

October 2017

Essays in International Economics: Decomposing Episodes of Large Growth in International Trade

Brandon K. Malloy

The University of Western Ontario

Supervisor

Dr. James MacGee

The University of Western Ontario

Joint Supervisor

Dr. Ananth Ramanarayanan

The University of Western Ontario

Graduate Program in Economics

A thesis submitted in partial fulfillment of the requirements for the degree in Doctor of Philosophy

© Brandon K. Malloy 2017

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>

 Part of the [International Economics Commons](#), and the [Macroeconomics Commons](#)

Recommended Citation

Malloy, Brandon K., "Essays in International Economics: Decomposing Episodes of Large Growth in International Trade" (2017). *Electronic Thesis and Dissertation Repository*. 4927.

<https://ir.lib.uwo.ca/etd/4927>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact tadam@uwo.ca.

Abstract

My thesis consists of three chapters relating to topics in International Economics.

In the first essay, I use bilateral trade data from Canada, Germany, Japan, Mexico, the U.S. and the U.K. to decompose the patterns of trade growth across various goods classifications during episodes of rapid growth in bilateral trade. I find that bilateral trade growth during these episodes is granular — less than 5% of goods classifications account for over 65% of overall bilateral trade growth. I quantitatively assess whether “Melitz-style” trade models, with heterogeneous productivity firms, CES demand and fixed and variable costs of exporting, can match the observed granularity of bilateral trade growth. I find the standard model generates only 10% of the observed granularity in the data, as measured by the share of total trade growth accounted for by various quantiles of goods classifications. However, by incorporating heterogeneous productivity changes and tariff reductions imputed from the U.S. production and export data, I find that the model generates roughly 70% as much granularity of trade growth across goods as in the data.

When firms export their goods to foreign markets, they often choose between multiple distribution technologies in transporting their goods to their final destination. The second essay extends the standard trade model by incorporating a choice among two distribution technologies in the exporting process

— one low-fixed, high-variable cost method, and one high-fixed, low-variable cost method — and assessing the implications for trade growth across goods. In this model, I find that heterogeneous productivity or tariff changes may lead firms to “switch” their optimal distribution method — from not-traded to traded, or from the low-fixed cost to the high-fixed cost technology. This results in disproportionately larger trade growth for these types of firms, since they benefit from a double reduction in the variable costs of exporting — the direct effect of the fall in trade costs, and the indirect effect of switching to a lower variable cost distribution method. Calibrating this model to bilateral trade flows, I find that model simulations with multiple distribution technologies generate up to 90–95% of the granularity in trade growth observed in the data.

The third essay examines the role of variation in transportation options — what I denote the “supply network” — on observed price differences between locations for a specific good, retail gasoline. I use a unique data set of weekly gasoline prices across 44 Canadian cities to analyze how the existence of variation in the available modes of transportation for gasoline between cities (via pipeline, marine tanker, rail or truck) accounts for observed price differences across locations. I find that the supply network is significant — cities connected by lower per-unit cost methods like pipelines or seaports exhibit smaller mean- and weekly-price differences than those connected only by road or rail, after controlling for variables such as distance, regional effects and market size.

Acknowledgements

I cannot adequately express my gratitude to my supervisors, with particular thanks to Jim MacGee, without whose wisdom and expertise I would never have had the ability to conduct this work. Their patient support, feedback and critiques provided countless instances of guidance along the path of learning that helped me complete this manuscript. I am truly indebted to their service.

I must also thank a large number of people whose comments and suggestions have helped shape my work at various stages. First and foremost, I wish to express my deepest thanks to the members and participants in the Economics Workshops at the University of Western Ontario, particularly the Macroeconomics seminar and “Brown Bag lunch” groups. Jim MacGee, Igor Livshits, Ananth Ramanarayanan, Simona Cociuba, Betsy Caucutts and Jacob Short, among others, provided invaluable insights and comments that helped shape my work and bring attention to critical issues that might provide fruitful extensions. I further wish to extend my thanks to other members of the faculty, particularly Jim Davies, Charles Saunders and Al Slivisnki, for similarly helpful comments and ideas along the way.

I must also acknowledge the valuable feedback provided by numerous attendees of the 2015-2017 Canadian Economics Association’s Annual Meetings, who allowed me to discuss my work with them and offered concise but welcome suggestions and comments that helped shape each of my chapters.

Personal Acknowledgments:

I would be remiss not to acknowledge the contribution of my colleagues and classmates over the years, who have help me along my academic path. I am forever grateful to my cohort, Peter Foltin, Zinaida Foltin, Youngmin Park, Hiroaki Mori, Jin Zhou, Selina Gonzalez, James Partridge, Andrew Agopsowicz, Aaron Gertz, Sunghoon Cho, Jon Rosborough and Zeynep Ozkok, among others, for their counsel, academic assistance and most importantly their friendship and support along the journey. I cannot imagine a better environment to approach the graduate experience than with friends and colleagues like these.

I must similarly acknowledge the amazing academic environment that is the Economics Department at the University of Western Ontario. From administrators Yvonne Adams, Sandra Augustine, Jennifer Hope and Sharon Phillips, to the entire teaching Faculty, particularly Igor Livshits, Ananth Ramanarayanan, Peter Streufert, Maria Goltsman and John Knight, everyone involved in the department served admirably to honor the tradition of learning and excellence exemplified by the University of Western Ontario.

On a personal note, I wish to extend my eternal thanks to Jim MacGee, who served as my supervisor, advisor, role model, and voice of reason in the pursuit of an academic career. I simply cannot fathom undertaking such an endeavor without his unwavering expertise and guidance, in both my work and my approach to an academic career. Finally, I am forever thankful for the love and support of my parents, Don and Jean Malloy, who have encouraged me in all my pursuits and provided that most supportive family environment I could have ever asked for. I could not have completed this manuscript without them.

London, Ontario - 2017

Brandon Malloy

To my parents, Don and Jean Malloy

Table of Contents

Certificate of Examination	ii
Abstract	iii
Acknowledgements	v
Table of Contents	viii
List of Figures	x
1 Introduction	1
1.1 References	5
2 Decomposing Episodes of Large Growth in International Trade	6
2.1 Introduction	6
2.2 Related Literature	14
2.3 Data	17
2.3.1 Stylized Facts	17
2.3.2 Data Sources	19
2.3.3 Granularity of Trade Growth	23
2.3.4 Tariff Rates	30
2.3.5 Production Growth and Trade Intensity	34
2.3.6 Extensive Margin Growth	40
2.4 Model	41
2.4.1 Model Set-up	44
2.4.2 Characterizing Equilibrium	45
2.4.3 Trade Liberalization	47
2.5 Numerical Exercises: Trade Growth	51
2.5.1 Parameterization and Calibration	52

2.5.2	Standard Model with Uniform Tariff Reductions	54
2.5.3	Standard Model with Heterogeneous Productivity Changes and Tariff Reductions	56
2.5.4	Quantitative Analysis	60
2.6	Conclusions	63
2.7	References	64
3	The Role of Multiple Distribution Technologies in Accounting for Trade Growth	67
3.1	Introduction	67
3.2	Related Literature	75
3.3	Data	78
3.3.1	Bilateral Trade Data	79
3.3.2	Production Data	82
3.3.3	Tariff Data	84
3.4	Model	85
3.4.1	Model Set-up	87
3.4.2	Characterizing Equilibrium	89
3.4.3	Comparative Statics: Tariff and Productivity Changes . .	92
3.5	Quantitative Analysis	95
3.5.1	Parameterization and Calibration	97
3.5.2	Initial Period	98
3.5.3	Final Period	99
3.6	Conclusions	103
3.7	References	104
4	The Supply Network and Price Dispersion in the Canadian Gasoline Market	108
4.1	Introduction	108
4.2	Related Literature	112
4.3	Supply Network	115
4.4	Data	121
4.4.1	Descriptive Statistics from Pre-Tax Data	122
4.5	Empirical Analysis	129
4.5.1	Mean Price differences	133
4.5.2	Weekly relative prices	135
4.5.3	Robustness	136
4.5.4	Summary: Impact of Supply Network on Price Dispersion	139
4.6	Case Studies of Supply Shocks	140

4.6.1	Stylized Facts from the data	142
4.6.2	Retail prices	149
4.7	Conclusions	150
4.8	References	152
5	Conclusion	155
5.1	References	157
A	Chapter 2 Appendix	159
B	Chapter 3 Appendix	164
C	Chapter 4 Appendix	168
C.1	Data Sources and Implementation	168
C.2	Regression Tables	170
	Curriculum Vitae	176

List of Figures

2.1	Growth by Good: 5-digit SITC bilateral trade	25
2.2	Growth by Good: 5-digit SITC bilateral trade	26
2.3	Large-growth goods: Imports vs. Tariff changes	32
2.4	Large tariff decrease goods: Imports vs. Tariff changes	33
2.5	Decile Growth (6-digit NAICS) bilateral trade	36
2.6	Trade Intensity Share vs. Share of Trade Growth	38
2.7	Production Growth Share vs. Share of Trade Growth	39
2.8	Decile Growth: 5-digit SITC bilateral trade	42
2.9	Standard model with uniform tariff reductions	48
2.10	Standard Model with uniform tariff reductions (Can-Mex)	55
2.11	Productivity Distributions: Productivity Shocks	57
2.12	Standard Model with Heterogeneous Productivity Changes and Tariff Reduction	59
2.13	Histograms of Trade Growth by Goods	61
3.1	Growth by Goods Category: US Exports to Canada	82
3.2	Productivity Distributions: Implied Productivities	83
3.3	Heterogeneous Tariff Changes: Canada	84
3.4	Initial level of export sales	91
3.5	Final level of export sales	93
3.6	Model Simulation: Optimal Distribution Choice	102
4.1	North American Oil Pipeline Infrastructure	117

4.2	Weekly Price Correlations: Levels	123
4.3	Weekly Price Correlations: Weekly Changes	124
4.4	Retail Price Correlations vs. Mean Price Differences: National . .	125
4.5	Coefficient of Variation: National	127
4.6	Max/Min histogram	128
4.7	Min/Max spread and Mean Price	128
4.8	Calgary Planned shutdown	144
4.9	Nanticoke Unplanned shutdown	145
4.10	Refinery Inventories	147
4.11	Refinery Net Imports	148
A.1	Worldwide Trade: 1960-2015	159
A.2	Growth by Good: Least-Traded Goods (Canada-Mexico)	160
A.3	Growth by Good: Mid-Traded Goods (Canada-Mexico)	160
A.4	Growth by Good: Most-Traded Goods (Canada-Mexico)	161
A.5	Correlation: Imports vs. Tariff changes (U.S.-Canada)	161
A.6	Correlation: Imports vs. Tariff changes (U.S.-Mexico)	162
A.7	Correlation: Imports vs. Tariff changes (U.S.-U.K.)	162
A.8	Growth by Good: 6-digit NAICS bilateral trade	163
B.1	Growth by Goods Category: US Exports to Mexico	164
B.2	Growth by Goods Category: US Exports to Japan	165
B.3	Productivity Distributions: Heterogeneous Productivity Changes	165
B.4	Heterogeneous Tariff Changes: Mexico	166
B.5	Heterogeneous Tariff Changes: Canada	166
B.6	Heterogeneous Tariff Changes: Japan	167

Chapter 1

Introduction

My thesis consists of three essays in the field of international economics. The first two essays examine trade growth during periods of large growth in international trade over the past 60 years. I use bilateral trade data to decompose trade growth across goods classifications during rapid growth episodes and find that trade growth is granular — the majority of overall trade growth between given country-pairs is accounted for by a small number of goods classifications exhibiting large growth in trade, while most goods exhibit small or negligible growth. In the first essay (Chapter 2), I ask whether standard “Melitz-style” international trade models, when supplemented with productivity and tariff changes imputed from trade data, can generate the observed level of granularity in bilateral trade data. In the second essay (Chapter 3), I add a choice among multiple distribution technologies for exporting firms, with differing fixed and variable costs, to quantify the model’s ability to match the observed granularity of trade growth. The third essay (Chapter 4) examines how variation in transportation methods, which I denote “the supply network”, impacts price dispersion across locations. I use a unique data set on weekly gasoline

prices across 44 Canadian cities to quantify how differences the available methods of transporting gasoline between locations (by pipeline, marine tanker, rail or truck) account for observed mean and weekly price differences across cities.

International trade has grown nearly twice as fast as world GDP over the past 60 years. For most countries, this growth in trade has occurred in a small number of rapid growth episodes. Standard international trade models have significantly under-predicted this large growth in international trade during these periods of rapid growth. As such, recent Melitz-style trade models have investigated the role of heterogeneity in trade growth across various industries and goods. Chapter 2 analyzes bilateral trade data, at the 5-digit SITC classification level for Canada, Germany, Mexico, Japan, U.S., and U.K. (accounting for 20–25% of global trade) between 1989 and 1999, in order to decompose the patterns of trade growth across various goods classifications. I find that bilateral trade growth during these rapid growth episodes is granular — less than 5% of goods classifications account for over 65% of overall bilateral trade growth while the majority of goods exhibit little to no growth in trade.

Trade theory suggests that disproportionately large growth in this small number of goods categories may be accounted for by the goods experiencing the largest reductions in tariff rates. However, I find that tariff reductions cannot account for the granularity of bilateral trade growth. Further, I find that production growth and increases in trade intensity (defined as the share of domestic production that is exported) are both significant factors in these large-growth goods classifications.

I use a Melitz-style trade model, featuring heterogeneous productivity firms

facing CES demand under monopolistic competition, calibrated to match bilateral trade flows, to quantify the predictions for trade growth across goods classifications during growth episodes. I find that for reasonable parameter values, the model predicts less granular trade growth, generating only about 10% of the granularity in the data, as measured by the share of total trade growth accounted for by various quantiles of goods classifications. To better match the observed granularity, I augment this standard model by incorporating heterogeneous productivity changes and tariff reductions imputed from the U.S. production and export data. Doing so, I find that the model predicts roughly 70% of the granularity across goods observed in the trade data.

To account for the remaining heterogeneity of trade growth in the bilateral trade data, I examine the role of heterogeneity in exporting methods in accounting for differences in observed levels of trade and trade growth across goods. Firms often choose between alternate methods of exporting their goods to foreign countries. Generally, these can be grouped into two broad categories — methods with high fixed costs and low variable costs, and those with low fixed costs and high variable costs.

In Chapter 3, I use the standard Melitz-style model from Chapter 2, and add a choice among multiple distribution technologies for exporting firms — one low-fixed, high-variable cost option, and one high-fixed, low-variable cost option. Solving the model, I find that, following productivity or tariff changes, firms that switch from not-traded to traded, or from the low-fixed to high-fixed cost method exhibit disproportionately larger growth than non-switching firms, generating higher granularity in trade growth in the model. To quantify the effect of this mechanism on trade growth granularity, I calibrate and simulate the model, incorporating heterogeneous productivity and tariff changes

imputed from trade data, to match data on bilateral trade flows and trade growth. I find that the model with multiple distribution technologies increases the predicted granularity of trade growth across goods to 90–95% of the level observed in the data, as opposed to the 70% generated by a model with a single distribution technology in Chapter 2.

Chapter 4 investigates how variation in available transportation methods impacts observed price differences across locations in the Canadian gasoline market. Price dispersion is often attributed to transportation costs — the larger the costs of transporting goods between locations, the larger the price gaps that can be sustained over time. Many studies use geographic distance as a proxy for transportation costs.¹ However, little is known about the quantitative impact of variation in the methods used to transport goods between locations on these relative price differences.

I use a unique data set on weekly average gasoline prices in 44 Canadian cities between 2001 and 2017, as well as differences in the four main modes of transporting gasoline — via pipeline, marine tanker, rail or truck — to quantify the impact of the supply network on relative price differences between locations. Controlling for distance, regional and market effects, I find that the supply network has a significant impact on price dispersion in the Canadian gasoline market. City-pairs connected via pipeline — a faster, lower cost-per-unit method — exhibit 3.5% less mean-price dispersion than those connected by higher cost-per-unit methods like rail or truck. Further, the existence of pipelines connecting cities has the effect of reducing weekly price differences by the equivalent of a 53% reduction in geographical distance, while a seaport connection between cities reduces the effective distance by 38%, compared to

¹See, for example, Burdett and Judd (1983), Crucini, Telmer and Zachariadis (2003), or Engel and Rogers (1996).

land-route alternatives.

A brief case study of supply disruptions at the refinery level indicates that unplanned refinery shut-downs result in price spikes that are higher in regions closest to these supply disruptions, and that retail price shocks have lower variation in cities that share a pipeline connection than in those that do not. These results reinforce the finding that the structure of the supply network is significant in accounting for observed price dispersion across locations.

1.1 References

- BURDETT, K. AND K.L. JUDD (1983). “Equilibrium Price Dispersion”, *Econometrica*, Vol. 51, No. 4, pp. 955–969
- CHANEY, T. (2008). “Distorted Gravity: The Intensive and Extensive Margins of International Trade”. *American Economic Review*, 98:4, 1707–1721.
- CRUCINI, M., C. TELMER AND M. ZACHARIADIS (2003). “Price Dispersion: The Role of Borders, Distance and Location”, *Tepper School of Business*, Paper 490
- ENGEL, C. AND J. ROGERS (1996). “How Wide Is the Border?”, *American Economics Review*, Vol 86 No 5, pp. 1112–1125
- MELITZ, M.J. (2003). “The Impact of Trade on Intraindustry Re-allocations and Aggregate Industry Productivity.” *Econometrica*, Vol. 71.6 (2003): 1695–1725.

Chapter 2

Decomposing Episodes of Large Growth in International Trade

2.1 Introduction

Over the past 50 years, international trade has grown nearly twice as fast as world GDP.¹ For most countries, this growth has not been smooth and constant over time, but rather occurred over a small number of rapid growth episodes.² For some countries, episodes of large growth follow implementations of trade liberalization, such as the North American Free Trade Agreement (NAFTA) in the early 1990s. However, episodes of large trade growth also occur in the absence of formal trade liberalizations.³

It is well-documented that in many instances, trade models significantly

¹Source: IMF WEO 2012
available at: <http://www.imf.org/external/pubs/ft/weo/2012/02/weodata/index.aspx>.

²Figure A.1 shows bilateral trade for several countries over the past 30–60 years, where an average of roughly 70–90% of overall growth over this period is accounted for by periods spanning only 10% of the time span (i.e. 5 years).

³For example, between 1990 and 1999, Mexican exports to Canada quintupled, from \$0.4 billion to \$2.1 billion. By contrast, Mexico's GDP grew by only 70% during this same period. Over the same period, German exports to the U.S. grew from \$25 billion to \$49 billion while German GDP grew by only 50%.

under-predicted the magnitude of trade growth during these episodes.⁴ More recently, papers like Melitz (2003) and Chaney (2008) have focused on the role of heterogeneity in accounting for trade growth across industries and goods. As a result, trade literature has examined responses to trade liberalization across various types of goods and industries: previously-traded goods that become traded in larger values — the *intensive* margin — and previously not-traded goods that become traded — the *extensive* margin.⁵

Several questions remain to account for the large growth during these rapid growth episodes:

1. **What is the distribution of growth across goods during growth episodes?** What proportion of overall trade growth is attributable to intensive margin growth as opposed to extensive margin growth? Kehoe and Ruhl (2013) find evidence that the extensive margin is in fact large and significant in accounting for overall trade growth. Further, within these margins, is trade growth widely dispersed across a large number of different goods, or is trade growth *granular* — i.e. is trade growth concentrated in a small number of goods accounting for the majority of overall trade growth?⁶

⁴For example, see Kehoe(2005) for a detailed examination of the erroneous predictions of trade models following NAFTA.

⁵Although the terms “goods” and “industries” may have different connotations in other economic literature, for brevity I will hereafter use the term “goods” to refer to various industries and the goods they produce and trade.

⁶The use of the term “granular” in the literature is relatively recent and sparse — papers like Gabaix (2011), DiGiovanni and Levchenko (2014) and DiGiovanni, Levchenko and Mejean (2017) use granularity to refer to the “incompressible grains of economic activity” that result from the idiosyncratic shocks to the upper end of the fat-tailed distribution of firms within an industry, that pervade to aggregate economy-level shocks. For the purposes of this dissertation, I extend my connotation of granularity to include two main properties: (1) the level of disaggregation in cross-sectional trade that identifies idiosyncratic behavior across industries in a given period, similar to Gabaix and others (as opposed to homogeneous behaviors across industries or focusing on aggregate values, and (2) idiosyncratic behavior in trade *growth* across industries

2. What accounts for heterogeneity of growth across different goods?

Do tariff reductions account for overall trade growth for each good? Is trade growth a function of production growth or increases in trade intensity (or both)? That is, does trade growth for each good coincide with increases in domestic production, with a constant share of output being traded, or does the share of domestic output being traded increase?

3. Are the predictions of Melitz-style models consistent with the level of granularity observed in international trade data? Do the mechanisms for tariff reductions and productivity changes in these models deliver similar patterns of trade growth across goods as observed in the data?

In this paper, I document new data facts on bilateral trade growth across goods during rapid growth episodes. I analyze 5-digit Standard International Trade Classification (SITC) bilateral trade data, between 1989 and 1999, for 6 countries (Canada, Germany, Japan, Mexico, USA and UK) accounting for 20–25% of global trade during this period. Although data limitations restrict the sample to this 10-year period, it does include the implementation of the North American Free Trade Agreement (NAFTA) and the Canada-US Free Trade Agreement (CUSFTA) in the early-to-mid 1990s, thus allowing a comparison of growth for trade-liberalization and non-liberalization country-pairs.⁷ I supplement this data with 8-digit Harmonized Tariff Schedule (HTS) data

— industries react differently across the distribution of trade, even from similar previous levels of trade. The latter property contrasts with the connotation of terms like “lumpiness” of trade (as in Armenter and Koren (2014), among others) which denotes the concentration of trade and trade growth at the upper end of the distribution, where the majority of trade growth would similarly be accounted for by those industries accounting for the majority of previous trade.

⁷Numerous papers, such as Caliendo and Parra (2014), Gould (1998), Romalis (2007), etc., examine the overall impact of NAFTA and CUSFTA, providing a large, diverse literature.

on U.S. tariff rates and 6-digit North American Industry Classification System (NAICS) data on U.S. manufacturing production. The U.S. tariff data is matched to U.S. import data to analyze the impact of changes in tariff rates on the growth in trade across goods classifications. U.S. manufacturing data is matched to U.S. export data to decompose trade growth into changes in production and changes in trade intensity.

I document four key facts during these rapid trade growth episodes:

1. Bilateral trade growth is granular across time and across goods — over 70% of total trade growth over time is accounted for by 10% of the total time period (i.e. 5 of 50 years), and less than 5% of goods classifications account for 65–110% of total bilateral growth across country-pairs⁸
2. Changes in tariffs do not account for episodes of large trade growth
3. Increases in trade intensity and domestic production each account for large growth in trade across goods
4. The extensive margin is significant for overall bilateral trade, but is driven by a relatively small number of extensive margin goods (<20%) that account for the majority of overall extensive margin growth

Across all bilateral country-pairs, trade growth is granular across time and goods. Over the past 50 years, most of the total growth in bilateral trade is accounted for by a small number of rapid growth episodes. Examining growth for each year, the 5 largest-growth years (i.e 10% of the total time span) account for over 70% of the total growth in bilateral trade over the past 50 years.

⁸These goods can account for over 100% of trade growth due to negative growth in some goods classifications.

Cross-sectionally, the majority of total bilateral trade growth is accounted for by a small number of goods classifications exhibiting disproportionately large growth.⁹ Nearly all these large-growth goods had non-zero levels of trade prior to rapid growth. Large growth from goods categories with zero previous trade is rarely observed. Further, among this set of large-growth goods, there is large variation in their level of trade prior to rapid growth.

Reductions in trade barriers alone are unable to account for these large-growth goods categories which comprise the majority of trade growth across country-pairs. Many episodes of trade growth occur in goods whose tariff rates do not change, while most changes in tariff rates do not result in large trade growth for those goods. As a result, there is no statistically significant correlation between reductions in U.S. tariffs rates and U.S. bilateral trade growth across goods categories. Of the 100 goods contributing the largest shares of bilateral trade growth, over half exhibit no reduction in their ad valorem equivalent (AVE) tariff rate. Further, of the goods with the largest reductions in tariff rates, few (10–20%) exhibit substantial growth in trade.

Examining U.S. manufacturing data, I find that episodes of large growth in bilateral trade coincide with increases in both trade intensity and production.¹⁰ Controlling for changes in domestic production, observed increases in trade

⁹On average, across country-pairs, roughly 1% of goods classifications account for approximately 60% of total trade growth.

¹⁰To clarify this distinction, global growth of new technologies like smart-phones and laptop computers, employed on a global scale, necessitates increased production of semi-conductors, the key foundation of internal circuitry for modern electronics. A large increase in exports for goods like semi-conductors may therefore be proportional to increases in their overall production, with trade intensity (the share of domestic production that is exported) remaining relatively constant. Alternatively, episodes of trade liberalization resulting in decreased barriers to trade may lead to a larger share of domestic production being exported to a given destination, representing an increase in trade intensity, even in the absence of large increases in gross production.

intensity account for 35–90% of total bilateral trade growth across country-pairs.¹¹ Conversely, holding trade intensity constant, production growth accounts for 30–40% of observed trade growth for some country-pairs, while for others it accounts for virtually none of the overall bilateral trade growth.

Extensive margin growth is measured as changes in trade accounted for by the goods classifications representing the bottom 10% of initial trade.¹² I find extensive margin growth is significant to overall bilateral growth for all country-pairs. Additionally, within the extensive margin, large growth in a small number of goods accounts for the majority of overall extensive margin growth, mirroring the granularity of trade growth across goods in the intensive margin, for trade-liberalizing and non-liberalizing countries alike.

I quantitatively compare these empirical findings on trade growth across goods to the predictions of Melitz-style trade models. As in Melitz(2003), Helpman, Melitz and Yeaple (2004) or Chaney (2008), these models typically feature heterogeneous productivity firms, monopolistic competition pricing, and CES demand, as well as fixed and variable costs of exporting. Since I am interested in the quantitative results, I use reasonable parameter values and calibrate this “standard model” to match observed trade flows and quantitatively assess the model’s ability to match the observed granularity in bilateral trade data for two main cases:

¹¹Due to the data limitations of only having U.S. production data, I only consider 5 country pairs — U.S. exports to Canada, Mexico, Japan, Germany and the U.K.

¹²This paper adopts the Kehoe and Ruhl (2013) definition of the extensive margin—the set of goods that account for the bottom 10% of initial trade. This is due to reporting issues in international trade—shipments with sufficiently low values need not be declared on Customs reports for many countries, so distinguishing true zero-trade goods from small trade goods can cause issues. For similar reasons, I follow the Kehoe and Ruhl approach of using the bottom decile of goods, according to initial trade value to represent the extensive margin.

1. **Standard model with uniform tariff reductions.** A standard Melitz-style model featuring firms with fixed productivities drawn from a Pareto distribution, where trade liberalization takes the form of a uniform reduction in tariffs across all goods, a common approach in the trade literature.
2. **Standard model with heterogeneous productivity changes and tariff reductions.** A standard Melitz-style which I augment by incorporating heterogeneous changes in productivity across goods imputed from U.S. production data, as well as heterogeneous tariff reductions across goods matched from U.S. import data.

I find that the predictions of the standard model with uniform tariff reductions do not match the stylized facts observed in the data. With a Pareto distribution of productivities, and a fixed and variable cost of exporting, this model produces much less granularity of trade growth than found in the data. In equilibrium, goods are stratified into traded goods (those with sufficiently high productivity to cover the fixed cost of exporting) and goods that are not-traded (i.e. those where the fixed cost exceeds the potential profits from exporting). Following a uniform reduction in trade costs for all goods, the increase in exports of previously traded is proportional to their productivities. As a result, growth in the intensive margin of trade is smooth across previously-traded goods, not granular. Extensive margin growth arises from a shift in the productivity cut-off, resulting in some goods that were previously not-traded becoming traded, as the reduction in variable cost makes it profitable to pay the fixed cost of exporting.

The key result delivered by this standard mechanism is the smoothness of trade growth — exports grow proportionally to the uniform reduction in

trade costs and proportionally to their productivity. The only disproportionate growth in the standard model comes from extensive margin goods, jumping from zero trade to non-zero levels. However, these goods become traded in such small values (as a result of their relatively low productivities), that they account for a small proportion (<20%) of overall trade growth. This implies that across goods, goods with similar initial levels of trade are predicted to grow in similar magnitudes, and trade growth is only slightly more granular than cross-sectional trade in any given period. Conversely, in the data trade growth is highly granular and trade growth is substantially more granular than cross-sectional trade.

