the argument that firms may have other approaches to product development like staff training that are not reported in a separate cost category like R&D. (If the regression was run without the $I(R&D)$ dummy, the coefficient on the $Log(R&D)$ variable remains positive and statistically significant). In summary, the evidence suggests a role for R&D in product innovation, with indications that other factors also play a part in product innovation. Column 4 mimics the pattern in Column 1 for the product innovation dummy.

Columns 2 and 5 suggest that having asset investments predicts slightly higher levels product innovation. This is true after controlling for firm size and whether the firm made any investments at all. (Asset investments occur for only about a third of the observations in the data). The estimated coefficients on this variable are statistically significant, even if the coefficients of the dummy variable are not as distinguishable from zero. Large firms in terms of assets and employees do not appear to be most likely to introduce product innovations, but new products are a greater share of output for these large firms.

¹⁶ As these are industrial firms, for outlays to be recorded as investments, they must be for equipment or equally durable productive assets.

Table 2.7: R&D and Asset Purchases Increase Product Innovation

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	Log(New Product)			I(Product Innovation)		
Log(R&D)	0.005*** (0.000)		0.005*** (0.000)	0.011*** (0.001)		0.010*** (0.001)
I(R&D > 0)	0.007*** (0.002)		0.007*** (0.002)	0.003 (0.003)		0.003 (0.003)
Log(Asset Investments)		0.001*** (0.000)	0.001*** (0.000)		0.002*** (0.000)	0.001*** (0.000)
I(Asset Investments)		0.004*** (0.001)	0.003** (0.001)		0.005* (0.003)	0.004 (0.003)
Log(Total Assets)	0.002*** (0.000)	0.002*** (0.000)	0.001*** (0.000)	0.001 (0.001)	0.005*** (0.001)	0.001 (0.001)
Log(Employees)	0.001*** (0.000)	0.001*** (0.000)	0.001*** (0.000)	-0.000 (0.000)	0.004*** (0.001)	-0.000 (0.000)
Constant	-0.004 (0.030)	-0.011 (0.031)	-0.006 (0.030)	0.091 (0.058)	0.042 (0.062)	0.090 (0.058)
Observations	760,777	760,777	760,777	762,883	762,883	762,883
R-squared	0.788	0.787	0.788	0.754	0.754	0.754
Prov. FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The product innovation measure follows the definition in the rest of the paper: new products as a share of total output. Columns 4-6 use a product innovation dummy as the dependent variable. Logged values of R&D and Asset Investments for each firm year are the key control variables (I include a dummy variable if the reported value is zero - the $I()$ items. In those cases, the logged variable is set to zero). The dummy variables that indicate when these variables are zero address possible concern about bias due to the zeros in the RHS variables. I also include size controls – the log of total assets and the log of total employee numbers. To interpret these OLS estimates, it helps to know that the mean values of the Log(New Product), I(Product Innovation), Log(R&D) and Log(Assets Investments) variables are 0.92, 0.09, 0.71 and 5.78 respectively. 90 and 60 % of the observations had a value of zero for R&D and Asset Purchases.

Columns 3 and 6 combine the two key variables and their dummy indicators and yield estimates that are consistent with the other specifications in Table 2.7. Firms that undertake R&D tend to do product innovation and investing in assets predicts product innovation, although the coefficients of the size variables are not statistically significant in the Column 6 - suggesting that the incidence of product innovation may not be linked to size after one controls for inputs like asset investments and R&D.

Table 2.8 links exporting and foreign ownership to R&D and asset investments. Column 1 shows R&D, while Column 2 shows asset investments. The annual survey dataset reports both the depreciated and original or purchase values of fixed assets. Therefore, it is possible to track net asset purchases using their reported original values for fixed assets in 2006-2007. I estimate equation (2.3) with the net purchases data.¹⁷

Firms' patterns of spending on innovation inputs change in the period that they start exporting. The Tobit specification in Table 2.8 shows that R&D for the average firm increases by almost an order of magnitude when a firm starts exporting and continues to increase for firms that remain exporters. (Using the average value for the R&D variable 455 and the coefficients in column 1, calculated as $\frac{3,347.525-455}{455}$ and $\frac{2,875.744-455}{455}$). That group in turn, invests more than firms that stop exporting. The firms that stop exporting still invest more than firms with no export record. That positive difference suggests a positive effect from their past export experience. (Their R&D spending is more than 200% above average). Asset investment patterns do not follow the trend exactly, but remain broadly consistent: firms that start exporting invest more than the average non-exporter, while those that stop exporting reduce their investments. Experienced exporters invest in equipment at higher rates than new exporters or non-exporters, arguably because exporting provides new growth opportunities. That said, the three-year panel data can only provide limited support for claims about investment trends in the long run. The size controls behave as expected; firms that are larger in terms of assets or employees also undertake more R&D and investments in fixed assets.

In contrast, firms that start FDI do not spend on R&D more than the average firm. Their level of asset investments actually falls below the average domestically owned firm

¹⁷To use net asset purchases in place of gross asset purchases, I only need to assume that asset sales are small relative to purchases. (Note that the paper argues for a link between product innovation and new equipment purchases rather than net investments). Negative asset purchases are rare in the data, consistent with this assumption. Using net purchases, i.e. gross asset purchases minus disposals, biases the dependent variable towards zero, which implies that my estimates are conservative. If foreign-owned or exporting firms were upgrading equipment, which requires the disposal of old assets, one expects those firm categories to have higher-than-average disposals, hence more conservative estimates of asset purchases.

Table 2.8: Innovation Drivers by Stage of Export/FDI Participation

Dependent Variables:	(1) R&D	(2) Asset Purchases	(3) R&D	(4) Asset Purchases
Started_Exports	3,347.525*** (363.363)	26,332.042*** (6,106.260)		
Stayed_Exports	(182.941)	(2,935.440)	2,875.744***	46,722.274***
Stopped_Exports	1,254.035*** (401.223)	-12,574.655** (6,322.127)		
Started_FDI			-4,895.070*** (640.465)	-14,372.539 (9,450.938)
Stayed_FDI			-5,016.336*** (278.767)	8,941.144** (4,322.030)
Stopped_FDI			-4,222.461*** (677.583)	-34,973.959*** (10,105.218)
R&D, Lagged	1.126*** (0.002)		1.125*** (0.002)	
Asset Investments, Lagged		0.101*** (0.003)		0.101*** (0.003)
Constant	-95,297.229*** (923.096)	-1449025.224*** (14,876.781)	-97,851.472*** (922.512)	-1469633.267*** (14,851.725)
σ	27,249.671*** (53.514)	521,751.713*** (1,167.716)	27,241.946*** (53.489)	521,782.902*** (1,167.825)
Observations	437,768	437,768	437,768	437,768
Prov. FE	Y	Y	Y	Y
Ind.-Year FE	Y	Y	Y	Y
Size Controls	Y	Y	Y	Y

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The dependent variable is the value of R&D and asset investments undertaken in each firm-year. The main explanatory variables are firms' foreign ownership or export status. I also use lagged values of the dependent variable to mitigate concerns about endogeneity. The σ captures the Tobit specification's equivalent of the square root of residual variance.

after controlling for other factors. The estimated effect of the change in ownership status is statistically significant, and supports a narrative in which foreign investors ship the R&D arms of their Chinese affiliates elsewhere. Similar patterns obtain for remaining majority foreign-owned or reverting from foreign-owned to domestic ownership. Asset investments have a slightly different pattern, firms in the first year of foreign ownership on are not distinguishable from the average firm on this measure. However, firms that remain foreign-owned past the first year have asset investments more than twice the average of 4,660 yuan, after controlling for other factors. As expected, firms that revert to domestic ownership from foreign investors decrease their investment outlays.

The pattern of lower R&D spending by foreign-owned entities is consistent with the literature - multinationals generally prefer to keep R&D centralized where they have stronger intellectual property protection (Fernandes and Tang, 2012; Branstetter et al., 2006). In contrast, locally owned exporters generally do not have the option to outsource their R&D. Their spending on R&D and assets therefore reflects their efforts to update production processes as they compete in global markets.

These differences in patterns of R&D growth experienced by exporters may lend some credence to the Economist magazine's claim: Exporters in developing economies are staking their claim on the innovation terrain.

2.3.6 Learning with Corrected Biases

To address concerns that the regression coefficients in Table 2.8 are biased upwards because of self-selection, the next two paragraphs present the results of tests that match treated observations to counterfactuals with similar observed attributes (propensity score matching).

Table 2.9 presents results consistent with the findings of the OLS step. Each 'treatment condition' is tested separately. In the spirit of matching propensities for the treatment and control groups, the control group was selected to match each treatment: Observations with the *Start* treatment were matched to others who were similarly not exporters in the previous period. *Stay* was matched against new exporters and those that had stopped exporting, while those with the *Stop* treatment were compared with firms that had no exporting history.

As in Table 2.8, firms that start exporting invest more in R&D and fixed production capital than their non-exporting peers. (774 for new exporters vs. 668 for their non-exporter peers). The estimate in Table 2.9 is not statistically significant for first-year exporters. For firms in first year as foreign-owned, the sign is opposite and the coefficient is statistically

Table 2.9: Matching Estimates by Stage of Export Participation

VARIABLES	Log(R&D)	Log(Asset Investments)
Started_Exports	156.244 (144.959)	1762.503*** (789.824)
Stayed_Exports	602.194*** (170.850)	3383.571*** (1645.369)
Stopped_Exports	-143.509** (78.417)	-885.979 (843.667)
Started_FDI	-210.047*** (93.668)	-1168.146 (2087.223)
Stayed_FDI	77.412 (181.583)	4999.749 (4548.513)
Stopped_FDI	-113.867 (130.826)	-1355.626 (1380.328)

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: The reported effects are the estimated average treatment effects on treated observations (ATT). For these propensity score matching exercises, the counterfactual for each row was limited to comparable firm-years as follows: *Started_** was matched to observations not foreign-owned or an exporter, and *Stayed_** to observations with a history of exporting or foreign ownership, but currently not in a second consecutive year in that status. *Stopped_** was matched to either non-exporters or firms with no foreign ownership in that year. The dependent variables are the reported values of R&D and asset investments. The number of treated observations were 15,724, 110,260 and 17,100 respectively for Columns 1 of Exports. The FDI segment had 6,757, 35,394 and 5,636 treated observations. The numbers vary by segment because the match was limited to items on the common support. The matching variables include firm size, output per assets and employee, as well as 4-digit industry dummies. Further detail on the mean outcomes for treated and untreated items, the control items on common support and balancing tests for the matching variables are available on request from the author.

significant. In the first year that a firm becomes majority foreign-owned it invests less in R&D than comparable Chinese-owned firms. This is consistent with Table 2.8, and supports the suggestion that when firms start exporting, they learn to do R&D. Chinese firms that become foreign-owned may actually reduce their R&D efforts if the parent company opts to locate R&D efforts elsewhere, to maintain better control over intellectual property rights. While firms that start exporting out-invest peers in terms of production assets, the difference is not statistically significant for firms in the first year of foreign-ownership status.

The estimates in Table 2.9 were not intended to measure the effects of each transition, but to show how firm spending changes with each transition in export or FDI status. Firms that remained as exporters spend more on R&D than new exporters and firms that stopped exporting – the comparison group for this exercise. Similarly, spending on assets investments is higher for firms that stay as exporters in a statistically significant sense. Firms that remained majority foreign-owned, compared with new or formerly foreign-owned firms do not register any statistically significant difference in their spending on R&D and asset purchases.

Firms that stop exporting invest less in R&D and new capital than other non-exporting peers. The difference is small enough that it is not statistically significant at the 99% level, however. This may imply that characteristics like output per employee or other matching variables drive the learning suggested by Table 2.8. It does not invalidate the claim altogether, just how it is interpreted. Firms that changed from majority foreign ownership report lower outlays on both measures, but not with any level of statistical significance. These firms spend roughly the same on R&D and asset investments as comparable Chinese-owned firms.

2.3.7 Other Empirical Specifications

The definition of foreign capital excluded funds from Hong Kong, Macau and Taiwan (HMT) throughout this paper. This definition was motivated by the similarity of business cultures, technology and connections in the region.

Nevertheless, I show below in Table 2.10 that the coefficients of the OLS tests in Tables 2.2 and 2.8 would remain mostly unchanged if foreign capital were redefined to include funds from Hong Kong, Macau and Taiwan. (The implication is that the two categories of foreign capital sources in the data are not inherently associated with different propensities for product innovation). For the PSM tests, matching estimates for both versions of the model are broadly similar, showing that firms increase R&D and asset purchases when they enter the export market, invest more as they remain exporters, and reduce the pattern if they stop exporting, but not to the level of firms that never exported.

Table 2.10 only indicates that the conclusions of this paper should not change, even if the definition of foreign capital had been more expansive from the start. In fact, I expect any other definition of foreign capital to enhance the contrast between the effects of trade and foreign investment presented in Tables 2.4 and 2.8.

Table 2.10: Comparing coefficients for FDI with HMT

	(1)	(2)	(3)	(4)	(5)	(6)
	Product Innovation			Product Innovation > 0		
Exporter	0.039*** (0.000)		0.044*** (0.000)	0.126*** (0.001)		0.145*** (0.001)
FDI with HMT		-0.001 (0.001)	-0.019*** (0.001)		-0.007*** (0.001)	-0.065*** (0.001)
Constant	0.029*** (0.000)	0.040*** (0.000)	0.031*** (0.000)	0.065*** (0.000)	0.101*** (0.000)	0.070*** (0.000)
Observations	760,777	760,777	760,777	762,883	762,883	762,883
R-squared	0.081	0.072	0.082	0.141	0.110	0.146

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Note: The dependent variable in columns 1 - 3 is new product's share of total output, while columns 4 - 6 use a dummy that is 1 if new products represent a positive share of outputs. FDI with HMT is a categorical variable that switches from zero to 1 if more than 50% of ownership is from outside mainland China. Foreign capital in this table is defined to include Hong Kong, Macau and Taiwan (HMT). In sign and significance, the results are comparable to Table 2.2

The appendix includes tests of the match quality for all the propensity score-based tests in the previous section.

2.4 Conclusions

This paper compares the direct impacts of exporting and foreign ownership (FDI) on product innovation. FDI and export promotion are the two main channels that developing economies have adopted to lead private sector growth; hence the motivation to evaluate their relative merits in promoting product innovation (Harding and Javorcik, 2011; ?). Firms with an interest in stimulating product innovation may also consider the same question as a matter of strategy.

Using propensity score matching methods and rich firm level data, this paper shows that exporting causes firms to engage in greater levels of product innovation. The finding lends support to the ‘learning-by-exporting’ hypothesis presented in Bratti and Felice (2012), Damijan et al. (2010) and De Loecker (2007). FDI does not yield the same impact on product innovation, in terms of either incidence or intensity. In some specifications, foreign ownership actually leads to less innovation and less spending on items like R&D. In a developing economy like China, the absence of a positive relationship between FDI and innovation may be due to foreign owners’ efforts to protect intellectual capital by moving R&D abroad (Fernandes and Tang, 2012; Branstetter et al., 2006). Those firms could also be reducing innovation efforts in the developing-economy subsidiary to avoid effort duplication.

I further explore potential pathways for the relationship between exporting and innovation, using innovation inputs like R&D and asset purchases, as R&D predicts product innovation in this context. Exporting or foreign ownership may drive the decision to create new products, but the act of creating new products must require measurable changes to these or other innovation inputs. Estimates from that exercise indicate that firms that start exporting undertake more R&D and invest more in new production assets. These results also suggest that firms learn from exporting – firms that stop exporting spend more on R&D and new assets than the average non-exporter, even if less than new or continuing exporters. In all specifications, firms that change from Chinese to foreign ownership reduce R&D spending on average. Their asset purchases are higher than average, but less than the comparable number for new exporters.

The contrast between these findings and Commander and Svejnar (2011) suggest that context may matter for whether foreign investment leads to product innovation. On the other hand, exporting consistently predicts higher levels of innovation efforts like R&D and better product innovation outcomes. In a context where a foreign owner only wants the low production cost of a location like China, foreign ownership may actually lead to lower levels of product innovation. The owners’ priorities determine whether the firm undertakes costly innovation efforts.

Relating these findings to papers like Commander and Svejnar (2011) and Guadalupe et al. (2012) that find a positive relationship between product innovation and foreign ownership in European contexts signals the potential for additional work on how context, property rights and economic development influence technology transfer through ownership.

Bibliography

- Abadie, A., Imbens, G. W., 2009. Matching on the Estimated Propensity Score. Tech. rep., National Bureau of Economic Research.
- Aitken, B., Hanson, G. H., Harrison, A. E., 1997. Spillovers, Foreign Investment, and Export Behavior. *Journal of International Economics* 43 (1), 103–132.
- Aitken, B., Harrison, A., 1999. Do Domestic Firms Benefit from Direct Foreign Investment? Evidence from Venezuela. *American Economic Review* 89 (3), 605–618, cited By 889.
- Amiti, M., Freund, C., 2010. The anatomy of china’s export growth. In: *China’s Growing Role in World Trade*. University of Chicago Press, pp. 35–56.
- Arnold, J. M., Javorcik, B. S., 2009. Gifted kids or pushy parents? foreign direct investment and plant productivity in indonesia. *Journal of International Economics* 79 (1), 42–53.
- Augier, P., Cadot, O., DAVIS, M., 2013. Imports and TFP at the firm level: The role of absorptive capacity. *Canadian Journal of Economics/Revue canadienne d’économique* 46 (3), 956–981.
- Aw, B. Y., Chung, S., Roberts, M. J., 2000. Productivity and Turnover in the Export Market: Micro-level Evidence from the Republic of Korea and Taiwan (China). *World Bank Economic Review* 14 (1), 65–90.
- Aw, B. Y., Roberts, M. J., Xu, D. Y., 2008. R&D Investments, Exporting, and the Evolution of Firm Productivity. *American Economic Review*, 451–456.
- Benhabib, J., Perla, J., Tonetti, C., 2014. Catch-up and Fall-back through Innovation and Imitation. *Journal of Economic Growth* 19 (1).
- Bernard, A. B., Bradford Jensen, J., 1999. Exceptional Exporter Performance: Cause, Effect, or Both? *Journal of International Economics* 47 (1), 1–25.
- Bernard, A. B., Jensen, J. B., 2004. Why Some Firms Export. *Review of Economics and Statistics* 86 (2), 561–569.
- Bernard, A. B., Redding, S. J., Schott, P. K., 2011. Multiproduct Firms and Trade Liberalization. *Quarterly Journal of Economics* 126 (3), 1271–1318.
- Brahmbhatt, M., Hu, A., 2010. Ideas and Innovation in East Asia. *World Bank Research Observer* 25 (2), 177–207.
- Branstetter, L. G., Fisman, R., Foley, C. F., 2006. Do stronger intellectual property rights increase international technology transfer? empirical evidence from us firm-level panel data. *Quarterly Journal of Economics* 121, 321–349.