Matching both the levels of cross sectional bilateral trade and the observed level of granularity from the data is important for several reasons. For policy analysis, it is necessary to identify the disparity in trade growth across industries so as to assess the welfare impact of trade policy across income distributions. It is important that the model capture the granularity of observed trade to provide insight into heterogeneous responses across industries to policy proposals such as trade liberalization. Matching the granularity of trade across goods may also be important for quantifying the impact of shocks (both cyclical and secular) on trade flows, sectoral output dynamics and the reallocations of inputs across industries.

To attempt to better match the level of granularity observed in the data, I introduce into the model heterogeneous tariff and productivity changes imputed from U.S. manufacturing and export data. Heterogeneous productivity changes across goods allow for large jumps in productivity for some goods classifications that generates disproportionately large trade growth for these goods. Heterogeneous tariff reductions similarly allow for disproportionately large growth

in the goods exhibiting the largest tariff reductions, while goods experiencing smaller reductions in trade costs exhibit much smaller growth.

To quantify this mechanism, I calibrate the standard model including heterogeneous productivity and tariff changes to match U.S. export data and compare the granularity of trade growth across goods in the model to that in the data. I find that incorporating these productivity and tariff changes significantly increases the granularity of trade growth predicted in the model. Measuring the proportion of overall bilateral trade growth accounted for by various quantiles of large-growth goods, I find that this augmented model generates roughly 70% of the granularity observed in the trade data, a marked improvement from the 10% generated by the standard model with uniform tariff reductions.

2.2 Related Literature

Several empirical papers document the “lumpiness” of trade — a majority of international trade in a given period is accounted for by a relatively small number of goods traded in large volumes, while a large proportion of domestically-produced goods are not exported. Armenter and Koren (2010) suggest that the concentration of large amounts of total bilateral trade in a small number of goods categories may be systematic of the sparseness of trade data — in a given period, the fact that there are relatively few international shipments (balls) compared to the large number of potential goods classifications (bins) necessarily results in this inherent “lumpiness” of trade across goods categories, with a large number of goods being not-traded or traded in very small amounts, while the majority of bilateral trade is concentrated in a small number of large-trade

goods. However, Blum, Claro and Horstmann (2016) argue that a key underlying assumption driving the balls and bins model — that export shipment size is randomly allocated across good categories, and is independent of firm size — is inconsistent with shipment data and invalidates the Armenter and Koren model’s findings. Alternatively, Fernandes et al (2015) argue that the granularity of trade growth, particularly on the intensive margin, may arise from a log-normal productivity distribution, as opposed to the Pareto distribution commonly assumed in Melitz-style trade models, which causes firms to react disproportionately in response to trade liberalization that lowers variable costs of exporting. Alessandria, Kaboski and Midrigan (2010) attribute the lumpiness of trade to economies of scale in transportation and delivery lags. They find that large costs associated with international transportation of goods leads firms “stock up” with larger and less frequent shipments, accounting for the lumpiness of trade across goods and time.

While these papers focus on the patterns of cross-sectional trade, a central contribution of this chapter is to document the fact that trade *growth* is granular. Further, I find that the set of large-growth goods accounting for the majority of trade growth is uncorrelated with the set of goods in which per-period trade is concentrated. Finally, I find that trade growth is more granular than cross-sectional trade, and that the level of trade growth across goods is uncorrelated with initial levels of trade.

A large literature (e.g. Krugman (1979), Lancaster (1980), Deardorff (1984), etc) documents recent growth in international trade. Most germane to this chapter is the documented fact that the ratio of trade to GDP has increased over the past 50 years. Bergoeing and Kehoe (2001) document key trade growth

facts and investigate the ability of standard trade models to quantitatively capture the patterns observed in international trade data. This chapter documents that this trade growth is highly concentrated in a small number of rapid growth episodes over this period.

This chapter is also related to empirical work identifying the effects of trade liberalizations on bilateral trade flows. Romalis (2007) uses a difference-in-difference approach to exploit variation between trade liberalizing and non-trade liberalizing countries to estimate elasticities of traded goods. He reports that trade liberalization has a significant impact on volume, but small impact on prices and welfare. This chapter extends these data facts by investigating trade growth variation across individual goods, and develops a model that can deliver the patterns of trade observed in the data, which standard models do not produce.

Kehoe and Ruhl (2013) find that the extensive margin plays a significant role in accounting for bilateral trade growth, with a 10% increase in trade between country-pairs being accompanied by a 36% increase in the extensive margin, on average.¹³ While I find similar results on the impact of the overall extensive margin, this chapter adds to Kehoe and Ruhl by decomposing extensive margin growth between goods that were previously not traded and those that were traded in very small amounts. I document that extensive margin growth is quite granular, with less than 5% of extensive margin goods accounting for over 25% of growth in the extensive margin.

Many papers have used standard international trade models to examine the impact of trade liberalizations (and other structural change) on trade volumes

¹³Kehoe and Ruhl's interpretation of the extensive margin varies from the theoretical definition identified in Chaney (2008) — Kehoe and Ruhl classify the extensive margin as the growth in trade among the set of *least-traded goods*, classified as the bottom decile of goods sorted by initial trade value.

and the patterns of trade growth. However, these models do not produce the high level of granularity of trade growth observed in the data. Melitz (2003) introduces heterogeneous productivity across firms with a fixed cost of entry to a model with monopolistic competition, which predicts that intensive margin growth is smooth across goods. Chaney (2008) builds on the Melitz framework to isolate the role of the extensive margin in trade growth, identifying the set of goods that enter the exporting market following reductions in trade barriers. However, in this framework, extensive margin growth is smooth across extensive margin goods. Arkolakis (2010) builds a model of market penetration, in which firms essentially choose their fixed cost in order to penetrate a market and then face increasing marginal costs to reach additional consumers. However these models, whether employing a single fixed-cost export technology (as in Melitz, Chaney), or a continuum of fixed cost options (as in Arkolakis), do not produce the non-convexity of trade growth across goods observed in the data. This chapter extends the literature by adding heterogeneous productivity and tariff changes to investigate whether this standard framework can generate the level of granularity of trade growth observed in the data across both the intensive and extensive margins.

2.3 Data

2.3.1 Stylized Facts

I use data on bilateral trade values, ad valorem tariff rates, and manufacturing data to decompose trade growth between the intensive and extensive margins across goods classifications. Four key facts emerge from the data:

- 1. Trade growth is highly granular across time and across goods**

classifications.

- a. Across countries, bilateral trade growth over time occurs in short periods of rapid growth. Rather than smooth, consistent growth over time, growth is concentrated in a small number of rapid growth episodes. Over the past 50 years, the 5 largest-growth years account for roughly 75% of overall trade growth.
- b. Across goods classifications, trade growth is concentrated in a small number of goods classifications, with 5% of classifications accounting for roughly 65–110% of total bilateral trade growth by country-pair. In all bilateral pairs, there are many goods classifications that begin with zero trade, and remain so over time.¹⁴ Goods that switch from zero reported trade to positive trade remain at low values. The majority of goods that are initially traded grow very little — it is only a small number of goods, growing from low levels of initial trade to high levels of final trade, or from high levels of initial trade to very high levels of final trade, that account for the majority of the growth in bilateral trade.

2. Reductions in tariffs coincide with only 10–20% of the large-growth goods.

Tariff reductions do not account for the granularity in trade growth. Among the subset of large-growth goods which account for the majority of trade, large decreases in tariff rates accompany only a small number (10–20%) of these goods, while the remaining goods exhibit no change or small increases in their ad valorem equivalent tariff rate.

¹⁴While this property is consistent across country-pairs, it is not true that it is the *same* set of goods with zero trade across all country pairs. One good with zero trade for one country-pair may be highly traded across other country-pairs, and vice-versa.

3. **Trade intensity and production growth both account for large growth in bilateral trade**

Controlling for growth in production across goods classifications, increases in trade intensity account for 35–90% of total observed trade growth across various U.S. bilateral pairs. Conversely, fixing trade intensity at initial levels, growth in domestic production accounts for 30–40% of U.S. export growth to certain countries (such as Canada and Japan), whereas production growth accounts for virtually none of the observed bilateral trade growth in others (Germany, Mexico, U.K.).

4. **Extensive margin growth is granular and a significant factor in total bilateral trade growth** The rise in trade growth of the least-traded goods accounting for the bottom decile of initial bilateral trade is significant in accounting for overall trade growth. A 10% increase in bilateral trade is accompanied by a 24% increase in the extensive margin, on average across country-pairs. Extensive margin growth is granular, with a small number of extensive margin goods (<5%) accounting for roughly 25% of total extensive margin growth, but less granular than total bilateral growth.

2.3.2 Data Sources

To study and decompose bilateral trade growth, I use annual UN Comtrade data on bilateral trade from 1989–1999, at the 5-digit Standard International Trade Classification (SITC) level, for 6 countries — Canada, Germany, Japan, Mexico, U.S.A, and U.K. Trade between these country-pairs accounts for roughly 20–25% of global trade during this period. At the 5-digit SITC code level there

are 1836 different goods classifications.¹⁵ Since the 1989–1999 period includes the implementation of NAFTA and CUSFTA — episodes of trade liberalization for some country-pairs (Canada, Mexico, U.S.A.) but not for others (Germany, Japan, U.K.) — I am able to examine the differential effects of trade liberalization on trade growth across country-pairs.¹⁶

I use data on U.S. tariff rates and U.S. imports from Canada, Germany, Japan, Mexico, and the U.K. to investigate the impact of trade liberalization, and heterogeneous changes in tariff on trade growth across goods. The U.S. tariff rate data is from the NBER database on international trade and reports the ad valorem equivalent (AVE) tariff rates on U.S. imports between 1989–1999 at the Harmonized Tariff Schedule (HTS) 8-digit level.¹⁷ This includes estimates of AVE rates for all Most-Favored-Nation-status (MFN) trading partners, such as Germany, Japan and U.K., as well as the tariff schedules for Canada and Mexico following the implementation of NAFTA. I match this tariff data, with data on U.S. imports, resulting in approximately 8000–10000 different goods categories, depending on the import source country.¹⁸

¹⁵The UN uses a lexicographic ordering of SITC codes, with a parent category divided into up to 10 subcategories at each stage of disaggregation — thus a 3 digit code 782 may be subdivided into 7821 and 7822, and 7822 may be subdivided into 78221, 78223, 78225, 78227 and 78229, etc. This may invoke concerns of endogeneity in the classifications — goods with larger trade values may enable or necessitate more subdivisions, while goods with less trade may be lumped into one broader parent category. However, the SITC codes underwent their most recent rounds of revisions in 1986 and 2006 respectively — since the data in this work focuses on the time periods 1989–1999, this concern should be reasonably mitigated. This is similarly true for the NAICS and HTS classification systems. For a more detailed methodology, consult the United Nations Statistical Divisions reports, available at <http://unstats.un.org/unsd/family/default.asp>.

¹⁶This trade data is particularly appropriate as it represents a subset of the Kehoe and Ruhl (2013) data set, allowing for a comparison of the results on trade liberalizations and extensive margin growth found in that paper.

¹⁷This data set is compiled by Feenstra, Romalis and Schott (NBER working paper), available at <http://www.johnromalis.com/publications/>.

¹⁸U.S. import data comes from a dataset on U.S imports and exports compiled by Peter Schott, containing data on U.S. imports at the 10-digit HTS level, which is then aggregated up to the 8-digit level over the same time period, 1989–1999, available at http://faculty.som.yale.edu/peterschott/sub_international.htm.

To identify the relationship between trade and gross domestic production, I include matched data on U.S. manufacturing and U.S. exports. I compile production data from the NBER-CES Manufacturing Industry database to provide data on 473 goods classifications at the 6-digit North American Industry Classification System (NAICS) level between 1989 and 1999. I match the production data to U.S. export data at the 6-digit NAICS level for the same destination countries as the tariff and bilateral trade data.¹⁹

Since international trade data often suffers from issues of sparseness and lumpiness in shipments, I use a 3-yr average to control for reporting errors in the timing of shipments, or instances of shipment and Customs reports being dispersed over multiple years. I classify the initial level of trade as the 1989–1991 average and the final level of trade as the 1997–1999 average for each good classification. The difference in these 3-yr averages therefore represents the growth in trade over this period.

Descriptive Statistics

Tables 2.1 and 2.2 present summary statistics, for NAFTA country-pairs and non-NAFTA country-pairs respectively, of the 5-digit SITC bilateral trade data. There are three key facts to note. First, on average, the initial level of trade for NAFTA and non-NAFTA pairs is approximately the same (roughly \$36B), in total shipment value in \$US, while the average final level of trade is higher in NAFTA country-pairs (\$75B) than in non-NAFTA pairs (\$57B). This reflects the larger average growth in total trade for NAFTA country-pairs of 148%, while non-NAFTA pairs average 63% growth over the same period. Second, the bottom decile of goods, ordered by initial trade value, grow to account for 18%

¹⁹The NBER data comes from a database compiled by Becker, Gray and Marvakov, available at <http://www.nber.org/nberces/>, while the export data comes from the Schott database (Ibid.).

of final trade in NAFTA country-pairs, while accounting for 13% of final trade in non-NAFTA pairs. This implies a slightly larger role for the extensive margin in trade-liberalizing countries. Third, the 10 largest-growth goods in each country pair (out of 1836 total good classifications) accounts for approximately 50% of total trade growth in NAFTA and non-NAFTA country-pairs alike. Similarly, the 100 largest-growth goods account for 101% of total trade growth in non-NAFTA countries, and 90% of total trade growth in NAFTA country-pairs. Together, these three facts suggest that trade growth is highly concentrated in a small number of goods classifications. This pattern is consistent across all country-pairs, although it is slightly more prominent in trade-liberalization country-pairs.

Table 2.1: Summary Statistics: NAFTA country-pairs

Bilateral Trade Statistics: 5-digit SITC codes							
Variable	Can-Mex	Mex-Can	Can-Usa	Usa-Can	Mex-Usa	Usa-Mex	NAFTA Avg
Initial trade	452451	1711096	83278894	71313846	28298716	25796112	35141852
Final trade	880883	5537384	160144152	123700777	93734628	70447948	75740962
Total trade growth	428432	3826288	76865257	52386931	65435912	44651836	40599109
% Δ in trade	94.7%	223.7%	92.3%	73.5%	231.2%	173.1%	148%
Share of trade growth from 10 largest growers	94.7%	61.3%	41.1%	29.4%	42.5%	29.6	51.9%
Share of trade growth from 100 largest growers	136.5%	95.0%	77.6%	68.5%	85.2%	71.2%	91.8%
Extensive margin share of final trade	25.7%	26.4%	15.6%	11.5%	17.5%	13.3%	18.3%
Number of goods	1836						
(All trade values in thousands of \$US)							

Table 2.2: Summary Statistics: Non-NAFTA country-pairs

Bilateral Trade Statistics: 5-digit SITC codes							
Variable	Ger-Usa	Usa-Ger	Jpn-Usa	Usa-Jpn	UK-Usa	Usa-UK	Non-NAFTA Avg
Initial trade	25828198	21430441	93940987	48749684	18021847	19532210	37917228
Final trade	47977038	34737504	124383525	67346115	33471956	34286200	57033723
Total trade growth	22148840	13307063	30442537	18596431	15450108	14753990	19116495
% Δ trade growth	85.7%	62.1%	32.4%	38.1%	85.7%	75.5%	63.3%
Share of trade growth from 10 largest growers	52.0%	56.1%	57.8%	58.1%	41.1	38.7%	50.6%
Share of trade growth from 100 largest growers	85.3	107.9%	121.1%	118.7%	87.2%	85.9%	101.1%
Extensive margin share of final trade	13.0%	12.9%	12.3%	12.4%	11.9%	14.5%	12.8%
Number of goods	1836						
(All trade values in thousands of \$US)							

2.3.3 Granularity of Trade Growth

Granularity Across Time

Over the past 50–60 years, trade growth across NAFTA and non-NAFTA country-pairs exhibit similar patterns.²⁰ All country-pairs show large overall growth over this period, but exhibit large variation in the rate of growth between periods. In most cases, the majority of the overall trade growth over time is accounted for by a small number of episodes of rapid growth.

To quantify the granularity in trade growth over time, I calculate the share of total trade growth accounted for by each year, and calculate the proportion of overall trade growth accounted for by the top 5, 10 and 15 years of growth. Table 2.3 reports that the top 5 years of growth (representing only 10% of the total time frame) account for roughly 75% of overall trade growth across time.²¹ The top 10 and 15 years of growth account for roughly 150% and 200%, respectively, of total bilateral trade over this period, suggesting that most of the overall growth over time is concentrated in a small number of rapid growth episodes.

Table 2.3: Proportion of total bilateral trade from “X” top-growth years

Country-Pair	Proportion of trade growth from top X% years (5-digit SITC codes)								
	1-yr intervals			2-yr intervals			3-yr intervals		
	5 yrs (10%)	10 yrs (20%)	15 yrs (30%)	3 yrs (10%)	6yrs (20%)	9yrs (30%)	2 yrs (10%)	4 yrs (20%)	6 yrs (30%)
Mex-Can	0.6856	0.9931	1.1836	0.7257	1.0360	1.1356	0.7146	0.9552	1.0612
Usa-Jpn	0.6781	1.0553	1.3400	0.6000	0.9775	1.2586	0.5747	0.8556	1.0261
Jpn-Usa	0.9172	1.5070	1.9293	0.8442	1.2970	1.5322	0.6426	0.9776	1.2874
Mean	0.7603	1.1852	1.4843	0.7233	1.1035	1.3088	0.6440	0.9295	1.1249

²⁰Figure A.1 plots total bilateral trade values for exports over the past 50 years for Mexican exports to the U.S. and Canada, as well as U.S. exports to Japan and Japanese exports to the U.K. For some countries in the data set, such as Mexico, the data only extends back to 1980. For brevity, only 4 countries are plotted here, but other the country-pairs exhibit similar growth patterns over time.

²¹It should be noted that these need not be concurrent years, but rather the 5 individual years demonstrating the largest growth across the entire time period.

Due to the inherent lumpiness and variability in reporting of international shipments, using single year intervals to measure growth may potentially overestimate trade growth in any given year. I therefore repeat the exercise by using 2-year and 3-year intervals for the length of each period in calculating trade growth. I find similar results to those using 1-year intervals, with a high degree of granularity across time periods, with the top 10% of time periods accounting for 70–75% of total trade growth. This reinforces the finding that trade growth is not smooth and uniform across time, but highly concentrated in a small number of episodes of rapid growth.


Granularity across goods

To identify the properties of trade growth, I examine the share of total trade growth accounted for by each good for each country pair. I arrange goods, in ascending order, by initial value of exports for each bilateral pair and calculate each good’s corresponding share of the total growth in bilateral trade. For ease of exposition, I separate goods into 3 groupings — the least-traded goods (comprising the bottom decile of total initial exports, by 5-digit SITC code); the mid-traded goods (comprising the second through fifth deciles); and the most-traded goods (comprising the top 50% of initial exports).. Across the various bilateral country-pairs, the least-traded goods make up 75–90% of the goods classifications. Roughly half of these categories are goods with zero reported trade.²² The set of most-traded goods comprises a small number (5–40 of 1836) of good classifications. This reflects the “lumpiness” in cross-sectional trade data — many goods are not-traded or traded in small quantities, while a small number of highly-traded goods accounts for a large portion of total bilateral

²²Due to the fact that most countries exclude very small shipments (i.e. less than \$2000) from customs duties and reports, the term “not-traded” is used in this paper to refer to goods with no recorded trade value in these customs reports.

trade in any given period.

Figure 2.1: Growth by Good: 5-digit SITC bilateral trade

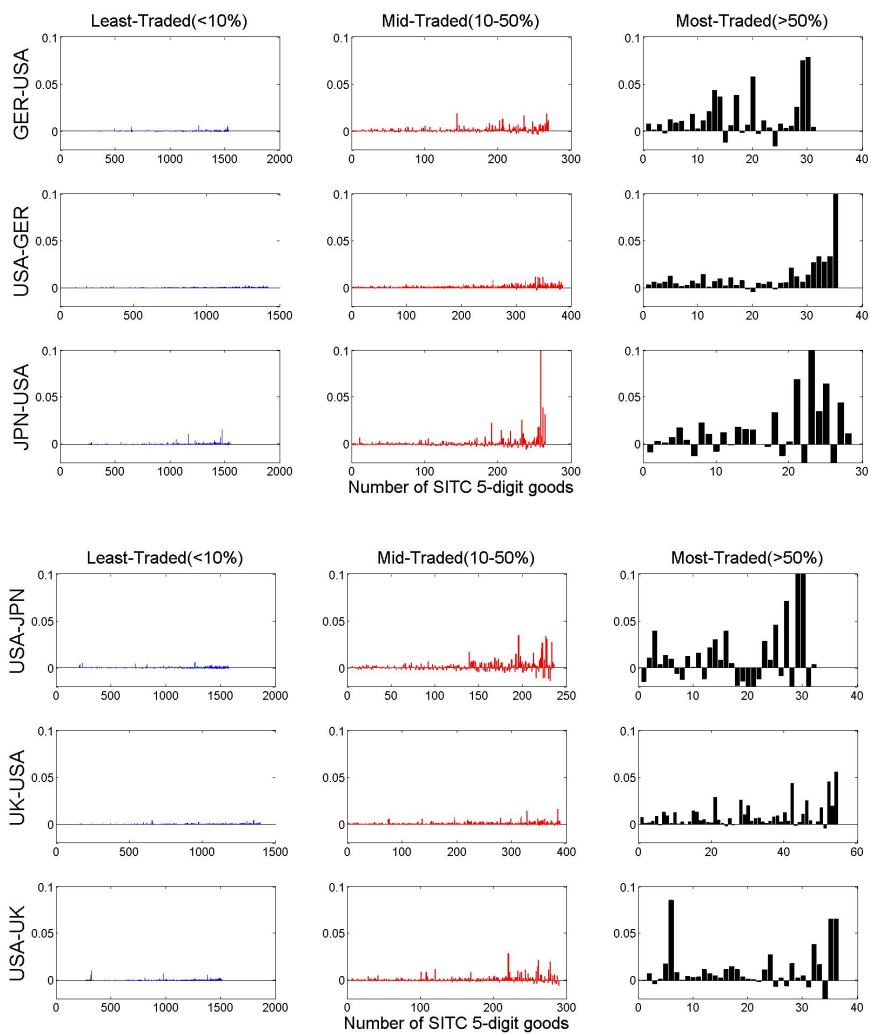


(b) NAFTA county-pairs

To decompose trade growth across goods, I plot the trade growth accounted for by each good as a share of total bilateral trade growth in Figures 2.1–2.2.²³ The height of each bar represents that good’s contribution to the total growth

²³Formally, the y-axis of Fig. A.2–2.2 is measured, for each good i , as: $\frac{\Delta Exports(i)}{\sum_j \Delta Exports(j)}$.

Figure 2.2: Growth by Good: 5-digit SITC bilateral trade



(b) Non-NAFTA county-pairs

in trade for that country-pair.²⁴

Bilateral trade growth is highly granular, with a small number of goods exhibiting large increases in trade values accounting for the vast majority of total trade growth. Most of these goods originate from the mid- or most-traded categories, and almost never from the least-traded category. The remaining goods (90–95% of all goods classifications) exhibit small (<0.05% of total trade), zero, or negative growth.

This finding can be summarized as follows: many goods begin and remain not-traded; some not-traded goods become traded in small amounts, but almost never go from zero to large amounts of trade; most traded goods grow very little; a small number of goods grow from small amounts of trade to large amounts, or from large amounts to even larger amounts, the two cases that account for the majority of trade growth between country-pairs.

The goods exhibiting large growth that account for the majority of trade growth come from varying levels of initial trade across bilateral pairs. That is, large growth goods (“growers”) do not all begin with similar levels of initial trade, and goods with similar initial levels of trade do not necessarily grow in similar proportions.²⁵ Further, goods that begin with zero trade do not typically become traded in large values, if they become traded at all.

To quantify this relationship, I calculate the share of total bilateral trade growth accounted for by these high-growth goods. Table 2.4 lists the proportion of overall trade growth accounted for by the top 1% of large-growth goods, as

²⁴To more clearly illustrate the patterns of trade growth, I also plot groupings of goods individually, as seen in Figures A.2–A.4. These figures represent the case for Canadian exports to Mexico; Figures 2.1–2.2 show similar results for all country-pairs.

²⁵One possible exception to this might be the U.S.-Canada pairing, which seems to exhibit less concentration of trade growth and smoother growth in trade, according to initial value of trade by good.

well as the top 2%, 5% and 10% of goods.²⁶ The mean values, across all country-pairs, identify the granularity in trade growth, with the top 1% of goods accounting for 62% of total trade growth, while the top 2%, 5% and 10% of goods accounting for 76%, 93% and 105% of total bilateral trade growth, respectively. There is variability in the degree of granularity across country-pairs — the top 1% of Canadian exports to the U.S. account for 108% of total growth in trade, while the top 1% of U.S. exports to Mexico account for only 39% of total trade growth. However, all country-pairs exhibit granularity of trade growth across goods, with a disproportionately high concentration of trade growth in a small number of goods.

Table 2.4: Proportion of total bilateral trade concentrated in top “X”% of goods
Proportion of growth from top X% of goods
(5-digit SITC codes)

Country-Pair	1% (18 goods)	2% (36 goods)	5% (91 goods)	10% (183 goods)
Can-Usa	1.081	1.223	1.355	1.410
Can-Mex	0.491	0.606	0.760	0.884
Ger-Usa	0.589	0.692	0.837	0.955
Jpn-Usa	0.722	0.940	1.189	1.313
Mex-Can	0.718	0.818	0.940	1.004
Mex-Usa	0.551	0.681	0.836	0.935
Usa-Can	0.373	0.478	0.665	0.815
Usa-Ger	0.710	0.882	1.062	1.167
Usa-Jpn	0.753	0.928	1.165	1.311
Usa-Mex	0.396	0.518	0.692	0.834
Usa-UK	0.492	0.647	0.841	0.972
UK-Usa	0.516	0.653	0.853	0.974
Mean	0.616	0.756	0.933	1.048

While this pattern of granularity of trade growth is pervasive across all country-pairs, it is not the same goods categories that account for the majority of trade growth across different country-pairs. Rather, there is significant

²⁶As there are 1836 goods classifications, the top 1% of goods consists of the top 18 goods classifications, the top 2% consisting of 36 goods, etc.

Table 2.5: Common large-growth goods across all countries

Number of Goods with $X\%$ of Trade Growth 5-digit SITC codes (1836 goods)				
Countries	>3%	>1%	>0.5%	>0.1%
All	0	1	2	9
		Pass. Autos	Switches/Fuses Other auto parts	Unhardened rubber Pipe valves Metal mountings Peripheral units TV transmitters Static converters Elec. microcircuits Gas instruments Polarizing lenses

heterogeneity in which goods account for the majority of trade growth across country-pairs. Table 2.5 shows that of the 1836 SITC 5-digit classifications, none account for at least 3% of total trade growth in every country-pairs. Only one category, passenger automobiles (excluding buses), accounts for at least 1% of trade growth in every country-pair. Further, only a dozen of the 1836 goods classifications account for at least 0.1% of total trade growth in every country-pair. Examining bilateral pairs, of the small number of goods categories accounting for the majority of trade growth, only a small proportion (roughly 5–20%) are common to both countries' trade growth. Table 2.6 demonstrates that for each bilateral pair, of the 1–6 goods categories accounting for at least 5% of bilateral trade growth for each country, none or only 1 of them is common to both countries. Similarly, of the roughly 20–30 goods each accounting for at least 1% of total bilateral trade growth, only 4–8 of these goods are common to both countries.