- Bratti, M., Felice, G., 2012. Are Exporters More Likely to Introduce Product Innovations? *The World Economy* 35 (11), 1559–1598.
- Broda, C., Weinstein, D. E., 2006. Globalization and the Gains from Variety. *Quarterly Journal of Economics* 121 (2), 541–585.
- Caliendo, M., Kopeinig, S., 2008. Some Practical Guidance for the Implementation of Propensity Score Matching. *Journal of economic surveys* 22 (1), 31–72.
- Castellani, D., Zanfei, A., 2006. Multinational Firms, Innovation and Productivity. *Multinational Firms, Innovation and Productivity*.
- Clerides, S. K., Lach, S., Tybout, J. R., 1998. Is Learning by Exporting Important? Microdynamic Evidence from Colombia, Mexico, and Morocco. *Quarterly Journal of Economics* 113 (3), 903–947.
- Commander, S., Svejnar, J., 2011. Business Environment, Exports, Ownership, and Firm Performance. *Review of Economics and Statistics* 93 (1), 309–337.
- Damijan, J. P., Kostevc, C., Polanec, S., 2010. From Innovation to Exporting or Vice Versa? *World Economy* 33 (3), 374–398.
- De Loecker, J., 2007. Do Exports Generate Higher Productivity? Evidence from Slovenia. *Journal of International Economics* 73 (1), 69–98.
- De Loecker, J., 2013. Detecting Learning by Exporting. *American Economic Journal: Microeconomics* 5 (3), 1–21.
- Desai, M. A., Foley, C. F., Hines Jr, J. R., 2004. The Costs of Shared Ownership: Evidence from International Joint Ventures. *Journal of Financial Economics* 73 (2), 323–374.
- Djankov, S., Hoekman, B., 2000. Foreign Investment and Productivity Growth in Czech Enterprises. *World Bank Economic Review* 14 (1), 49–64.
- Eckel, C., Neary, J. P., 2010. Multi-Product Firms and Flexible Manufacturing in the Global Economy. *Review of Economic Studies* 77 (1), 188–217.
- Fernandes, A. P., Tang, H., 2012. Determinants of vertical integration in export processing: Theory and evidence from China. *Journal of Development Economics* 99 (2), 396–414.
- Girma, S., Gong, Y., Görg, H., 2008. Foreign Direct Investment, Access to Finance, and Innovation Activity in Chinese Enterprises. *The World Bank Economic Review* 22 (2), 367–382.
- Girma, S., Gong, Y., Görg, H., Lancheros, S., 2012. Foreign Ownership Structure, Technology Upgrading and Exports: Evidence from Chinese Firms. *Kiel Working Paper*.

- Gorodnichenko, Y., Schnitzer, M., 2013. Financial Constraints and Innovation: Why Poor Countries Don't Catch Up. *Journal of the European Economic association* 11 (5), 1115–1152.
- Gorodnichenko, Y., Svejnar, J., Terrell, K., 2010. Globalization and Innovation in Emerging Markets. *American Economic Journal-Macroeconomics* 2 (2), 194–226.
- Greenaway, D., Kneller, R., 2007. Firm Heterogeneity, Exporting and Foreign Direct Investment*. *The Economic Journal* 117 (517), F134–F161.
- Guadalupe, M., Kuzmina, O., Thomas, C., 2012. Innovation and Foreign Ownership. *American Economic Review* 102 (7), 3594–3627.
- Harding, T., Javorcik, B. S., 2011. Roll Out the Red Carpet and They Will Come: Investment Promotion and FDI Inflows. *The Economic Journal* 121 (557), 1445–1476.
- Hering, L., Poncet, S., 2010. Market access and Individual wages: Evidence from China. *Review of Economics and Statistics* 92 (1), 145–159.
- Hu, A. G., Jefferson, G. H., 2002. FDI Impact and Spillover: Evidence from China's Electronic and Textile Industries. *World Economy* 25 (8), 1063–1076.
- Hummels, D., Klenow, P., 2005. The Quality and Variety of a Nations Trade. *American Economics Review* 95 (3), 704–723.
- Iacovone, L., Javorcik, B., Keller, W., Tybout, J., 2009. Walmart in Mexico: The Impact of FDI on Innovation and Industry Productivity. University of Colorado Working Paper.
- Javorcik, B. S., 2004. Does Foreign Direct Investment Increase the Productivity of Domestic Firms? In Search of Spillovers through Backward Linkages. *American Economic Review*, 605–627.
- Keller, W., 2004. International Technology Diffusion. *Journal of Economic Literature* 42 (3), 752–782.
- Kneller, R., Pisu, M., 2007. Industrial Linkages and Export Spillovers from FDI. *World Economy* 30 (1), 105–134.
- Krugman, P., 1980. Scale Economies, Product Differentiation, and the Pattern of Trade. *The American Economic Review* 70 (5), 950–959.
- Lai, E. L.-C., 1998. International Intellectual Property Rights Protection and the Rate of Product Innovation. *Journal of Development economics* 55 (1), 133–153.
- Leuven, E., Sianesi, B., 2012. PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing. *Statistical Software Components*.

- Lileeva, A., Treffer, D., 2010. Improved access to foreign markets raises plant-level productivity for some plants. *The Quarterly Journal of Economics* 125 (3), 1051–1099.
- Madsen, J. B., 2008. Semi-endogenous versus Schumpeterian growth models: testing the knowledge production function using international data. *Journal of Economic growth* 13 (1), 1–26.
- Schott, P. K., 2008. The Relative Sophistication of Chinese Exports. *Economic Policy* 23 (53), 5–49.
- Schumpeter, J. A., 1942. *Capitalism, Socialism and Democracy*. Routledge.
- Segerstrom, P. S., 1991. Innovation, Imitation, and Economic Growth. *Journal of Political Economy*, 807–827.
- Sun, S., 2009. How Does FDI Affect Domestic Firms' Exports? Industrial Evidence. *World Economy* 32 (8), 1203–1222.
- Woo, J., 2012. Technological Upgrading in China and India: What Do We Know? OECD Development Centre Working Papers (308).
- Xiao, G., 2004. Peoples Republic of Chinas round-tripping FDI: Scale, causes and implications. Asia Development Bank Institute Discussion Paper 7.

Appendix

II.A Covariate Balancing and Common Support

The key variables are summarized in Table II.A.1.

Table II.A.2 reports the standardized bias before and after matching for the results reported in Table 2.4. The group averages for the variables used to predict exporting and FDI are generally within the 5% bias range that is considered reasonable (Caliendo and Kopeinig, 2008). This ranges in absolute terms from -0.2% for variable the output asset ratio to -8.8% for the foreign share of ownership. (The corresponding range for FDI is -0.4% and -4.4% for Employees and Assets respectively).

The quality of common support for covariates can also be shown as a histogram of propensity scores for each of the firm categories. Figure II.A.1 represents both Exporting and FDI categories' propensity scores. The upper histogram in red shows the distribution of propensity scores for exporters (or foreign-owned firms). As expected, this histogram falls to right of the lower or blue histogram of untreated observations. The firms that exports, or those that are foreign owned tend to have higher predicted probabilities of being exporters (or foreign-owned).

Tables II.A.3 and II.A.4 report the standardized bias before and after matching for the exporting and FDI results reported in Table 2.9. The matching between the control and treatment groups is excellent, with bias being less than 5% in all cases except for Output per employee and Foreign Share of Ownership for the Stay_Exporting variable and for 9 of the 28 tests for FDI.

II.B Product Innovation Before and After Exporting or FDI

Table II.B.1 provides more details findings in support of the summary in Table 2.6 that product innovation increases in the first year of exporting for firms that become exporters.

Table II.A.1: Summary of Key Variables

Variable	Mean	Std. Dev.	Min.	Max.	N
year	2006.08	0.813	2005	2007	763036
Product Innovation	0.04	0.163	0	1	760930
Exporting	0.274	0.446	0	1	763036
FDI	0.08	0.272	0	1	763036
FDI with HMT Capital	0.163	0.37	0	1	763036
Export Share of Sales	0.167	0.337	0	1	760992
Foreign Share of Ownership	0.085	0.262	0	1	763036
Asset Purchases Index	0.365	0.291	0	1	763036
State-Owned Dummy	0.088	0.283	0	1	763036
Started_Exporting	0.036	0.186	0	1	437841
Stayed_Exporting	0.252	0.434	0	1	437841
Stopped_Exporting	0.039	0.194	0	1	437841
Started_FDI	0.013	0.113	0	1	437841
Stayed_FDI	0.07	0.255	0	1	437841
Stop_FDI	0.012	0.109	0	1	437841
Age	9.282	9.104	1	126	763036
Log(Age)	1.898	0.806	0	4.836	763036
Employees	192.794	810.633	1	188151	763036
Log(Employees)	4.337	1.288	0	12.145	763036
R&D Expenses	454.727	16747.735	0	7142497	763036
Equipment Vintage	0.698	0.208	0	1	763036
Equipment (Original Value)	39782.204	532832.116	1	157000000	763036
Equipment (Current Value)	29484.101	355024.556	1	76589209	763036
Total Assets	82449.921	775616.848	1	154000000	763036
Log(Total Assets)	9.776	1.397	0	18.852	763036
Output	102839.952	908177.022	0	186000000	763036
New Product Value	12634.049	376061.876	0	110000000	763036
Sales	100776.006	898002.239	0	187000000	763036
Exports	22308.734	441338.881	0	181000000	763036
Paid up Capital	19685.688	156928.073	0	17512000	763036

Table II.A.2: Balancing Test for Propensity Score Matching Variables

Predictors of Exporting					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.5738	1.5765	-0.2	-0.29	0.771
Output/Employees	6.1047	6.1443	-3.5	-4.48	0.000
Foreign Share of Ownership	.17867	.20366	-8.8	-8.93	0.000
State-Owned Dummy	.08919	.08761	0.5	0.72	0.471
Log(Total Assets)	11.54	11.522	1.1	1.34	0.181
Log(Employees)	5.5823	5.5187	4.6	5.86	0.000
Equipment Vintage	.65905	.66529	-3.3	-4.34	0.000
Log(R&D)	6.1893	6.1547	1.6	1.93	0.054
Employee Training	.01378	.01403	-0.4	-0.71	0.479
Log(Age)	2.2552	2.2174	4.6	5.97	0.000

Predictors of Majority Foreign Ownership					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.5869	1.6055	-1.5	-0.94	0.348
Output/Employees	6.4284	6.4654	-3.2	-1.85	0.064
Export Share of Sales	.38399	.37957	1.3	0.69	0.489
Log(Total Assets)	11.615	11.684	-4.4	-2.67	0.008
Log(Employees)	5.2847	5.2905	-0.4	-0.24	0.808
Equipment Vintage	.64691	.64471	1.2	0.72	0.473
Log(R&D)	6.1767	6.2185	-1.8	-1.10	0.271
Employee Training	.01099	.01165	-1.1	-1.22	0.221
Log(Age)	1.9825	1.9722	1.4	0.93	0.353

Please see descriptions of each variable at the beginning of this section of the appendix

Table II.A.3: Balancing Test for Propensity Score Variables: Transition Test I

Predictors of Started Exporting					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.5281	1.5178	0.8	0.24	0.810
Output/Employees	5.925	5.9263	-0.1	-0.04	0.972
Foreign Share of Ownership	.13155	.13907	-3.0	-0.67	0.505
State-Owned Dummy	.09321	.09587	-0.9	-0.25	0.803
Log(Total Assets)	11.622	11.607	1.0	0.26	0.798
Log(Employees)	5.7027	5.6865	1.3	0.36	0.721
Equipment Vintage	.6641	.66567	-0.8	-0.24	0.811
Employee Training	.0156	.01577	-0.3	-0.10	0.918
Log(Age)	2.2718	2.2646	0.9	0.26	0.798

Predictors of Stayed Exporting					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.5474	1.579	-2.7	-2.25	0.024
Output/Employees	5.9308	6.0291	-10.2	-7.75	0.000
Foreign Share of Ownership	.17727	.20722	-10.0	-6.44	0.000
State-Owned Dummy	.10083	.10293	-0.7	-0.55	0.584
Log(Total Assets)	11.995	12.026	-1.9	-1.45	0.146
Log(Employees)	6.1769	6.1311	3.6	2.76	0.006
Equipment Vintage	.62876	.62543	1.8	1.50	0.135
Employee Training	.01259	.01286	-0.7	-0.74	0.457
Log(Age)	2.4545	2.3972	7.5	6.07	0.000

Predictors of Stopped Exporting					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.5485	1.5306	1.4	0.35	0.729
Output/Employees	5.8648	5.8697	-0.5	-0.12	0.908
Foreign Share of Ownership	.0759	.08364	-3.6	-0.75	0.453
State-Owned Dummy	.1368	.12554	3.3	0.80	0.423
Log(Total Assets)	11.375	11.417	-2.7	-0.64	0.519
Log(Employees)	5.5437	5.564	-1.7	-0.41	0.683
Equipment Vintage	.656	.65473	0.7	0.16	0.870
Employee Training	.01713	.01743	-0.6	-0.17	0.866
Log(Age)	2.3614	2.3585	0.4	0.09	0.928

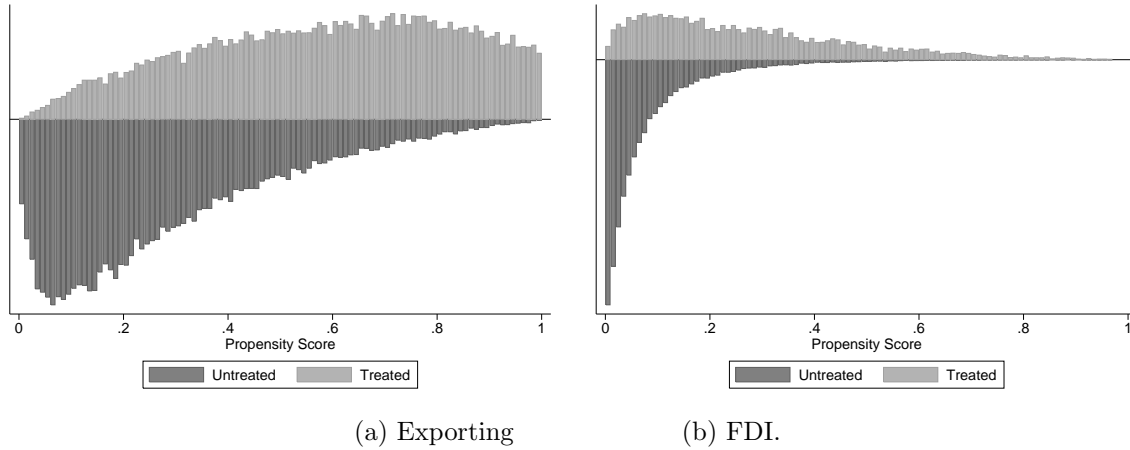
Table II.A.4: Balancing Test for Propensity Score Variables: Transition Test II

Predictors of Started FDI					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.6976	1.6848	1.1	0.15	0.877
Output/Employees	6.1857	6.2	-1.6	-0.22	0.826
Export Share of Sales	.343	.3436	-0.2	-0.02	0.982
State-Owned Dummy	.07126	.08551	-4.4	-0.77	0.442
Log(Total Assets)	11.923	11.895	1.8	0.26	0.794
Log(Employees)	5.8305	5.8129	1.3	0.20	0.842
Equipment Vintage	.64612	.63734	4.8	0.70	0.481
Employee Training	.01038	.00969	1.6	0.61	0.544
Log(Age)	2.1074	2.175	-9.8	-1.60	0.110

Predictors of Stayed FDI					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.6572	1.7246	-5.8	-2.06	0.040
Output/Employees	6.3125	6.4117	-10.1	-3.09	0.002
Export Share of Sales	.37451	.37824	-0.9	-0.31	0.755
State-Owned Dummy	.08289	.0575	8.5	3.21	0.001
Log(Total Assets)	11.947	12.094	-10.2	-3.26	0.001
Log(Employees)	5.7729	5.8381	-5.1	-1.68	0.093
Equipment Vintage	.60634	.60716	-0.4	-0.15	0.882
Employee Training	.01157	.00977	4.5	1.49	0.135
Log(Age)	2.1783	2.1604	3.3	1.06	0.288

Predictors of Stopped FDI					
Variable	Means			t-test	
	Treated	Control	%bias	t	$p > t $
Output/Total Assets	1.63	1.6816	-4.3	-0.55	0.581
Output/Employees	5.9498	5.9698	-2.2	-0.28	0.783
Export Share of Sales	.33614	.38017	-12.6	-1.44	0.152
State-Owned Dummy	.11799	.14749	-8.4	-1.13	0.258
Log(Total Assets)	11.711	11.73	-1.2	-0.16	0.876
Log(Employees)	5.8332	5.8516	-1.4	-0.18	0.856
Equipment Vintage	.61518	.60384	6.0	0.77	0.444
Employee Training	.01034	.01078	-0.9	-0.22	0.825
Log(Age)	2.1231	2.1793	-8.4	-1.13	0.257

Figure II.A.1: Graphing Covariate Match Quality



(FDI also leads to increased product innovation in the first year, though not on the same scale). Each ‘treatment condition’ is tested separately. As with Table 2.9, the control groups were selected to match each treatment: Observations with the *Start* treatment were matched to others who were similarly not exporters in the previous period. *Stay* was matched against new exporters and those that had stopped exporting, while those with the *Stop* treatment were compared with firms that had no exporting history.