Table 2.6: Common large-growth goods across country-pairs

Number of Goods with $X\%$ of Trade Growth 5-digit SITC codes (1836 goods)										
Country	>5%		>3%		>1%		>0.5%		>0.1%	
	Same		Same		Same		Same		Same	
Canada	6		9		21		40		103	
Mexico	4	1	17	1	18	4	26	6	107	27
Canada	1		4		14		34		176	
USA	1	0	12	1	15	5	32	15	202	108
Mexico	2		6		21		35		138	
USA	1	0	16	1	18	8	36	15	190	76
Germany	2		3		12		31		165	
USA	4	1	21	1	26	8	44	17	139	72
Japan	3		5		31		57		146	
USA	3	1	21	3	25	8	53	19	178	62
UK	3		6		19		47		159	
USA	3	0	14	2	22	6	42	18	182	91

2.3.4 Tariff Rates

Trade theory suggests that reductions in import tariffs should lower barriers to trade and increase trade for those goods. To examine the relationship between changes in tariff rates and corresponding trade growth across goods, I match data on U.S. tariff rates to U.S. import data at the 8-digit HTS code level.²⁷ I calculate the change in the tariff rate for each good as the difference between the initial and final ad valorem equivalent (AVE) tariff values calculated by Romalis (2007). For each bilateral U.S. trade partner, I match the AVE tariff change to the corresponding growth share in U.S. imports for each good. Across all goods, the tariff changes show a wide range of values, from large increases

²⁷To examine tariff data at the most disaggregated level available, I use 8-digit HTS data, which has significantly more goods classifications than the 5-digit SITC data (>8000 vs. 1836) — however, I find the pattern of trade growth granularity, while slightly more pronounced in the HTS classifications, displays similar patterns of the concentration of trade growth as the 5-digit SITC code data. As in the 5-digit SITC data, trade growth is also slightly more granular in NAFTA countries (Canada and Mexico) than it is in non-liberalizing countries (U.K.).

to no change to large decreases. Tariff changes are generally larger and more negative in the trade liberalizing countries of Canada and Mexico than those for Most Favored Nation (MFN) status countries like the U.K. However, for each country-pair, the correlation between import growth and the change in tariff rates across all goods is not significantly different from zero. Isolating each set of least-, mid- and most-traded goods, this lack of correlation holds.²⁸

Due to the granularity of trade growth across goods, I isolate the subset of largest-growth goods contributing the majority of total trade growth and examine their corresponding tariff changes. Figure 2.3 plots the 100 goods classifications that contribute the most to import growth (accounting for $\approx 80\%$ of total growth). These large-growth goods display a wide range of tariff changes — depending on the country-pair. Only 10 to 30 of these 100 largest “growers” exhibit notable decreases in the AVE tariff rate. The remaining goods reflect no change or an increase in the AVE rate, suggesting that tariff decreases alone cannot account for the granularity in trade growth across goods.

Alternatively, in Figure 2.4, I isolate the subset of goods that exhibit the largest decreases in tariff rates to examine their corresponding share of trade growth. Roughly one-third of the goods experiencing the largest tariff reductions exhibit positive growth in trade. However, few of these goods ($< 10\text{--}15\%$) contribute significant shares to the overall increase in imports. The majority of goods with the largest tariff reductions exhibit no growth or slight decreases in imports from initial to final trade levels, coinciding with the lack of statistically significant correlation between tariff changes and import growth.

²⁸Refer to Figures A.5–A.7.

Figure 2.3: Large-growth goods: Imports vs. Tariff changes

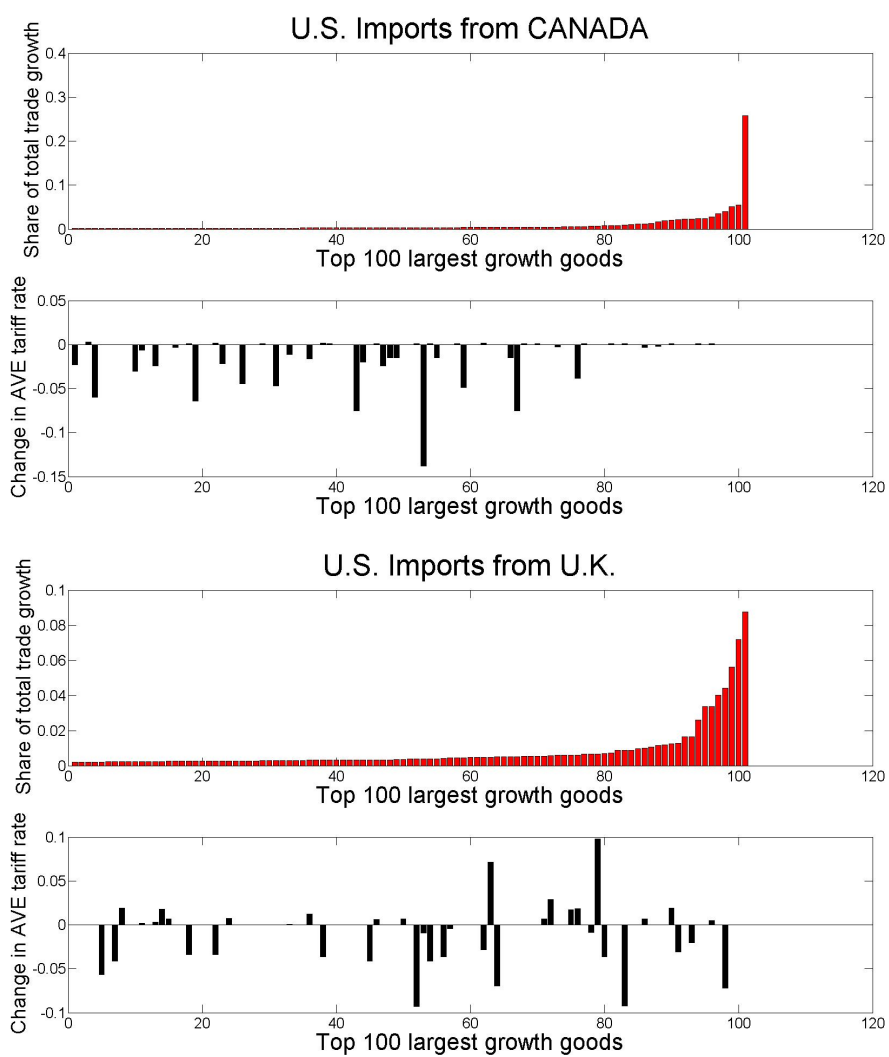
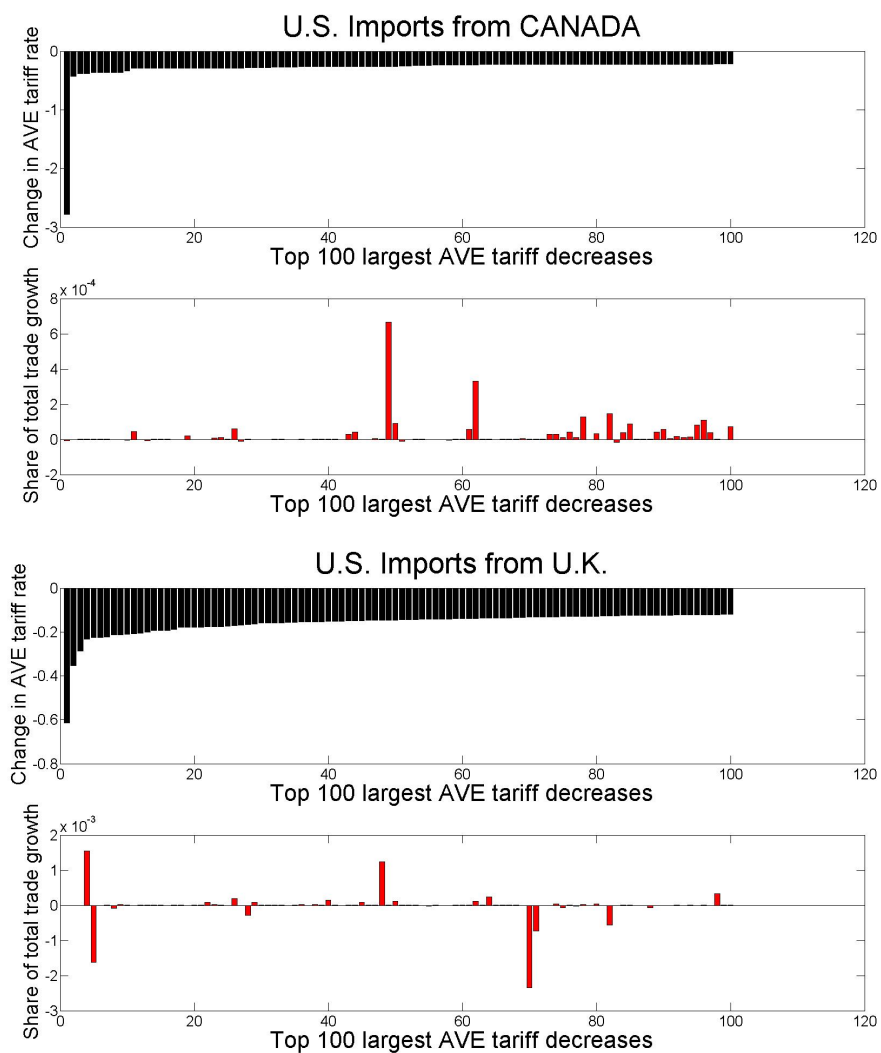


Figure 2.4: Large tariff decrease goods: Imports vs. Tariff changes



2.3.5 Production Growth and Trade Intensity

Trade theory predicts that increases in trade across goods may be related to increases in overall domestic production, with the share of domestic production being exported staying constant. Alternatively, increases in exports may be driven by an increase in trade intensity — exporting a proportionally larger share of domestic production, independent of the level of domestic production.

To examine the relationship between bilateral trade growth and domestic production, I focus on U.S. data, due to the availability of U.S. production data from the NBER manufacturing database at the 6-digit NAICS code level. The trade data is therefore confined to U.S. exports to Canada, Germany, Japan, Mexico, and U.K. I concord the U.S. export data from the 5-digit SITC codes to the 6-digit NAICS classifications and match it to the corresponding production data. This less disaggregated 6-digit NAICS code level results in less classifications (473 goods) than the 5-digit SITC level (1836 goods).²⁹

One possibility is that these large-growth goods that account for the majority of trade are simply due to large increases in domestic production, with a constant share of domestic production being exported. To determine whether these goods exhibit similar granularity in their domestic production, I calculate the share of domestic production growth accounted for by each good. To be consistent, I group the production data into the same categories of least-, mid- and most-traded goods, that are arranged in the same fashion as the export data. Specifically, the goods are grouped according to their *initial trade* levels, and then analyzed for their growth in *production*. Thus, the “bottom decile”

²⁹As seen in Figures A.9(a)–2.2 and 2.5. U.S. export data exhibits similar granularity at the less disaggregated 6-digit NAICS code level as at the 5-digit SITC code level. Extensive margin growth, represented by the set of least-traded goods, is still significant; however, due to the less disaggregated data, the extensive margin accounts for a smaller fraction of final bilateral trade ($\approx 14\%$) than it does in the SITC data.

of production goods account for the bottom 10% of initial trade, and may not account for the bottom 10% of initial domestic production.

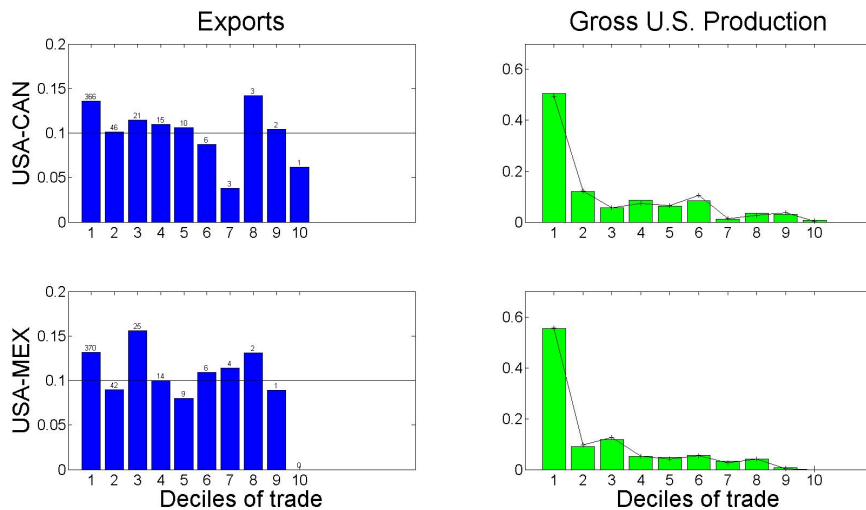
Production growth is concentrated in a small number of goods for each country-pair, but is less granular in the production data than in the export data, particularly among the mid- and most-traded goods sets.³⁰ Examining the deciles of U.S. production, each decile primarily remains at its initial levels in the final production data. In Figure 2.5, the value of final production appears as the height of each bar, with the initial value plotted as the “+” in each decile. Most deciles, across all country-pairs, demonstrate small to no change in their fraction of total production from initial to final levels. Also, the bottom decile in each country-pair accounts for a large share of overall production, both in initial and final levels, reflecting the fact that many domestically-produced goods are not traded between each country-pair.

The production data suggests that the granularity in trade growth is not merely a function of domestic production. To examine the effect of changes in total production and changes in trade intensity in accounting for bilateral trade growth, I construct the exports-to-production ratio for each good, for both initial and final levels of trade. I then create two separate measures to decompose the effects of changes in trade intensity and production growth in contributing to trade growth:


1. Holding production constant at initial levels, and multiplying by the final level of trade intensity, I calculate the share of trade growth that can be attributed to growth in trade intensity
2. Holding the trade intensity fixed at the initial level, and multiplying by the final level of production, I calculate the share of trade growth that can

³⁰See Figures A.9(a)–A.9(b).

Figure 2.5: Decile Growth (6-digit NAICS) bilateral trade



(a) NAFTA country-pairs



(b) Non-NAFTA country-pairs

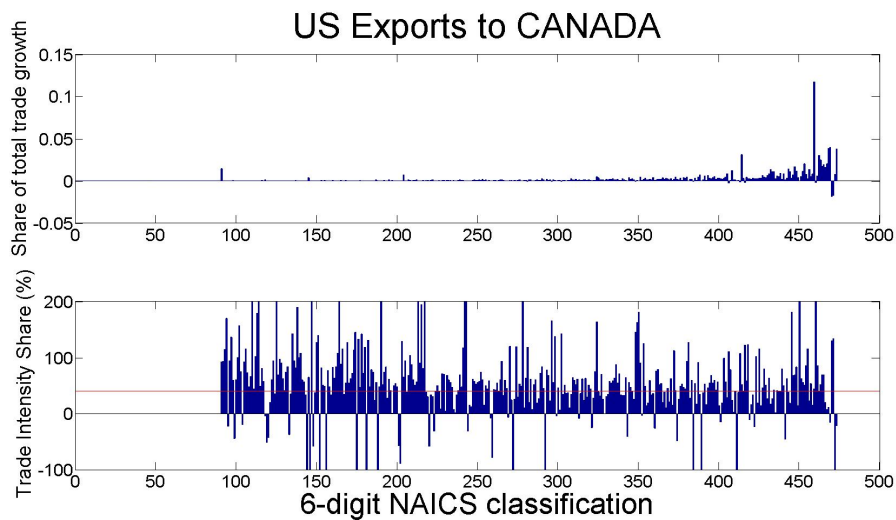
be attributed to production growth

To illustrate this decomposition, Figure 2.6 plots the share of total trade growth (in the top row of each figure) along with the corresponding trade intensity share for each of the 473 6-digit NAICS goods classifications. Across all goods, increases in trade intensity account for, on average, 30% to 72% of the observed trade growth across country-pairs. However, there is little correlation between the share of trade growth accounted for by each good and the corresponding trade intensity share, as reported in Table 2.7. There is similarly no significant correlations between initial trade values and the resultant trade intensity share across goods.


Similarly, increases in production account for, on average, -5% to 33% of the observed trade growth across country-pairs. For some destination countries (Canada and Japan), production growth accounts for large portions of total trade growth (30–40%), while for others (Germany, Mexico, U.K.) production growth accounts for virtually none of the overall growth in trade, on average, as seen in Figure 2.7. Across goods there is little correlation between the share of trade growth and the corresponding production growth share for each good, as reported in Table 2.7. However, unlike the trade intensity shares, goods with higher initial trade values generally have higher production growth shares.

On average, increases in trade intensity are significant factors in accounting for episodes of large growth in bilateral trade across all country-pairs. Production growth shares vary more widely across destination countries in their significance in accounting for trade growth. There is no discernible difference between trade-liberalizing destination countries and non-liberalizing destination countries for both the mean trade intensity and production growth shares.

Figure 2.6: Trade Intensity Share vs. Share of Trade Growth




(a) U.S.-Canada




(b) U.S.-U.K.

Figure 2.7: Production Growth Share vs. Share of Trade Growth



(a) U.S.-Canada



(b) U.S.-U.K.

Table 2.7: Correlations: Trade Intensity & Production Growth Shares vs. Share of Trade Growth

Correlation	US-Can	US-Ger	US-Jpn	US-Mex	US-UK
Trade Intensity Share vs Share of Trade Growth	-0.007	-0.052	-0.010	-0.015	-0.006
Production Growth Share vs Share of Trade Growth	0.023	0.053	0.015	0.032	0.023

2.3.6 Extensive Margin Growth

Kehoe and Ruhl (2013) find that the extensive margin is significant in accounting for overall bilateral trade growth. Due to inconsistencies in the reporting of small shipment values in international trade data, Kehoe and Ruhl examine the set of *least-traded goods* in each bilateral country pair, represented by the subset of goods accounting for the bottom decile of initial trade. Examining the growth in the share of final trade accounted for by this set of least-traded goods then gives a proxy for extensive margin growth in bilateral trade data.

To quantify the contribution of new or previously not-traded goods to overall trade growth, I calculate extensive margin growth in bilateral trade data following the methodology of Kehoe and Ruhl (2013). Due to the large number of goods that are not traded (30–50% of the 1836 classifications for most countries) between any given country-pair, it takes 75–90% of all goods classifications to account for 10% of initial trade. Mirroring Kehoe and Ruhl’s findings, I find the extensive margin is significant in accounting for bilateral trade growth across country-pairs in the 5-digit SITC trade data. Across all country-pairs, the set of least-traded goods accounts for an average of approximately 16% of

final trade. A 10% growth in total trade between country-pairs is accompanied by a 24% increase in the share of final trade accounted for by these goods. However, there is also significant growth in other deciles by country-pair. For example, the 6th decile of Canadian exports to Mexico grows to account for over 20% of final trade, and the 5th decile of U.S. exports to Germany grows to account for 18% of final trade, as seen in Figure 2.8.

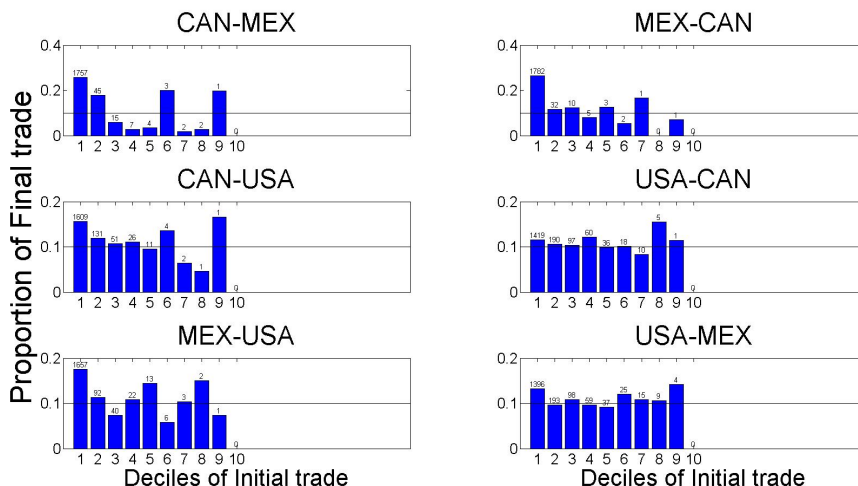
Kehoe and Ruhl (2013) examine the overall role of the extensive margin in contributing to total bilateral trade growth. I extend their findings by identifying that within the set of least-traded goods, the majority of extensive margin growth comes from a small number of goods growing from small values to slightly larger values of trade, not from not-traded goods becoming traded in large amounts. Nearly half of the goods classifications contained within this set of least-traded goods begin as not-traded.³¹ The majority of these remain not-traded, while those that become traded only do so in small amounts. A small number of goods (<10% of all extensive margin goods classifications) in each bilateral pair exhibit relatively large growth, from a small value of initial trade to large values of final trade. This result shows the granularity of trade growth within the extensive margin that mirrors that in the intensive margin. However, due to the exceedingly large number of goods classifications exhibiting small growth, these large-growth extensive margin goods only contribute 5–25% of total extensive margin growth.

2.4 Model


I now ask whether the predictions of standard Melitz-style trade models are consistent with the empirical findings in the data. The term “standard” here

³¹Refer to Figure A.2.

Figure 2.8: Decile Growth: 5-digit SITC bilateral trade



(a) NAFTA country-pairs



(b) Non-NAFTA country-pairs

refers to recently used models where firms differ in productivity, face fixed and variable costs of exporting (as in Melitz (2003), Chaney (2008), etc.) and CES demand under monopolistic competition. Specifically, I ask whether these models, for reasonable parameter values, generate the high level of granularity of trade growth across goods classifications observed in the data, in response to both uniform and heterogeneous tariff reductions, and when incorporating heterogeneous productivity changes imputed from the data.

I analyze a two-country model with firms indexed by heterogeneous productivities, facing CES demand under monopolistic competition, with fixed and variable costs of exporting. I characterize the equilibrium objects for profit-maximizing firms in their export decisions, and identify the stratification of goods as exported or not-traded. I then introduce trade liberalization, represented as a reduction in the variable costs of exporting, for two main cases, in order to determine the implications for bilateral trade growth across goods classifications, to compare to the empirical findings. First, I follow an approach common to many trade models, with firms exhibiting heterogeneous productivities that are fixed across periods, where trade growth is generated by trade liberalization that takes the form of a uniform proportional reduction in tariff rates across goods. Second, I augment this standard approach by incorporating heterogeneous tariff reductions, matched from global tariff data, as well as heterogeneous productivity changes, imputed from U.S. production data, across goods.

I find that the standard model with uniform tariff reductions does not capture the granularity of trade growth — the model predicts trade growth is widely dispersed across a large number of goods classifications. Additionally, growth in the extensive margin, as represented by the set of least-traded goods,

does not match the pattern of growth in previously not-traded goods or the small subset of goods growing from small values of trade to much larger values. Adding heterogeneous productivity changes and tariff reductions allows the model to better match the lumpiness of initial trade values across goods seen in the trade data, as well as predicting significantly more granularity in trade growth across goods. This version of the model also better captures the granularity of growth within the extensive margin, as well as better capturing the contribution of the extensive margin to overall trade growth.

2.4.1 Model Set-up

The model builds on the standard trade model framework, with 2 symmetric countries, Home(H) and Foreign(F).³² Firms are heterogeneous in their labor productivity, with each firm producing a differentiated good using only labor as an input. Firms are indexed by their productivity, $\frac{1}{a}$, where a is the amount of labor required to produce one unit of output. Pricing follows the Dixon-Stiglitz (1977) model of monopolistic pricing with firms facing CES demand in both a home and foreign country. Firms may sell in the domestic market only, or may choose to also export to the foreign market via a common distribution technology that requires a fixed cost of exporting, and an iceberg variable cost per unit shipped to the destination market. Following the literature, episodes of trade liberalizations can be represented by a uniform reduction in the variable cost of exporting in the second period.

³²For brevity, I will consider the problem of Home's exports to Foreign — by symmetry, the Foreign country faces the same problem as the Home country in its exporting decisions.

Consumers

Preferences of consumers in country $i \in \{H, F\}$ can be represented by the CES utility function

$$U^i = \left(\int_{a \in A} c(a)^{\frac{\epsilon-1}{\epsilon}} da \right)^{\frac{\epsilon}{\epsilon-1}} \quad (2.1)$$

where $c(a)$ is consumption of good variety a , ϵ is the elasticity of substitution across goods ($\epsilon > 1$) and A is the set of available goods.

Firms

I assume firms maximize profits under monopolistic competition pricing, as in Dixon-Stiglitz (1977) while facing this CES demand. Firms first draw their labor requirement from a Pareto distribution, $G(a)$, which they hire at domestic wage rate, ω^i in country i . Firms then choose whether or not to service the foreign market by paying a fixed cost of exporting, f , and a variable cost of exporting, τ in the form of an iceberg transportation cost.

2.4.2 Characterizing Equilibrium

Demand

Due to the symmetry of the two-country model, I consider only the demand from Foreign for product varieties produced in Home, and set aside the demand from the domestic market. For utility-maximizing consumers with this common CES utility specification, the demand in Foreign for Home variety a is:

$$c^F(a) = E^F \frac{p(a)^{-\epsilon}}{P_F^{1-\epsilon}} \quad (2.2)$$

where E^F is national income in Foreign and P_F is the Foreign aggregate price level.

Firm's Decision

Under monopolistic competition, firms price at a constant mark-up over marginal cost. In the Home market, the marginal cost of production is $\omega^H a$, while the marginal cost for firm a to service the foreign market is $\omega^H \tau a$, as each exporting firm must pay the additional variable export cost τ per unit. This results in a price in Foreign for variety a produced in Home of

$$p^F(a) = \left(\frac{\epsilon}{\epsilon - 1} \right) \omega^H \tau a \quad (2.3)$$

For a firm with labour requirement a , domestic profits for profit maximizing firms are:³³

$$\pi_D(a) = (\omega^H a)^{1-\epsilon} D^H \quad (2.4)$$

Under CES demand and monopolistic competition pricing, these profits are proportional to the firm's productivity.

For a firm with labour requirement a , the *additional* profits from exporting are:

$$\pi_X(a) = (\omega^H \tau a)^{1-\epsilon} D^F - \omega^H f \quad (2.5)$$

That is, firms generate sales that are proportional to their productivity and the variable transportation costs τ , but must pay the fixed cost, f , in order to access the foreign market.³⁴

Productivity thresholds

These profit functions result in productivity thresholds dictating which firms will choose to export and which will only be produced domestically. With no

³³ $D^H = \frac{E^H \epsilon^{-\epsilon}}{[P^H(\epsilon-1)]^{1-\epsilon}}$ and is a constant w.r.t. a .

³⁴ $D^F = \frac{E^F \epsilon^{-\epsilon}}{[P^F(\epsilon-1)]^{1-\epsilon}}$.

fixed cost of entry, domestic profits are always non-negative, so all firms choose to produce domestically.³⁵ To export, additional profits must be non-negative: $\pi_X(a) \geq 0$. This leads to productivity cut-off, $\frac{1}{\bar{a}}$, that satisfies:

$$\frac{1}{\bar{a}} = \left[\frac{\omega^H \epsilon f \tau^{\epsilon-1}}{D^F} \right]^{\frac{1}{\epsilon-1}} \quad (2.6)$$

The productivity cut-off for exporting is increasing in both f and τ .

The top panel of Figure 2.9 shows the stratification of goods according to the exporting choice. Arranging goods by increasing productivity $\frac{1}{a}$, the least productive goods will only be produced for domestic consumption. Firms producing good a with productivity higher than $\frac{1}{\bar{a}}$ will choose to export.

2.4.3 Trade Liberalization


In order to examine whether the standard model can match observed patterns of trade growth in the data, I characterize equilibrium in the static standard model in two cases: an “initial” period and a “final” period. I categorize the difference in the equilibrium trade values for each good across periods as the corresponding growth in trade for each good. Following the trade growth literature, I consider the counter-factual of an episode of trade liberalization between periods, represented as a reduction in the variable cost of exporting, to determine the patterns of trade growth across goods.³⁶

A fall in iceberg transportation costs lower the marginal cost of servicing the foreign market. This results in a lower price offered to the foreign market.

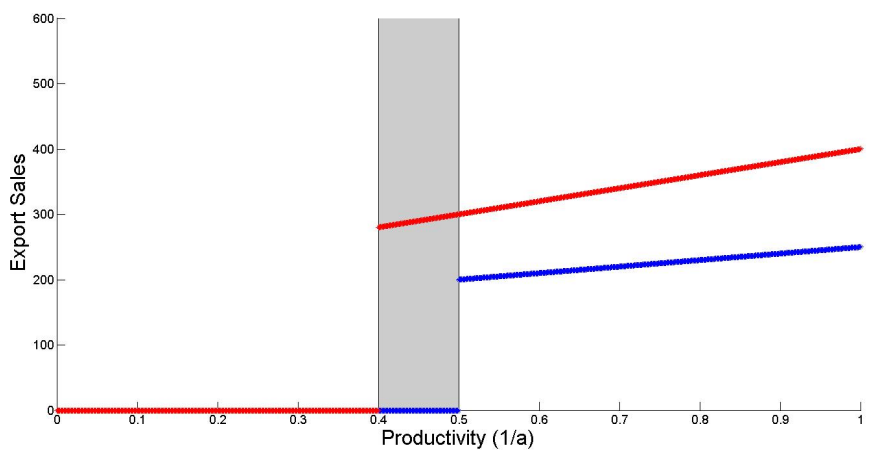
³⁵ $\pi_D(a) = (\omega^H a)^{1-\epsilon} D^H \geq 0$ implies productivity cut-off $(\frac{1}{\bar{a}}) = 0$.