The results show, as expected, that new products as a share of output increase by an average of 3.8% for firms in their first year of exporting, compared to firms with similar observable traits. Firms that remain exporters undertake more product innovation than those that are new or that stopped exporting. On the other hand, firms that stopped exporting undertake slightly more product innovation than non-exporters, but this difference is not statistically significant.

The same effect is not obvious for foreign investment. There is no statistically significant difference between firms in their first year of foreign ownership and comparable peers with matched propensities for FDI. On the contrary, new products as a share of output is less for firms with more than one year of foreign ownership, compared with other firms that are newly or formerly foreign-owned.

II.C More Tests with FDI Defined to Include HMT

Table II.C.1 shows that the coefficients in Table 2.9 of section 2.3.6 should remain largely unchanged if foreign capital was redefined to include funds from Hong Kong, Macau and

Table II.B.1: Matching Estimates by Stage of Export Participation

VARIABLES	Product Innovation
Started_Exports	0.038*** (0.002)
Stayed_Exports	0.010*** (0.002)
Stopped_Exports	0.001 (0.002)
Started_FDI	0.0003 (0.003)
Stayed_FDI	-0.024*** (0.005)
Stopped_FDI	-0.009*** (0.003)

Standard errors in parentheses
 *** p<0.01, ** p<0.05, * p<0.1

Note: The reported effects are the estimated average treatment effects on treated observations (ATT). For these exercises, the counterfactual for each row was limited to comparable firm-years as follows: *Started_** was matched to observations not foreign-owned or an exporter, and *Stayed_** to observations with a history of exporting or foreign ownership, but currently not in a second consecutive year in that status. *Stopped_** was matched to either non-exporters or firms with no foreign ownership in that year. The dependent variables is product innovation – new products’ share of the firm’s output. The number of treated observations were 15,724, 110,260 and 17,100 respectively for Columns 1 of Exports. Columns 2 of that segment had 6,757, 35,394 and 5,636. The numbers vary by column because the match was limited to items on the common support. The matching variables include firm size, output per assets and employee, as well as 4-digit industries. Further detail on the mean outcomes for treated and untreated items, the control items on common support and balancing tests for the matching variables are available on request.

Taiwan. In showing how firms change their investments in R&D and fixed assets as they become foreign-owned, keep that status or leave it, the estimates remain remarkably consistent with those of Table 2.9. Columns 1 and 2 reflect the values in Table 2.9 while the last two columns use the new definition of foreign ownership. Note that there are twice as many observations that are foreign-owned by this new definition, compared to the old. That is, firms in the first year of majority foreign ownership by this definition invest less in R&D than comparable Chinese-owned firms. One reason Chinese firms that become foreign-owned may reduce their R&D efforts is to avoid duplication of efforts by the foreign parent. It appears that foreign owners from Hong Kong, Taiwan and Macau are no more inclined to keep R&D in mainland China than other foreign owners. Similarly, firms that start majority foreign-ownership out-invest their peers in terms of fixed production capital. The estimates are not statistically significant, just like in Table 2.9.

For the firms that remain or stop being majority foreign-owned, the estimates in Table II.C.1 are remarkably similar to estimates with the original foreign ownership definition. The one exception is that firms that leave the foreign-owned status in the new definition under-invest in assets relative to comparable foreign-owned entities and the difference is statistically significant. In other words, the new definition does not help the argument in the literature that foreign ownership promotes innovation and the asset investments associated with product innovation.

In sum, Table II.C.1 suggests that the conclusions of this paper are robust to many definitions of foreign capital.

Table II.C.1: Comparing coefficients for FDI with and without HMT

FDI definition: VARIABLES	(without HMT)		(with HMT)	
	R&D	Asset Investments	R&D	Asset Investments
Started_FDI	-210.047*** (93.668)	-1168.146 (2087.223)	-171.522* (97.164)	-4018.628 (2673.713)
Stayed_FDI	77.412 (181.583)	4999.749 (4548.513)	265.349 (99.501)	1070.389 (2890.818)
Stopped_FDI	-113.867 (130.826)	-1355.626 (1380.328)	-391.639 (149.841)	-7710.459*** (2599.643)

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Note: HMT stands for Hong Kong, Macau and Taiwan. The estimates in Columns 3 and 4 use foreign ownership definitions that include capital from these sources outside the Chinese mainland. Columns 1 and 2 replicate the results in Table 2.9. The propensity score matching approach also follows the pattern in that table, the counterfactual for each row was limited to comparable firm-years as follows: *Started_FDI* was matched to observations not foreign-owned, *Stayed_FDI* to observations in the first year of foreign ownership or that was majority foreign-owned in the previous year, *Stopped_FDI* was matched to firms no foreign-owned in that year. The dependent variables are R&D and asset investments. The number of treated observations were 6,753, 66,902 and 5,635 respectively. The numbers vary by column because the match was limited to items on the common support.

CHAPTER III

Bridging the Enforcement Gap in International Trade: Explaining Participation in the New York Convention on Arbitration¹

Chapter Abstract

International trade transactions come with a peculiar hazard – there are no common courts or legal jurisdictions to enforce contracts if disputes arise, unlike commerce inside national borders. Distance, language and cultural disparities further complicate contract enforcement for international trade. A 1958 United Nations treaty, the New York Convention on the Recognition and Enforcement of Foreign Arbitral Awards (NYC) addresses this enforcement gap. NYC signatory countries commit to enforce the outcomes of private arbitration between importers and exporters. This makes private international arbitration a viable legal alternative, supported with state enforcement like national courts. Surprisingly, many countries waited decades to join the NYC despite the obvious benefits. This puzzle presents an opportunity to examine treaty membership as a policy adoption problem - how the costs and benefits of membership lead to delays in participation. The model of policy adoption in this paper makes it possible to compare the roles of trade interactions, social contagion or peer effects and internal barriers in the timing of membership decisions. Tests of the model's predictions indicate that countries are more likely to join when their regional trade partners do - trade by itself does not fully explain this pattern. I also estimate significant effects for internal barriers to policy adoption, using systems of government and legal origins as proxy. Democracies on average join earlier, and countries with socialist legal origins have the shortest average delays to joining the treaty.

¹Versions of this chapter have been circulated as coauthored with Prof. Scott Masten. The chapter started as a research project under his supervision: the original idea arose from our conversations and the project has progressed with his guidance. He is not listed as a co-author here at his own suggestion.

3.1 Introduction

International trade transactions come with a peculiar hazard – there are no common courts or legal jurisdictions to enforce contracts if disputes arise, unlike commerce inside national borders. Distance, language and cultural disparities further complicate the challenges of contracting for international trade. Even if the parties agree to arbitration or the courts of one party’s country, the problem of enforcement remains; it may be difficult or impossible to secure performance or damages if the losing party resides outside the controlling jurisdiction. In sum, the legal options present a dilemma: Dim prospects of recovery for awards made in domestic courts that the counter-party may not recognize, versus the potential for biased judgments in the courts of the partner’s host country.

To bridge this enforcement gap, a 1958 United Nations conference produced the Convention on the Recognition and Enforcement of Foreign Arbitral Awards, known as the New York Convention (NYC). Countries that sign the NYC treaty agree to enforce the judgments issued by arbitrators to resolve private commercial disputes arising from international trade. The benefits to signatories are several; businesses get relatively fast and high-quality adjudication from arbitrators of their choice and governments have fewer cases congesting the courts. The convention also provides the opportunity to send a credible signal that a national government supports honest trade. Surprisingly, many countries waited decades to join the treaty, despite the obvious benefits. Large trading economies like the U.S. and the U.K. joined in 1970 and 1975, more than a decade after the first members. Treaty membership was only 51 out 155 possible countries in 1978, two decades after the initial signatures. By comparison, in the first two decades of GATT, the World Trade Organization’s predecessor founded in 1948, 75 countries became members, despite GATT’s more stringent entry requirements. (Figure III.B.2 in the appendix illustrates the comparison of WTO and NYC membership).

This paper asks why countries delayed entry into the treaty, given the apparent benefits of NYC membership. The exploration follows a model of treaty membership that recognizes multiple sources of incentives for countries to participate. The model is rich enough to embody many of the traits of seminal models of product adoption like Bass (1969) and models of influence or diffusion in the microeconomics and policy literature (Golub and Jackson, 2010; Young, 2009; Shipan and Volden, 2008; Simmons and Elkins, 2004).²

²The rich body of work on peer effects, contagion and diffusion spans many fields – from economics to political science and beyond. Space constrains this paper to citing only a limited number.

The paper has two main objectives: showing that treaty participation is endogenous, and highlighting the role of peer effects. In showing that participation in the NYC is endogenous, the paper enables better estimates of the effects of participation on trade. I show that participation and the timing of membership are not random – they are related to countries’ trade relationships, legal origins and the nature of traded goods. This complements papers that measure the impact of NYC membership on trade. (See Berkowitz et al. (2006), for example). This study may also help other efforts to understand the impact of membership for other treaties, e.g. the WTO, although differences in the WTO’s entry procedures call for additional model features.³


In the context of international governance, identifying peer effects is important for many reasons. If one considers peer effects, initiators of international treaties may recognize that the success of a treaty depend as much on whether influential countries are the first to join as on the benefits of the treaty to the average country. Furthermore, governments adopt policies with significant economic impact like smoking bans or motorcycle helmet laws, largely in response to their neighbors’ actions Lee and Lee (2011); Bramoullé et al. (2009); Lee (2007). This chapter aims to describe the drivers of membership in a trade treaty with the legal significance of the NYC, where peer effects are a possibility. (I use contagion and peer effects interchangeably in this context; the first term is less common in the economics literature except for papers on financial crisis, but is preferred by related papers like Young (2009) and Simmons and Elkins (2004).)

Specifically, I ask whether peer effects play a role in countries’ decision to join the NYC. The data suggests that the answer is affirmative. Following the basic structure in Young (2009), contagion produces participation patterns shaped like an S-curve. The intuition behind the structure is straightforward: in a model of contagion or peer effects, the hazard of participation is proportional to interactions between members and non-members. That hazard is lowest when most countries fall in one category, and it peaks when members and non-members are in equal proportion.

In Figure 3.1, a relatively featureless model of contagion explains 90% of the variation in total NYC membership over time. Participation follows the S-curve predicted by a model of contagion. The model behind Figure 3.1 implies that a country joins only because others have joined. That model identifies no specific incentives faced by potential treaty members,

³Membership in the NYC requires only a signature by an eligible country. (See discussion of eligibility section 3.2.1). GATT/WTO membership on the other hand, requires all WTO members to agree on the terms of each new member’s accession through negotiations.

Figure 3.1: NYC Members over Time: Model of Contagion vs. Data



The plot shows actual NYC membership in year t as circles and model predictions as hollow squares. The model uses only the number of eligible countries N , current members n and time t to predict changes to membership. The predictions represent the best-fit line from estimating: $\log\left(\frac{n_t}{N_t - n_t}\right) = \text{constant} + \gamma t + \varepsilon_t$. (I do not show N_t , which changed as new countries were formed, or former colonies gained political independence).

Data Source: UNCITRAL (2014)

nor does it describe costs or frictions that delay entry, if any. The model is agnostic with respect to the specific nature of interactions that promote membership and how interactions of any kind translate to treaty participation. Nevertheless, the level of fit between this basic model of contagion and the data is noteworthy.

Interpreting the level of fit in Figure 3.1 depends on how one defines interactions between treaty members, or whether interactions contribute to the hazard of joining the treaty. First, the skeptical position that interactions do not contribute: If it was true that countries join at an exogenous rate that does not depend on interactions, the rate was constant over time, the number of countries is fixed, and all countries were equally susceptible, Figure 3.1 would be a linear plot. If membership was strictly exogenous but countries were not equally susceptible, the curve would be consistently concave. The most susceptible countries would be first to join and the slope of the curve would be decreasing (or at least non-increasing). Neither of these scenarios involve a specific explanation for an exogenous constant driver of participation, if any exists. In either case, if an exogenous membership process was part of a mix of factors driving membership, the linear or concave plot they would generate should contribute to the degree of fit that one observes in the figure.

Second, if treaty participation was only due to interactions between countries, Figure 3.1 offers a narrow set of assumptions about how those interactions cause countries to join the treaty. The model behind the figure does not identify whether the interactions responsible for treaty participation is having shared borders, a common language or the simple act of observing other countries' membership status. Interactions are assumed equally likely between any two parties, and all interactions contribute equally to the likely of a non-member joining the treaty. Justifying these assumptions is difficult – bilateral relations between government can hardly be described as uniform.


Figure 3.2 extends the basic model of contagion behind the previous figure. It shows the predictions of a model in which treaty members' share of trade drives the likelihood of non-members joining the treaty. The regressions underlying the figure suggest that NYC members' share of trade explains more than 95% of changes in membership over time. (Section 3.3 describes this model). While interactions between countries drive membership in both models, trade determines the weight given to interactions in Figure 3.2. By this reasoning, countries with no NYC trade partners have zero incentive to join. In Figure 3.1, the share of a country's trade with NYC members is irrelevant to the treaty-joining decision, only their numbers matter. The model in Figure 3.2 looks less like an S-curve because trade between NYC members and non-members should be constantly increasing – trade between any pair of transactors is equally likely in the model.

The model represented in Figure 3.2 allows for the possibility that countries that join the treaty because it lowers trade costs or provides benefits in the nature of a network good. One could crudely describe the benefit of NYC membership in this model as free contract enforcement insurance for firms in member countries. Any non-member that trades with NYC members has an incentive to join for the 'free insurance' that it represents to its trading firms. If that country joins, all its trade partners that are not NYC members in turn have an incentive to join. However, there is no logical reason why narrowly defined economic incentives must be the only driver of treaty membership.⁴

The two figures make the case for a hybrid model of treaty participation. Earlier models of decision-making with more than one driving mechanism include the relatively successful Bass (1969) model of product adoption. That model proposed that purchases of durable

⁴There is a testable difference between membership motivated by 'free contract insurance' on the one hand and social learning, contagion or imitation on the other hand. Membership motivated by the insurance or cost incentive will not respond to changes in the membership of countries that are not active trade partners of a potential members. Membership driven by interactions reflects the influence of other countries, even if they are marginally trade partners or not trade partners at all.

Figure 3.2: NYC Members over Time: Network Good Model vs. Data



The plot shows the number of NYC members in each year t as circles. The squares show the prediction of a model that uses only the number of eligible countries N , NYC members' share of global trade, and the number of current members n to predict changes to membership. The predictions represent the best-fit line from estimating $\log\left(\frac{n_t}{N_t - n_t}\right) = \lambda(\text{NYC} \text{tradeshare}_t * t) + \varepsilon_t$. Trade data is only available up to 2006. Section 3.3 explains the assumptions and estimates further.

Data Source: UNCITRAL (2014), Head et al. (2010)

goods reflect an *innovation rate* due to product advertising and a *diffusion rate* due to interactions between current and potential product users. This paper's hybrid model differs in how it proposes richer interactions between decision makers, given that trade interactions are primarily between businesses, not the governments that ultimately make the membership decision. It also recognizes that trade is not the only reason behind interactions between national governments, even if the NYC treaty is limited to trade.⁵

In exploring how countries adoption policies, treaties or innovations, this paper considers three approaches: [1] Countries join to allow their firms the economic benefits of trade with

⁵ Aral et al. (2013) provides another model in this hybrid category, one that emphasizes peer influence, while allowing higher levels of influence for peers with more similar traits. Papers like Young (2009) and Shipan and Volden (2008) emphasize the diffusion mechanism, based on mechanisms that include learning and social contagion.

I also considered the possibility that countries join the NYC as a requirement for joining other groups, motivated by interactions of the coercive sort described in Shipan and Volden (2008). A simple test of whether NYC participation is linked to WTO membership does not support this argument. (The WTO is the only better-known trade agreement with a comparable number of signatories). A correlation of entry sequence rank for the two agreements yields a coefficient of 32% for the 128 countries that signed up to both organizations. Rank is based on GATT dates for countries that joined before 1995. It is not clear that countries join the NYC because they need to join the WTO/GATT.

partners from NYC member countries. [2] Countries join because of social contagion or interactions with NYC members outside of trade between firms. [3] Countries join when the benefits of participation sufficiently exceed the political costs or internal barriers to participation. The costs of political action should be related to system of government and the legal system. These broad responses can be explored in detail while taking advantage of the remarkable fit between the aggregate membership trend and models driven by interactions between countries.

Other non-trade factors can also be identified or included in the estimation framework that flows from these answers. For example, considering political costs or internal barriers has two benefits: it recognizes that country-specific propensities may vary, and it addresses the skeptical argument that interactions play no part in membership from earlier in this section. If interactions played no part in membership, then internal barriers to participation should explain all the variation in the data. In other words, with a model this flexible, an observer has more room for trying to learn what best explains the data.

The empirical section of this paper uses a hazard model to estimate the relative contribution of the aforementioned drivers to the timing of countries' decision to join the NYC. The main findings from those exercises are:

- Countries are more likely to join the treaty when their regional trade partners are members. The findings are consistent with the predictions of a model of peer influence.⁶
- Surprisingly, the share of trade with NYC members by itself is not a significant driver of membership after controlling for the NYC membership of RTA partners.
- Country-specific traits influence the hazard of treaty participation.
 - Countries with differentiated goods as a greater share of trade join sooner
 - There is weak evidence that democracies join sooner on average,
 - Countries of socialist legal origin join the treaty sooner, on average.

⁶The argument for peer influence comes with the caveat that further work is required on how countries' policy choices depend on those of the parent-countries or colonial metropolises. Though more than 190 countries were eligible to join the treaty in 2014, only 87 had that privilege in 1958 - most had not even been founded. The argument for peer influence in this paper rests largely on an empirical model that does not consider the NYC status of parent-countries like Yugoslavia on the choices of countries like Serbia that emerged from its split, nor on the NYC status of colonial metropolises - in part because only four countries account for most colonial relationships. Table III.A.4 clearly suggests further research into how the NYC status of 'parent-countries' influence the choices of the countries they spawned.

The findings are relevant to several questions in the literature. First, Masten and Prüfer (2014) show that formal adjudication is a substitute for reputation-based enforcement, especially when trade is expanding. This paper provides an opportunity to take this hypothesis to the data. Specifically, one could test the hypothesis that the propensity to join the treaty is higher for countries with higher shares of trade with new or distant partners, or countries with trade situations that are generally contract-intensive.

The model in the paper also contributes to the product/policy adoption literature. Its hybrid nature includes elements of the Bass (1969) model of product adoption – choices are private even if the treaty is a network good because every country that joins increases the value of participation for current and potential members. However, it is flexible enough to recognize that countries may join for reasons other than taking trade contracts beyond the bounds of reputational enforcement. In that way, it complements the papers on contagion or policy diffusion (Aral et al., 2013; Young, 2009; Shipan and Volden, 2008).

The next section presents related literature and some background on the treaty. It is followed by sections that model treaty participation and empirics that test the model’s predictions.

3.2 Background and Related Literature

3.2.1 Background on the New York Convention

It is surprising that many countries took decades to join the NYC.⁷ There are no real barriers to entry, unlike the WTO where current members may hold up would-be joiners in negotiations (Bagwell and Staiger, 2011). (See additional notes on the WTO in Section 3.2.2). Any eligible state may ratify or accede to the NYC, usually through an act of Senate or Parliament. A geographic unit may become eligible for NYC membership at formation, independence or when it otherwise obtains treaty-signing powers. The responsible UN secretariat notifies current members when new members join.⁸

Membership in the NYC directly addresses the issue of legal jurisdiction for international trade contracts. Signatory countries agree to enforce awards made by arbitrators against the party to a contract dispute that belongs in their jurisdiction. Trading parties only need to

⁷Ultimately, the NYC became one of the most successful international treaties, with 140 members by December 2010. On membership, it is one of the top seven treaties of the 1313 multilateral agreements in the (The treaty database shows memberships up to the year 2000; at that time the NYC had 124 signatories).

⁸Appendix Section III.A.1 in the appendix explains membership eligibility as defined by the NYC charter. It also provides data sources for countries’ membership and eligibility dates.

agree on arbitration as the dispute resolution mechanism.⁹ The threat of state enforcement of arbitral outcomes represents an implicit guarantee that is usually sufficient for effective contracting between international trade partners based in signatory countries.¹⁰

The NYC treaty effectively allows the use of formal arbitration in place of self-enforcement based on reputation. Self-enforcement based on reputation is the default mode for international trade in the absence of court-enforceable contracts. (The history of international trade expansion on the basis of reputation or self-enforcement extends to the Law Merchant or *Lex Mercatoria* in Western Europe's Middle Ages (Masten and Prüfer, 2014), and to Maghrebi traders in North Africa and the Middle East (Goldberg, 2005)). However, the ability of reputation and trader networks to provide security relative to a formal court system is likely to decrease with increasing geographic and social distance. The scope of trade under reputation is limited, and with it, the realization of scale and specialization economies.¹¹

Independent arbitration panels offer the formality of a court, faster average case resolution times, and uniform legal procedures that generally do not depend on a contracting party's country. Effectively, they reduce favoritism and procedural disparities in the resolution of contract disputes. Several scholars claim that most international trade contracts refer possible disputes to arbitration.¹² Nevertheless, international arbitration by itself does not address two critical issues: how to force a recalcitrant party to participate in arbitration,

⁹ Quoting the United Nations body that oversees the Convention:

The Convention's principal aim is that foreign and non-domestic arbitral awards will not be discriminated against and it obliges Parties to ensure such awards are recognized and generally capable of enforcement in their jurisdiction in the same way as domestic awards. An ancillary aim of the Convention is to require courts of Parties to give full effect to arbitration agreements by requiring courts to deny the parties access to court in contravention of their agreement to refer the matter to an arbitral tribunal.

The treaty allows specific limited exceptions: e.g. contravention of public policy in the jurisdiction of enforcement. Governments generally have no sustained incentive to protect private citizens that violated an international contract.

¹⁰ The NYC applies strictly to disputes between private parties. It is different from the International Centre for Settlement of Investment Disputes (ICSID), which administers disputes between private parties and governments.

¹¹ The alternative formal arrangement for contract enforcement is reliance on the domestic courts of one party to a trade contract, but that puts the other party at a disadvantage in terms of distance, familiarity with the legal system, and potential local bias. It should be clear that participation in the NYC by a government only gives firms the option of dispute resolution through private arbitration, no imposition is made on importers and exporters.

¹² Leeson (2008) claims that more than 70% of contracts have an arbitration clause, citing a survey by Casella (1996). The arbitration fora to which these contracts refer disputes are easily accessible. They are in more than forty-seven countries, with the most popular being the International Chamber of Commerce (Paris) and the London Court of International Arbitration (Chartered Institute of Arbitrators, 2011).

and how to enforce arbitral decisions against a losing party.

The foregoing shows why the NYC addresses the enforcement gap, a gap that becomes more important as international trade expands beyond firms' social and business circles. (The same principle applies if a country changed its trade portfolio from a standard commodity like wheat to differentiated goods like automobiles). Article II of the Convention requires courts in member states to recognize and enforce agreements to submit disputes to arbitration, and Article III obliges courts to recognize and enforce awards resulting from arbitrations occurring in other states. In principle, the New York Convention provides the necessary institutional support to secure international bargains. This favors businesses looking to expand the scale or scope of imports and exports, especially when trade is near the limits of self-enforcement. (Masten and Prüfer, 2014; Dixit, 2003).

The treaty generally finds support with the private sector. Appendix Section III.B.4 shows that at least in the US, the private sector lobbied for the treaty. In a sense, the treaty's offers *contract enforcement insurance* for all firms that agree to use arbitration in international trade disputes. The cost of arbitration is the analog of the deductible expense for firms. Countries act as insurers for their resident firms' trade partners. This analogy illustrates how this treaty favors trading firms. National governments should also favor the treaty, given how more private sector trade could create higher government revenue.

Nevertheless, arbitration under the auspices of the NYC costs signatory countries a modicum of sovereignty. By signing the treaty, a country agrees to enforce agreements and awards made possibly outside its jurisdiction, using its own legal resources. The willingness of countries to do this may vary, given the different systems of laws that exist, the frictions that exist within different systems of government and the attitudes of governments to international cooperation.

20 governments signed the treaty in June 1958 – the NYC came into force for those first signatories in June 1959. These founding members were a diverse array of countries that included Germany, France, India and the Philippines. Appendix Table III.A.1 shows membership dates by country, it also shows the dates countries first became eligible so that one can measure the delay to membership. Membership in the treaty increased to 140 in 2010 (144 in 2014).

3.2.2 Related Literature

3.2.2.1 Modelling Treaty Participation (Policy or Product Adoption)

The literature on the adoption of policies and products goes back more than a half-century. These include studies of how Iowa farmers adopted hybrid corn seeds (Ryan and Gross, 1943), to more recent papers in economics that consider herd behavior in financial decision making (Banerjee, 1992) or households' decision to use personal computers (Goolsbee and Klenow, 2002). Earlier in the business literature, Bass (1969) described the purchase of durable goods with an empirically successful imitation-based model.¹³

Much has been published on the subject in the political science literature. The studies extend from cities' adoption of smoking bans (Shipan and Volden, 2008), through states' decision to run lotteries (Berry and Berry, 1990), to the international adoption of financial liberalization policies (Simmons and Elkins, 2004).

On policies related to international trade, a few papers have studied factors like past colonial relationships that facilitate or obstruct WTO membership (Copelovitch and Ohls, 2012; Cattaneo and Braga, 2008). Wagner (2010) provides a robust model of accession to the Montreal Protocol, which emphasizes spillover effects of early joiners on other countries. Others like Baldwin and Jaimovich (2012) suggest the possibility of contagion in countries participation in regional trade agreements.¹⁴

This paper builds on these works by considering treaty adoption as a special case. Examining contagion in the context of international relations provides the opportunity to examine

¹³Exploring the difference between rational learning and 'unsophisticated imitation' holds great research potential, but that is not the goal of this paper. The chapter attempts to stay close to the script in Masten and Prüfer (2014) in addressing the factors that motivate countries' adoption of trade policies like this treaty. Whether the potential drivers of peer influence are rational or not is not explored.

¹⁴Given the WTO's popularity, one must make an effort to distinguish the NYC from the WTO. The WTO is about reducing trade costs in the form of tariff and non-tariff barriers, while the NYC focuses on facilitating arms-length trade by making legal recourse enforceable. The WTO always had more members than the NYC, in part because it was initiated ten years earlier. (Figure III.B.1 plots membership over time for the two treaties). I am not aware of a paper with the explicit objective of evaluating factors that drive membership in the WTO, like this paper does for the NYC. As mentioned earlier in this paper, modelling membership in the WTO requires consideration of current members' holding up would-be members. Modelling NYC participation gets at the question of peer effects without this wrinkle. Nevertheless, such an exercise would be interesting given the volume of work debating whether GATT/WTO membership increases trade, (notable examples in this literature include Tomz et al. (2007); Subramanian and Wei (2007); Rose (2004)).

Furthermore, the dispute resolution mechanism of the WTO should not be confused with that of the NYC. The NYC is strictly for contractual disputes between private parties, while the WTO is for disputes between national governments. (Even if a WTO dispute is initiated by a government in the interest of a particular firm, the dispute is usually about trade policy and not about a particular contract between two private firms).

how countries influence one another, and whether the role of international institutions in economic processes like trade.

3.2.2.2 Estimation Model Choice

The literature on contagion in the context of policy or production adoption has always leaned on research describing biological diffusion or the spread of diseases (Aral et al., 2013; Young, 2009; Jackson and Yariv, 2007). Young (2009) describes this linkage in building a formal model of contagion from the basic elements of epidemiology’s susceptible-infected-recovered (SIR) framework.¹⁵ Problems of this form lend themselves to estimation using logit models for the simplest cases, or duration models more generally.

This paper follows the majority of the related literature in specifying duration models (also known as hazard models). A duration model helps to address the concern that the data is right-censored. It is not sufficient to know the membership status of countries in 2014. Estimates that ignore the lag to participation will err in comparing countries like Venezuela and Kazakhstan – both joined in 1995, but Venezuela had been eligible since the treaty’s founding, while Kazakhstan did not become a country until the nineties. Other papers that make a case for applying duration models to similar research questions include Simmons and Elkins (2004) and Bikhchandani et al. (1992).

In considering a duration model, I also recognize that several variables that motivate membership are time-varying; the number of NYC members being the primary example. Several papers discuss estimating parameters for the class of duration models with time-varying factors, (see Simmons and Elkins (2004) and Jenkins (1995), for example). Furthermore, interactions between countries may not be uniform, such that the incentive created for country B by country A’s participation depends on the characteristics of the A-B country-pair. These considerations guided the choice of an empirical model.

The next section presents a formal, if simple, model of countries adopting a trade treaty like the NYC.

3.3 Model

Explaining NYC membership decisions is challenging, not least because it implicates organizations and governance at multiple levels: businesses involved in international trade,

¹⁵Carroll (2001) provides another notable usage of the SIR model in economics, using it to represent the diffusion of macroeconomic expectations.

private arbitration tribunals, state courts, international organizations (the UN), and for ratifying treaties, national governments. I focus on two of those levels: businesses and governments.

Consider a set of N countries, each populated by F_i firms, ($i \in 1, \dots, N$). At time t , n_t countries belong to a treaty that benefits only firms in member-countries. For each non-member country, the potential benefit of joining is proportional to the sum of each of its resident firms' private benefits from joining the treaty. For the firms, treaty benefits represent the net cost of an insurance policy that guarantees the execution of contracts with trade partners in NYC member countries. (If one calls NYC membership free contract enforcement insurance for firms, the expected cost of appearing before arbitration is the deductible expense). National governments that are responsive to the private sector's interests will want to join, as long as there are gains from trade and trade with NYC members is significant. The hazard of joining will be proportional to the value of NYC's benefits to each country's importers and exporters.

National governments may also want to join the treaty in response to peer effects or contagion. Like individuals, governments may be susceptible to peer influence, so that treaty adoption behaves like contagion. (The mechanism mimics the workings of a susceptible-infected (SI) model of infectious disease transmission). The rationale behind conforming to a norm could be the intangible cost of being a 'pariah' state, or the desire to be one of the crowd, in order to gain the social influence that comes with being part of a trend (Aral et al., 2013; Golub and Jackson, 2010; Young, 2009).

The timing of entry depends on the combination of these incentives, as well as the responsiveness of the decision-makers to the incentives. One could also describe responsiveness in terms of barriers. Governments have competing priorities – one expects high priority items will take less time and low priority items to encounter long delays. In the language of disease transmission, if the barrier to transmission is low, then most countries will join quickly. Some countries may have a high barrier before taking action on treaty membership, so that the decision is delayed, despite the benefits conforming to a global norm or the economic benefits of contract enforcement guarantees to firms.

The next subsections explain these mechanisms further.

3.3.1 Contagion

The data argues for treaty adoption through contagion. The participation curve visually resembles an S-curve, which is typical of systems where changes in status are driven by in-

teractions between non-members and members in a given population. The hazard is initially low when the proportion of members is low, increases to a maximum when the relative proportions of members and non-members are equal, and decreases again as the proportion of non-members falls.¹⁶

Contagion can be explained as the outcome of a polling process for countries that are either members or non-members of a treaty. (Say there are n_t members out of N countries at time t). Hypothetically, each country polls every other country on their membership status in each period. It decides to join with a non-negative probability proportional to the fraction of members in the population $P(t) = n_t/N$. The fraction of countries joining in each period or the hazard rate is initially low because few non-members meet members in the early rounds of polling, given that there are few members. Eventually, the hazard diminishes as non-members become few.

Formally:

$$\frac{dP}{dt} = \gamma P(1 - P) \tag{3.1}$$

All countries are equally influential in the diffusion process, and mixing between countries is perfect. Each interaction between a member and a non-member increases the probability of the non-member switching status. γ represents that marginal contribution to the probability of joining the treaty.¹⁷

¹⁶Appendix Section III.D presents a formal test that the data plot in Figure 3.1 is indeed S-shaped.

¹⁷ This model is equivalent to the basic epidemiological model of contagion. As countries never revoke the treaty, the susceptible-infected-recovered (SIR) model simplifies to an SI model.

The trivial baseline for the model is one where membership is not driven by interactions, but occurs stochastically at some rate γ . Young (2009) motivates his paper with this baseline, which helps in defining γ later in this paper. It also sets up a notable feature of this model i.e. membership decisions are not guaranteed, but occur with a non-negative probability. In the baseline, all countries join eventually. If γ is heterogeneous, the more responsive countries with high γ s should have shorter average entry delays.

Equation (3.1) also mirrors the form presented by the Bass (1969) model of product adoption, $\frac{dP}{dt} = \gamma[P(1 - P)] + \hat{p}(1 - P)$, i.e. if one sets the innovation rate \hat{p} to zero. That is, the Bass Model is an SI model with an additional term that depends strictly on the uninfected fraction of the population.

Solving (3.1) yields:

$$\ln\left(\frac{P}{1-P}\right) = \gamma t + A \quad (3.2)$$

where A is the odds ratio at time $t = 0$,

$$A = \ln[P_0/(1 - P_0)].$$

$$P = \frac{Ae^{\gamma t}}{1 + Ae^{\gamma t}} \quad (3.3)$$

Figure 3.1 represents equation (3.3), which clearly follows a logit form. Section 3.4 shows that this naive model of contagion with a constant γ explains about 90% of the variation in treaty membership over time. To continue the polling analogy, $P(1 - P)$ in equation (3.1) is the share of poll questions asked across the membership divide.¹⁸

The γ parameter may be country-specific, such that some countries have high γ_i and are more responsive to interactions with NYC members. National attributes that may determine the responsiveness of government, or the propensity to join the treaty may include the system of government and trade as a share of the national economy. For example, a country in constant political gridlock, or dictators wary of commitments to the international community will generally be less responsive to participation incentives like interactions with NYC members. In general, those countries will take longer to join the treaty. With heterogeneous γ_i s, one must define country-specific participation hazards dP_{it}/dt :

$$\frac{dP_{it}}{dt} = \gamma_i P_t \quad (3.4)$$

In the aggregate:

$$\frac{dP_t}{dt} = \frac{1}{N} \sum_{i=1}^{N-n} \gamma_i P_t \quad (3.5)$$

With N countries and n members, so that $(1 - P) = (N - n)/N$, equation (3.5) reverts to (3.1) if γ_i was constant.

Furthermore, interactions between countries are not uniform. Countries belong to different regional agreements, commonwealths and other fora for negotiation that make it possible

¹⁸ N^2 questions are asked in each period, of which $NP * (N - NP)$ are between the $n = NP$ members and $N - n$ nonmembers.

to observe other countries' NYC status, or to relate on the basis of membership in the treaty. Regional trade agreements represent a notable forum for negotiating and observing trade policy between countries. That countries in these agreements belong to the same geographic region also means that members often share boundaries, language and other traits that could facilitate contagion. For example, it is more likely that Argentina would emulate Brazil, its MERCOSUR partner, rather than a distant country like China, even if the latter were large.

One cannot claim that equation (3.5) only represents contagion through imitation. The interactions in the equation could represent trade. Trade by itself captures interactions between private parties - importers and exporters, not necessarily governments. Yet, one expects that the level of trade affects government choices on policy, either directly because governments want the benefits of trade, or indirectly, because trade prompts governments to come together at forums like regional trade agreements to discuss policy.

The next paragraphs shift the focus to trade as the primary mechanism that promotes membership in the NYC treaty.

3.3.2 Treaty Membership Motivated by Trade

Treaty participation driven by trade will follow the form described by contagion, with weighted interactions. By definition, trade is a form of interaction between countries. Therefore, to estimate the trade benefits of NYC membership, one must recognize that those benefits derive from interactions, much like the hazard of participation in the polling analogy from the previous subsection.

Consider the N countries mentioned at the start of this section, each with F_i firms, $i \in [1, N]$. If w_{ij} is country j 's known share of country i 's trade, $\sum_{j \neq i} NYC_j * w_{ij} \tau$ is country i 's hazard of joining. $NYC_j = 1$ indicates that country j is a member and τ captures the value of NYC membership benefit per dollar of trade.