³⁶ The work here addresses a period of trade liberalization in the standard context of reducing τ . However, it is noted that τ , the variable cost of exporting, can be decomposed into $\tau = MC + tar$, where MC represents the distribution costs of production and transportation to reach the foreign market, and tar is the tariff rate. Thus examining a decrease in τ provides insight into cases of pure trade liberalization, captured by a decrease in tar or cases like technological improvements in the distribution network, MC .

Figure 2.9: Standard model with uniform tariff reductions



(a) Initial export sales



(b) Final export sales

With an elasticity of substitution greater than 1, this generates an increase in export sales for all exported goods that is proportional to the fall in tariffs and the firm's productivity — this is growth in the intensive margin. Additionally, the decrease in variable costs lowers the productivity threshold, as the increase in variable profits resulting from the lower marginal costs of reaching the foreign market allow a subset of goods that were previously not-traded to overcome the fixed cost of exporting and enter the foreign market — this is growth in the extensive margin.

The bottom panel of Figure 2.9 demonstrates the change in export sales resulting from a decrease in trade costs for the standard model with uniform tariff reductions. First, all previously exported goods increase their export sales due to the direct decrease in variable costs, proportional to their productivity. The model predicts that the largest intensive margin growth will occur among goods with the largest initial level of trade, with goods traded at lower levels exhibiting smaller growth in trade. Second, as the productivity threshold shifts to the left, lower productivity firms enter the export market, representing the extensive margin. As these were the marginally excluded firms before the trade liberalization, the model predicts the final level of trade for these goods to be similar to that of the goods previously least-traded. As these goods all previously accounted for zero export sales, the model predicts the growth in trade for the extensive margin goods to be relatively larger than that of other previously traded goods with similar productivities, i.e. those previously just above the cut-off.

Implications for trade growth

The standard model predictions for growth in bilateral trade are inconsistent with the empirical findings. Growth is smooth across all previously traded goods, and neither the intensive nor extensive margins display the high level of granularity of a small number of goods accounting for the majority of trade observed in the data.

First, the standard model delivers much less granularity of trade growth than observed in the data. The model predicts that trade growth for previously traded goods will be much less concentrated, as the growth in trade is proportional to the initial level of trade, reflecting each good's productivity. With a common fixed and variable cost of exporting, goods with similar productivities, and thus similar initial levels of trade, will have similar levels of trade following homogeneous reductions in trade costs, and thus contribute similar shares of total trade growth. Additionally, no goods with low levels of initial trade exhibit disproportionately large growth, as seen in the data. Further, all goods with high initial levels of trade grow similarly, with none of these goods exhibiting zero or negligible trade growth as seen in the data.

Second, growth in the extensive margin does not match the stylized facts from the data. The model predicts that most extensive margin growth, in terms of the Kehoe-Ruhl representation of the set of least-traded goods, comes from large increases in previously not-traded goods becoming traded in relatively high values. All previously-traded goods within the extensive margin grow proportionally to their productivity and the reduction in tariffs, similar to intensive margin goods. The standard model predicts that due to the relatively large growth in the intensive margin, the extensive margin accounts for a small portion of overall total trade growth.

Finally, standard models often use the simplifying assumption that trade liberalization takes the form of a decrease in the marginal cost of exporting that is constant and proportional across all goods. The tariff data shows that periods of trade liberalization tend to exhibit high levels of heterogeneity in the magnitude of AVE tariff reductions across goods, which will not be captured by a homogeneous reduction of the marginal costs of exporting across all goods in the model.

2.5 Numerical Exercises: Trade Growth

Characterizing equilibrium in the standard model generates smooth growth across goods, which is proportional to each firm's productivity and the reduction in variable costs from trade liberalization. The standard model does not generate the granularity of trade growth observed in the data. To quantify the level of granularity generated by the standard model, I use reasonable parameters from trade literature to calibrate the standard model. I compare both the level of granularity in trade growth and the correlations of trade growth with production growth and trade intensity predicted in the model against those observed in the data.

I simulate the model for two main cases:

1. **The standard model with uniform tariff reductions.**

Firms in the model exhibit productivities drawn from a Pareto distribution, which are fixed for each good category across the initial and final periods. As is common in many standard trade models, trade growth is driven by trade liberalization between the initial and final periods that takes the form of a uniform proportional reduction in tariffs across all

goods.

2. The standard model with heterogeneous productivity changes and tariff reductions.

Firms again take Pareto productivities in the initial period. However, final period productivities are calculated by incorporating heterogeneous productivity changes to the Pareto distribution as calculated by the imputed changes in productivity backed out from the U.S. production data. This simulation also introduces trade liberalization between periods by incorporating heterogeneous tariff reductions across goods, imputed from tariff data.

2.5.1 Parameterization and Calibration

The numerical strategy is to select reasonable parameter values for the elasticity of substitution across goods classifications, productivity parameters, and fixed and variable costs of exporting, within the estimated ranges of the leading trade literature.

The model predictions are compared to those of the 5-digit SITC bilateral trade data with 1836 goods classifications, as well as the 6-digit NAICS data, with 473 classifications, in order to match production data. To represent these goods in the model, I draw each firm's productivity from a Pareto distribution, and map each firm to the production of one good.³⁷ I arrange firms according

³⁷The Pareto distribution is commonly used in a vast literature on international, (i.e. Melitz (2003), Chaney(2008), Gabaix (2008), Helpman, Melitz and Rubenstein (2008), among many others), both for its computational expediency and its ability to match certain stylized facts, such as the upper-right tail of the firm distribution, which motivates the choice of a Pareto distribution employed here. However, although not employed in this chapter, recent work has begun to explore alternate distributions, such as log-normal (Eeckhout (2004), Fernandes et al(2014), etc) or a mixed distribution of the two (Nigai (2017) that match other salient properties of firm-level productivity distributions.

to their heterogeneous productivity levels, in ascending order.³⁸ The iceberg cost, representing the variable cost of exporting, is set at 2.0. The elasticity of substitution across goods in the CES demand function is set toward the high end of estimated values in the literature, at 6.0, in order to give the model a better chance to match the granularity of growth across traded goods observed in the data.³⁹

The fixed cost in the model determines the productivity threshold beyond which firms will enter into the exporting market. As such, the fixed cost is calibrated to match the number of non-exported goods classifications in each bilateral trade pair. Finally, D^F , the constant which contains foreign income and price levels as well as other constants in the model set-up, is calibrated to match the total level of initial bilateral trade for each country-pair.

For the standard model with uniform tariff reductions, I calibrate the proportional reduction in tariffs (as represented by a reduction in the variable costs of exporting) necessary to match the total growth in exports for each country.⁴⁰

For the standard model with heterogeneous productivity and tariff changes, I first determine the distribution of productivity changes across goods classifications imputed from U.S. production data. I use the 6-digit NAICS data to

³⁸Thus “good 1” will be least productive firm, “good 2” will be the next least-productive firm, etc.

³⁹See, for example, Imbs and Mejean(2010), who estimate elasticities for more than 30 countries between 1 and 7.5, and supported by McDaniel and Balistreri (2003) who summarize the literature’s findings that estimated elasticities are higher with more highly disaggregated data.

⁴⁰Specifically, I take the observed change in AVE tariff rates from the data, and scale them by the factor necessary to match the observed change in exports for each bilateral country pair. This method implicitly assumes that D^F is constant across periods, thus all growth is generated via the uniform tariff reduction. However, since I am concerned with the share of trade growth accounted for by each good category, this method is computationally equivalent to choosing a set reduction in tariffs (i.e. 10%) and calibrating D^F to match the overall level of trade growth.

back out the implied productivities for each good in the U.S. production data, in both the initial and final periods. I then calculate the growth in productivity, and apply the imputed productivity change onto the initial Pareto distribution of firms in the model, in order to calculate the new productivities for the final period. I then determine the distribution of tariff changes across goods imputed from tariff data from U.S. export destinations (e.g. Canada, Germany, Mexico, Japan and the U.K.). I normalize this distribution of tariff changes and calibrate it to match the overall level of bilateral trade growth for each country-pair, as in the case of the uniform tariff distribution, given the productivity changes already imposed.⁴¹

2.5.2 Standard Model with Uniform Tariff Reductions

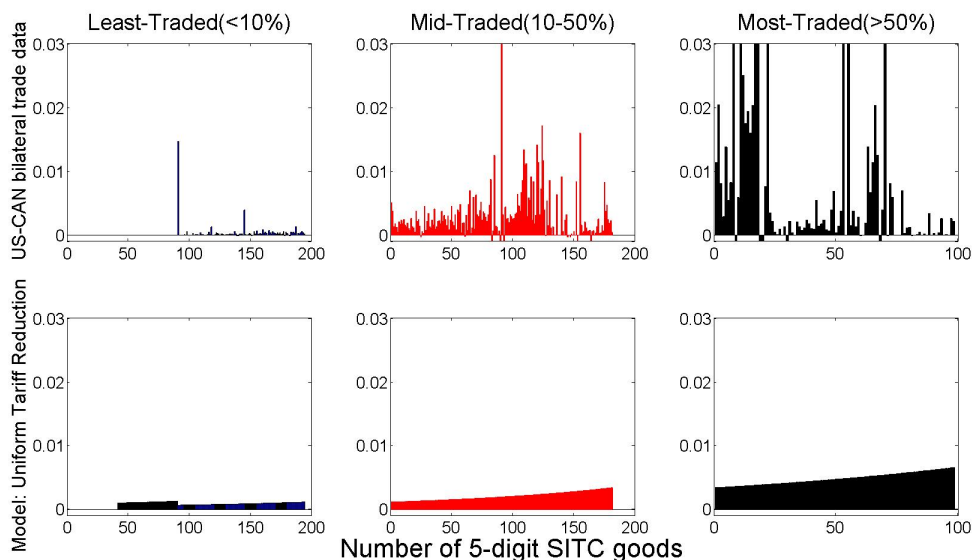
Figure 2.10 presents the results of the model simulation for Canadian exports to Mexico using the 5-digit SITC data.⁴² In the bottom row of each column, the standard model specification is unable to match the patterns of trade observed in the actual bilateral trade data presented in the top row. The overall granularity of trade growth, with the majority of growth concentrated in a small number of goods is not produced in the standard model with uniform tariff reductions. The model delivers smooth growth across goods, with each previously traded goods experiencing trade growth proportional to its initial level of trade, reflective of its productivity. Further, growth in the extensive margin, represented by the bottom decile of initial trade, occurs solely in previously traded goods (represented in blue), with virtually none of the overall trade growth

⁴¹Similar to the case of uniform tariff reductions, for the sake of calculating the share of trade growth across goods, this method is computationally equivalent to calibrating D^F to match the overall level of trade growth, or scaling the productivity changes proportionally to match the observed trade growth.

⁴²All other country-pairs present similar results, and are omitted here for brevity.

attributable to previously not-traded goods becoming traded.

Figure 2.10: Standard Model with uniform tariff reductions (Can-Mex)



To quantify the degree of granularity of trade growth, Table 2.8 reports the proportion of overall trade growth accounted for by various quantiles of goods classifications. On average, across country-pairs, the top 1% of goods classifications accounts for 32% of overall trade growth in the data, but only 3% of overall trade growth in this model simulation.⁴³ The standard model with uniform tariff reductions delivers only roughly 10% of the granularity observed in the data, as measured by the fraction of overall trade growth accounted for by the various quantiles of largest-growth goods.

This lack of granularity is not surprising — by construction, with all firms exhibiting fixed productivities and a uniform reduction in tariffs across goods, the majority of trade growth is evenly dispersed across previously traded goods.

⁴³It should be noted that this level of concentration is less than the 65% of trade growth accounted for by the top 1% of goods in the 5-digit SITC data. This is mainly due to the necessity of switching to the less disaggregated NAICS data (with only 473 goods) in order to impute productivities from U.S. production data, where 1% of goods is only 5 goods, as opposed to 18 goods in the SITC classification data.

Table 2.8: Standard Model with Uniform Tariff Reductions

Proportion of growth from top $X\%$ of goods (6-digit NAICS codes)					
Model	1% (5 goods)	2% (10 goods)	5% (24 goods)	10% (47 goods)	20% (94 goods)
Data	0.318	0.473	0.696	0.852	0.988
Standard	0.032	0.062	0.143	0.259	0.4490

The only disproportionate growth comes from the extensive margin — firms which cross the productivity threshold switch from not-traded to trade, but become traded only in relatively small volumes. Trade growth from these extensive margin goods is outweighed in the share of total trade growth by the increases in goods at the upper ends of the productivity distribution. In all cases, growth is smooth and proportional to productivity, and fails to produce the granularity of trade growth observed in the data.

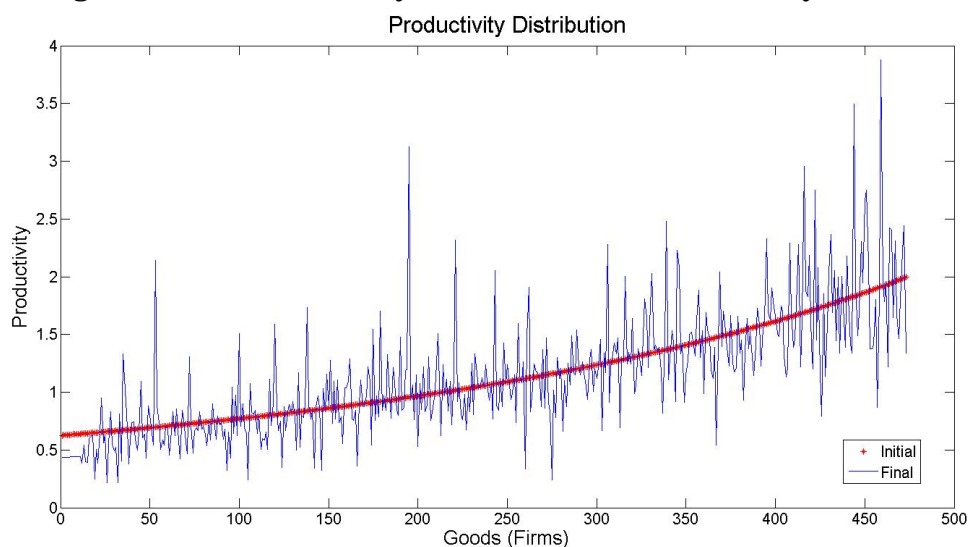
2.5.3 Standard Model with Heterogeneous Productivity Changes and Tariff Reductions

The standard model with uniform tariff reductions fails to generate the granularity in trade growth found in the trade data due to the smoothness in the productivity distribution, and the uniformity of tariff reductions with are common across all goods. A possible solution is adding heterogeneity in the productivity distribution and in the reductions in tariff rates across goods. If some goods categories experience large increases in productivity, it may lead to disproportionately larger growth in production and trade for those goods. Similarly, large decreases in tariff rates for some goods may lead to disproportionately larger growth in those categories. Combining these effects may be able to generate the level of granularity observed in the data, due to the large growth in a small subset of goods experiencing the complementarity of these effects. I re-simulate

the standard model incorporating heterogeneous productivity changes and tariff reductions imputed from U.S. production and export data, to determine if this augmented standard model can produce the documented patterns of trade growth.

Figure 2.11 demonstrates the distribution of productivity changes imputed from U.S. production data. There is a great deal of heterogeneity across the final distribution (in blue), once applied to the initial Pareto distribution (in red). This suggests there may be some ability for the model to capture a higher degree of granularity of trade growth as in the data, due to the relatively small number of goods exhibiting disproportionately large productivity shocks, arising from varying levels of initial trade.

Figure 2.11: Productivity Distributions: Productivity Shocks

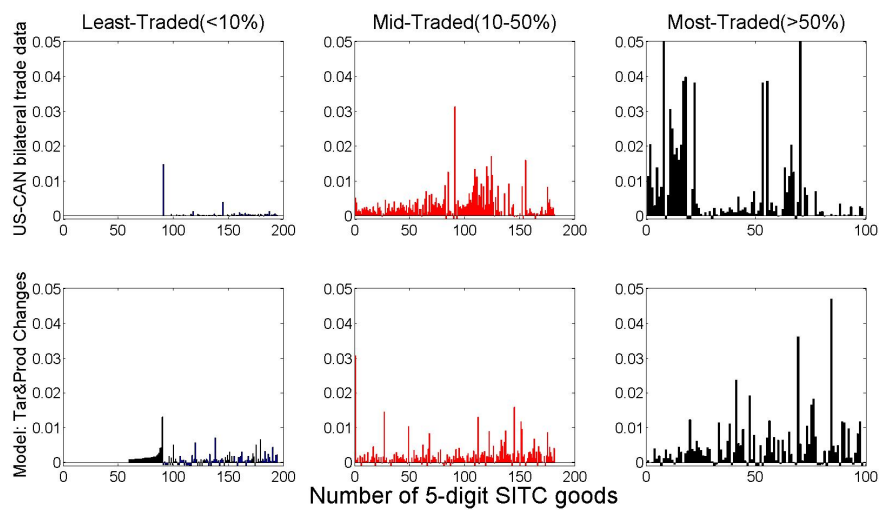


Once applied to the standard model, the heterogeneous productivity changes and tariff reductions generate a larger degree of granularity in trade growth, similar to that in the data. As can be seen in the bottom row of Figures 2.12(a)–2.12(b), the standard model with productivity shocks generates much

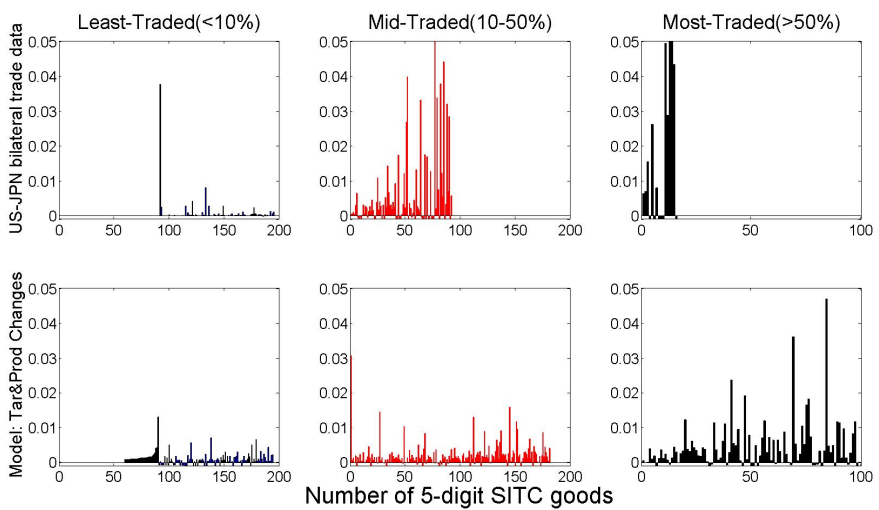
more heterogeneity across goods in the level of trade growth than the standard model, due primarily to the large degree of heterogeneity in the productivity shocks. Contrasting with the observed patterns of trade growth in the bilateral trade data in the top row, there are typically a larger number of large growth goods in the most-traded goods, and fewer large spikes among the mid-traded goods in the model than what is observed in the data. This model simulation does appear to generate similar patterns of trade growth among the least-traded goods, although there is a larger role of the traditional extensive margin goods — those switching from not-traded to traded — in the model than is observed in the data. Overall, there appear to be a larger number of goods exhibiting relatively large growth in the model, such that the degree of concentration of a large proportion of overall trade growth in a small number of goods is slightly less than in the data, and thus the magnitude of growth concentrated in each large-growth good is less in this model simulation.

Table 2.9 quantifies the degree of granularity generated in this model simulation. Across various country-pairs, including heterogeneous productivity changes and tariff reductions greatly improves the standard model's ability to generate the granularity observed in the data. Analyzing the share of trade growth accounted for by various quantiles of large-growth goods, the model now generates roughly 70% as much granularity as in the data, a marked improvement from the 10% generated in the standard model with uniform tariff reductions. In many cases, this improves to 80–90% of the granularity observed in the data as more goods are included (i.e. the top 10% or 20% of largest-growth goods).

Figure 2.12: Standard Model with Heterogeneous Productivity Changes and Tariff Reduction



(a) U.S.-Canada



(b) U.S.-Japan

Table 2.9: Heterogeneous Productivity Changes and Tariff Reductions

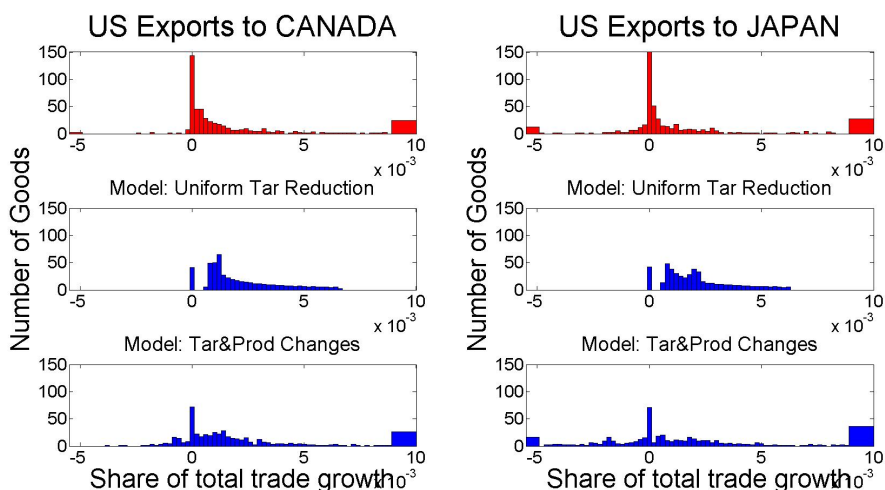
Share of growth from top $X\%$ of goods (6-digit NAICS codes)					
US Exports to Mexico					
Model	1%	2%	5%	10%	20%
Data	25%	36%	57%	74%	89%
Productivity & Tariff Changes	18%	26%	39%	53%	71%
US Exports to Canada					
Model	1%	2%	5%	10%	20%
Data	26%	38%	56%	71%	86%
Productivity & Tariff Changes	14%	22%	39%	58%	80%
US Exports to Japan					
Model	1%	2%	5%	10%	20%
Data	37%	56%	86%	101%	114%
Productivity & Tariff Changes	25%	37%	60%	82%	105%

2.5.4 Quantitative Analysis

To further illustrate the granularity generated by the standard model, I plot the histograms of the share of total trade growth accounted for by each good, for each of the various iterations of the model, in Figure 2.13. For all country-pairs, the high level of granularity in the trade data results in a bimodal distribution, with large spikes at small positive values, and at the top end of the distribution. This reflects the empirical findings that most goods grow very little, and the majority of trade growth is concentrated in a small subset of goods exhibiting large growth.

The standard model with uniform tariff reductions produces a flatter, smoother distribution of trade growth shares, reflecting much less granular growth across

Figure 2.13: Histograms of Trade Growth by Goods



goods, and much less growth at the top end of the distribution. The standard model with heterogeneous productivity changes and tariff reductions produces comparable growth at the top end of the distribution, reflecting the large growth among the small subset of goods accounting for the majority of trade growth. However, the remainder of the distribution is less granular than in the data, with growth more evenly dispersed across the remaining goods.

To quantify this dispersion of growth across goods, Table 2.10 provides summary statistics of trade growth across goods classifications for the various iterations of the model. The standard deviation gives an idea of the dispersion across goods in the share of trade growth accounted for by each good, while the kurtosis provides insight into the steepness of the distribution, both of which reflect the granularity of trade growth. Notably, the variation across goods in their share of total trade growth is higher in the data than in any of the model iterations, and the kurtosis is significantly higher. The standard model with heterogeneous productivity changes and tariff reductions performs best, with a standard deviation 65% as large and a kurtosis 45% of that of the data. The

Table 2.10: Summary Statistics: Growth by Goods — Data vs. Models

		Standard Deviation and Kurtosis				
Model		US/Can	US/Ger	US/Jpn	US/Mex	US/UK
Data	Std	0.0072	0.0094	0.0107	0.0065	0.0092
	Kurt	142.2410	34.0410	89.1526	62.6012	127.3663
Standard	Std	0.0016	0.0015	0.0015	0.0017	0.0016
	Kur	3.0950	3.0340	3.0825	3.0545	3.1281
Prod. shock	Std	0.0043	0.0077	0.0072	0.0037	0.0052
	Kur	39.9845	42.2558	42.2955	37.3227	41.5337

standard model with uniform tariff reductions generates standard deviations and kurtosis that are much lower than those observed in the data.

Table 2.11 reports correlation coefficients between the initial level of trade for each good and the resultant share of total trade growth accounted for by each good. In the data, this correlation is statistically insignificant, ranging from -0.02 to 0.05 across country-pairs, signifying that the high degree of concentration of growth in a small subset of goods is independent of those goods' initial trade values. In the standard model with uniform tariff reductions, in which goods have fixed productivities, this correlation is very high, greater than 0.85 for all country-pairs. This characterizes the “smooth” productivity distribution in the model, which results in all goods growing proportionally to the tariff decrease and their initial productivity. The model with heterogeneous productivity changes perform better, though the correlation coefficients are still statistically non-zero, and in some countries much higher than observed in the data. Again, the inherent smoothness in the initial productivity distribution passes through in at least some degree to the levels of initial and final trade, even when including heterogeneous productivity changes, generating stronger correlation between initial trade levels and trade growth across goods.

Table 2.11: Correlations: Initial Trade vs Share of Total Trade Growth

Model	US/Can	US/Ger	US/Jpn	US/Mex	US/UK
Data	-0.0113	-0.0063	-0.0193	0.0453	-0.0155
Standard	0.9795	0.9114	0.9255	0.9869	0.9674
Prod. shock	0.3620	0.1715	0.1886	0.4405	0.2903

2.6 Conclusions

This chapter identifies key facts of trade growth during episodes of large growth in bilateral trade. Bilateral trade growth is granular across goods classifications, with less than 5% of goods categories accounting for over 65% of total trade growth during these periods. Tariff changes do not account for the large growth in this subset of large growth goods, and many goods experiencing large drops in their AVE tariff rates exhibit little to no growth. For all U.S. exporting partners, increases in trade intensity, on average, account for a large share of total trade growth, while production growth, on average, accounts for significant portions of trade growth for some country-pairs, but virtually none of the overall trade growth for others.

Characterizing the predictions of the a Melitz-style model with uniform tariff reductions shows that it does not generate the granularity observed in the data. Trade growth is much less granular in this model, generating only 10% as much granularity as in the data, as measured by the share of total trade growth accounted for by various quantiles of goods classifications. Adding heterogeneous productivity changes and tariff reductions to the model generates a higher degree of granularity in trade growth in the model, capturing approximately 70% of the granularity observed in the data.

Further research is needed to match the model to the level of granularity

generated by the data. More detailed data on tariff and production levels, for a larger collection of countries, at this level of disaggregation across goods may improve the model's ability to match the observed level of granularity. Similarly, examining non-tariff trade barriers may help to account for the remaining 25–30% of trade growth concentrated in a small subset of large-growth goods. Finally, examining variation in the methods of transportation and distribution of exported goods may help account for the granularity of trade growth, as goods exported via different distribution methods may respond differently to trade liberalization, leading to a larger degree of granularity generated when these factors are considered in a Melitz-style model environment.

2.7 References

- ABEL-KOCH, J. (2013) "Who Uses Intermediaries in International Trade? Evidence from Firm-level Survey Data," *The World Economy*, 36(8),1041–1064.
- AHN, J., A. KHANDELWAL, AND S. WEI (2011). "The Role of Intermediaries in Facilitating Trade," *Journal of International Economics*, 84(1), 73–85.
- ALESSANDRIA, G., J. KABOSKI AND V. MIDRIGAN (2010). "Inventories, Lumpy Trade, and Large Devaluations". *American Economic Review*, 58(Dec), 2304–39.
- ANDERSON, J.E. AND E. VAN WINCOOP (2003). "Gravity with Gravitas: A Solution to the Border Puzzle". *American Economic Review*. 93(1). 170–192
- ARKOLAKIS, C. (2010). "Market Penetration Costs and the New Consumers Margin in International Trade". *Journal of Political Economy*, 118(Dec), 1151–99.
- ARMENTER, R., AND M. KOREN (2010). "A Balls-and-Bins Model of Trade". *American Economics Review* (forthcoming).
- BERGOEING, R., AND T. KEHOE (2001). "Trade Theory and Trade Facts". *Federal Reserve Bank of Minneapolis, Research Department*.
- CALIENDO, L. AND F. PARRA (2014). "Estimates of the Trade and Welfare Effects of NAFTA". *The Review of Economics Studies*. rdu035.
- CHANEY, T. (2008). "Distorted Gravity: The Intensive and Extensive Margins of International Trade". *American Economic Review*, 98:4, 1707–1721.