That hazard is not distinguishable from the polling analogy in equation (3.5). In the polling analogy, the trade share w_{ij} could be used as weights for each bilateral interaction. The weights will range from zero to one, so every country is polled, but the opinion of non-trade partners are completely discounted. In this scenario, the country-specific hazard of participation will be $\sum_{j \neq i} NYC_j * w_{ij} \gamma$, and the right hand side of equation (3.5) will be $\hat{\gamma} * \sum_{i=1}^N w_j \sum_{j \neq i} NYC_j * w_{ij}$. The modified marginal probability of joining is $\hat{\gamma}$ and w_j is country j 's share of global trade. Replacing τ with γ in the previous paragraph shows that the two approaches are equivalent.

The model can be extended to include contagion effects due to trade directly, as well as

other interactions. If only one of these two set of incentives is at work, then the estimation step should yield a zero parameter for the other.

Defining the hazard formally:

$$\frac{dP_{it}}{dt} = \gamma_i \left[\sum_{j \neq i} NYC_j (w_{ij} \tau + \kappa) \right] \quad (3.6)$$

κ is the relative contribution of factors not directly proportional to the level of trade.

Figure 3.2 represents equation (3.6). Setting non-trade weights $\kappa = 0$ and the enforcement gap parameter $\tau = 1$ yields equation (3.5).

The rest of the paper will focus on estimating γ , κ and τ . It is important to check whether κ is greater than zero – this addresses the question of whether state actors are subject to peer influence. This line of research was championed by Shipan and Volden (2008); Simmons and Elkins (2004).

The responsiveness of governments is captured by the γ term - with responsiveness to economic interests as a special case. The value of the term should be influenced by attributes of national attributes like legal origins or system of government. If for example, all democracies equally favor the treaty but parliamentary democracies are more efficient, then their average γ s should be higher. Similarly, if democracies are on average more *pro*-trade than monarchies, then monarchies will have lower γ s. (The political costs of treaty membership are also couched in this term, so if monarchies and democracies are equally *pro*-trade, but the political costs are less in monarchies, then democracies will have lower γ s).¹⁹ Accommodating systems of law to new requirements can be a complicated and time-consuming undertaking. Governments have competing demands on their time and will act on a proposal, even when no opposition exists, only when pressures reach a level sufficient to move the issue ahead of others.

The τ term represents the enforcement gap – the benefits of treaty membership, conditional on trade, or more precisely, the net benefits of formal adjudication over informal, reputation-based enforcement. One expects this term to be influenced by the nature of trade – trade with proximate longstanding partners needs courts or arbitration less than trade with new distant partners. Similarly, trade in undifferentiated goods like wheat on global exchanges is expected to require less contract adjudication than trade in differentiated goods like clothing and airplanes. This follows the arguments in Nunn (2007) and Berkowitz et al.

¹⁹III.B.4 explains this point further.

(2006) that some goods are more contract-intensive.

The next section takes equation (3.6) to the data, first in the aggregate, then later with country-specific variables.

3.4 Empirics

In describing how and when countries join the NYC, this paper will consider:

[1] Internal frictions to policy adoption

[2] The enforcement gap, or factors that make reputational enforcement less valuable than formal arbitration

[3] Interactions through trade and at other fora, e.g. regional trade agreements (RTAs).²⁰

There are three parts to this section: [1] I estimate equation (3.2) on the assumption that trade with NYC members does not contribute to countries' NYC decision (i.e. setting $\kappa = 1$ and $\tau = 0$ in equation (3.6)); this provides baseline estimates for the responsiveness of governments. [2] I allow non-zero values for τ , while forcing all countries to have the same value of γ and τ ; this allows claims about whether non-trade drivers of membership are statistically significant, and [3] I allow some country-level heterogeneity in the τ parameter.

First, I describe the data used for these estimates. The main variable of interest is the delay to membership. I define this as the number of years between when a country became eligible to join the treaty and when it actually joined –for the countries that joined before December 2010. (The empirics that do not require trade data use membership through 2014). A longer lag or delay implies a lower hazard of treaty participation. For the first set of estimates, I simply compare the number of member countries with eligible countries in each year following 1958. The UNCITRAL website provided membership dates. Appendix Section III.A.1 explains the variable in detail, and describes sources for eligibility dates. Table 3.1 presents a summarized version of the data in Figures 3.1 and 3.2, but includes the related trade volumes and trade features.

NYC members increased from less than a quarter of eligible countries and a third of global trade in 1958 to nearly three quarters of all countries in 2006. At that time, imports and exports from NYC members represented 99% of the global total. These numbers indicate that NYC members have always had larger than average trade shares, an early sign

²⁰ Figure III.B.1 in Appendix Section III.B.2 shows that the increase in the number of RTAs was largely contemporaneous with the growth of NYC membership. Comparing the efforts required to establish RTAs, relative to joining the NYC, which only requires a unilateral signature, it is reasonable to accept that countries do not form RTAs because of NYC membership.

Table 3.1: NYC Members over Time

YEAR	Eligible Countries	NYC Members	Total Trade ('000 USD)	NYC Trade Share	Rauch Diff'd Share	Colonial Trade Share
1958	81	20	204,593	0.32		0.175
1968	124	39	489,892	0.46	0.48	0.122
1978	149	55	2,720,496	0.79	0.50	0.083
1988	158	73	5,853,784	0.93	0.61	0.084
1998	186	119	1.18e+07	0.97	0.66	0.078
2006	187	139	2.58e+07	0.99	0.60	0.067

The Eligible Countries column represents the number of member and non-members states that qualify to join the treaty. (Appendix Section III.A.1 defines eligibility in detail). NYC Share represents the fraction of global imports and exports that originated or terminated in NYC member countries. Diff'd Share represents the Rauch differentiated goods' share of trade. This data was derived from Feenstra et al. (2005), which starts in 1962. The table stops at 2006, the last year of trade for the Head et al. (2010) data.

Data Sources: UNCITRAL, Head et al. (2010), Feenstra et al. (2005).

that membership is endogenous. Furthermore, the nature of trade changed in this period, as the last two columns of Table 3.1 show. Differentiated goods, which tend to be more contract-intensive, increased from less than half to 60% of world trade, using the product categories defined by Rauch (1999). The share of trade between countries with colonial ties also decreased significantly, so the business relationships behind trade transactions now more frequently cross the barriers of language, legal origin and cultural affiliation.

Bilateral trade data came from Head et al. (2010). These include the US dollar value of trade between country-pairs in all years between 1948 and 2006, as well as bilateral variables like colonial linkages, distance and shared languages. I derive NYC members' share of global trade using the data on membership dates. The sources and descriptions for other data items like regional trade agreement memberships are introduced with each relevant variable. Regional trade agreements as a form of interaction between countries are briefly discussed after the next baseline estimates on aggregate membership trends.

The 20 founding members of the treaty provide the initial conditions that I take as given. This leaves the timing of membership decisions for about 120 economies available for analysis, as well as the lags for the 45 non-members that I consider censored observations.

3.4.1 Baseline Estimates of Government Responsiveness γ

The γ term in equation (3.2) should be zero if interactions between governments play no role in their treaty membership decisions. This motivates my first stage of analysis, which

ignores all country-level attributes and trade. To do this, I estimate equation (3.2) fitting data on the number of NYC members for each year between 1958 and 2014 to the equation $\ln[P/(1 - P)] = \gamma t + A$.

The estimates yield a γ of 0.033 (using GMM). (The parameter’s z-statistic of 34.3 strongly suggests that it is not zero). With OLS, one gets similar estimate of 0.035, and an R^2 of 0.90 with an A term constructed to match the number of treaty members in 1958. The fit between this naive model and countries’ accession to the NYC is itself remarkable. This estimate yields the model predictions in Figure 3.1.

Having established that interactions between countries play a role in the NYC membership decision, the next steps are to disentangle the direct effects of trade from other interactions, as those other interactions will represent contagion - the adoption of a policy because others did the same, without .

3.4.2 Trade-Driven and Interaction-Driven Membership τ vs. κ

Here I allow interactions between countries to be weighted by trade, while keeping all countries homogeneous as shown in equation (3.5). The data for this illustrative exercise is limited to n , the number of eligible countries, N , the number of NYC members and $NYC_{tradeshare}$, members’ share of trade in each year between 1958 and 2010.²¹

Formally, I redefine equation (3.5) as follows:

$$\ln\left(\frac{P}{1 - P}\right) = (\tau NYC_{tradeshare})t + A \tag{3.7}$$

Estimating equation (3.7) yields a τ of 0.037 (using GMM). (The parameter’s z-statistic of 45.0 indicates that it is not zero). With OLS, one gets similar estimates of 0.038, and an R^2 of 0.98. This specification fits the aggregate trend better, which is not surprising – for a trade treaty, the relevant form of interaction between countries should be trade. Estimating (3.7) yields the model predictions in Figure 3.2. It is encouraging that more than 95% of the variation in membership over time is explained by time and the share of trade by NYC members, as shown by the fit of the model to the data in the graph.

These estimation exercises on aggregate membership trends are only suggestive, as they

²¹ To justify the form in (3.7), one must emphasize that the $NYC_{tradeshare}$ term must be interacted with the t term, in (3.1), the prevalence of NYC membership is one more factor that drives the ‘infection’ of non-members. It cannot be separated from time. For this exercise, I extended the Head et al. (2010) data using the UN COMTRADE database from 2007 to 2010.

Table 3.2: Summary Statistics for Model Variables

VARIABLES	Mean	Std. Dev.	Minimum	Maximum	N
NYC Share of Trade	.66	.33	0	1	8835
NYC Share of RTAs	.37	.34	0	1	10441
Rauch Diff'd Goods	.44	.15	0	1	7692
Colonial Share	.14	.16	0	1	8534

The Rauch differentiated goods variable only becomes available from 1962. The table reports values for 1968 to 2006. NYC membership increases from 20 in 1958 to 139 in 2006. It is 140 in 2010.

do not consider the specific countries that adopted the treaty, nor the number and size of their trade partners.

The rest of this section explores these findings in detail, using bilateral and country-specific attributes.

3.4.3 Country-Level Differences in Trade-Driven Membership

Table 3.2 summarizes the variables used in the duration models that follow. At the country-year level NYC partners generally represented more than half of imports and exports for the average country between 1958 and 2006. As Table 3.1 suggests, this value increased over time. The broadly consistent trend was stronger for some countries than others.

Interactions with NYC members in the context of regional trade agreements also changed in this period. The average share of each country's partnerships through RTAs that were due to NYC members was less than four in ten, but the average depended on the country and time. For some countries, all their RTA partners were NYC members, others had none. In addition, I measure the nature of interactions between countries with the *Colonial Share* and *Rauch Diff'd Goods* variables: the first is the fraction of a country's trade that involved partners with colonial ties and the second represents contract-intensive goods as a share of each country's trade. In the data, *Colonial Share* was generally in decline while differentiated goods were gaining global trade share.

I will emphasize RTA membership as a forum for interaction between countries for several reasons: It is a reasonable measure of peer linkage between countries that is not a direct measure of trade. Countries with strong trade ties tend to have regional trade agreements, but this correlation is not perfect. By definition, RTAs are tied to regional geography. Therefore RTA membership is a good proxy for interactions between countries on the basis

of geographic proximity and the commonalities that usually come with such proximity. The difference between the share of a country's trade that is due to NYC members, and the share of a country's RTA partnerships that is due to NYC members provides a useful means to identify whether countries join primarily because of trade, or because they want to be NYC members like their peers.

Other variables that I use include legal origins and systems of government. There are categorical by nature. The legal origins of the 197 countries in the analysis are: French Common Law (76), German/Scandinavian (10), Socialist (33) and U.K. Common Law (63). Government systems vary by year for several countries. Countries like Saudi Arabia and the United Kingdom remained a monarchy or parliamentary democracy for all the years they were in the data, while others changed systems of government. This category breaks down into country-years as: Parliamentary Democracy (2013), Semi-Presidential Democracy(670), Presidential Democracy (1056), Military Dictatorship (2327), Civilian Dictatorship (1573), Monarchies (686). Cheibub et al. (2010) provided the data on systems of government between 1960 and 2008.

Table 3.3 presents the contributions of trade with NYC members to countries' decision to join. I assume an exponential form for the hazard model, following the distinctive structure of equation (3.6). This form implies that countries join only because of interactions between members and non-members over time. This also requires the explicit assumption that the baseline participation hazard is constant. I argue that this is the case for NYC treaty participation using an analogy. A person standing outdoors during a hailstorm faces a non-negative hazard of getting hit by hail. If the said person had a perfectly sturdy umbrella, one could say that hazard was constant (i.e. zero). However, carving a hole in the umbrella returns the hazard to positive values. The hazard with a 'holey' umbrella is always the constant value of zero plus the hazard due to exposure, just as the hazard of joining the NYC is a baseline of zero added to the various forms of exposure to NYC participation through direct benefits or peer effects.

Column 1 of Table 3.3 is consistent with the graph in Figure 3.1; the number of countries that are current NYC members is a reasonable predictor of the participation hazard for non-members. The estimates suggest that the hazard of non-members joining increases by 32% [i.e. $\exp(0.15 * 1.888) - 1$] if the fraction of NYC members increased by 0.15 (one standard deviation). That said, it is not plausible to assume that South Africa joining the treaty would have the same effect on China as South Korea joining the treaty. The regional and cultural links make South Korea a more proximate source of contagion for China. Column

Table 3.3: NYC Membership Hazard Estimates: Contagion

VARIABLES	(1)	(2)	(3)	(4)
NYC Countries	1.888*** (0.00291)			1.027 (0.206)
NYC Share of RTAs		1.081*** (0.00121)		0.968** (0.0136)
NYC Share of Trade			1.098** (0.0152)	0.611* (0.0872)
Constant	-4.491*** (0)	-4.001*** (0)	-4.372*** (0)	-4.873*** (0)
Observations	4,238	4,238	4,015	4,015

The table reports estimates from a duration model, where the duration of interest is the lag between when a country is eligible to when it joins the NYC or the data is censored. The time-varying covariates represent country-year observations, except for *NYC Share*, the fraction of all countries in a particular year that belong to the NYC. *NYC Trade Share* is the fraction of imports and exports for a country-year traded with NYC members. *NYC RTAs* is the fraction of a country's RTA partners that belong to the NYC in each year. The specification assumed an exponential form, i.e. constant baseline hazard.

Robust standard errors in parentheses, clustered by country. ***, **, and * denote statistical significance at the 1, 5, and 10% levels.

2 explores this idea.

By 2010 most countries belonged to a regional trade agreement. These regional fora for negotiating trade policy are one of the defining features of the last half-century (Crawford and Fiorentino, 2005). It is also reasonable to expect that information about the benefits of NYC membership, or negotiations to influence non-members to join may happen in the course of RTA meetings. Column 2 estimates the effects of belonging to RTAs with NYC members on the hazard of participation: the hazard is estimated to be three times higher if all of a country's RTA partners were NYC members, compared with belonging to no RTAs or RTAs with no NYC members.²²

Column 3 of Table 3.3 mildly suggests that trade between members and non-members promotes NYC participation. It is consistent with the Figure 3.2, with estimates that suggest a 43% increased likelihood of participation for a non-member whose share of trade with NYC members increases by 0.33 (one standard deviation). However, the level of statistical significance for this estimate is less than the comparable estimates for contagion outside the

²²If a country belongs to multiple RTAs, each linkage to partners through an RTA is counted separately. Data on RTA memberships by year from (D'úr et al., 2014).

direct value of trade.

Combining all three measures of interactions between NYC members and non-members in Column 4 confirms the intuition that interactions through trade and RTAs can explain away the purported effects of contagion through observing the global trend in participation. The estimates that follow will exclude the global NYC share measure in favor of the country-year measures of interaction with NYC members.

Table 3.4 builds on the argument in Nunn (2007) that contract enforcement may be more vital for some goods. Higher shares of contract-intensive goods in a country's trade portfolio should motivate participation in the NYC, which facilitates formal dispute resolutions through arbitration. I measure the share of a country's imports and exports that are contract-intensive as those in the Rauch (1999) differentiated goods category. For an effective contract, these differentiated goods not sold on a global exchange or reference-priced require at the minimum that trading parties agree on the quality of the traded item. The category covers a broad range of products that range from tee-shirts to airplanes. (The average share of differentiated goods increased from 38% in 1958 to about 50% in 2010).

The estimates suggest that increasing differentiated goods as a share of a country's trade portfolio also increases the likelihood of NYC participation. The hazard of participation according to column 1 increases by 47% if one increased this index by 0.152 (one standard deviation). [Estimates computed as $\exp(2.545 * 0.152) - 1$]. The estimated effect decreases slightly in column 2 when one accounts for the contributions of RTA memberships and trade with NYC memberships.

Table 3.4 also tests a hypothesis from Masten and Prüfer (2014) – that traders will favor formal adjudication as trade expands beyond the scope of past established relationships. In this case, colonial history is the proxy for established relationships. Head et al. (2010) shows that the erosion of colonial linkages is one of the most significant changes to trade patterns over the last half-century. By this paper's estimates, trade between countries with historical colonial ties fell from above 20% in the 1950's to about 7% in 2010. Historical colonial linkages as a proxy also capture some of the effects of having similar languages, legal institutions and trade treaties.

Column 3 of the table implies that the hazard of joining the NYC decreases by 12% for an increase in the share of a country's trade with past colonial partners of 0.145 (one standard deviation). This agrees with the hypotheses in Masten and Prüfer (2014) that established relationships with self-enforcement are a substitute for formal arbitration. However, the estimates are not statistically significant.

Table 3.4: NYC Membership Hazard Estimates: The Nature of Trade and Goods

VARIABLES	(1)	(2)	(3)	(4)	(5)
Rauch Diff'd Share	2.545*** (0.00001)	1.992*** (0.000800)			1.951*** (0.00103)
Colonial Share			-0.890 (0.153)	-0.705 (0.300)	-0.670 (0.360)
NYC Share of RTAs		1.032*** (0.00354)		1.176*** (0.000545)	0.983*** (0.00509)
NYC Share of Trade		0.765* (0.0646)		0.665* (0.0950)	0.767* (0.0651)
Constant	-4.581*** (0)	-5.289*** (0)	-3.423*** (0)	-4.394*** (0)	-5.151*** (0)
Observations	3,426	3,373	3,994	3,994	3,363

The table reports estimates from a duration model, where the duration of interest is the lag between when a country is eligible to when it joins the NYC or the data is censored. The time-varying covariates represent country-year observations. *Rauch Share*, is the fraction of a country's trade in a particular year that belong to the Rauch differentiated goods category. I use the liberal definition. *Colonial Share* is the fraction of imports and exports for a country-year traded with former colonies or colonizers. *NYC Trade Share* is the fraction of imports and exports for a country-year traded with NYC members. *NYC RTAs* is the fraction of a country's RTA partners that belong to the NYC in each year. The specification assumed an exponential form, i.e. constant baseline hazard.

Robust standard errors in parentheses, clustered by country. ***, **, and * denote statistical significance at the 1, 5, and 10% levels.

Combining all the foregoing in Column 5 of Table 3.4 does not change any of the predictions substantially. Having more differentiated goods is still associated with a statistically significant hazard of NYC participation, as is having NYC trade and RTA partners. The share of trade with colonial partners remains a negative contributor to the hazard, but is still not statistically significant.

One must consider countries' systems government or legal origins, given earlier discussions of country-specific factors not related to trade. Table 3.5 examine how these two factors contribute to the hazard of participating in the NYC. For example, one may consider democracies more likely to agree to an international covenant. Among democracies, it may also be more likely that parliamentary democracies approve treaties faster – the executive and legislature are unlikely to be politically opposed, by definition. There is room for within- and between- variation in the system of government variable. 90 of the 188 countries for which data is available do not change systems of government. 40 of the remaining 98 change more than once, with transitions to democracy becoming more common over time.

Columns 1 and 2 of Table 3.5 do not support the view that the system of government matters. The estimates suggest that compared to monarchies, democracies are more likely to join, although this is not statistically significant when one considers the nature of traded goods and the prevalence of NYC members in trade relationships. In Column 5, the estimates that control for countries' legal origins indicate that dictatorships are less likely to join the treaty, but this is also not statistically significant.

Legal origins provide another measure of countries' propensity to sign a treaty about using legal resources to support trade. Some legal systems may be more disposed to cooperative agreements of this nature, which should reduce the national government's perceived political costs of signing the treaty. Civil Law (76 countries) and Common Law (63 countries) are the most common categories. These categories are time-invariant.

Columns 3 and 4 of the Table show that countries with socialist legal origins are more likely to join the NYC convention. Even after controlling for interactions with NYC members in column 4, the participation hazard for socialist legal origin countries is 1.9 times that of Civil Law countries, [calculated as $\exp(0.65)$]. The difference is statistically significant, while the difference between Common Law and Civil Law countries is not. German/Scandinavian legal origin countries have a higher estimated participation hazards, but this is not robust to controlling for other interactions with NYC members and government systems in Column 5.

The socialist legal system was built on Civil Law and shares many features, for example, in deciding cases judges must follow the Legal Code and not judicial precedent (as Common

Table 3.5: NYC Membership Hazard Estimates: Country Features

VARIABLES	(1)	(2)	(3)	(4)	(5)
Government System					
Parliamentary	0.532 (0.136)	0.387 (0.331)			0.270 (0.496)
Semi-Presidential	1.146*** (0.00898)	0.755 (0.142)			0.356 (0.504)
Presidential	0.315 (0.387)	-0.0213 (0.951)			-0.0546 (0.886)
Military Dictator	0.369 (0.267)	0.199 (0.564)			-0.0864 (0.808)
Civilian Dictator	0.0755 (0.840)	-0.154 (0.687)			-0.354 (0.389)
Legal Origin					
Germ./Scandinavian			1.086** (0.0157)	0.877*** (0.127)	0.523 (0.431)
Socialist			0.835*** (0.00402)	0.651** (0.0379)	0.684** (0.0246)
Common Law			-0.0917 (0.620)	0.0209 (0.923)	-0.111 (0.648)
NYC Share of RTAs		0.862** (0.0214)		0.996*** (0.00250)	0.852** (0.0117)
NYC Share of Trade		0.742* (0.0863)		0.648 (0.127)	0.774* (0.0763)
Rauch Diff'd Share		1.740*** (0.00639)		1.948*** (0.00222)	1.696** (0.0107)
Colonial Share		-0.885*** (0.257)		-0.242 (0.734)	-0.305 (0.699)
Constant	-3.963*** (0)	-5.113*** (0)	-3.701*** (0)	-5.257*** (0)	-5.093*** (0)
Observations	4,152	3,342	4,077	3,293	3,289

The table reports estimates from a duration model, where the duration of interest is the lag between when a country is eligible to when it joins the NYC or the data is censored. The time-varying covariates represent country-year observations. Monarchy is the comparison category for government type that is not shown. The reported categories are democracies – semi-Presidential and Presidential democracies, and dictatorships – military and civilian. Similarly, Civil Law as a legal origin category is the comparison category. *Rauch Share*, is the fraction of a country's trade in a particular year that belong to the Rauch differentiated goods category. I use the liberal definition. *Colonial Share* is the fraction of imports and exports for a country-year traded with former colonies or colonizers. *NYC Trade Share* is the fraction of imports and exports for a country-year traded with NYC members. *NYC RTAs* is the fraction of a country's RTA partners that belong to the NYC in each year. The specification assumed an exponential form, i.e. constant baseline hazard.

Robust standard errors in parentheses, clustered by country. ***, **, and * denote statistical significance at the 1, 5, and 10% levels.

Law countries do). However, the socialist legal system puts more emphasis on power of the state over property. Countries with this legal origin may have been inclined to join the NYC in order to signal a willingness to trade with other countries, given the notable differences in legal principles and procedures.

Combining all the factors in Column 5 gives a picture that is largely consistent with previous columns and tables. Having either more differentiated goods or a socialist legal origin increases the hazard of NYC participation, and so does having NYC members as RTA partners. The share of trade with NYC partners is marginally significant.

Overall, the results reported in Table 3.5 suggest that both internal factors like legal origins and external incentives like the membership of regional trade partners influence country's decisions to join a treaty like the NYC. This lends support to other papers that find evidence of peer effects in policy adoption, e.g. Shipan and Volden (2008) and Simmons and Elkins (2004). The findings also support claims that some goods are more contract-intensive (Nunn, 2007; Berkowitz et al., 2006), and the hypothesis in Masten and Prüfer (2014) that the preference for formal adjudication will increase as the nature and scope of trade expands.²³

3.5 Conclusion

I evaluate factors that influence the hazard of participation in the New York Convention (NYC) – a treaty that supports trade across legal jurisdictions through the commitment of its signatories to the outcomes of international contract dispute arbitration. The significance of countries' commitment to support private dispute settlement with the power of the state has led some to call the NYC a modern day *Lex Mercatoria*. The study is motivated by a puzzling fact: despite benefits of joining the treaty, several large trading economies like the United States, Canada and the United Kingdom waited more than a decade to join. (See a discussion of the US case in Appendix Section III.B.4).

I find that participation in the treaty is driven by interactions with countries that are members. Regional trade agreements are consistently statistically significant as a forum for interaction that promotes faster adoption of the treaty. Consistent with this finding, countries like Canada and the United States, who had no trade agreements in the 1960s, or

²³ Masten and Prüfer (2014) argues that trade becomes more contract-intensive with geographic distance, though the same principle applies to changes in the composition of trade. Distance provides little usable variation for my tests: the global average for trade distance has remained unchanged near 5000km since 1985 (Berthelon and Freund, 2008).

whose trade agreement partners did not belong to the NYC were less likely to join in the first decade, compared with members of the European Free Trade Area, which was established in 1960, and whose members include the first NYC signatories.

Peer effects matter in this context of international policy. Trade with NYC members in itself does not consistently predict a higher hazard of joining the NYC. This implies that in making the decision to join, peer effects or contagion through fora like RTAs may be as influential as the direct benefits from trading with NYC members. This is consistent with some of the work on policy diffusion and suggests further inquiry into the factors that make such peer effects more or less effective. Early works in this area include (Golub and Jackson, 2010; Young, 2009; Shipan and Volden, 2008; Simmons and Elkins, 2004). Table III.A.4 in the appendix suggests that much of what is measured as peer effects could be from the relationships of countries to their former colonial metropolises, or the countries from which they broke away.²⁴

Country-level factors like legal origins also affected the hazard of treaty participation, with the estimated hazard of participation being almost twice as high for socialist legal origin countries than Common Law or Civil Law countries. Other country-level factors like the system of government do not yield statistically significant predictions. It is not clear that democracies are more likely to join the treaty than dictatorships or monarchies.

The nature of traded goods also influences the hazard of participation. Countries that specialize in contract-intensive goods have a higher hazard of joining the treaty compared with those that trade commodities like cotton and coffee. The latter products are sold on regular exchanges and presumably require little legal intervention for a successful transaction. This finding supports the hypothesis in Masten and Prüfer (2014) that the preference for formal adjudication will increase as trade expands beyond the scope of reputational enforcement.

These findings come from estimating duration models that reflects the nature of the problem, and is flexible enough to accommodate country-level categories, as well as time-varying covariates like trade that are driven by interactions between countries.

The relevance of peer effects through RTA membership contributes to the literature on policy adoption. First, the presence of peer effects suggests that rational choice by national

²⁴In cases like Serbia and Slovenia, the speed with which Slovenia joined after seceding from Yugoslavia may reflect the fact that Yugoslavia was a member of the NYC, just as Serbia succeeded to Yugoslavia's NYC status after the split. The main body of the paper interprets Slovenia's entry as due to peer effects from Serbia, for simplicity. Table III.A.4 nevertheless suggests further research into how the choices of 'parent-countries' influence those of the countries they spawned.

governments may not be completely divorced from social influence or contagion. Second, the finding suggests that inter-governmental organizations may be important to the process of policy adoption because of others' influence. Whether this is due to learning, or homophily as suggested by Jackson (2008), I leave for later work.

These findings also encourage future research on the estimated effect of NYC participation on trade, after controlling for endogeneity in participation. Such work would contribute to research on the subject by Leeson (2008) and Berkowitz et al. (2006).

Bibliography

- Aral, S., Muchnik, L., Sundararajan, A., 2013. Engineering Social Contagions: Optimal Network Seeding in the Presence of Homophily. *Network Science* 1 (02), 125–153.
- Bagwell, K., Staiger, R. W., 2011. What do Trade Negotiators Negotiate About? Empirical Evidence from the World Trade Organization. *The American Economic Review*, 1238–1273.
- Baldwin, R., Jaimovich, D., 2012. Are free trade agreements contagious? *Journal of International Economics* 88 (1), 1–16.
- Banerjee, A. V., 1992. A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 797–817.
- Bass, F. M., 1969. A New Product Growth for Model Consumer Durables. *Management Science* 15 (5).
- Berkowitz, D., Moenius, J., Pistor, K., 2006. Trade, Law, and Product Complexity. *The Review of Economics and Statistics* 88 (2), 363–373.
- Berry, F. S., Berry, W. D., 1990. State lottery adoptions as policy innovations: An event history analysis. *American Political Science Review* 84 (02), 395–415.
- Berthelon, M., Freund, C., 2008. On the Conservation of Distance in International Trade. *Journal of International Economics* 75 (2), 310–320.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 992–1026.
- Bramoullé, Y., Djebbari, H., Fortin, B., 2009. Identification of Peer Effects through Social Networks. *Journal of Econometrics* 150 (1), 41–55.
- Carroll, C. D., 2001. The Epidemiology of Macroeconomic Expectations. National Bureau of Economic Research.
- Casella, A., 1996. On Market Integration and the Development of Institutions: The Case of International Commercial Arbitration. *European Economic Review* 40 (1), 155–186.
- Cattaneo, O., Braga, C. P., 2008. Everything You Wanted to Know about WTO Accession (but Were Afraid to Ask). World Bank Policy Research Working Paper (5116).
- Chartered Institute of Arbitrators, 2011. International arbitration survey.
- Cheibub, J. A., Gandhi, J., Vreeland, J. R., 2010. Democracy and Dictatorship Revisited. *Public Choice* 143 (1-2), 67–101.

- Copelovitch, M. S., Ohls, D., 2012. Trade, Institutions, and the Timing of GATT/WTO Accession in Post-colonial States. *The Review of International Organizations* 7 (1), 81–107.
- Crawford, J.-A., Fiorentino, R. V., 2005. The Changing Landscape of Regional Trade Agreements. World Trade Organization.
- Dúr, A., Baccini, L., Elsig, M., 2014. The Design of International Trade Agreements: Introducing a New Dataset. *The Review of International Organizations* 9 (3), 353–375.
- Dixit, A., 2003. Trade Expansion and Contract Enforcement. *Journal of Political Economy* 111 (6), 1293–1317.
- Feenstra, R. C., Lipsey, R. E., Deng, H., Ma, A. C., Mo, H., 2005. World Trade Flows: 1962-2000. National Bureau of Economic Research.
- Goldberg, J., 2005. Geographies of Trade and Traders in the Eleventh-Century Mediterranean: A Study Based on Documents from the Cairo Geniza. Ph.D. thesis.
- Golub, B., Jackson, M. O., 2010. Naive Learning in Social Networks and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 112–149.
- Goolsbee, A., Klenow, P. J., 2002. Evidence Of Learning And Network Externalities In The Diffusion of Home Computers. *Journal of Law and Economics* 45 (2), 317–343.
- Head, K., Mayer, T., Ries, J., 2010. The Erosion of Colonial Trade Linkages after Independence. *Journal of International Economics* 81 (1), 1–14.
- Jackson, M., 2008. *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O., Yariv, L., 2007. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 92–98.
- Jenkins, S. P., 1995. Easy Estimation Methods for Discrete-time Duration Models. *Oxford Bulletin of Economics and Statistics* 57 (1), 129–136.
- Lee, J., Lee, Y., 2011. A Grouped Mixed Proportional Hazard Model with Social Interactions: The Passage of the Motorcycle-Helmet-Use Law.
- Lee, L.-f., 2007. Identification and Estimation of Econometric Models with Group Interactions, Contextual Factors and Fixed Effects. *Journal of Econometrics* 140 (2), 333–374.
- Leeson, P. T., 2008. How Important is State Enforcement for Trade? *American Law and Economics Review* 10 (1), 61–89.
- Masten, S. E., Prüfer, J., 2014. On the Evolution of Collective Enforcement Institutions: Communities and Courts. *Journal of Legal Studies* 43 (2), 359–400.

- Nunn, N., 2007. Relationship-specificity, incomplete contracts, and the pattern of trade. *Quarterly Journal of Economics*, 569–600.
- Rauch, J. E., 1999. Networks versus Markets in International Trade. *Journal of International Economics* 48 (1), 7–35.
- Robins, J., Tsiatis, A. A., 1992. Semiparametric Estimation of an Accelerated Failure Time Model With Time-Dependent Covariates. *Biometrika* 79 (2), 311–319.
- Rose, A. K., 2004. Do We Really Know That the WTO Increases Trade? *American Economic Review* 94 (1), 98–114.
- Ryan, B., Gross, N. C., 1943. The Diffusion of Hybrid Seed Corn in Two Iowa Communities. *Rural Sociology* 8 (1), 15–24.
- Shipan, C. R., Volden, C., 2008. The Mechanisms of Policy Diffusion. *American Journal of Political Science* 52 (4), 840–857.
- Simmons, B. A., Elkins, Z., 2004. The Globalization of Liberalization: Policy Diffusion in the International Political Economy. *American Political Science Review* 98 (01), 171–189.
- Subramanian, A., Wei, S.-J., 2007. The WTO Promotes Trade, Strongly but Unevenly. *Journal of International Economics* 72 (1), 151–175.
- Tomz, M., Goldstein, J. L., Rivers, D., 2007. Do We Really Know that the WTO Increases Trade? Comment. *American Economic Review*, 2005–2018.
- UNCITRAL, 2014. Status: Convention on the Recognition and Enforcement of Foreign Arbitral Awards (New York, 1958).
- Wagner, U. J., 2010. Estimating a Strategic Model of Treaty Formation: The Case of the Montreal Protocol. SSRN Working Papers (1730183).
- Young, H. P., 2009. Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review* 99 (5), pp. 1899–1924.

Appendix

III.A NYC Treaty Membership and Delays to Entry

III.A.1 Entry Date: Signature, Ratification, Accession and In-Force Dates

Countries join the New York Convention by ratification, accession or succession. All dates for these actions by members are at the UNCITRAL website: http://www.uncitral.org/uncitral/en/uncitral_texts/arbitration/NYConvention_status.html. Tables III.A.1 and III.A.2 report these dates and the delay to membership for countries that had joined by October 2014.

The original treaty was signed by 20 countries on June 10, 1958 and came into force on June 7, 1959. (Four countries signed later in 1958). All 24 countries that signed in 1958 are considered founding signatories of the treaty. The treaty only took effect after ratification by national legislatures; both dates are shown in the table at the link. From January 1959 countries acceded to the treaty by submitting a statement to that effect. Ratification was necessary when national laws had to be amended to give the treaty effect. The joining country's legislature records the state's obligations under the treaty, and a signed copy of the treaty is submitted to the other contracting parties. Signatories that have acceded to, but not ratified the treaty are obliged to support the treaty's objectives in good faith. Countries like Serbia that were carved out of an NYC member could succeed to the treaty on the terms of the former member. The ratification or accession document defines when the treaty comes into force, usually a few months after the ratification or accession date.

III.A.2 Entry Delay: Lag from Eligible Date to Signature Date

The delay to participation reported in the paper is the lag between when a country first becomes eligible to when it signed the treaty, in years (with months as fractional years). Participation is defined broadly to cover signature, accession, ratification or succession as de-

scribed in the preceding paragraphs. The condition for eligibility, as stated by the governing UN secretariat is:

The Convention is open to accession by any Member State of the United Nations, any other State which is a member of any specialized agency of the United Nations, or is a Party to the Statute of the International Court of Justice (articles VIII and IX).

84 countries met these criteria in 1958. A wave of countries achieved independence from colonialism in the 1960s, increasing the pool of potential participants. By 2014, the number of eligible countries had increased to roughly 200. Table III.A.1 shows members of the treaty sorted by the lag between eligibility and participation. The eligibility date is June 10, 1958 for countries with treaty status when the NYC was created. For others, I used the dates of membership in the UN, the International Court of Justice (ICJ), and three selected specialized UN agencies.²⁵ The earliest membership date for these organizations formed the basis for eligibility.