- DORNBUSCH, R., S. FISCHER, AND P. SAMUELSON (1977). “Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods”. *American Economic Review*, 67(5)823–39.
- FEENSTRA, R., J. ROMALIS, AND P. SCHOTT “U.S. Imports, Exports and Tariff Data, 1989–2001”, NBER Working Paper 9387.
- FERNANDES, A.M., P. KLENOW, S. MELESHCHUK, M.D. PIEROLA, AND A. RODRIGUEZ-CLARE (2015) “The Intensive Margin in Trade: Moving Beyond Pareto”, Working Paper.
- GABAIX, X. (2011) “The Granular Origins of Aggregate fluctuations.” *Econometrica* 79(3): 733–772.
- GOULD, D.M. (1998). “Has NAFTA Changed North American Trade?” *Economic Review- Federal Reserve Bank of Dallas* 12–23.
- HELPMAN, E., M.J. MELITZ, AND Y. RUBINSTEIN (2008). “Estimating Trade Flows: Trading Partners and Trading Volumes.” *The Quarterly Journal of Economics* 123(2): 441–487.
- HELPMAN, E., M.J. MELITZ AND S. YEAPLE (2004). “Export Versus FDI with Heterogeneous Firms”. *American Economic Review*, 94(1), 300–316.
- HORNOK, C., AND M. KOREN (2015). “Per-shipment Costs and the Lumpiness of International Trade.” *Review of Economics and Statistics*. 97(2), 525–530.
- IMBS, J., AND I. MJEAN (2010) “Trade Elasticities” Paris School of Economics, unpublished mimeo.
- KEHOE, T. (2005). “An Evaluation of the Performance of Applied General Equilibrium Models on the Impact of NAFTA”, in *Frontiers in Applied General Equilibrium Modeling: Essays in Honor of Herbert Scarf*, edited by Kehoe, Srinivasan and Whalley, 341–77. New York: Cambridge University Press.
- KEHOE, T. AND K. RUHL (2013). “How Important is the New Goods Margin in International Trade?”. *Journal of Political Economy*, 121 (2013), 358–92.
- KRUGMAN, P. (1980). “Scale economies, product differentiation, and the pattern of trade.” *The American Economic Review*, 950–959.
- MCDANIEL, C., AND E. BALISTRERI (2003) “A Review of Armington Trade Substitution Elasticities.” *Economie internationale* (2) 301–313.
- MELITZ, M.J. (2003). “The Impact of Trade on Intraindustry Reallocations and Aggregate Industry Productivity.” *Econometrica*, Vol. 71.6 (2003): 1695–1725.
- NAGURNEY, A. (2010). “Optimal Supply Chain Network Design and Redesign at Minimal Total Cost and with Demand Satisfaction”. *International Journal*

of Production Economics. 128, 200–208.

NIGAI, S. (2017). “A Tale of Two Tails: Productivity Distribution and the Gains from Trade.” *Journal of International Economics* (104): 44–62.

PIERCE, J. AND P. SCHOTT (2012). “A Concordance Between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries”. *Journal of Economic and Social Measurement*, 37(1–2), 61–96.

ROMALIS, J. (2007). NAFTAs and CUSFTAs Impact on International Trade, *Review of Economics and Statistics*, 89(3), pp.416–435.

SANTOSO, T., S. AHMED, M. GOETSCHALCKX, AND A. SHAPIRO (2005). “A Stochastic Programming Approach for Supply Chain Design Under Uncertainty”. *European Journal of Operation Research*, 167, 95–115.

SCHOTT, P. (2008). “The Relative Sophistication of Chinese Exports ”. *Economic Policy*, 53, 5–49.

Chapter 3

The Role of Multiple Distribution Technologies in Accounting for Trade Growth

3.1 Introduction

Fixed and variable costs of exporting have been found to play a significant role in accounting for international trade flows.¹ A wide literature establishes the role of geographic barriers, international borders, and tariffs and quotas in accounting for variation in bilateral trade growth across countries.² It is well-documented, in papers like Anderson (1979), or Parsley and Wei (1996), that trade flows generally decrease with distance — the larger the geographic distance, the higher the presumed transportation costs of shipping goods to foreign locations, resulting in smaller trade flows, *ceteris paribus*.

While a large literature examines the impact of transportation costs in accounting for bilateral trade flows across countries, many studies assume a

¹See, for example Melitz (2003), Chaney (2008), or Bernard et al (2007) for prominent examples.

²See Eaton, Kortum (2002), Engel, Rogers (1996), among others.

singular exporting technology, with a common fixed and variable cost. However, when exporters ship their goods internationally, they typically compare the costs of multiple available technologies to choose their optimal distribution networks.³ In many cases, these can be divided into two broad categories: (1) methods with low-fixed and high-variable costs, and (2) methods with high-fixed and low-variable costs. For example, firms exporting large volumes may opt to build their own infrastructure, complete with marine or air transport and offices at home and abroad to handle shipments, incurring large fixed costs, but relatively low per-unit costs. Conversely firms with relatively small or infrequent shipments may choose to export using an intermediary like FedEx, with minimal fixed costs, but higher per-unit costs incurred.

One of the main findings of Chapter 2 is that trade growth is granular — that the majority of bilateral trade growth is accounted for by a small number of large-growth goods classifications. I find that using a standard Melitz-style trade model, and including heterogeneous productivity and tariff changes imputed from U.S. production and tariff data, generates roughly 70% as much granularity as in bilateral trade data. Can incorporating heterogeneity in available exporting technologies help account for the remainder of this granularity? How do the differences in available transportation options, and their associated costs, affect suppliers exporting decisions and account for increased granularity in trade growth across goods?

To answer these questions, I build a model of bilateral trade with standard features such as heterogeneous productivities across firms, CES preferences and monopolistic competition pricing similar to Chapter 2. In this chapter, I

³For a general overview of recent developments in the strategic supply network formation literature, see Mills, Schmitz and Frizelle (2004), or work on optimal supply network design, such as Nagurney, Dong and Zhang (2004).

add a choice for exporters among multiple distribution technologies — one low-fixed, high variable cost method, and one high-fixed, low-variable cost method — to examine how variation in the available distribution technologies impacts bilateral trade flows and trade growth across goods classifications. This model framework generates three main channels for generating large trade growth across goods:

1. Large increases in productivity that lower the marginal costs of production
2. Large reductions in tariff rates that lower the marginal costs of servicing the foreign market
3. The ability for firms to switch distribution technologies, from a high variable cost method to a lower variable cost method

The contribution of this chapter is to examine the impact of this third channel. In Chapter 2, I find that a Melitz-style generates only 70% as much granularity in trade growth as bilateral trade data. Including this “switching” mechanism may allow the model to bridge this gap, which requires a higher concentration of overall trade growth in a smaller number of goods exhibiting disproportionately larger growth.

With CES demand and monopolistic competition pricing, the fixed costs of exporting generate productivity thresholds — the least productive firms do not export, firms with sufficiently high productivities export via the low-fixed cost method, and the most productive firms find it profitable to export via the higher-fixed cost method. As trade barriers fall (commonly represented by reductions in the variable costs of exporting), the productivity thresholds shift — some previously not-traded goods now become exported, and some previously

traded goods switch from the low-fixed/high-variable cost method, to the high-fixed/low-variable cost method. These “switchers” may exhibit disproportionately larger trade growth, due to double effect of the original reduction in trade costs, combined with the switch to a lower variable cost exporting method.

To quantitatively assess the ability of this channel to generate granularity in trade growth, I calibrate the model to match observed trade shares across goods in bilateral trade data. As in Chapter 2, I use 6-digit North American Industry Classification System (NAICS) data on U.S. exports to Canada, Japan, and Mexico between 1989 and 1999, to determine the share of cross-sectional trade and trade growth accounted for by each good category.⁴ I target the level of granularity in trade growth, which I quantify by calculating the share of bilateral trade growth accounted for by each of the 473 NAICS goods classifications. I find that trade growth is highly granular in the data, with the top 5% of large-growth goods accounting for, on average, 66% of bilateral trade growth.⁵ Most goods account for very little of the overall growth in trade — it is only a small number of goods, growing from small initial levels of trade to large levels of trade, or from large initial levels to much larger levels, that account for the majority of growth.

⁴The analysis in this chapter is confined to these three destination countries due to data limitations — in order to include heterogeneous productivity changes at this disaggregated level of goods classifications, I use 6-digit NAICS production data for the United States, making the U.S. the only source country for bilateral trade; in order to include heterogeneous tariff reductions at this level of disaggregation, I use 8-digit Harmonized Tariff Schedule (HTS) import tariff data, recorded to 6-digit NAICS classifications, that is only available during this period for Canada, Japan and Mexico.

⁵Further, I find that trade growth is more granular than cross-sectional trade — that is, the top 5% of large-trade goods categories accounts for 52% of trade in a given period, but the top 5% of large-growth categories accounts for, on average, 66% of trade *growth* across countries. Note: these top 5% categories need not be the same — the top 5% of traded goods in one period are not necessarily the largest growth goods in the next period.

I quantitatively assess whether this model with multiple distribution technologies can generate the high degree of granularity observed in the data. As in Chapter 2, each firm produces a single, unique good using labour as the sole input, indexed by its productivity, $\frac{1}{a}$, where a is the per-unit labor requirement, facing CES demand with monopolistically competitive markets. In this chapter, I add a discrete choice among multiple distribution technologies for exporting firms — X_1 : a low-fixed, high-variable cost method (f_1, τ_1) , and X_2 : a high-fixed, low-variable cost method (f_2, τ_2) . Given a productivity drawn from a Pareto distribution, firms profit maximize by deciding whether to export, and if so, choosing their optimal distribution technology.

Characterizing equilibrium, I find that with CES preferences, foreign demand for each good produced in the home country is inversely proportional to that good's price offered in the foreign market. Monopolistic competition between firms implies that each good is priced in the foreign market at a constant mark-up over marginal cost. This marginal cost has multiple components — the marginal cost of domestic production, a , which is a function of each firm's productivity; and the marginal cost of reaching the foreign market, which itself has two components — the variable costs of the distribution method employed, τ , and any tariffs imposed by the destination country, which is treated similar to a change in the variable transportation cost. However, with multiple distribution technologies, the variable distribution cost can take two values — τ_1 or τ_2 , allowing for greater heterogeneity across goods in the growth in trade generated by productivity changes or trade liberalization.

Exporting profits are thus a function of the firm's productivity and the fixed and variable costs of the chosen distribution method. The existence of fixed

costs of exporting generates productivity thresholds — only goods with sufficiently high productivities will find it profitable to export, and only goods with productivities high enough to cover the higher fixed cost will find it profitable to export via method X_2 , while all other exported goods will be exported via the lower-fixed cost method X_1 . Goods are stratified in the model according to their optimal distribution choice, and trade flows are a function of the relative productivities and export costs across goods.

To quantify the granularity of trade growth generated by this model, I first solve the model for given values of fixed and variable costs, productivities and tariff rates, to determine export sales and optimal distribution choices across goods. To analyze growth in this static model, I follow a common approach of representing trade liberalization as a fall in trade costs which reduce the variable costs of exporting, and then re-solve the model to determine changes in the model's predictions for export sales and distribution choice.

With CES demand, monopolistic competition pricing, and a single distribution technology, a fall in trade costs results in a proportional increase in exports for each previously-traded good. However, with multiple distribution technologies, the fall in variable costs also shifts the productivity thresholds for the high- and low-fixed cost distribution methods. Some previously not-traded goods become traded, as the reduction in variable costs makes paying the fixed cost to export profitable. This is extensive margin growth, and leads to disproportionately larger growth than the size of the tariff reduction for these goods, as they grow from zero to larger share of total trade by crossing the exporting threshold. Further, some goods previously traded via the low-fixed cost method now profitably switch to the higher-fixed cost method, accessing the lower variable cost. These “switchers” exhibit disproportionately larger export growth

due to a double effect — the direct effect of a reduction in the variable costs due to tariff decreases, and the indirect effect of switching to a distribution method with a lower variable cost.

To compare the predictions of the multiple distribution technologies model to the data, I calibrate and simulate the model and calculate the share of total trade growth accounted for by each good. I first simulate the model with each good classification drawing a heterogeneous productivity from a Pareto distribution, and determine the share of trade accounted for by each good. I re-simulate the model, adding heterogeneous productivity changes for each good that I impute from U.S. production data, and heterogeneous tariff reductions concorded from destination-country tariff data. The changes in the model's predictions for each good's level of exports between the two simulations are then compared to the changes in trade for each good in the data.

A key issue is how to calibrate the fixed and variable costs, f_1 , f_2 and τ_1 , τ_2 . Since data on direct measures of shipping and distribution costs is limited at this level of disaggregation, I calibrate the fixed and variable costs in the initial simulation to match the granularity of cross-sectional trade in the data for the 1989–1991 average for each country-pair. For a given τ_1 , the lower fixed costs, f_1 is calibrated to match the number of not-traded goods categories.⁶ The higher fixed cost, f_2 and lower variable cost τ_2 are calibrated to best match the granularity in cross-sectional trade, as measured by the share of total trade accounted for by top quantiles of goods categories (i.e. the top 2, 5, and 10%) in a given period. I retain these values for the fixed and variable costs in the second simulation.

Simulating the model, I find the multiple distribution technologies model

⁶This is determined by the equation for the lower productivity cut-off, Eq.(5).

closely matches the degree of granularity in cross-sectional trade. On average across country-pairs, the top 2, 5 and 10% of goods categories, respectively, account for 33, 56 and 64% of cross-sectional trade, as compared to 35, 52 and 67% of cross-sectional trade in the data. Achieving this granularity requires a relatively high ratio of fixed costs, with $\frac{f_2}{f_1}$ calibrated anywhere from 670 for U.S. exports to Mexico, up to 1840 for U.S. exports to Japan. The ratio of variable costs for the two methods is fairly consistent, calibrated near $\frac{\tau_2}{\tau_1}=0.60$ for all country-pairs.

I find that the two-method model generates a higher degree of granularity, with the top quantiles of goods categories accounting for roughly 90–95% of the share of total trade growth as in the corresponding trade data. Specifically, the top 2, 5 and 10% of goods categories account for 43, 57 and 81%, respectively, of total trade growth in the two-method model, compared to 43, 66 and 82% in the data. Conversely, performing similar model simulations with only a single distribution technology, as in Chapter 2, accounts for only 23, 46 and 64% of total trade growth, roughly 60–70% of the granularity observed in the data.

Examining the choices of distribution methods in the simulations, I find that switching behaviour generates increased granularity in trade growth in the model. Goods that begin and remain traded via method X_1 account for an average of less than 0.01% of total trade growth. Goods that begin and remain traded via method X_2 contribute a larger share of total trade growth, on average 0.40% of total trade growth for each good in this group. However, goods that switch from not-traded to traded via either method contribute an average of 0.39% of total trade growth, while goods that switch distribution methods from X_1 to X_2 contribute an average of 0.60% of total trade growth. This supports the theory that the ability for firms to choose among multiple distribution

technologies in the exporting process increases the granularity predicted by the model, more closely matching the level of granularity observed in bilateral trade data.

3.2 Related Literature

As in Chapter 2, this chapter contributes to the “trade lumpiness” literature, in multiple ways.⁷ First, this chapter provides a novel mechanism of incorporating multiple exporting technologies into a standard Melitz-style trade model, that is capable of matching a high degree (roughly 90–95%) of the granularity of cross-sectional trade observed in the data. Second, a central contribution of Chapter 2 was to document the fact that trade growth is more granular than cross-sectional trade, and that the set of large-growth goods is uncorrelated with the set of goods that were previously most highly traded. Including multiple exporting technologies in the model generates a mechanism whereby firms may switch distribution technologies, leading to disproportionately larger trade-growth than non-switching firms, and helping to account for the granularity of both cross-sectional trade and trade growth over time.

This chapter builds a framework similar to that of Helpman, Melitz and Yeaple (2004), in which firms face a choice of exporting versus foreign direct investment (FDI). In their model, firms face different fixed costs of exporting and FDI, and become stratified: domestic-only producers, firms productive enough to pay the relatively lower fixed costs of exporting, and firms that are most productive and can pay the relatively higher fixed costs of FDI to profit from the lower marginal costs of production in FDI rather than exporting. However,

⁷See Armenter, Koren (2010), Hornok, Koren (2015) or Alessandria, Kaboski and Midraggan (2010), for example.

since FDI is not observed in trade data as exports, this mechanism would not be present in accounting for the granularity of export growth. This chapter introduces a similar mechanism, with goods stratified according to their optimal distribution technology, but which allows for the model to account for the granularity of trade growth observed in export data.

A wide literature investigates optimal supply network design, to identify how firms choose their distribution networks when confronted with options across multiple distribution technologies. Nagurney (2010), and subsequent papers related to this work, investigate the formation, and potential re-optimization, of firms' optimal supply networks. This work focuses on channels such as the role of excess capacity — firms strategically choose distribution network capacity and usage, to dynamically optimize around potential changes in the economy. Similarly, Santoso et al (2005) propose an algorithmic approach, similar to Nagurney, for solving optimal supply network design across firms. While this chapter abstracts from innovation in new supply network design or capacity constraints, it contributes to the supply network literature by quantitatively assessing the impact of distribution options in accounting for observed patterns of trade growth.

Another literature related to optimal distribution networks focuses on the role of wholesalers as intermediaries in international trade. Abel-Koch (2013) uses firm-level data to empirically examine the relationship of firm size and production to the use of trade intermediaries in Turkish exporting firms, and find that intermediary use is decreasing in firm size, and that newly traded goods are more likely to use trade intermediaries to export their products. Ahn, Khandlwal and Wei (2011) incorporate an intermediary sector into a model with heterogeneous firms, and find that firms are stratified into non-traded,

use of trade intermediary, and direct exporting firms, according to productivity. This chapter finds a similar result of the stratification across the multiple technologies in firms' distribution choice; however, these previous works largely seek to match data on cross-sectional trade, not on trade growth. Further, they primarily focus on matching which firms employ each method of trade, and do not examine the possible change in these optimal choices following a reductions in trade barriers and the resultant impact on trade growth across firms. This chapter extends this literature by allowing for possible changes in distribution choice across multiple exporting technologies as trade barriers fall or other structural changes may take place.

A large literature highlights the role of heterogeneity in accounting for trade growth in international trade models. However, the focus of these models is typically in matching overall trade flows and growth in trade. This chapter focuses not only on matching overall bilateral trade growth, but also the distribution of growth across goods, in matching the observed granularity of bilateral trade data. Melitz (2003) and Chaney (2008), both pre-eminent works in this literature, document the role of heterogeneous productivities across firms, as well as fixed and variable costs of exporting, to identify the roles of the intensive and extensive margins in accounting for overall trade growth. Freund and Pierola (2015) empirically determine that the majority of cross-sectional trade can be attributed to a small number of the largest firms in a given sector, which can account for variation in the sectoral distribution of exports relative to income across countries. Arkolakis (2010) focuses on market penetration, building a model in which firms essentially choose their fixed cost in order to access a foreign market. Depending on this choice of fixed costs, they then face varying options of increasing marginal costs to reach additional consumers and increase sales. However, as documented in Chapter 1, even adding heterogeneous

productivity and tariff changes to these types of models does not generate sufficient granularity of trade growth. By adding a discrete choice across multiple distribution technologies in the exporting process, allowing some firms to switch technologies and leading to disproportionately larger growth for these goods, the model in this chapter is better able to match the granularity of trade growth across goods observed in the data.

3.3 Data

To decompose trade growth across goods, I use 6-digit North American Industry Classification System (NAICS) data for U.S. exports to Canada, Mexico and Japan, from 1989 to 1999. Unlike Chapter 2, the bilateral trade data is confined to this subset of country-pairs by data limitations. In order to include heterogeneous productivity changes imputed from production data, I confine the analysis to U.S. exports, as data at this level of disaggregation is only available for U.S. production. Similarly, in order to include heterogeneous tariff changes across goods, I confine the analysis to U.S. exports to Canada, Mexico and Japan, as tariff data at this level of disaggregation is only available for this subset of the countries used in Chapter 2 for this time period.

To combine these disparate data sources, I map all data to the 6-digit NAICS classification system. While these limitations restrict the analysis to a smaller set of countries, I am still able to analyze trade growth across goods at a relatively high level of disaggregation, and for countries that experience formal trade liberalization (U.S.-Canada, U.S.-Mexico) as well as those that do not (U.S.-Japan).

In Chapter 2, I used 5-digit Standard International Trade Classification

(SITC) data to document the granularity of trade growth across goods. One key distinction between the two goods classification systems is the level of disaggregation — there are 1836 goods categories in the 5-digit SITC data, while the NAICS data is less disaggregated, with only 473 goods classifications. Since the level of disaggregation has a potential impact on the level of granularity, I re-calculate the shares of total trade growth accounted for by top quantiles of goods categories for each country-pair in the 6-digit NAICS data.

As in Chapter 2, I use a 3-year average for measuring trade flows, to account for issues such as shipping delays and customs reporting irregularities. For each 6-digit NAICS good, I calculate the average export value from 1989–1991, which I label as the “initial” trade value, and similarly calculate the average export value from 1997–1999, which I label as the “final” trade value. Trade growth is thus measured as the difference in these three-year averages for each good.

3.3.1 Bilateral Trade Data

Descriptive Statistics

Table 3.1 presents summary statistics for U.S. exports to Canada, Mexico and Japan, between 1989 and 1999. Countries experiencing formal trade liberalization (U.S./Canada, U.S./Mexico around NAFTA) exhibit larger total growth in trade than those non-liberalizing countries (U.S./Japan). Additionally, this trade growth represents a larger percentage increase in trade over initial levels, with U.S. exports to Canada and Mexico growing by 115% and 189% respectively, while U.S. exports to Japan grew by 44% over this period.

Table 3.1: Summary Statistics: Bilateral Trade Data
(6-digit NAICS codes)

Variable	USA-CAN	USA-MEX	USA-JPN
Initial Trade	65196.99	24536.51	34976.65
Final Trade	139880.12	70818.53	50320.17
Trade Growth	74683.13	46282.01	15343.51
% Δ in Trade	115%	189%	44%
Share of initial trade from top $X\%$ of goods			
2%	37%	27%	41%
5%	53%	45%	59%
10%	66%	62%	74%
Share of trade growth from top $X\%$ of goods			
2%	38%	36%	56%
5%	56%	57%	86%
10%	72%	74%	102%
Median share of trade growth	0.05%	0.03%	0.03%
Number of goods: 473			
(All trade values in thousands of \$US)			

Trade Growth Granularity

I re-calculate the share of total trade growth between country-pairs that is accounted for by each 6-digit NAICS classification. Table 3.1 reports the share of total trade growth accounted for by the top 2, 5 and 10% of goods. At this lower level of disaggregation, trade growth remains granular, with a small number of goods (roughly 10, 25 and 50 goods) accounting for 43%, 66% and 82% of total trade growth, respectively, but is not as pronounced as in the 5-digit SITC data from Chapter 2.⁸ However, trade growth is still highly concentrated in a small number of goods, while the majority of goods account for very little of the overall growth in trade. Further, for each country-pair, the median share of trade growth per good (0.03–0.05%) is well below the mean share of trade growth (2.1% for each of the 473 goods categories), signifying a highly skewed distribution of trade growth across goods. Finally, trade growth remains more granular than cross-sectional trade — the top 2, 5 and 10% of goods categories account for 43%, 66%, and 82% of trade growth, but only 35%, 52% and 67% of trade in the initial period.⁹

Figure 3.1 illustrates the granularity of trade growth across goods at this lower level of disaggregation. As in Chapter 2, this is shown by the small number of spikes representing goods exhibiting large growth shares, while most goods exhibit little negligible growth shares.¹⁰ Further, trade growth shares remain uncorrelated with initial levels of trade — the small number of goods

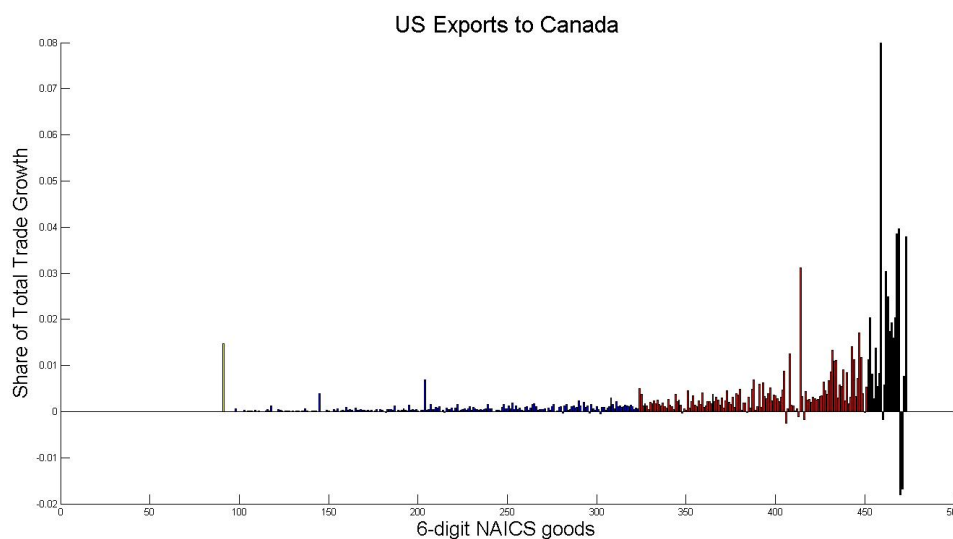
⁸For example, in the SITC data, the top 1% of goods categories accounted for roughly 60% of total trade growth, on average across country-pairs — accounting for 60% of trade growth in the NAICS data would require 2–5% of largest-growth goods categories.

⁹These 2, 5 and 10% of goods categories need not be the same in accounting for trade growth versus cross-sectional trade — that is, the most-traded goods in the initial period are not necessarily the same goods that exhibit the largest growth in trade.

¹⁰Figure 3.1 illustrates this case for U.S. exports to Canada — similar examples for U.S. exports to Mexico and Japan can be found in Appendix B.

accounting for the majority of overall trade are not necessarily just the largest, or smallest, initially traded-goods, but rather arise from varying levels of initial trade.

Figure 3.1: Growth by Goods Category: US Exports to Canada



3.3.2 Production Data

To quantitatively assess the the model’s predictions for the trade growth across goods, I include heterogeneous productivity changes, imputed from U.S. production data, and compare the granularity of trade growth predicted by the model to that observed in the bilateral trade data. I use 6-digit NAICS data on U.S. production from the NBER-CES Manufacturing Industry Database, for the initial period (1989–1991 average) and the final period (1997–1999 average), to calculate the changes in imputed productivity between periods. I match the changes in productivity to the data on bilateral trade growth for each 6-digit NAICS good category.

Figure 3.2: Productivity Distributions: Implied Productivities

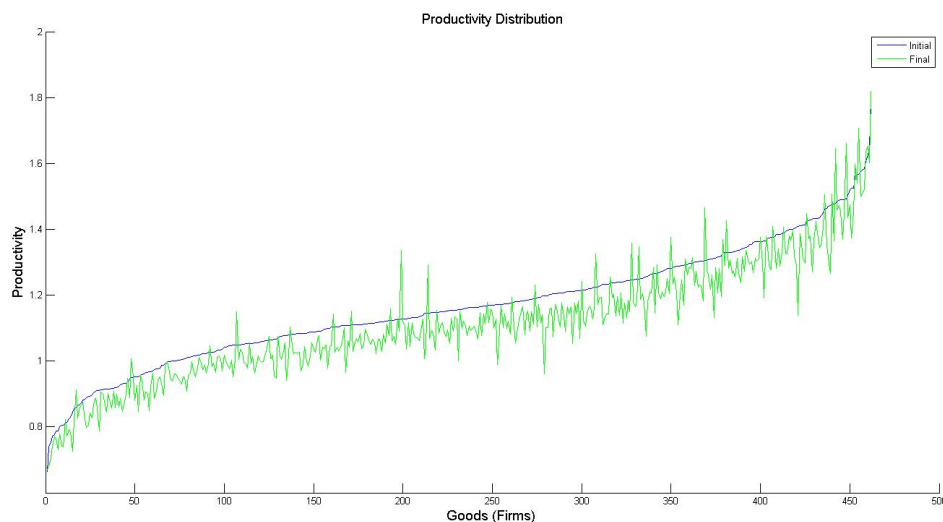


Figure 3.2 demonstrates the changes in the imputed productivities for U.S. production. For each good category, I use the observed level of domestic production to back out the implied productivity from the production function for the initial period and arrange goods in order of ascending productivity. I repeat the process for the final period and map the change in imputed productivity to each corresponding good from the initial period productivity distribution.¹¹ There is a large degree of heterogeneity of productivity changes across goods. Further, the changes in productivity are uncorrelated with the initial productivities — that is, it is not just the most productive firms that become even more productive, or vice-versa.

¹¹I also normalize the changes in productivity to mean zero to highlight the heterogeneity in productivity changes across goods, rather than merely an increase in mean productivity.

3.3.3 Tariff Data

As in Chapter 2, I analyze the role of heterogeneous tariff changes in accounting for trade growth granularity in a standard trade model. To do so, I map changes in the ad valorem equivalent (AVE) tariff rates for U.S. export destinations to each good category in the model. I use 8-digit Harmonized Tariff Schedule data from the World Bank’s World Integrated Trade Solution (WITS) database on AVE import tariff rates on U.S. exports for Canada, Mexico and Japan, from 1989–1991 and 1997–1999 to match the initial and final periods in the bilateral trade data. I use concordances provided by Pierce and Schott (2012) to map the 8-digit HTS tariff rates to the 6-digit NAICS classifications.

Figure 3.3: Heterogeneous Tariff Changes: Canada

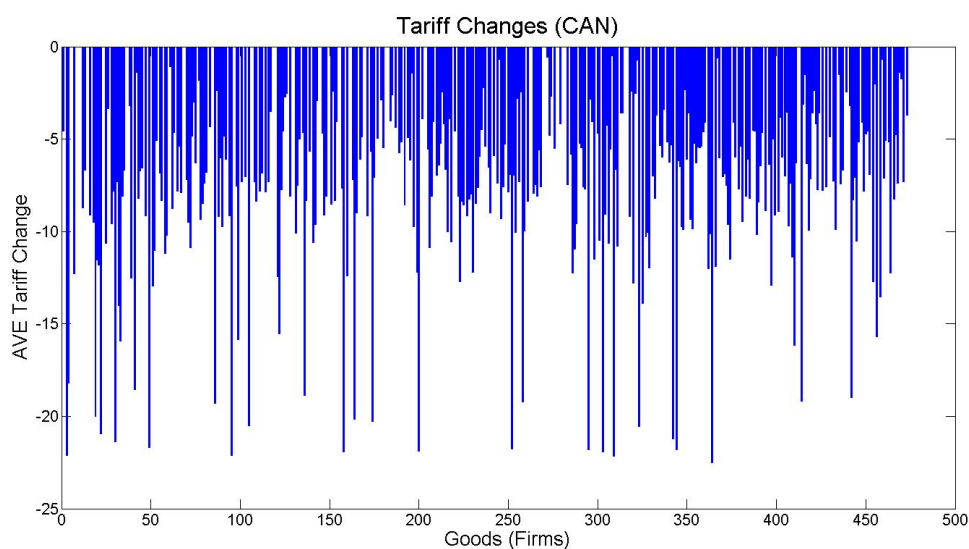


Figure 3.3 shows the changes in the AVE tariff rates between the initial and final periods for each concorded 6-digit NAICS good category for Canada.¹²

¹²Similar tariff changes for Mexico and Japan are shown in Figures B.4–B.6.

There is a large degree of heterogeneity across goods in the AVE tariff changes, with most goods falling somewhere in the range between zero change and a 25% reduction tariffs.¹³ The mean reduction in tariffs are largest in Canada, followed by Mexico and finally Japan, reinforcing that countries engaged in formal trade liberalization policy over this period (U.S./Canada, U.S./Mexico) have larger overall reductions in tariffs than those countries that lack a formal trade liberalization agreement (U.S./Japan). Further, there are a larger proportion of goods categories exhibiting large decreases in tariffs for Canada and Mexico, and relatively more goods with small or zero reductions in tariffs for Japan.

3.4 Model

To account for the granularity of trade growth in the data, I extend the Melitz-style model framework from Chapter 2 by including a choice among two exporting technologies — one low-fixed, high-variable cost method, and one high-fixed, low-variable cost method. The model includes heterogeneous productivities across firms, monopolistic competition, and CES preferences for consumers. I characterize consumer demand and firms' profit maximization in equilibrium, including firms' export decisions and distribution method choice.

As in Chapter 2, I use a static two-country model, and characterize partial equilibrium to predict export sales of home-produced goods in the foreign market. To analyze growth in this static environment, I solve the model for two cases: first, I solve the model with a Pareto distribution of firms facing CES demand, under monopolistic competition, given a choice among the two

¹³There is a notable exception, with an outlier of Mexican tariffs on U.S. exports of NAICS 311213 — “Malt Manufacturing, from barley, rye or other grains” rising by over 130%.

distribution technologies— this will serve to represent the “Initial period”; second, I re-solve the model, incorporating heterogeneous productivity and tariff changes across goods, under the same choice of distribution methods for exporting firms — this will serve as the “Final period”. Trade growth for each good is calculated as the changes in the model’s predictions for exports between the two cases.

In the Initial period, I find that the model generates productivity thresholds for exporting, and goods are stratified according to their initial productivities. The least productive goods are not exported, as the fixed costs of exporting make exporting unprofitable regardless of distribution method. Goods with sufficiently high productivity find it profitable to export via the low-fixed, high variable cost method. The most productive goods exhibit the largest export sales, finding it more profitable to export the high-fixed, low-variable cost method.

In the Final period, I find that adding productivity and tariff changes across goods generates a large degree of heterogeneity in trade growth across goods. The model produces three channels for disproportionately larger growth in certain goods, generating the granularity of trade growth:

1. Goods with large increases in productivity
2. Goods with large decreases in tariffs, represented by a reduction in the variable costs of exporting
3. The ability for firms to switch distribution technologies, from a high-variable to a low-variable cost method

With productivity increases or tariff reductions (or a combination of the two), some goods that were previously not-traded now find exporting profitable, via

the low-fixed cost method, or via the high-fixed cost method if these shocks are sufficiently large. Similarly, some goods that were previously exported via the low-fixed cost method now find it profitable to switch to the high-fixed cost method, with a corresponding reduction in variable costs. These “switching” goods exhibit disproportionately larger growth, as they benefit directly from the productivity increase or tariff reduction, but also indirectly from the choice of switching to a lower-variable cost method, leading to larger export sales. This “switching” mechanism generates more granularity than the standard model with a single distribution technology, as a larger share of total trade growth is concentrated in these disproportionately larger-growth goods.

3.4.1 Model Set-up

The model builds on the standard Melitz-style framework, adding a discrete choice among multiple exporting technologies. I use a static, two-country model with symmetric countries, Home (H) and Foreign (F). As in Chapter 2, firms are heterogeneous in their productivities, $\frac{1}{a}$ where a is the per-unit labour requirement. Firms engage in monopolistic competition pricing, facing CES preferences generating consumer demand from each country. Firms choose whether to sell to the foreign market by utilizing one of two distribution methods — one with a lower fixed and higher variable (iceberg) cost of exporting, and the other with a higher fixed and lower variable cost of exporting.

Firms

I assume each firm produces a differentiated good, so that each good category in the model is synonymous with one firm, indexed by its per-unit labour requirement, a . Firms use labour as the single input, paying the domestic wage

rate ω^i to produce in country $i \in \{H, F\}$. For simplicity, I consider the case of Home exports to Foreign, though the symmetric problem is equivalent. Firms draw their productivity, $\frac{1}{a}$, from a Pareto distribution, $G(\frac{1}{a})$. Firms engage in monopolistic competition pricing, as in Dixit-Stiglitz (1977), facing CES demand from each country. Given foreign demand for their product, $c^F(a)$, firms choose whether or not to service the foreign market. If they choose to export, firms must choose their optimal distribution technology among:

1. Method $X_1(f_1, \tau_1)$: a low-fixed, high-variable cost option
2. Method $X_2(f_2, \tau_2)$: a high-fixed, low-variable cost option

For a firm with labour requirement a , given foreign demand $c^F(a)$, the firm's profit maximization problem is:

$$\max_{X_j, j \in \{1, 2\}} \{ \max_{p(a)} \{ p(a)c^F(a) - h(a, X_j)c^F(a) \} \} \quad (3.1)$$

where $h(a, X_j)$ is the cost function, depending on the choice of distribution method $X_j, j \in \{1, 2\}$. These costs include the fixed cost of exporting, f_j , as well as the marginal cost of exporting, which embeds the marginal cost of production, $\omega^H a$, and the variable cost of exporting, τ_j . The variable cost can be further decomposed into two components: the shipping costs of physically transporting goods between locations, and trade barrier costs, such as tariffs. By assumption, there is no fixed cost for domestic production.

Consumers

Preferences of consumers in country $i \in \{H, F\}$ are represented by the CES utility function,

$$U^i = \left(\int_{a \in A} c^i(a)^{\frac{\epsilon-1}{\epsilon}} da \right)^{\frac{\epsilon}{\epsilon-1}} \quad (3.2)$$

where $c^i(a)$ is consumption of good variety a , ϵ is the elasticity of substitution across goods ($\epsilon > 1$) and A is the set of available goods. Since the two countries are symmetric, I solve the case of Home exports to Foreign, and therefore only consider the utility function U^F to derive demand for Home-produced goods in the Foreign market, $c^F(a)$.

3.4.2 Characterizing Equilibrium

I solve the model to characterize equilibrium export flows for all goods varieties, $a \in A$. With no fixed cost of production, all goods are produced in Home, and consumers in Foreign have demand for these varieties that is given by the well-known demand function for CES preferences:

$$c^F(a) = E^F \frac{p(a)^{-\epsilon}}{P_F^{1-\epsilon}} \quad (3.3)$$

where E^F is national income in Foreign and P_F is the aggregate price level in Foreign. With $\epsilon > 1$, demand in the Foreign market, $c^F(a)$, is inversely proportional to the price offered in the Foreign market for each Home-produced good.

Monopolistic competition implies that firms price at a constant mark-up over marginal cost. The marginal cost of servicing the foreign market has multiple components. The marginal cost of domestic production is $\omega^H a$ — however, in order to have a full unit reach the foreign market, the firm must ship $\tau > 1$ units (where τ is commonly referred to as the iceberg transportation cost). Thus the marginal cost of exporting is $\omega^H \tau_j a$, where $j \in \{1, 2\}$ depending on the distribution method employed by the firm. This results in a price in Foreign for Home-produced variety a of:

$$p^F(a, X_j) = \left(\frac{\epsilon}{\epsilon - 1} \right) \omega^H \tau_j a \quad (3.4)$$

Export sales, $p^F(a, X_j)c^F(a)$, are proportional to productivity and inversely proportional to the variable cost of exporting:

$$ExSales(a, X_j) = \left(\frac{\epsilon}{\epsilon - 1} \right) \frac{E^F}{P_F^{1-\epsilon}} [\omega^H \tau_j a]^{1-\epsilon} \quad (3.5)$$

The exporting profits for a firm with labour requirement a using distribution technology $X_j, j \in \{1, 2\}$ with associated fixed cost f_j and variable cost τ_j are:

$$\pi(a, X_j) = \left[\frac{1}{\epsilon - 1} \right] \frac{E^F}{P_F^{1-\epsilon}} [\omega^H \tau_j a]^{1-\epsilon} - \omega^H f_j \quad (3.6)$$

Productivity thresholds

To export via either distribution method, exporting profits must be non-negative:

$\pi(a, X_j) \geq 0$. Re-arranging the profit function yields a productivity threshold $\frac{1}{\bar{a}_1}$, that satisfies:

$$\left(\frac{1}{\bar{a}_1} \right) = \left[\frac{(\omega^H)^\epsilon f_1 \tau_1^{\epsilon-1}}{D^F} \right]^{\frac{1}{\epsilon-1}} \quad (3.7)$$

where D^F is a constant with respect to a .¹⁴ This implies that firms must be sufficiently productive to cover at least the lower fixed cost of exporting, f_1 , to make exporting profitable.¹⁵

There is also a threshold that determines which firms will choose to export via method X_2 rather than method X_1 , where $\pi(a, X_2) > \pi(a, X_1)$. Solving this inequality yields the productivity threshold $\frac{1}{\bar{a}_2}$, that satisfies:

$$\left(\frac{1}{\bar{a}_2} \right) = \left[\frac{\omega^H (f_2 - f_1)}{D^F (\omega^H \tau_2)^{1-\epsilon} - (\omega^H \tau_1)^{1-\epsilon}} \right]^{\frac{1}{\epsilon-1}} \quad (3.8)$$

¹⁴ $D^F = \left[\frac{1}{\epsilon-1} \right] \frac{E^F}{P_F^{1-\epsilon}}$

¹⁵It is assumed that (f_1, τ_1) and (f_2, τ_2) are such that $\frac{1}{\bar{a}_1}$ is lower for method X_1 than method X_2 . Thus, by assumption, $\frac{f_1}{f_2} < \left(\frac{\tau_2}{\tau_1} \right)^{\epsilon-1}$.

Firms with sufficiently high productivity will choose to export via the higher-fixed cost method, X_2 , while firms with lower productivity will choose to export via method X_1 , as long as their productivity is such that paying the lower-fixed cost of exporting still yields positive exporting profits.¹⁶

Figure 3.4: Initial level of export sales

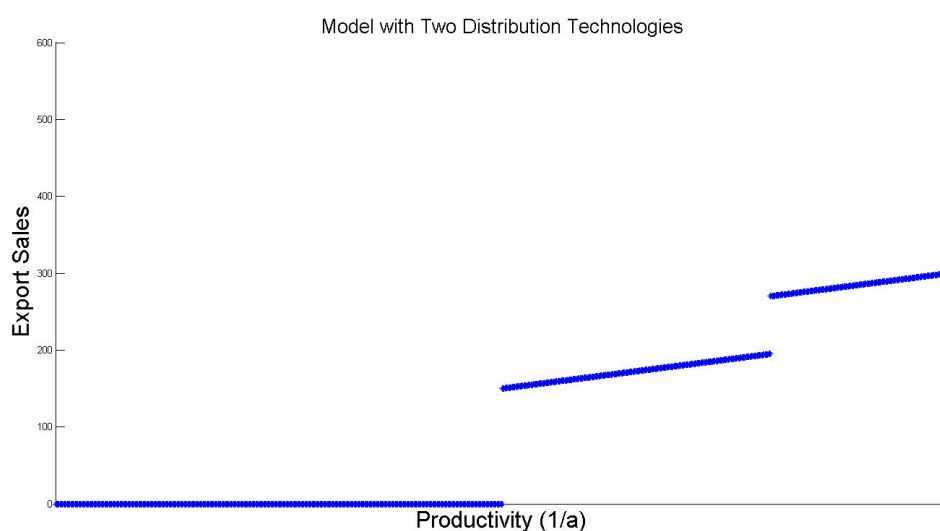


Figure 3.4 demonstrates the stratification of export sales across goods, by productivity. Firms with low productivity draws do not export. Firms with sufficiently high productivity find exporting profitable after paying the fixed cost, with sales increasing in productivity. The less productive of these firms will export via method X_1 , while the more productive firms will export via method X_2 . There is a discontinuous jump in export sales at the upper productivity threshold $\frac{1}{a_2}$, with the more productive firms benefiting from the lower variable cost, leading to higher export sales.

¹⁶It is assumed that $\frac{1}{a_2} > \frac{1}{a_1}$, which further disciplines the relationship of the fixed and variable costs: $f_2 > f_1 \left[1 - \left(\frac{\tau_1}{\omega^H} \right)^{\epsilon-1} (\tau_2^{1-\epsilon} - \tau_1^{1-\epsilon}) \right]$.

3.4.3 Comparative Statics: Tariff and Productivity Changes


To analyze trade growth in the static model, I solve the model twice — for an “Initial” period and for a “Final” period. I calculate the difference in predicted export sales across the two cases to represent trade growth for each good. To determine the Initial period’s export sales, I solve the model with a random productivity draw from a Pareto distribution for each firm producing good variety a and solve for each firm’s profit maximization export decision and distribution method choice. To determine the Final period’s export sales, I re-solve the model while incorporating heterogeneous changes in productivity and tariffs across goods, imputed from bilateral trade data. I impute the changes in productivity from U.S. production data between 1989–1991 and 1997–1999, and apply the corresponding change directly to the productivity of each good, $\frac{1}{a}$ from the Initial period. Similarly, I impute changes in AVE tariff rates for each good from U.S. export destination tariff rates, which takes the form of a proportional reduction in the variable costs of exporting, τ_1 or τ_2 , relative to the Initial period’s values.

Across goods, an increase in productivity or a decrease in tariffs lowers the price offered in the foreign market, as given in Equation 3.4. The direct effect is a proportional increase in exports, as given by Equation 3.5 due to the heterogeneous changes in productivity and tariff rates across goods.

Changes in productivity and tariff rates also have an indirect effect on export sales by altering the productivity thresholds that dictate the optimal distribution choice for each good. Equations 3.7 and 3.8 show that a reduction in tariffs which decreases the variable costs of exporting, τ_1 and τ_2 , lowers the productivity thresholds, $\frac{1}{\bar{a}_1}$ and $\frac{1}{\bar{a}_2}$. This allows some goods that were below the exporting threshold, $\frac{1}{\bar{a}_1}$, to become exported, and some goods that were below

the upper threshold, $\frac{1}{a_2}$, to switch from exporting via method X_1 to X_2 . The former is categorized as extensive margin growth— previously not-traded goods become exported. For the latter, firms that switch from X_1 to X_2 also benefit from the reduction in variable costs by moving from τ_1 to τ_2 , bringing a secondary increase in export sales, as in Equation 3.5.¹⁷

Figure 3.5: Final level of export sales



With three distribution possibilities, $\{0, X_1, X_2\}$ in each period, there are 9 potential combinations for distribution method choice across the two periods. Figure 3.5 shows the stratification of these possibilities. With a reduction in the variable costs of exporting from a decrease in tariffs, export sales will increase proportionally across the productivity distribution for each method. Additionally, the productivity thresholds will shift, as represented by the shaded regions in Figure 3.5. Trade growth can be decomposed as follows: Goods that

¹⁷It should be noted that if productivity decreases or tariffs increase, the converse is true—the productivity thresholds increase and sales decrease, with firms potentially “switching down” in their distribution choices.

begin with and retain the lowest productivities will remain not-traded. Some previously traded goods will increase export sales proportionally to the reduction in tariffs, while continuing to use the same distribution method, either X_1 or X_2 . However, some previously not-traded goods may now cross the lower productivity threshold and become exported via method X_1 .¹⁸ Similarly, with the reduction in variable costs of exporting, some goods that were previously traded via method X_1 now find it profitable to switch to method X_2 (as shown in the blue shaded region). This increase in export sales is disproportionately larger than it would be at similar productivity levels where the choice of distribution method does not change (as in the white shaded regions to the left and right).

Adding changes in productivity moves firms across the productivity distribution, which potentially generates more granularity in trade growth. Firms with a sufficiently large increase in productivity may “jump” across a productivity threshold in Figure 3.5. For example, a firm with low productivity may choose not to export in the initial period, but an increase in productivity may move them into a region of the distribution where exporting becomes profitable via X_1 , or via X_2 if the increase in productivity is sufficiently large. Similarly, goods initially exported via X_1 may receive a sufficient increase in productivity to move to a portion of the distribution where switching to X_2 is profitable, bringing a disproportionately larger growth in trade than the direct effect of the productivity increase.¹⁹

The model now provides three channels for generating heterogeneity in

¹⁸While not shown explicitly in 3.5, it is theoretically possible for goods to jump from not-traded to exported via method X_2 if the tariff reductions are sufficiently large.

¹⁹Again, it should be noted that although the cases presented here deals with tariff decreases and productivity increases, the converse holds for cases of tariff increases and productivity decreases and the resultant decreases in export sales across goods.

trade growth across goods. First, heterogeneous tariff changes, represented as changes in the variable cost of exporting for both distribution technologies, changes the price offered in the foreign market and results in proportionally larger export sales for the goods exhibiting the largest tariff reductions. Second, heterogeneous productivity changes, represented as changes to the variable cost of domestic production, similarly result in heterogeneous changes in the price offered in the foreign market and proportionally larger export sales for the goods exhibiting the largest productivity increases. Third, these first two channels may lead to certain goods crossing the productivity thresholds, resulting in disproportionately larger growth for these “switchers” — goods that switch from not-traded to traded via either technology, or switching from method X_1 to X_2 . This occurs due to the combined effect of these goods exhibiting increases in export sales proportional to the growth in the first two channels being magnified by a switch to a lower-variable cost method, amplifying the growth in trade. The larger growth exhibited by this subset of goods increases the granularity of trade growth predicted by the model, compared to the model with a singular distribution technology.

3.5 Quantitative Analysis

To quantitatively assess the model’s ability to match the granularity in the data, I calibrate and simulate the model with multiple distribution technologies to match bilateral trade flows. I first simulate the model and calibrate the fixed and variable costs of exporting to match U.S. exports to Canada, Mexico and Japan, for the 3-year average from 1989–1991, representing the “Initial period”. I re-simulate the model, adding heterogeneous productivity and tariff changes imputed from U.S. data, to match bilateral trade data for each

country-pair over a three-year average from 1997–1999, representing the “Final period”. I calibrate the distribution of tariff reductions for the Final period to match the overall growth in trade for each country-pair between 1989–1991 and 1997–1999 averages.

One key issue in simulating the model is how to implement the fixed and variable costs $(f_1, \tau_1; f_2, \tau_2)$ associated with the different distribution methods available to exporters. Since direct data on variations in shipping rates across goods is highly limited at this level of disaggregation, I calibrate these parameters to match the granularity of cross-sectional trade flows. Specifically, I calibrate the relative fixed and variable costs for the two methods, X_1 and X_2 , to match the shares of total trade accounted for by the top 2, 5 and 10% of goods categories in the Initial period. I then take these costs as given when re-simulating the model for the Final period.

In Chapter 2 I found that including heterogeneous productivity and tariff changes in the Melitz-style model accounted for roughly 70% of the observed granularity of trade growth in bilateral trade data. Adding a choice among multiple distribution technologies to this model framework, I find that the model now captures roughly 90–95% of the granularity in the data. Further, I find evidence of switching behaviour in the model simulations, which increases the level of granularity in predicted trade growth. Although a relatively small number of firms choose to switch technologies, with the majority of overall trade growth occurring from goods retaining their distribution choice across periods, I find evidence that on average, each switching firm accounts for a greater share of trade growth (0.5–1.0% per good) than each non-switching firm (0–0.5% per good), increasing the granularity of trade growth across goods.

3.5.1 Parameterization and Calibration

The calibration strategy in this chapter is similar to that of Chapter 2, with one notable exception — I now have two additional free parameters to calibrate: f_2 and τ_2 , the costs associated with the second distribution option X_2 . As in Chapter 2, I first parameterize and calibrate the model to match bilateral trade data on U.S. exports to Canada, Mexico and Japan for the 3-year average from 1989–1991, which I classify as the Initial period. To analyze changes in trade flows across goods, I re-simulate the model, adding heterogeneous productivity and tariff changes imputed from U.S. data, and calibrating the model to match bilateral flows for U.S. exports to Canada, Mexico and Japan for 1997–1999, which I classify as the Final period. To incorporate these heterogeneous productivity and tariff changes, I calibrate the model to match data at the 6-digit North American Industry Classification System (NAICS) level, resulting in 473 distinct goods classifications.

For the Initial period, I draw each firm’s initial productivity from a Pareto distribution, $G(\frac{1}{\alpha})$, mapping each firm to a single good category, and set the elasticity of substitution, ϵ , at 6.0.²⁰ Since my variable of interest is the share of growth accounted for by each good category (not the absolute level of trade flows for each category), I begin with τ_1 , the higher variable cost, as a free parameter, as in Chapter 2. For a given τ_1 , the lower fixed cost, f_1 , which determines the lower productivity threshold for exporting, is calibrated to match the number of not-traded goods in the data. I then simultaneously calibrate the higher fixed cost f_2 , and lower variable cost, τ_2 , to match the cross-sectional trade flows accounted for by the top 2, 5 and 10% of goods in the Initial period

²⁰Refer to Chapter 2, Section 5.2.1 for a discussion on these parameter choices and implications.

(1989-1991 average for each destination country in the data).²¹

Similar to Chapter 2, I use U.S. manufacturing data to impute changes in productivity across goods between the Initial and Final periods, and apply these changes to the productivity distribution from the Initial period simulation.²² I use 8-digit Harmonized Tariff Schedule (HTS) data, which I concord to 6-digit NAICS classifications, to determine the changes in ad valorem equivalent (AVE) tariff changes for Canada, Mexico and Japan between the 1989–1991 and 1997–1999 periods.²³ I scale the distribution of tariff changes to match total trade growth for each U.S. export destination for the Final period, 1997–1999, retaining the same set of fixed and variable cost options calibrated in the Initial period, to determine exports, distribution choice, and share of total trade growth accounted for by each good category.²⁴

3.5.2 Initial Period

Table 3.2 presents the results of the model calibration for each destination country. For the top 2, 5, and 10% of goods categories, the model captures 93–95% of the overall granularity in cross-sectional trade, on average across

²¹Specifically I grid search over a range of all plausible values of f_2 and τ_2 consistent with the literature, to determine the parameter values that minimize the loss function over various quantiles of trade, $\|\vec{r}_{data} - \vec{r}_{model}\|$, where the \vec{r} 's are vectors of the proportion of total trade accounted for by the top 2, 5 and 10% of goods. Therefore, in keeping with the focus on matching the share of trade growth accounted for across goods in the data, the calibration determines the *ratio* of fixed and variable costs between methods X_1 and X_2 .

²²Figure B.3 shows the distribution of Initial productivities along with the applied productivity changes.

²³Figures B.4–B.6 show the distribution of tariff changes across goods classifications. Of note, the mean tariff decreases are larger in Canada and Mexico, due to the formal trade liberalization of NAFTA, as opposed to Japan. However, while there are many more goods exhibiting large AVE decreases for Canada and Mexico, there are still large AVE reductions for some goods categories in Japan, in the absence of formal trade liberalization agreements.

²⁴As in Chapter 2, this is computationally equivalent to calibrating D^F to match the overall level of trade growth.

country-pairs. For all country-pairs, the ratio of the two calibrated variable costs, $\frac{\tau_2}{\tau_1}$, is consistently between 0.55–0.65.²⁵ There is more variation in the ratio of calibrated fixed costs, $\frac{f_2}{f_1}$, across country-pairs. For U.S. exports to Mexico, where the top quantiles of goods account for lower shares of Initial trade, the ratio is lower, at 670. For U.S. exports to Canada and Japan, where the trade shares accounted for by the upper quantiles of goods is higher, the ratio of fixed costs is higher, at 1430 and 1840, respectively.²⁶ These results suggest that a higher fixed cost ratio creates a higher productivity threshold for firms using method X_2 , which, combined with a lower variable cost τ_2 , concentrates a larger share of Initial trade in a smaller proportion of goods categories at the top end of the distribution.

3.5.3 Final Period

Table 3.3 reports the share of trade growth accounted for by quantiles of largest-growth goods categories for each U.S. export destination. The share of trade growth accounted for by the top 2, 5 and 10% of goods is roughly 90–95% as high as that observed in the data, on average across country-pairs. This marks an improvement from the roughly 70% of granularity generated by the model with a single distribution technology, as in Chapter 2.

²⁵With the higher variable cost, τ_1 , set as a free parameter at 2.0, this results in values for τ_2 across countries of 1.12–1.32, consistent with the interpretation in the literature of the iceberg costs shipping, where τ represents the number of units necessary to ship to result in 1 unit arriving at the destination country.

²⁶Due to the scarcity of direct measures of the fixed costs of exporting, it is difficult to assess the appropriateness of these fixed cost ratios — the vast majority of the relatively small literature that seeks to directly measure distribution and transportation costs in international trade goes no deeper than identifying total costs, without decomposing the shares of fixed vs. variable costs, and virtually all assume a common distribution technology. However, Kropf and Saure (2014) estimate variation across Swiss firms in fixed costs per shipment, and find logged values that range from -3 to 9 — thus the fixed cost ratios I calibrate between 670 and 1840 correspond to logged differences of 6.5 to 7.5, which arguably fall within realistic ranges of those imputed by Kropf and Saure.