Figure III.A.1 represents a histogram of entry delays in years. (Each bin in the histogram is six months wide). As described earlier in the paper, a batch of countries joined within a year of becoming eligible; about 25% of all countries (33% of countries eligible in 1958). Others joined in a relatively steady trickle that continued for more than a half-century, e.g., Myanmar joined in 2013, fifty-five years after it became eligible at the treaty's formation in 1958. Transition economies like Croatia and Bosnia-Herzegovina joined in 1993, the next year after attaining eligibility in 1992. Despite its size and trade, the US delayed entry by more than 12 years while smaller economies like Jordan, the Philippines and Poland were founding members. Country size offers no readily obvious patterns to membership. The second panel of the figure show a similar pattern for all the countries that were eligible in 1958; the lag for these is not complicated by independence dates or the membership status of other countries.

However, the patterns in Figure III.A.1 are themselves informative, following the reasoning in Section 3.1. The trickle of countries that follow the first five years of membership is not increasing like a model with network effects would predict.

²⁵For UN membership, <http://www.un.org/en/members/>; for ICJ membership, <http://www.icj-cij.org/jurisdiction/?p1=5&p2=1&p3=1&sp3=a>; for International Labor Organization (ILO) membership, <http://www.ilo.org/dyn/normlex/en/f?p=1000:11003:0>; for the International Telephone Union membership (ITU) http://www.itu.int/online/mm/scripts/mm.list?_search=ITUstates&_languageid=1; and Food and Agriculture Organization (FAO) membership <http://www.fao.org/legal/home/fao-members/en/>.

Table III.A.1: Countries, NYC Membership Status and Entry Delay

Country	Code	Eligible	Joined	Delay	Country	Code	Eligible	Joined	Delay
Delay < 5 years					Delay >5 years <30				
Belgium	BEL	1958	1958	0	Botswana	BWA	1966	1971	5.2
Costa Rica	CRI	1958	1958	0	Djibouti	DJI	1977	1983	5.8
Germany	DEU	1958	1958	0	Armenia	ARM	1992	1997	5.8
India	IND	1958	1958	0	Republic of Moldova	MDA	1992	1998	6.5
Israel	ISR	1958	1958	0	Antigua and Barbuda	ATG	1981	1989	7.3
Jordan	JOR	1958	1958	0	Azerbaijan	AZE	1992	2000	7.9
Netherlands	NLD	1958	1958	0	Tunisia	TUN	1958	1967	9.1
Philippines	PHL	1958	1958	0	Nigeria	NGA	1960	1970	9.4
Poland	POL	1958	1958	0	Ghana	GHA	1958	1968	9.8
El Salvador	SLV	1958	1958	0	Dominica	DMA	1978	1988	9.8
Argentina	ARG	1958	1958	0.2	Italy	ITA	1958	1969	10.6
Slovenia	SVN	1992	1992	0.2	Brunei Darussalam	BRN	1984	1996	11.8
Slovakia	SVK	1993	1993	0.3	United States of America	USA	1958	1970	12.3
Serbia	SCG	2000	2001	0.3	Mexico	MEX	1958	1971	12.8
Montenegro	MNE	2006	2006	0.3	Benin	BEN	1960	1974	13.7
France	FRA	1958	1958	0.4	Zimbabwe	ZWE	1980	1994	14.3
Luxembourg	LUX	1958	1958	0.4	Denmark	DNK	1958	1972	14.5
Bulgaria	BGR	1958	1958	0.5	Republic of Korea	KOR	1958	1973	14.7
Belarus	BLR	1958	1958	0.5	Marshall Islands	MHL	1991	2006	15.3
Switzerland	CHE	1958	1958	0.5	Cuba	CUB	1958	1974	16.5
Ecuador	ECU	1958	1958	0.5	Bahrain	BHR	1971	1988	16.6
Finland	FIN	1958	1958	0.5	Australia	AUS	1958	1975	16.8
Sri Lanka	LKA	1958	1958	0.5	Holy See	VAT	1958	1975	16.9
Monaco	MCO	1958	1958	0.5	Chile	CHL	1958	1975	17.3
Pakistan	PAK	1958	1958	0.5	United Kingdom	GBR	1958	1975	17.3
Russia	RUS	1958	1958	0.5	South Africa	ZAF	1958	1976	17.9
Sweden	SWE	1958	1958	0.5	Kuwait	KWT	1959	1978	18.7
Ukraine	UKR	1958	1958	0.5	Spain	ESP	1958	1977	18.9
Latvia	LVA	1991	1992	0.6	Bangladesh	BGD	1972	1992	19.9
Morocco	MAR	1958	1959	0.7	St. Vincent & Grenadines	VCT	1980	2000	20
Czech Republic	CZE	1993	1993	0.7	Cyprus	CYP	1960	1980	20.3
Egypt	EGY	1958	1959	0.8	Tajikistan	TJK	1992	2012	20.5
Syrian Arab Republic	SYR	1958	1959	0.8	Singapore	SGP	1965	1986	20.9
FYR Macedonia	MKD	1993	1994	0.9	Colombia	COL	1958	1979	21.3
Croatia	HRV	1992	1993	1.2	Lesotho	LSO	1966	1989	22.7
Bosnia-Herzegovina	BIH	1992	1993	1.3	Mozambique	MOZ	1975	1998	22.8
Thailand	THA	1958	1959	1.5	Ireland	IRL	1958	1981	22.9
Cambodia	KHM	1958	1960	1.6	Indonesia	IDN	1958	1981	23.3
Madagascar	MDG	1960	1962	1.8	New Zealand	NZL	1958	1983	24.6
Estonia	EST	1991	1993	1.9	Cook Islands	COK	1984	2009	24.7
Georgia	GEO	1992	1994	1.9	Uruguay	URY	1958	1983	24.8
Central African Republic	CAF	1960	1962	2.1	Kenya	KEN	1963	1989	25.2
San Marino	SMR	1977	1979	2.2	Haiti	HTI	1958	1983	25.5
Norway	NOR	1958	1961	2.8	Guatemala	GTM	1958	1984	25.8
Tanzania	TZA	1961	1964	2.8	Barbados	BRB	1966	1993	26.3
Austria	AUT	1958	1961	2.9	Panama	PAN	1958	1984	26.3
Japan	JPN	1958	1961	3	Algeria	DZA	1962	1989	26.3
Romania	ROU	1958	1961	3.3	Burkina Faso	BFA	1960	1987	26.5
Trinidad and Tobago	TTO	1962	1966	3.4	Oman	OMN	1971	1999	27.3
Lithuania	LTU	1991	1995	3.5	Malaysia	MYS	1958	1985	27.4
Kazakhstan	KAZ	1992	1995	3.7	Canada	CAN	1958	1986	27.9
Hungary	HUN	1958	1962	3.8	Cameroon	CMR	1960	1988	27.9
Uzbekistan	UZB	1992	1996	3.9	Mauritius	MUS	1968	1996	28.3
Greece	GRC	1958	1962	4.1	China	CHN	1958	1987	28.6
Niger	NER	1960	1964	4.1	Uganda	UGA	1962	1992	29.3
Kyrgyzstan	KGZ	1992	1996	4.8					

*lag is the gap in months/12 rounded to 1.dp. between 10-Jun-1958 or the eligibility date, and when a country joined

Table III.A.2: Countries, NYC Membership Status and Entry Delay (continued)


Country	Code	Eligible	Joined	Delay	Country	Code	Eligible	Joined	Delay
	Delay > 30 years					Non-Members			
Peru	PER	1958	1988	30.1	Ethiopia	ETH	1958		
Côte d'Ivoire	CIV	1960	1991	30.4	Iraq	IRQ	1958		
Qatar	QAT	1971	2002	31.3	Libya	LBY	1958		
Guinea	GIN	1958	1991	32.1	Sudan	SDN	1958		
Mongolia	MNG	1961	1994	33	Yemen	YEM	1958		
Bahamas	BHS	1973	2006	33.3	Togo	TGO	1960		
Mali	MLI	1960	1994	34	D. R. Congo	COD	1960		
Turkey	TUR	1958	1992	34.1	Congo	COG	1960		
Senegal	SEN	1960	1994	34.1	Somalia	SOM	1960		
United Arab Emirates	ARE	1971	2006	34.7	Chad	TCD	1960		
Mauritania	MRT	1961	1997	35.6	Sierra Leone	SLE	1961		
Malta	MLT	1964	2000	35.7	Malawi	MWI	1964		
Saudi Arabia	SAU	1958	1994	35.8	Gambia	GMB	1965		
Portugal	PRT	1958	1994	36.3	Maldives	MDV	1965		
Venezuela	VEN	1958	1995	36.7	Swaziland	SWZ	1968		
Bolivia	BOL	1958	1995	36.8	Equatorial Guinea	GNQ	1968		
Sao Tome & Principe	STP	1975	2012	37.2	Nauru	NRU	1969		
Viet Nam	VNM	1958	1995	37.3	Tonga	TON	1972		
Zambia	ZMB	1964	2002	37.3	Guinea-Bissau	GNB	1973		
Paraguay	PRY	1958	1997	39.3	Grenada	GRD	1974		
Nepal	NPL	1958	1998	39.8	Cape Verde	CPV	1975		
Jamaica	JAM	1962	2002	39.8	DPR of Korea	PRK	1975		
Fiji	FJI	1970	2010	39.9	Papua New Guinea	PNG	1975		
Lao PDR	LAO	1958	1998	40	Comoros	COM	1975		
Lebanon	LBN	1958	1998	40.2	Suriname	SUR	1975		
Honduras	HND	1958	2000	42.3	Angola	AGO	1976		
Albania	ALB	1958	2001	43	Seychelles	SYC	1976		
Bhutan	BTN	1971	2014	43	Samoa	WSM	1976		
Iran	IRN	1958	2001	43.3	Namibia	NAM	1977		
Iceland	ISL	1958	2002	43.6	Solomon Islands	SLB	1978		
Dominican Republic	DOM	1958	2002	43.8	Saint Lucia	LCA	1979		
Brazil	BRA	1958	2002	44	Vanuatu	VUT	1981		
Nicaragua	NIC	1958	2003	45.3	Belize	BLZ	1981		
Rwanda	RWA	1962	2008	46.1	St. Kitts & Nevis	KNA	1983		
Gabon	GAB	1960	2006	46.3	Kiribati	KIR	1986		
Afghanistan	AFG	1958	2004	46.4	Micronesia	FSM	1991		
Liberia	LBR	1958	2005	47.3	Turkmenistan	TKM	1992		
Liechtenstein	LIE	1963	2011	48	Eritrea	ERI	1993		
Guyana	GUY	1966	2014	48.3	Andorra	AND	1993		
Burundi	BDI	1962	2014	51.8	Palau	PLW	1994		
Myanmar	MMR	1958	2013	54.8	Tuvalu	TUV	1996		
					Niue	NIU	1999		
					Timor-Leste	TLS	2002		
					Tokelau	TKL	2011		
					South Sudan	SSD	2011		

*lag is the gap in months/12 rounded to 1.dp. between 10-Jun-1958 or the eligibility date, and when a country joined

Table III.A.3: States Without Treaty-Powers Annexed to the Treaty

Country	Code	Added By	Joined
Bermuda	BMU	United Kingdom	1979
British Virgin Islands	VGB	United Kingdom	2014
Cayman Islands	CYM	United Kingdom	1980
Faeroe Islands	FRO	Denmark	1976
Gibraltar	GIB	United Kingdom	1975
Greenland	GRL	Denmark	1976
Guernsey	GGY	United Kingdom	1985
Hong Kong	HKG	China	1997
Isle of Man	IMN	United Kingdom	1979
Jersey	JEY	United Kingdom	2002
Macau	MAC	China	2005
Netherlands Antilles	ANT	Netherlands	1964

Figure III.A.1: NYC Entry Delay



(a) Histogram of Entry Delay for all NYC members
Data Source: UNCITRAL (2014)

(b) Histogram for only countries eligible in 1958

III.A.3 Results for Countries Eligible in 1958

Table III.A.4 replicates the main findings of the paper using only a subset of the data - the 87 countries that were eligible to sign the NYC treaty in 1958. These results address the concern that membership decisions of countries that became subsequently eligible may have been biased, as the tables do not control for the NYC membership of parent-countries or former colonial metropolises.

The first set of findings in Table III.A.4 are remarkably similar to those in Table 3.5. The systems of government generally do not yield statistically significant prediction. (The exception is the negative coefficient on the parliamentary democracy dummy in columns 1 and 2, but this goes away when other factors are considered in column 5). Legal origins also follow a similar pattern, though more systems showing a statistically significant difference in the propensity to join relative to the French Civil Law origin that represents the baseline. The German/Scandinavian and Socialist legal origins are the most likely to join in this assessment, as in the previous table.

The variables that measure interactions and peer influence reverse the findings in Table 3.5. This suggests that some of what the paper reports as peer influence may derive largely from the choices of 'parent-countries', or imitation of the patterns chosen by the former colonial metropolises of countries that achieved independence after 1958. The conclusion to the paper notes this caveat and suggests additional work in later work.

III.B Data

III.B.1 Data Sources

The bilateral aggregate trade data used by this paper is the CEPII gravity dataset. The trade data was compiled from the IFS Direction of Trade Statistics (DOTS) Database by Head et al (2010) and covers more than 200 countries and territories from 1948 to 2006. This data set contains information on GDP and Population from the World Bank's World Development Indicators (WDI) (and national sources for Taiwan). It also has information on legal origins, official languages, colonial relationships and regional trade agreements. The data appendix of the Head et al paper provides further details and references for this gravity dataset.

Table III.A.4: NYC Membership Hazard Estimates: Countries Eligible in 1958

VARIABLES	(1)	(2)	(3)	(4)	(5)
Government System					
Parliamentary	-1.420** (0.037)	-1.607** (0.048)			-1.360 (0.103)
Semi-Presidential	-1.019 (0.155)	-0.887 (0.329)			-0.181 (0.858)
Presidential	-0.847 (0.199)	-0.997 (0.213)			-0.907 (0.291)
Military Dictator	-0.718 (0.278)	-1.020 (0.218)			-0.839 (0.329)
Civilian Dictator	-0.458 (0.505)	-0.835 (0.324)			-0.707 (0.439)
Legal Origin					
Germ./Scandinavian			1.234** (0.018)	2.632*** (0.000)	2.646*** (0.000)
Socialist			0.312 (0.638)	1.965*** (0.001)	1.740*** (0.001)
Common Law			1.039* (0.580)	2.622*** (0.000)	2.605*** (0.000)
NYC Share of RTAs		-1.015** (0.012)		-1.183*** (0.005)	-1.087*** (0.008)
NYC Share of Trade		-0.728 (0.342)	(0.731)	-0.273 (0.322)	-0.784
Rauch Diff'd Share	(0.055)	-1.749* (0.007)	(0.007)	-2.456*** (0.044)	-2.006**
Colonial Share	(0.540)	0.695 (0.769)	(0.769)	0.280 (0.879)	0.178
Constant	4.143*** (0.000)	5.954*** (0.000)	2.303*** (0.000)	2.446*** (0.000)	3.426***
Observations	1,664	1,507	1,681	1,507	1,507

Robust pval in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The table reports estimates from a duration model, where the duration of interest is the lag between when a country is eligible to when it joins the NYC or the data is censored. The time-varying covariates represent country-year observations. Monarchy is the comparison category for government type that is not shown. The reported categories are democracies – semi-Presidential and Presidential democracies, and dictatorships – military and civilian. Similarly, Civil Law as a legal origin category is the comparison category. *Rauch Share*, is the fraction of a country's trade in a particular year that belong to the Rauch differentiated goods category. I use the liberal definition. *Colonial Share* is the fraction of imports and exports for a country-year traded with former colonies or colonizers. *NYC Trade Share* is the fraction of imports and exports for a country-year traded with NYC members. *NYC RTAs* is the fraction of a country's RTA partners that belong to the NYC in each year. The specification assumed an exponential form, i.e. constant baseline hazard.

III.B.2 Explaining NYC Membership with RTA Membership

The increase in NYC membership coincided with the proliferation of trade agreements, and the countries that joined those regional trade agreements (RTAs). Figure III.B.1 shows the twin trend.

To support the assertion in the introduction that take up of the NYC was surprisingly slow, given its benefits, Figure III.B.2 plots membership over time in the NYC side-by-side with the WTO. Despite the obvious lack of a barrier to entry for the NYC, it only had 55 members in 1978, after two decades of its initiation. The comparable number for the WTO is 75.

Given the stricter requirements for GATT/WTO membership, it should be clear from the gap between GATT/WTO membership and NYC membership that initial take-up of the latter was slow. At no time do the number of NYC members surpass the number of GATT/WTO members. The WTO shows an initial surge in membership and the second surge in membership in the sixties as more countries became treaty-eligible. The NYC does not reflect such a pattern.