Table 3.2: Cross-Sectional Granularity: Initial Period

Share of trade from top $X\%$ of goods (6-digit NAICS codes)						
	USA-CAN		USA-JPN		USA-MEX	
	Model	Data	Model	Data	Model	Data
2%	37	37	41	41	21	27
5%	57	53	63	59	49	45
10%	62	66	68	74	62	62
20%	71	81	75	88	71	79
40%	84	94	87	97	84	94
$\frac{f_2}{f_1}$	1430		1840		670	
$\frac{\tau_2}{\tau_1}$	0.59		0.56		0.66	

Additionally, I track the distribution choice for each good across the two simulations to quantitatively assess the ability of the model’s “switching” mechanism to account for trade growth granularity. Figure 3.6 plots the distribution choice across goods, arranged by ascending initial productivity. Firms are initially stratified according to productivity, with the least productive firms not exporting, the most productive firms exporting via method X_2 , and those in between the productivity thresholds exporting via method X_1 , represented by the starred levels. After including the productivity and tariff changes from the data, and simulating the model, the bars represent the new distribution choice for each good. Many previously not-traded goods remain not-traded, but some now become traded via X_1 or X_2 (as reflected by the bars ascending to 1 or 2). Similarly, a large proportion of goods previously traded via method X_1 remain traded via X_1 ; however, some now switch to method X_2 (as reflected by the bar ascending above the starred line) while others no longer export (as reflected by a blank space for that good). The same is true for goods initially traded via X_2 , where the majority remain traded via X_2 , but some goods switch to either X_1

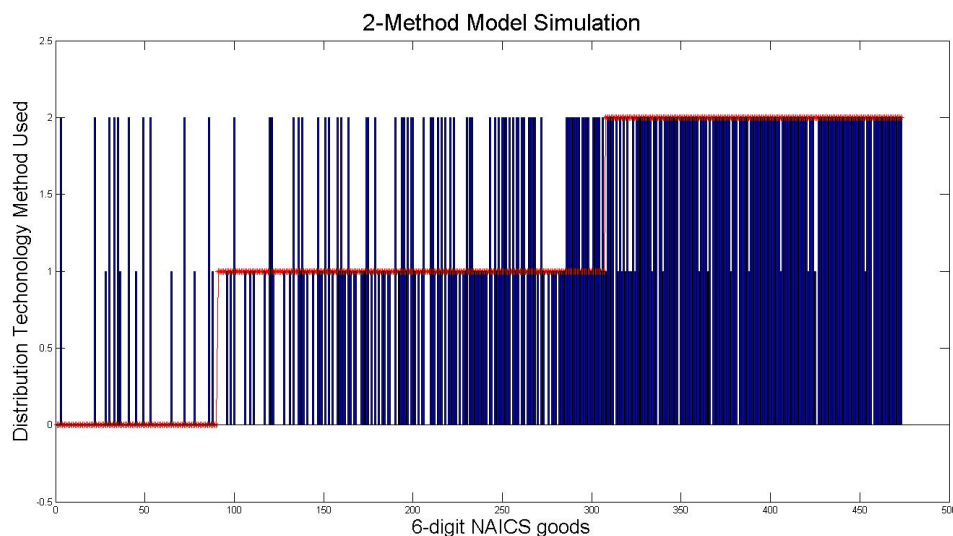
Table 3.3: Trade Growth Granularity: Second Period

Share of growth from top $X\%$ of goods (6-digit NAICS codes)						
	CAN-USA		JPN-USA		MEX-USA	
	Model	Data	Model	Data	Model	Data
2%	33	38	59	56	37	36
5%	48	56	74	86	50	57
10%	75	71	97	102	70	74
20%	92	86	116	114	87	89
40%	111	99	138	123	104	101
$\frac{f_2}{f_1}$	1430		1840		670	
$\frac{\tau_2}{\tau_1}$	0.59		0.56		0.66	

(a space down to level 1) or not-traded (a space down to 0).

Comparing across the two simulations, I find that firms switching from distribution technology 1 to 2 account for a larger shares of total trade growth than the traditional extensive margin of firms switching from non-traded to traded. Table 3.4 reports the shares of trade growth accounted for by each of the 9 possible combinations of distribution choices across the two simulations. While there are a smaller total number of firms switching from not-traded to traded, or from method X_1 to X_2 than there are firms that retain the same distribution choice, these switching firms account for relatively larger shares of trade growth, on average. Firms switching from not-traded to traded account for an average of roughly 0.3–0.7% of total trade growth per good, while firms switching from X_1 to X_2 account for an average of 0.5–1.0% of total trade growth per good, across the various country pairs. Conversely, firms that retain X_1 in both simulations account for virtually none of the total trade growth, and while firms that retain X_2 account for 50–60% of the total trade growth, the fact that there are relatively so many of these firms (140–150 per country)

Figure 3.6: Model Simulation: Optimal Distribution Choice



means that on average, they account for only 0.3–0.5% of total trade growth per good.

Table 3.4: Optimal Distribution Method Choice

Methods	Proportion of total trade growth								
	U.S.-Canada			U.S.-Japan			U.S.-Mexico		
	Tot. Share	# goods	Avg. Share	Tot. Share	# goods	Avg. Share	Tot. Share	# goods	Avg. Share
0→0	0%	74	0%	0%	78	0%	0%	74	0%
0→ X_1	0.03%	6	0.00%	0.03%	6	0.01%	0.05%	6	0.01%
0→ X_2	6.14%	10	0.61%	8.98%	7	1.28%	4.60%	10	0.46%
X_1 →0	-0.27%	77	0.00%	-0.50%	81	-0.01%	-0.36%	78	0.00%
X_1 → X_1	0.08%	74	0.00%	-0.08%	67	0.00%	0.32%	71	0.00%
X_1 → X_2	39.61%	66	0.60%	61.41%	62	0.99%	32.22%	69	0.47%
X_2 →0	-3.31%	7	-0.47%	-6.34%	5	-1.27%	-1.99%	7	-0.28%
X_2 → X_1	-7.06%	17	-0.42%	-15.88%	15	-1.06%	-3.95%	16	-0.25%
X_2 → X_2	64.78%	142	0.46%	52.39%	152	0.34%	69.11%	142	0.49%

It is also important to note that in Table 3.4 it is possible for firms to “switch down” in the model — firms with decreases in productivity or increases in tariffs may choose to switch from X_2 to X_1 following the inverse rationale of firms

that “switch up”, or may even decide to stop exporting altogether. While there are a smaller number of these types of down-switching firms than up-switching firms, their presence in the model serves to amplify the granularity of trade growth, as the up-switching firms account for an even larger share of total trade growth when these decreases in exports from down-switching firms are taken into consideration.

3.6 Conclusions

This chapter extends the literature by identifying a novel mechanism that helps account for the high degree of granularity of trade growth observed in bilateral trade data. I use data U.S. exports to Canada, Mexico and Japan, between 1989 and 1999, at the 6-digit NAICS level, to determine the distribution of trade growth shares across goods. I find that trade growth is granular — a small number of goods categories accounts for a majority of total bilateral trade growth- specifically, the top 5% of goods accounts for roughly 66% of overall trade growth, on average across country-pairs. Further, I find that trade growth is more granular than cross-sectional trade, and that trade growth is uncorrelated with previous levels of trade.

To match this level of granularity in the data, I use a standard trade model framework, and add a discrete choice among multiple distribution technologies for firms in the exporting process. I introduce a low-fixed, high-variable cost exporting technology, and a high-fixed, low-variable cost exporting technology, and include heterogeneous productivity and tariff changes, imputed from trade data, to analyze the predictions for trade growth across goods in the model. I calibrate the fixed and variable costs, and the distribution of productivity and tariff changes, to match overall bilateral trade flows, and simulate the model

to determine the distribution of growth in export sales predicted by the model.

I find that the model simulation generates roughly 90–95% of the observed granularity of trade growth in the bilateral trade data, as measured by the share of total trade growth accounted for by the top quantiles of goods categories. The model with multiple distribution technologies increases the granularity of trade growth in the model, as compared to a standard model with a single distribution technology, which generates only 60–70% as much granularity as the data. In the model, I find evidence that goods that switch their distribution technology in response to heterogeneous productivity changes and tariff reductions account for a relatively larger share of total trade growth than non-switching goods. This occurs as a result of the double effect of a direct reduction in variable costs of exporting from tariff reductions and productivity increases, combined with an indirect effect of switching to a lower variable cost method of exporting, resulting in disproportionately higher growth, accounting for the observed granularity of trade growth.

3.7 References

- ABELKOCH, J. (2013). “Who Uses Intermediaries in International Trade? Evidence from Firmlevel Survey Data”. *The World Economy*. 36(8), 1041–1064.
- AHN, J., A. KHANDELWAL, AND S. WEI (2011). “The Role of Intermediaries in Facilitating Trade,” *Journal of International Economics*. 84(1), 73–85.
- ALESSANDRIA, G., J. KABOSKI AND V. MIDRIGAN (2010). “Inventories, Lumpy Trade, and Large Devaluations”. *American Economic Review*. 58(Dec), 2304–39.
- ANDERSON, J.E. (1979). “A Theoretical Foundation for the Gravity Equation”. *The American Economic Review*. 69(1), 106–116.
- ANDERSON, J.E. AND E. VAN WINCOOP (2003). “Gravity with Gravitas: A Solution to the Border Puzzle”. *American Economic Review*. 93(1). 170–192
- ANDERSON, J.E. & Y.V. YOTOV (2010). “The Changing Incidence of Geography,” *American Economic Review, American Economic Association*. Vol. 100(5),

pages 2157–86

ARKOLAKIS, C. (2010). “Market Penetration Costs and the New Consumers Margin in International Trade”. *Journal of Political Economy*. 118(Dec), 1151–99.

ARMENTER, R. AND M. KOREN (2014). “A Balls-and-Bins Model of Trade”. *The American Economic Review*. 104(7), 2127–2151.

BERGOEING, R., AND T. KEHOE (2001). “Trade Theory and Trade Facts”. *Federal Reserve Bank of Minneapolis, Research Department*.

BERNARD, A.B., J. EATON, J.B. JENSEN, AND S. KORTUM (2003). “Plants and productivity in international trade”. *The American Economic Review*. 93(4), 1268–1290.

BLUM, B., S. CLARO, AND I. HORSTMANN (2016). “A Balls-and-Bins Model of Trade: Comment.” *American Economic Review* 106(3), 843–51.

CALIENDO, L. AND F. PARRA (2014). “Estimates of the Trade and Welfare Effects of NAFTA”. *The Review of Economics Studies*. rdu035.

CHANEY, T. (2008). “Distorted Gravity: The Intensive and Extensive Margins of International Trade”. *American Economic Review*. 98:4, 1707–1721.

CRUCINI, M., C. TELMER AND M. ZACHARIADIS (2003). “Price Dispersion: The role of borders, distance and location”, *Tepper School of Business*. Paper 490

DORNBUSCH, R., S. FISCHER, AND P. SAMUELSON (1977). “Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods”. *American Economic Review*. 67(5)823–39.

EATON, J. AND S. KORTUM (2002). “Technology, Geography, and Trade”. *Econometrica*, 70(5), 1741–1779.

ENGEL, C. AND J. ROGERS (1994). “How Wide is the Border?” *American Economic Review*. 86(5), 1112–1125.

FEENSTRA, R., J. ROMALIS, AND P. SCHOTT. (2002). “U.S. Imports, Exports and Tariff Data, 1989–2001”, NBER Working Paper 9387.

FERNANDES, A.M., P. KLENOW, S. MELESHCHUK, M. DENISSE PIEROLA AND A. RODRIGUEZ-CLARE (2015). “The Intensive Margin in Trade: Moving Beyond Pareto”, Working Paper.

FREUND, C. AND M. PIEROLA (2015). “Export Superstars” *Review of Economics and Statistics*, 97(5), pp. 1023–1032

GOULD, D. (1998). “Has NAFTA Changed North American Trade?” *Economic Review- Federal Reserve Bank of Dallas* 12–23.

- HELPMAN, E., M.J. MELITZ AND S. YEAPLE (2004). "Export Versus FDI with Heterogeneous Firms". *American Economic Review*, 94(1), 300–316.
- HORNOK, C., AND M. KOREN (2015). "Per-shipment Costs and the Lumpiness of International Trade." *Review of Economics and Statistics*. 97(2), 525–530.
- KEHOE, T. (2005). "An Evaluation of the Performance of Applied General Equilibrium Modes on the Impact of NAFTA", in *Frontiers in Applied General Equilibrium Modeling: Essays in Honor of Herbert Scarf*, edited by Kehoe, Srinivasan and Whalley, 341–77. New York: Cambridge University Press.
- KEHOE, T. AND K. RUHL (2013). "How Important is the New Goods Margin in International Trade?". *Journal of Political Economy*, 121 (2013), 358–92.
- KROPF, A., AND P. SAUR (2014). "Fixed costs per shipment." *Journal of International Economics* 92(1): 166–184.
- KRUGMAN, P. (1980). "Scale Economies, Product Differentiation, and the Pattern of Trade." *The American Economic Review*, 950–959.
- MELITZ, M.J. (2003). "The Impact of Trade on Intraindustry Reallocations and Aggregate Industry Productivity." *Econometrica*, Vol. 71.6 (2003): 1695–1725.
- MILLS, J., J. SCHMITZ, AND G. FRIZELLE (2004). "A Strategic Review of Supply Networks". *International Journal of Operations & Production Management*, 24(10), 1012–1036.
- NAGURNEY, A. (2010). "Optimal Supply Chain Network Design and Redesign at Minimal Total Cost and with Demand Satisfaction". *International Journal of Production Economics*. 128(1), 200–208.
- NAGURNEY, A., J. DONG AND D. ZHANG (2002). "A Supply Chain Network Equilibrium Model". *Transportation Research Part E: Logistics and Transportation Review*. 38(5), 281–303.
- PARSLEY, D.C., AND S.J. WEI (1996). "Convergence to the Law of One Price Without Trade Barriers or Currency Fluctuations". *The Quarterly Journal of Economics*, 111(4), 1211–1236.
- PIERCE, J. AND P. SCHOTT (2012). "A Concordance Between Ten-Digit U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries". *Journal of Economic and Social Measurement*, 37(1–2), 61–96.
- ROMALIS, J. (2007). NAFTAs and CUSFTAs Impact on International Trade, *Review of Economics and Statistics*, 89(3), pp.416–435.
- SANTOSO, T., S. AHMED, M. GOETSCHALCKX AND A. SHAPIRO (2005). "A Stochastic Programming Approach for Supply Chain Design Under Uncertainty". *European Journal of Operation Research*, 167, 95–115.

SCHOTT, P. (2008). “The Relative Sophistication of Chinese Exports ”. *Economic Policy*, 53, 5–49.

SHEN, Z.J.M. (2006). “A Profit-Maximizing Supply Chain Network Design Model with Demand Choice Flexibility”. *Operations Research Letters*, 34(6), 673–682.

Chapter 4

The Supply Network and Price Dispersion in the Canadian Gasoline Market

4.1 Introduction

Policymakers and consumers have long sought to understand why prices for identical goods differ across locations. Many studies use geographic distance, as a proxy for transportation costs, to account for large price differences of homogeneous goods across locations.¹ However, little is known about the impact of variation in the methods used to transport goods between locations on these relative price differences. In this chapter, the term “supply network” refers to the different modes of transporting products between locations. For suppliers servicing various locations, the associated transportation costs will be a function of not just the geographical distance covered, but also of the costs associated with the various methods of transportation for reaching each potential destination. The structure of the supply network therefore plays a key role in

¹See Burdett and Judd (1983), Crucini, Telmer and Zachariadis (2003), or Engel and Rogers (1996) among others.

limiting arbitrage opportunities and determining what level of price dispersion can be sustained across locations over time.

This chapter quantifies the impact of the supply network on relative price dispersion in the Canadian gasoline market. The Canadian gasoline industry is large — the average Canadian household spends over \$2000 per year on gasoline for transportation, accounting for over 3% of household spending.² Gasoline is a homogeneous good for which consumers make purchasing decisions based largely on price and accessibility, with varying brands being arguably indistinguishable in their physical composition.³ Further, there is substantial data on gasoline prices, at various levels of aggregation, provided by various data sources, allowing for a more thorough breakdown of demand-side and supply-side effects that are common to, and differ across, locations.

There are four methods of transportation for gasoline products employed across Canada: pipeline, marine tanker, rail and transport truck. Pipelines are generally the safest and most cost-effective means of transporting large volumes. Marine tankers similarly offer large capacity, but are limited to locations with access to seaports, while rail and truck offer access to a greater number of locations, but at much smaller scales.

While most studies use geographic distance as a proxy for transportation costs, few studies have quantified the impact of the structure of the supply network on price dispersion. Locations that are linked by fast, low-cost and large-scale methods may exhibit lower price dispersion than locations that are linked by more costly or smaller-scale technologies that can sustain larger price gaps. As such, the existence of pipelines or seaports between locations may be

²Statistics Canada, Survey of Household Spending, 2015.

³While there is an extensive literature examining the role of distance, location and competition on price dispersion *within* cities, this paper focuses on differences in city-wide average prices across cities. See Marvel (1976), Chandra, Tappata (2011), Lewis (2011), among others.

expected to decrease arbitrage opportunities for suppliers and result in smaller relative price dispersion than locations with only road or rail connections.

To quantify the impact of the supply network on price dispersion in the Canadian gasoline market, I use a unique data set, compiled from Kent Marketing Services, on weekly average gasoline prices from 44 Canadian cities, between 2001 and 2017. One source of price dispersion arises from differences in provincial and municipal gasoline taxes that may distort retail prices faced by consumers across Canada, even after controlling for transportation costs. I therefore use pre-tax prices for gasoline in each location, to control for variation in taxes and isolate the true prices received by suppliers that discipline the arbitrage conditions when price gaps arise between locations.

Examining the data provides several key facts about gasoline prices across locations over time. Price levels are highly correlated across locations, while changes in price are less highly correlated. Intra-regional prices are more highly correlated and exhibit smaller mean differences than inter-regional prices.⁴ The coefficient of variation for weekly prices is small and stable, both nationally and at the regional level. Finally, the location of the minimum and maximum weekly prices rarely changes within Canada.

Regressing measures of price dispersion on distance, region, market size and supply network variables, I find that that the supply network is significant in explaining observed price dispersion across Canadian cities. I find that cities connected by pipeline exhibit 3.5% less mean-price dispersion than cities only connected by road or train. Using weekly relative prices, I find that the existence of a pipeline connection between cities reduces weekly-price dispersion by 2.2%, while a maritime connection reduces weekly-price dispersion by 1.6%.

⁴Regions are broken down according to the natural supply orbits of the Canadian gasoline market: West, Ontario, Quebec and Atlantic.

A pipeline connection has the equivalent effect on weekly price dispersion as a 53% reduction in the geographical distance separating the two locations. Similarly, a seaport connection has the equivalent effect as a 38% reduction in geographic distance.

To put these findings into context, consider two cities such as London, ON and Halifax, NS, of roughly similar total population and density. These two cities are roughly 1850km apart, and at present share no pipeline or maritime connection, with a mean-price difference of approximately 1 ¢/L, with standard deviation of 3.5 ¢/L and a range of weekly-price differences between 0 and 14 ¢/L. The regression analysis suggests that a pipeline built between these two cities would effectively “move” Halifax to Quebec City — that is, the effect on weekly price differences between these two cities would be the equivalent of having the city of Halifax moved 50% closer in geographic distance to London. Ignoring this impact of the supply network would bias the evaluations of suppliers or policymakers considering potential infrastructure projects or policy assessments, if they considered only geographic distance as a proxy for transportation costs.

To check the robustness of these results, I consider two alternative cases — omitting city-pairs subject to price regulation, and omitting geographically remote cities. Some locations, like Quebec and the Atlantic provinces, impose regulations on weekly gasoline prices. Omitting all city-pairs containing a Quebec or Atlantic city, I find the supply network variables remain significant and become slightly larger in magnitude. Similarly, omitting Whitehorse and Yellowknife, which may be outliers due to their extreme geographic remoteness,

I find little change in the regression coefficients, with the supply network remaining significant in accounting for mean and weekly relative price dispersion across locations.

To further isolate the role of the supply network in accounting for observed price dispersion, I consider “supply shocks”, in the form of disruptions to the gasoline production and distribution process. I analyze several instances of refinery shut-downs to examine the effects on production volumes and retail prices across locations. I find that retail prices increase relatively more in regions closest to refinery shut-downs than in those further away, indicating that prices are most significantly impacted within the supply orbit of the affected refinery. Further, I find that price dispersion is lower across locations that share pipeline connections than those that do not, suggesting that arbitrage opportunities arising from supply disruptions are more constrained when locations are connected by faster, cheaper methods of transportation. Both of these findings reinforce the result that the structure of the supply network, not solely geographical distance between locations, is significant in determining the level of price dispersion that can be supported between locations.

4.2 Related Literature

A large and diverse literature examines violations of the Law-of-One-Price (LOP) to determine the causes of price dispersion across locations over time. Papers like Stahl (1982), Crucini, Telmer and Zachariadas (2003), and Crucini, Shintani and Tsurugu (2012) all examine various sources of price dispersion,

such as distance between locations (typically serving as a proxy for transportation costs), market effects and border effects.⁵ Recent work such as Crucini and Yilmazkuday (2014) integrates the role of productivity and wage differences, along with distance and border effects, in accounting for LOP violations. Kano, Kano and Takechi (2013) estimate a model of iceberg-type transportation costs to determine that geographic barriers are a significant contributor to failures of the LOP in Japanese wholesale agricultural markets. This chapter extends this literature by quantifying these types of effects on a particular market, retail gasoline, across Canadian cities. Further, this chapter considers not just geographic distance as a proxy for transportation costs, but also examines variation in transportation methods to account for observed price dispersion across locations in the Canadian gasoline market.

A separate literature examines supply networks formation and equilibrium structures.⁶ Shen (2006) investigates how firms strategically construct their supply networks, by choosing which groups of customers to serve in a profit-maximizing competitive environment. Nagurney, Dong and Zhang (2004) introduce a more general model of supply network equilibrium which is adaptable to different implementations of decision-makers and their independent behaviours. Although this chapter abstracts from strategic decisions regarding the formation of the gasoline supply network, it contributes to this literature by examining the effects of the existing supply network on retail price behaviour. This chapter is unique in its quantitative measure of the impact of the existing

⁵This literature ultimately branches back to the seminal work of Stigler (1961), which identifies the role of incomplete information and search in accounting for price dispersion across retailers, followed by works like Burdett and Judd's (1983) work on multiple equilibria in models with imperfect information, many of which can support long-run price dispersion, and provide a basis for the role of distance, as it relates to acquiring costly information, in accounting for relative price differences across locations.

⁶For a general overview of the strategic supply network formation literature, see Mills, Schmitz and Frizelle (2004).

supply network on retail price dispersion across locations.

A large literature has examined the oil and gasoline industry, at various levels of disaggregation, across a large number of locations, and across various shipping and retailing methods. Marvel's (1976) empirical analysis of the gasoline market examines the role of consumer responses to costly, imperfect information in explaining price dispersion across locations and price variability over time. Adams (1997) argues that, compared to other less homogeneous goods purchased in convenience stores, the relatively low search and information costs of gasoline explain a large degree of price dispersion in highly localized markets. Similarly, Pennerstorger et al (2015) use gas station-level data to test a model of costly information acquisition in the localized retail gasoline market and find that allowing for spatial variation in the share of informed consumers sampling gasoline prices along their commuting routes helps account for observed price dispersion across locations.

This chapter expands on the gasoline literature in multiple ways. First, these papers largely focus on demand-side effects, like consumer search costs, in accounting for price dispersion. While this chapter also considers demand-side market and regional effects, it extends the existing literature by adding a quantitative analysis of the impacts of supply-side effects in explaining price dispersion across locations. Second, this literature predominantly examines price dispersion within local markets, on a station-to-station basis. This chapter quantifies the effects of transportation costs and supply network variables on price dispersion on a larger geographic scale, accounting for city-level mean prices across more geographically diverse locations.

Eckert (2011) surveys the literature on gasoline retailing, including supply side effects related to gasoline pricing. Among these are numerous studies on

the relationship of crude oil prices to retail prices, and the asymmetries of price movements to increases and decreases in crude oil prices. Feyrer, Mansur and Sacerdote (2016) investigate the transmission of income shocks generated by the fracking revolution in the crude gasoline industry and find that wage and income shocks are most strongly transmitted to areas that are most closely connected, both geographically and along the supply network, to the fracking sites. One of the most closely related studies in the recent literature comes from Yilmazkuday and Yilmazkuday (2012), who attempt to attribute relative price differences between locations to difference stages of the production process. They determine that price dispersion across locations within a city can be decomposed as attributable to 50% from crude oil prices, 33% from refinery costs, 12% from taxes, 10% from mark-ups, and only 4% from spatial factors. While this chapter does not decompose the contributing factors of price dispersion in a manner similar to these papers, it does extend the price dispersion literature by quantifying the relationship of the supply network structure to observed price dispersion across locations.

4.3 Supply Network

A common approach in price dispersion literature is to use a measure of geographic distance as a proxy for the associated transportation costs that govern the arbitrage conditions sustaining price gaps between locations. However, if transportation costs differ across locations, or vary depending on factors like the method of transportation employed, geographic distance alone may be a biased measure. The potential profits for suppliers seeking to buy product at lower price locations and transport to higher price locations for resale depend

on the per-unit costs associated with the transportation methods available between those locations. Observed price dispersion will therefore be a function not only of geographic distance, but also of the characteristics of the transportation options that suppliers can employ between retail locations.

In this chapter the term “supply network” refers to the various methods available for transporting oil products, such as refined gasoline, between locations across Canada. There are four primary modes of transportation: pipeline, marine tanker, rail, and truck. Pipelines are the most prominently used method of transporting bulk quantities, with approximately 750 million barrels of refined oil products being shipped annually via pipeline. Alternatively, approximately 95 million barrels of refined products are transported by marine tanker, both domestically and via import, into Canadian ports annually. Put into perspective, the amount of oil products transported daily across Canada via pipeline would necessitate the equivalent of 4200 rail cars, or 15,000 trucks.

Pipelines are the most cost-effective means of transporting large quantities of gasoline products, followed in order by tanker, rail and truck.⁷ However, pipeline and sea transport face obvious limitations of geography and pre-existing infrastructure, as pipelines and sea routes are less prevalent than the extensive rail and highway networks throughout Canada. Therefore, locations that share existing pipeline or sea route connections may reflect different relative pricing patterns than locations that do not have access to these shipping options.

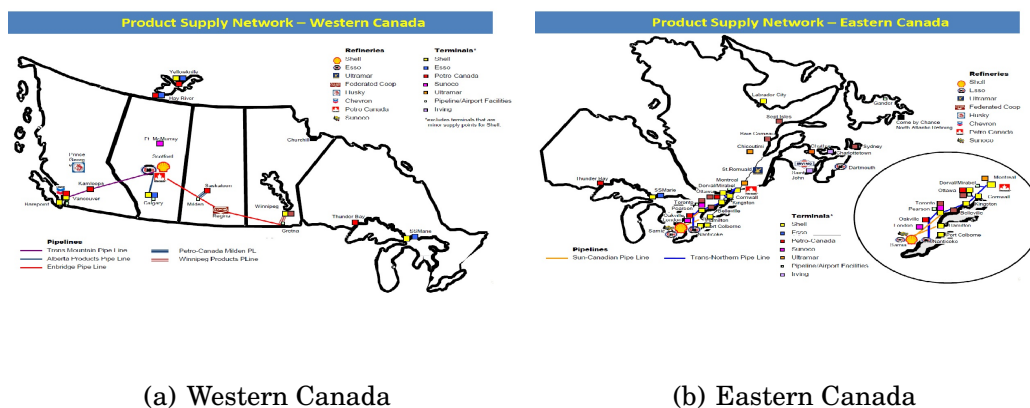
The Canadian gasoline supply network begins with crude oil reserves and imports. The vast majority of domestic crude production occurs in Western and

⁷For a detailed breakdown of the crude oil and petroleum products infrastructure in Canada, refer to Natural Resources Canada report, available at: <http://www.nrcan.gc.ca/energy/sources/infrastructure/1490>.