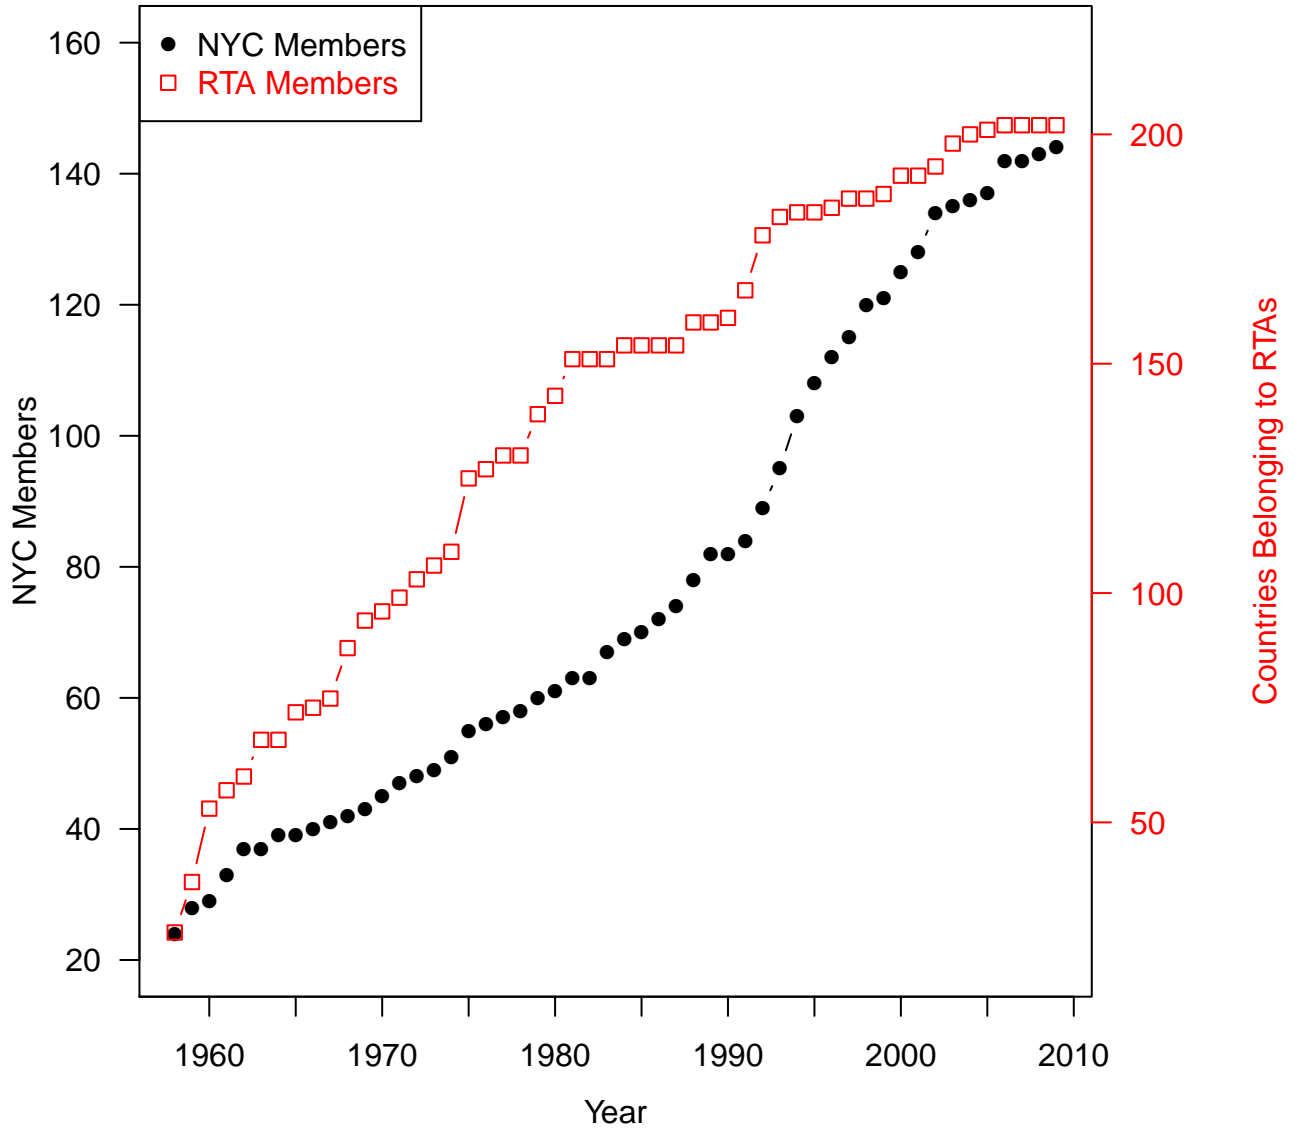
III.B.3 Trade with NYC members

Figure III.B.3 suggests that trade with other NYC members may not be the principal incentive for countries to join the NYC. The left panel in the graph shows the year in which each country joined the treaty on the horizontal axis, and the share of the joining country's trade that is with NYC members in the same year. Trade is the sum of exports and imports. For comparison, a blank circle in each period shows NYC members' share of trade for the average non-member. If countries were joining the treaty because their trade partners were in the treaty, then NYC members' trade shares for joiners in each year should be higher than the average for non-joiners.²⁶

This is not the case, as shown in the right panel of the figure. For example, when China (CHN) joined the NYC in 1987, members of the treaty provided about 65% of its imports and exports. In the same year, 80% of the average non-member's trade was with trade partners like the US that were NYC members. If membership was motivated by trade, some of the other non-members that traded more intensively with NYC partners should have joined before China. In general, the gap between the trade shares of countries that joined,


²⁶This is the intuition behind the semi-parametric linear rank estimator in the duration models of Robins and Tsiatis (1992).

Figure III.B.1: Number of NYC Members Over Time vs RTA Members



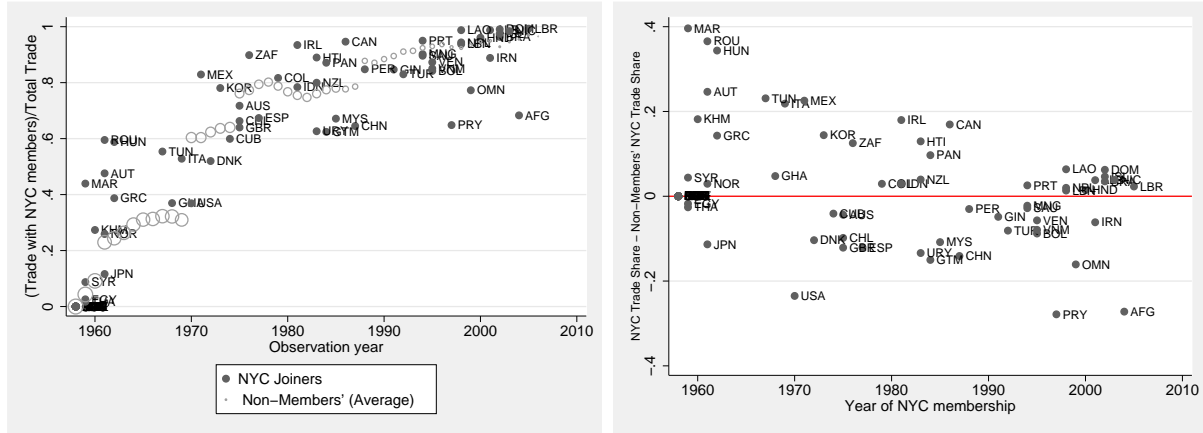
The plot shows the number of NYC members in each year t as circles. The squares show the number of corresponding number of countries belonging to at least one regional trade agreement (RTA).
Data Source: UNCITRAL (2014), D'úr et al. (2014)

Figure III.B.2: Number of NYC Members Over Time vs WTO Members



The plot shows the number of NYC members in each year t as circles. The triangles show the number of corresponding number of countries belonging to the WTO in the same year.
Data Source: UNCITRAL (2014),WTO

Figure III.B.3: NYC Trade Shares of Joiners and Non-Members



(a) The share of trade with NYC members
Data Source: UNCITRAL (2014)

(b) The gap between joiners in each year and other non-members

and the average for countries that remained outside the NYC are almost evenly distributed into negative and positive values.²⁷

Figure III.B.4 relays the same idea differently, while considering two mechanisms – distance and trade. The left panel plots global average minimum distances to the set of countries that are members of the NYC.²⁸ The right panel plots averages of exports to NYC members as a share of total exports.

The plot shows two averages in each year: one for countries that joined the NYC in a particular year, and another for non-members that maintained the status quo. The graph rests on the argument that if proximity-related incentives or mechanisms motivate joining the NYC, then it should spread outwards stepwise on a map from its initial locations. Therefore while the minimum distance to an NYC member should decrease globally over time, and


²⁷The same applies to the US (USA) when it joined in 1970, less than 40% of its trade was with NYC members, while trade with NYC members represented more than 60% for the average non-member. If trade with other NYC members motivated participation, then most countries should be above the red line on the right panel of the graph. Note that the figure is limited to countries that were eligible to join in 1958. (Countries that became eligible after 1958 are not shown because the delay to NYC membership is not comparable with those in the graph).

²⁸ Average distances computed as:

$$DistanceToNYC_t = \frac{1}{N_t} \sum_{i \in N_t} \sum_{j \neq i} distance_{ij} NYC_j$$

where N_t represents the number of joiners or non-joiners in year t , and NYC_j represents

Figure III.B.4: Proximity and Trade as Drivers of NYC Participation over Time



(a) Minimum Distance to NYC Members

(b) Share of Trade with NYC members

Note: The triangles represent average minimum distances to an NYC member (or exports to NYC members as a share of the total) for all countries that remain non-members in a given year, the circles is the average for the set of countries that joined in that year.
 Data Sources: Head et al. (2010), UNCITRAL (2014)

it should be on average lower for the countries that join in any given year. Similarly, the trade linkages of NYC members to countries that join should be stronger than linkages to the average non-member.

The plots do not suggest a dominant role for geographic or trade-driven participation. Of the 47 years between 1959 and 2006, joiners were closer on average and traded more with NYC members in 26 years and 17 years respectively. This is in comparison with other countries in the same that remained non-members. As expected, average minimum distance to an NYC member generally decreases over time, and average share of trade with NYC members gradually approaches 100%. These suggest further inquiry into what motivates participation in the treaty, or whether other variables moderate the effect of trade and distance on treaty participation.

III.B.4 The US case

As the value of state relative to private enforcement increases, one expects domestic business interests to increase pressure on the government for the extension of state enforcement to international agreements. Less clear is why states might resist that pressure or delay action.

One explanation may be a general reluctance to subject domestic legal systems to outside influences; as Trakman (1983: 42) observed, ‘National systems of law remain jealous of their jurisdiction over world trade and hesitate to lose such business to foreign systems.’ Delays may simple be due to the limited decision-making capacities of governments. Ratifying the treaty requires reconciling the NYC with the domestic legal system on matters like time limits for filing complaints, thresholds on the amount at stake, and grounds for refusal or deferral of award enforcement.

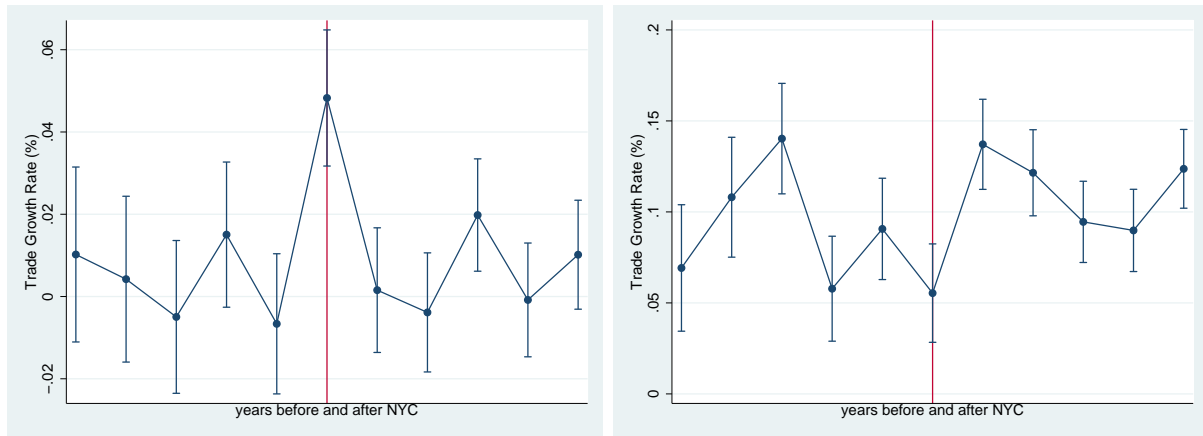
The U.S. experience is illustrative. The U.S. participated in the 1958 conference that established the New York Convention, but did not sign ‘because the American delegation felt that certain provisions were in conflict with some of our domestic laws’ (Congressional Report, 1970). Despite the urging of lawyers and businessmen for early and favorable action, support from the American Bar Association, the American Arbitration Association, the International Chamber of Commerce, the Departments of State and Justice, and the Bureau of the Budget; and assurances that ‘so far as is known, there is no opposition to the bill,’ the bill authorizing accession to the Convention did not arrive in the Senate until April 1968. By a vote of 57 to 0, the Senate approved the NYC in October 1969, but the decision was made to defer ‘depositing the instrument of accession’ until legislation modifying the chapter of the United States Code containing the United States Arbitration Act could be enacted. Asked why it took another four months (to February 1970) for the implementing legislation to be submitted, the State Department explained that an advisory committee on private international law in the State Department suggested changes to the approach for incorporating the convention into the US legal code. The treaty eventually went into effect in September 1970.

III.C NYC Membership’s Effect on Trade

As expected, Figure III.C.1 shows that trade growth on joining the New York Convention is mostly concentrated on the extensive margin. Trade tends to grow for most countries at average rates between 5 and 10%, with most of the growth coming from exports or imports that repeated product-country linkages from the previous year (i.e. the intensive margin). Growth due to any combination of new products or destinations is usually indistinguishable from zero, according to the left panel of the figure.

However the extensive margin experienced a statistically significant leap in the year that the NYC came into force, (for the countries that joined the NYC). Growth rates were

Figure III.C.1: NYC Membership and Trade Growth at the Intensive-Extensive Margins



(a) Growth Rate for Trade at the Extensive Margin

(b) Growth Rate for Trade at the Intensive Margin

The extensive margin represents trade in a combination of SITC product category and country that was new a country in a given year, e.g. in the first year the US exports computers to China, that trade contributes to the extensive margin. The intensive margin includes changes from the previous year's non-zero level of trade.

Data Source: UNCITRAL (2014);Feenstra et al. (2005)

averaged for all countries that joined, with each country's year of entry taken as the reference - or year zero. The graph was prepared using data from Feenstra et al. (2005), which shows trade between countries at the SITC4 product level. The extensive margin here represents a combination of a product and country that was new in a given year. In other words, several new trade relationships tend to coincide with NYC participation. The growth at the extensive margin does not dip significantly below zero in the following year, suggesting that this one-time jump in new export relationships adds to cumulative trade growth. The right panel of the figure confirms the stylized fact that most trade growth is on the intensive margin, but this margin does not react systematically to NYC membership in this assessment. That said, the data suggests that growth at the intensive margin is higher on average for the years after NYC accession.

III.D The NYC Membership Curve is S-Shaped

The next few paragraphs settle questions about the shape of the curve in Figures 3.1 and 3.2. Generally, the true curve could be a straight line, a convex form with monotonically increasing slope or a concave form with monotonically decreasing slope. Otherwise, the sign of the slope changes at some point like an S-curve, for which the slope first increases before

decreasing.

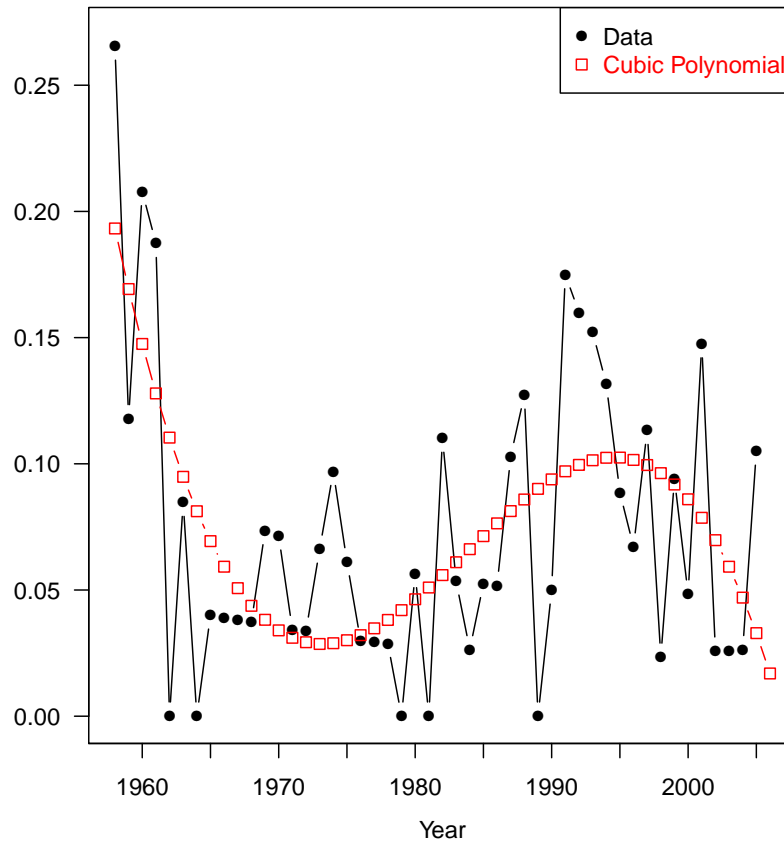
To formally show that the curve is S-shaped beyond the statistical margin of error, Figure III.D.1 plots the relative hazard rate of joining over time.²⁹ The relative hazard rate represents the number of countries that join in each period relative to the interaction between members and non-members, $\frac{(n_{t+1}-n_t)N}{n_t(N-n_t)}$. A plot of this ratio should be flat if the membership curve was a straight line, always sloping downwards with a concave membership curve and always sloping upwards with a convex membership curve. The figure shows a ratio that first dips and increases in later years; supporting the paper's claims of diffusion through interactions.³⁰

I add a cubic spline to the graph. The signs of the splines' first coefficients are in opposite. The statistical significance of these opposite-signed estimates, which fit the data closely, suggest that the curve is indeed S-shaped beyond the statistical margins of error.

²⁹Young (2009) takes the same approach to showing the pattern of hybrid corn adoption by farmers is S-shaped.

³⁰ The relative hazard rate would not change with time if membership was unrelated to the number of non-members or their interactions with members. It would be consistently non-decreasing(non-increasing) if membership was only motivated by the non-decreasing(non-increasing) fraction of countries that are members (non-members).

Figure III.D.1: S-Curve Evidence: Cubic Polynomial Fit of Relative Hazard Rate



The relative hazard rate $H_t = (n_{t+1} - n_t)N/n_t(N - n_t)$ describes the ratio of new members to interactions between members and non-members over time. The cubic spline equation is: $H_t = \alpha_0 + \alpha_1 t + \alpha_2 t^2 + \alpha_3 t^3 + \varepsilon_t$. Here are the estimated coefficients from fitting the spline with OLS:

	Coefficient	t-statistic
α_0	0.193	7.49
α_1	-0.025	-5.24
α_2	0.001	4.91
α_3	-0.000015	-4.52

Data Source: UNCITRAL (2014)