Northern Canada, with the remainder mainly occurring in offshore reserves in the North Atlantic Ocean.⁸ Crude oil is transported to refineries to be processed into finished petroleum products via two primary methods: in Western and Central Canada, where crude comes mainly from domestic production, this occurs via pipeline; in the East and Quebec, where most crude comes from offshore reserves or imports, this occurs primarily via marine tanker.⁹ This leads to a natural division of Canadian cities into distinct supply regions: Western Canada, Ontario, Quebec, and Atlantic Canada — in this chapter, I use the term “regions” to refer to these 4 distinct supply regions of the Canadian gasoline market.

Once the crude oil is transported from the source, it is converted into finished oil products at one of 19 Canadian refineries, 16 of which produce finished gasoline products, that are located in all Canadian provinces with the exceptions of PEI, Manitoba and the northern territories. The locations of these refineries are shown in Figure 4.1.

Figure 4.1: North American Oil Pipeline Infrastructure



⁸Energy Information Administration, United States Department of Energy, report in Oil and Gas Journal, December 2011, available at: www.eia.doe.gov/oiaf/ieo/index.html.

⁹There are an estimated 250,000km of liquids pipeline infrastructure across Canada. Source: Canadian Energy Pipeline Association.

Once refined, gasoline is shipped to regional terminals for distribution to retail outlets. Shipping from refineries to terminals can occur by pipeline, marine tanker, rail or truck. Once delivered to terminals, the refined gasoline may first receive additives to create unique blends that are specific to retail brands, or be shipped directly to retail outlets, which is always done by truck. To take advantage of economies of scale, retail brands often use the same terminals for distribution to retail outlets, through reciprocal purchase agreements.¹⁰ This minimizes total transportation costs for refined products, resulting in a relatively small number of terminals supplying large geographic regions for any number of distinct retail brands.

The supply network figures most prominently between the refinery and terminal stages, where methods of transportation are most varied. Pipelines, followed by ports, are more cost-effective than rail and truck, but are more limited by geography and terrain. Although pipelines have the lowest per-unit shipping cost, their construction is also the most expensive. Industry rule-of-thumb suggests that it requires approximately 15–20 years of pipeline operation to recoup the fixed costs of building a pipeline.¹¹

Since the methods of transportation and their associated costs are common to all locations for most links in the supply network (crude to refinery, terminal to final outlet), any variation in transportation costs between locations

¹⁰Source: Natural Resource Canada, *Ibid.*

¹¹Source: Canadian Energy Pipeline Association.

can be assumed to be a function of the methods employed between refinery-to-terminal across locations.¹² Therefore in quantifying the impact of the supply network on price dispersion across locations, I use variation in the supply network at the terminal-to-refinery stage, in conjunction with geographic distance, as a proxy for the transportation costs associated with supplying various locations.

The Impact of the Supply Network on Price Dispersion

For some intuition as to why the supply network may impact price dispersion in the gasoline market, consider two cities located X km apart that exhibit different pre-tax gasoline prices. Any observed price difference theoretically reflects the cost of arbitraging this price gap for potential suppliers. Suppliers would need to find it profitable to purchase gasoline at the lower price location and transport it to the higher price location for resale. However, the arbitrage opportunity depends on the existing supply network — the costs of transporting bulk quantities will depend not just on the distance between the two locations, but also on the available methods of transportation and their associated costs. In my quantitative analysis, I supplement data on traditional measures of geographic distance between locations with data on supply network variables to get an unbiased estimate of the quantitative impact of both distance and supply network variation on price dispersion across locations.

Pipelines are widely agreed to be the safest, quickest, and cheapest means of transporting large quantities of gasoline products between locations. Although

¹²Due to the mostly partitioned aspect of the crude-to-refinery stage of the production process, any difference between pipeline and seaport supply methods could be assumed to be captured by regional effects between the West, Ontario and Quebec/Atlantic regions — these divisions concur with the assertions made in the Natural Resources Canada report on gasoline supply infrastructure that define these regions as natural “supply orbits”.

cost may vary depending on location, transporting oil products by pipeline generally costs between \$3–6 per barrel, and move at speeds between 5 and 20km/hr. Since pipelines are generally laid underground, they tend to follow direct routes between locations and are largely indifferent to terrain. Once pipelines are constructed connecting locations, they can be used with relatively little cost of coordination, for large volumes of product — analogous to “flipping a switch” and sending the products to their desired location. However, one drawback of pipeline use is that the large capital and fixed costs associated with their construction mean the limited number of existing pipelines are typically run at or near capacity and accommodating large increases in demand is often difficult.

Conversely, consider the alternative of transportation by land, typically done by rail for large volumes of product. Although freight train shipping is typically faster once set in motion (roughly 30–35km/hr), the per-unit costs are higher, typically between \$10–15/barrel, and shipping routes may not be as direct, historically following roadways, and are often impeded by the surrounding terrain. Additionally, it is estimated that the amount of oil product shipped through pipelines on a daily basis in Canadian pipeline infrastructure would require approximately 4200 rail cars to transport. Thus, setting up large shipments by rail would be much more costly to arrange and would require a longer time frame to organize than with pipelines, potentially missing opportunities for arbitraging price differences. Similarly, shipping by marine tanker imposes geographical limitations, requiring ports for both locations, and faces higher variable costs and coordination time than shipping by pipeline.

Observed price gaps between two locations may therefore reflect differences in the supply network: locations connected via pipeline may exhibit smaller

ceteris paribus price differences than those that are not, with similar effects expected for locations connected by seaport, as opposed to land. For these reasons, the supply network may be significant in determining what price gaps can be sustained over short periods of time, reflected by weekly relative price differences. However, over time, using mean price differences, the supply network may not play as large a role as arbitrage opportunities may be reduced or eliminated over time.

4.4 Data

Gasoline Price Data This paper uses a unique data set compiled from gasoline price data collected by the Kent Group, a downstream data collection and marketing services firm.¹³ The Kent group performs weekly random surveys of gas stations to produce a snapshot of city-level average retail prices across Canada. The relevant data is compiled from weekly average price data for regular unleaded gasoline across a city-wide sampling of independent gas stations located in 44 Canadian cities between January 2001 and May 2017. Together they comprise 905 weeks of observations for a total of 39,820 independent city-average retail price observations.

Since this paper investigates the role of the supply network on price dispersion across locations, I examine pre-tax price data, controlling for variation in provincial and municipal gasoline taxes, to identify the price that would actually be received by suppliers, which governs the arbitrage condition. Summary statistics for the pre-tax price data at the national, as well as regional levels can be found in Table 4.1. The four regions correspond to the natural divisions suggested by the structure of the supply network. There is large variation at

¹³Available at <http://www.kentmarketingservices.com/dnn/PetroleumPriceData.aspx>.

the national level, with weekly pre-tax prices ranging from 14 ¢/L up to 128.5 ¢/L between 2001 and 2015, with a mean price of 67 ¢/L. The mean and standard deviation are consistent across the various regions, with the West typically exhibiting the highest mean prices and largest min-max spread, while the Atlantic region exhibits the smallest min-max spread.

Table 4.1: Pre-tax Weekly Unleaded Gasoline prices

Summary Statistics (¢/L)					
Region	Mean	Std. Dev.	Min	Max	Median
Canada	67.4	19.2	14.1	128.5	67.4
West	70.0	19.7	14.1	128.5	70.2
Ontario	65.8	19.2	20.9	113.5	66.4
Quebec	65.6	18.8	23.5	112.3	65.8
Atlantic	66.2	18.3	27.2	108.7	66.2

This retail price data is consistent with data collected by Statistics Canada, with the added benefit of more frequent observations (weekly rather than monthly) which allows for a more accurate accounting of responses to price gaps that arise (and may disappear) between locations across shorter periods of time. Further the 44 cities in this data set (as opposed to the 18 in the StatsCan data) cover a more geographically diverse set of locations across Canada, and include more variation in the available supply network between locations, allowing for a stronger analysis of the impact of the supply network on price dispersion in the Canadian gasoline market.

4.4.1 Descriptive Statistics from Pre-Tax Data


The data provides several key insights into pricing behaviour. Pre-tax prices are highly correlated across cities, both in price levels and, to a lesser extent,

changes in price. These correlations are strongest between cities within supply regions, which also exhibit smaller mean-price differences than inter-regional city-pairs. Relative prices change little across cities and variation across locations is small and stable.

1a. Pre-tax prices are highly correlated across locations

Over the 17 year period spanned by the data, pre-tax gasoline prices are highly correlated for all city-pairs. Figure 4.2 shows that most correlation coefficients for city-pairs are above 0.95, and all are above 0.90, indicating a high level of correlation in weekly prices over time. These prices are even more highly correlated across city-pairs than they are with crude oil prices, both with Brent crude, which is typically used as a price gauge in Eastern and Central Canada, as well as with West Texas Intermediate (WTI), which is used more in the United States and Western Canada.


Figure 4.2: Weekly Price Correlations: Levels



1b. Pre-tax price changes are less highly correlated across locations

Though the pre-tax prices exhibit high correlations across locations, the *changes* in weekly pre-tax prices are less highly correlated. Figure 4.3 plots the correlation coefficients for all city-pairs for the change in the pre-tax price from the previous week's value. These correlations, most of which are in the range of 0.15–0.60, are smaller than those of the price levels (0.90–0.99). This discrepancy between the high level of correlations between price levels and the lower level of correlation between price changes raises questions about the responses of retail prices to shocks across locations. Differences in market size, location-specific demand shocks, or variation in the available supply network may be significant in accounting for these retail price movements.

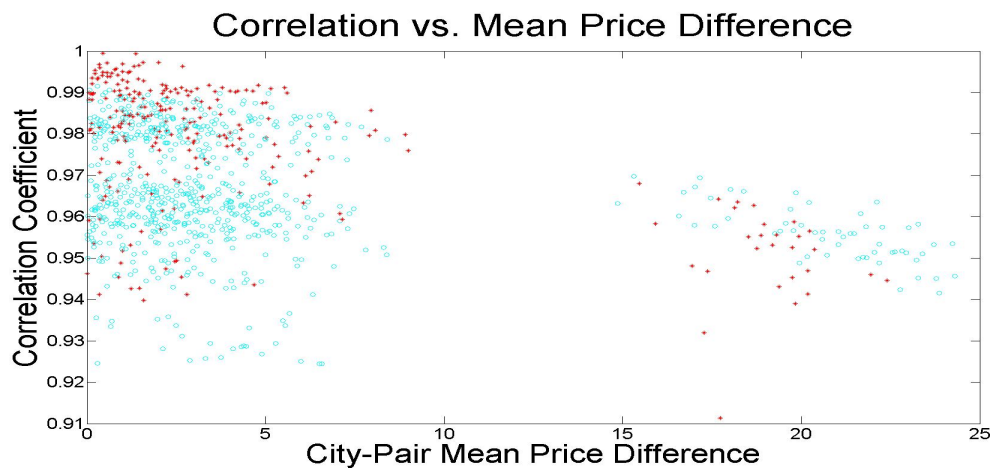
Figure 4.3: Weekly Price Correlations: Weekly Changes



2a. Intraregional pre-tax prices are more highly correlated

Pre-tax prices are more highly correlated for city-pairs within the same region than for interregional pairs. In Figure 4.4, all city-pairs display high correlation coefficients, above 0.90; however, while most intraregional city-pairs, seen in red, have coefficients between 0.96 and 0.99, the majority of interregional correlation coefficients, seen in blue, are generally lower, falling between 0.94 and 0.98. This property holds within each of the four regions across Canada, with intraregional city-pairs displaying higher correlations than interregional pairs.

Figure 4.4: Retail Price Correlations vs. Mean Price Differences: National



2b. Intraregional pre-tax prices exhibit smaller mean-price differences

Intraregional city-pairs also exhibit a smaller mean price difference than interregional pairs, regardless of the Canadian region in which they occur. With the notable exception of a grouping of points from city-pairs involving Whitehorse or Yellowknife (which may potentially be outliers due to their northern

isolation), most intraregional city-pairs exhibit lower mean-price differences, in the range of 0–5 ¢/L differences. Although many interregional city-pairs also exhibit low mean-price differences in this range, a larger proportion of interregional pairings display differences in the range of 5–10 ¢/L. This property is true within each region as well. These findings reinforce the relevance of regional effects in accounting for retail price dispersion.

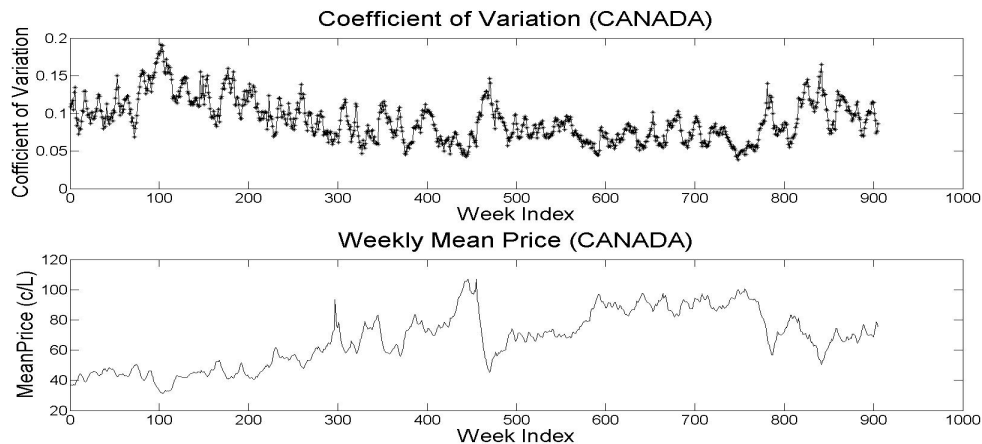
3. Coefficient of variation for pre-tax prices is small and stable

Pre-tax prices display a level of variation across Canada that remains small and stable over time. The coefficient of variation for weekly pre-tax prices lies in the range from 0.04 to 0.20 over the span of January 2001 to May 2017, the majority of which occur below 0.10, and exhibits few large changes. Figure 4.5 shows that as the mean national price changes and trends higher over time, the coefficient of variation remains small, reflecting a fairly stable relationship of prices across the country over time. Notable is the increase in mean prices during the crisis of 2008 (around week 400), with correspondingly low variation, and the ensuing spike in the coefficient of variation as the mean prices finally decreased in late 2008, indicative of national prices rising symmetrically, but declining at different rates in the ensuing periods. Across regions, the coefficient of variation is consistently low and stable.

4. Minimum and maximum prices' location rarely changes

Although weekly gasoline prices across Canada change frequently for each city, the location of the minimum and maximum weekly prices across Canada rarely changes. Over the 905 week period spanned by the data, only 3 of the 44 different cities assume the place of maximum pre-tax price, with the vast majority

Figure 4.5: Coefficient of Variation: National



of these occurrences being split between the two most geographically isolated cities, Yellowknife, NWT., and Whitehorse, YK.¹⁴ The location of the minimum pre-tax price changes more frequently than the maximum, but a small number of cities, such as Windsor, Quebec City, Montreal, Ottawa, Edmonton and Vancouver, account for the majority of minimum weekly prices. The histogram in Figure 4.6 plots the frequency of these occurrences.

The minimum and maximum weekly pre-tax prices also follow a stable relationship with the mean weekly price. The maximum price typically falls in a range of 120–150% of that of the mean national price, while the minimum price falls in the range of 75–90% of the mean national price, as seen in the top panel of Figure 4.7. This leads to a steady min-max spread that resides mainly between 30–70% of the mean price, as shown in the bottom panel of Figure 4.7.

¹⁴Corner Brook, NL accounts for the maximum price in only a few weeks.

Figure 4.6: Max/Min histogram

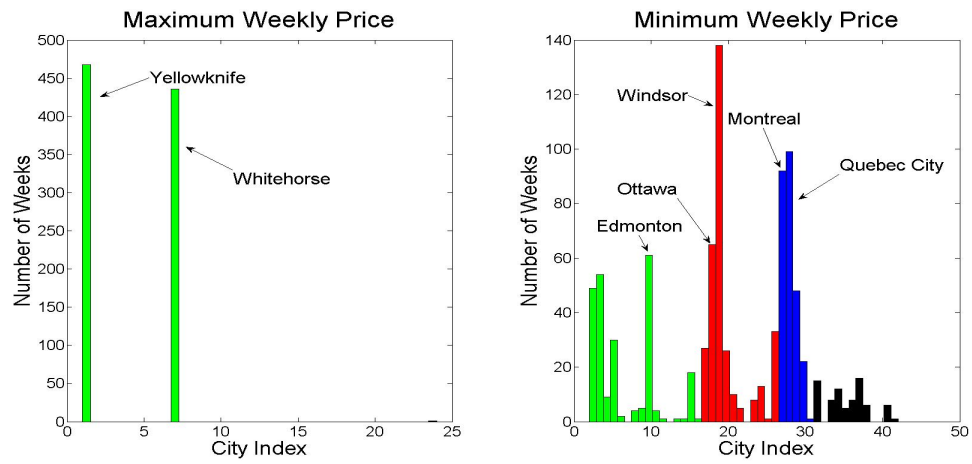
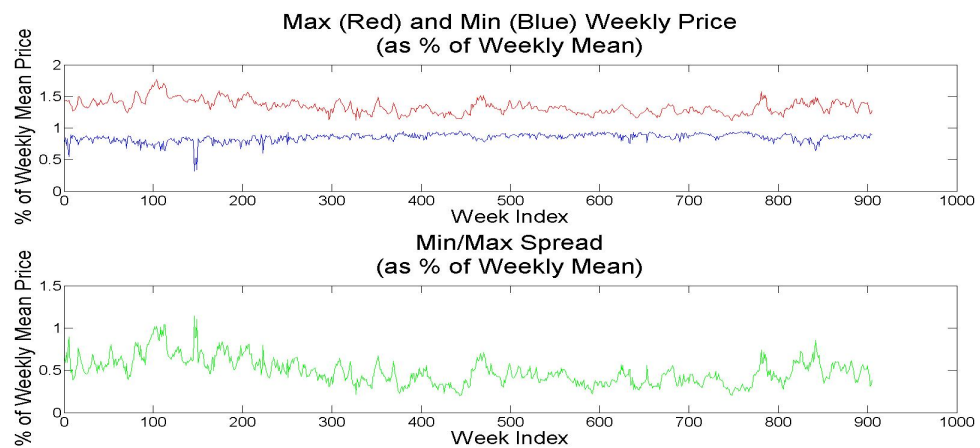


Figure 4.7: Min/Max spread and Mean Price



4.5 Empirical Analysis

I use regression analysis to quantify the impact of the structure of the supply network, as well as other explanatory variables, on price dispersion across Canadian cities. I regress measures of mean and weekly relative prices across Canadian cities on explanatory variables such as distance, region and market size, and also include supply network variables — specifically, I use dummy variables for the existence of a pipeline or seaport connection linking locations, as well as distances from each city to the closest supply terminal.¹⁵

Testing for Stationarity

To determine how the supply network and other explanatory variables impact price dispersion across locations over time, I follow the literature on Law-of-One-Price (LOP) deviations and test for stationarity in the weekly price data for each city, as well as relative prices between city-pairs, to rule out long-run convergence to parity of prices across locations.¹⁶ I use a standard Dickey-Fuller test of the AR(1) process with constant and time trend, of the form

$$\Delta P_{jt} = \alpha_0 + \alpha_1 t + \delta P_{jt-1} + u_t \quad (4.1)$$

Table 4.2 shows that the null hypothesis of a unit root ($\delta = 0$) can be rejected at the 95% confidence level for 19 of the 44 Canadian cities. I find that the more geographically isolated a city, the more likely is its weekly price data to be non-stationary — for example, northern cities such as Yellowknife, Whitehorse, many prairie cities and all of the Atlantic cities appear to be non-stationary,

¹⁵Source for terminal distance data: <http://www.essomaps.ca/terminals>.

¹⁶For example, see Ceglowski (2003) or Parsley and Wei(1996).

while most of Ontario and Quebec, as well as large Western cities like Vancouver and Victoria, appear to be stationary.

However, since the arbitrage opportunity is a function of the relative price between cities, I also test the time series of weekly relative prices across all city-pairs for stationarity. For each city-pair, ij , I test the AR(1) process of the form

$$\Delta|\log(P_{it}) - \log(P_{jt})| = \alpha_0 + \alpha_1 t + \delta|\log(P_{it-1}) - \log(P_{jt-1})| + u_t \quad (4.2)$$

With 44 Canadian city-pairs, I find that all 946 possible relative price series can reject the null hypothesis of a unit root at the 99% confidence level. This indicates that while each individual city's price series may or may not be stationary over time, the high degree of correlation between city-pair prices generates a stationary time series for each potential weekly relative price series. This result suggests that OLS regression, using weekly relative prices as the dependent variable is suitable for quantifying the impact of the supply network on price dispersion across locations.

Table 4.2: Testing for Stationarity of Time-Series Data

Stationary at the $X\%$ confidence level?					
	99%	95%		99%	95%
WEST Region					
Whitehorse	No	No	Red Deer	No	No
Vancouver	Yes	Yes	Edmonton	No	Yes
Victoria	Yes	Yes	Lethbridge	No	Yes
Prince George	No	No	Regina	No	No
Kamloops	No	Yes	Saskatoon	No	No
Kelowna	No	No	Prince George	No	No
Yellowknife	No	No	Winnipeg	No	No
Calgary	No	Yes	Brandon	No	No
ONTARIO Region					
Toronto	Yes	Yes	Thunder Bay	No	Yes
Ottawa	No	Yes	North Bay	Yes	Yes
Windsor	Yes	Yes	Timmons	No	No
London	No	Yes	Hamilton	No	Yes
Sudbury	Yes	Yes	St. Catherines	No	Yes
Sault Ste. Marie	No	Yes			
QUEBEC Region					
Montreal	Yes	Yes	Gaspe	No	No
Quebec City	No	Yes	Chicoutimi	No	Yes
Sherbrooke	No	No			
ATLANTIC Region					
Saint John	No	No	Yarmouth	No	No
Fredericton	No	No	Truro	No	No
Moncton	No	No	Charlottetown	No	No
Bathurst	No	No	St. Johns	No	No
Halifax	No	No	Gander	No	No
Sydney	No	No	Cornerbrook	No	No

Empirical Approach

To quantify the impact of the supply network on pre-tax gasoline price dispersion across Canadian cities, I use OLS regressions with two different measures of price dispersion between city i and city j as the dependent variable:

1. The absolute mean-price difference between city-pairs, $|\log(\bar{P}_i) - \log(\bar{P}_j)|$.
2. The weekly relative price between cities, $|\log(P_{it}) - \log(P_{jt})|$.

These measures capture different aspects of price dispersion across locations: mean city-pair differences examine systemic price variation across cities, while weekly relative prices examine how prices move and react to shocks differently across locations.

The explanatory variables include measures of distance, market size, and supply network variables. For each city-pair, distance is measured as the shortest driving route (in thousands of kilometres) between the two cities, $dist_{ij}$.¹⁷ This measure highlights the distinction between transporting products by pipeline and sea versus the default alternative of land routes used by train and truck, which may impose geographical limitations. Market size variables include a regional dummy, Reg_{ij} , which takes the value of 1 if both cities are in the same region (as defined by the supply network regions) and zero otherwise, and two population measures, pop_{ij} . The first measure is the difference in population between city-pairs, $|totpop_i - totpop_j|$.¹⁸ The second measure is the difference in population density between city-pairs, $|dens_i - dens_j|$, measured as population per square kilometer, in order to control for rural and urban

¹⁷See Appendix C.1 for a detailed discussion of data measures.

¹⁸The population measured used here is the population within city limits, as measured by Statistics Canada.

differences such as land prices that may be passed-through to pre-tax gasoline prices. The supply network variables include two dummy variables for pipeline connection, $Pipe_{ij}$, and seaport connection, Sea_{ij} , which are set to 1 if the two cities share a pipeline connection or traversable seaports, respectively, and zero otherwise.¹⁹ Finally, the variable $term_{ij}$ represents the difference between the distance to the closest supply terminal for the two cities.²⁰

With y representing one of the two dependent variables, the regression specification is then:

$$y = \alpha + \beta \log(dist_{ij}) + \gamma Reg_{ij} + \delta Pipe_{ij} + \lambda Sea_{ij} + \theta \log(pop_{ij}) + \psi \log(term_{ij}) + \epsilon_{ij} \quad (4.3)$$

All non-dummy variables are log-transformed to compensate for potential heteroskedasticity issues with prices, distance and population differences between pair-wise locations. For both mean- and weekly-prices, clustered standard errors are used to control for correlation between observations of relative prices that share at least one city in common.²¹

4.5.1 Mean Price differences

The coefficient estimates with city-pair mean price differences as the dependent variable can be found in Table C.1. The supply network has a significant impact on pre-tax mean-price differences across locations. The coefficients for

¹⁹Potential problems with endogeneity may arise with the pipeline variable, as it may be argued that the choice of locations for pipeline construction may be influenced by relative prices across those locations. However, since the pipelines employed in this study were constructed well before the time span of the retail price data, it is assumed in this chapter that these pipeline connections are predetermined as explanatory variables for pre-tax price dispersion.

²⁰For example, if city i is 100km from its nearest terminal and city j is 75km from its nearest terminal, $term_{ij}$ would be 25km.

²¹Specifically, errors are clustered in groups $g = 1, \dots, G$, where all elements of group g contain a relative price that includes city i in $|\log(\bar{P}_i) - \log(\bar{P}_j)|$ for mean-price differences, or city i in period t in $|\log(P_{it}) - \log(P_{jt})|$ for weekly-price differences.

the dummy variables on pipeline and seaport connections both have the expected negative sign, implying that retail prices are less dispersed across locations sharing these connections. However, the estimates for seaport are not significantly different from zero at a 90% confidence interval, as opposed to a 99% confidence level for the pipeline dummy. The coefficient on pipeline ($\delta = -0.0353$) in the preferred specification in regression 1, indicates that, *ceteris paribus*, cities connected by a pipeline exhibit 3.5% less dispersion in their mean prices. I also find that while statistically significant, doubling the difference in distances from the nearest terminal would increase mean-price dispersion by only 0.3%.

Distance and region have statistically significant coefficients. Interpreting these coefficients, with ($\beta = 0.0374$), two cities that are 100% further apart (in km) exhibit 3.7% more price dispersion in their mean prices. Intuitively, this means that if City X and City Y are identical in every way, but City X happens to be twice as far away from City Z as is City Y, then we would expect the difference in mean prices between City X and City Z to be 3.7% larger than that between City Y and City Z, *ceteris paribus*.²² The region coefficient is positive, predicting that prices are 6.3% more dispersed for city-pairs that are *within* the same region than for those that are not. This may potentially be due in part to demand side effects, indicating that demand shocks are highly localized over time, and not necessarily dispersed across entire regions.

Finally, the population coefficients are both found to be significant, but small. City-pairs with 100% larger population differences exhibit 0.5% more

²²While this result may seem small in magnitude relative to the percentage increase in distance between locations, a country as geographically scattered as Canada suggests such large variations in distance, with a coefficient of variation for the distance variable of 0.6817, compared to coefficient of variations for pre-tax prices Canada-wide with a mean value of 0.0943.

mean-price dispersion, while cities with 100% larger differences in population density also exhibit 0.5% more mean-price dispersion. While distance, region and pipeline are significant factors in predicting mean price differences, the value of the estimated coefficients changes very little across regressions as other explanatory variables such as seaports and population measures are added.

4.5.2 Weekly relative prices

The supply network also impacts relative weekly-price differences. Table C.2 presents the results of the regression analysis with weekly relative prices as the dependent variable. Controlling for distance and market effects, the supply network variables for pipeline and seaport connects are statistically significant and negative. In the baseline specification in Reg. 1, the existence of a pipeline connecting two cities ($\delta = -0.0221$) amounts to a 2.2% lower weekly relative price difference, while locations that share a seaport connection ($\lambda = -0.0160$) display to a 1.6% reduction in relative weekly price dispersion. Distance, region and market size variables are also all significant explanatory variables in this specification. A 100% increase in distance between two locations increases the weekly relative price gap by 4.2%, while a 100% increase in either total population differences or population density differences amounts to a 0.5% greater spread in relative prices. Again, intraregional relative prices are more dispersed than interregional prices, by approximately 5.6%. Using these parameter estimates, the existence of a pipeline between two locations produces an analogous effect on relative prices as would a 53% reduction in distance between them, while the existence of a maritime shipping link produces the effect of a 38% reduction in land route distance.

The distance coefficient estimates are stable across most regression specifications, although the coefficients suggest that omitting supply network characteristics from any regression analysis of these weekly relative prices may bias the predicted results. The omitted supply network variables may bias the distance coefficients upwards by roughly 5% across specifications, and ignoring regional effects (where region boundaries are defined with supply network arcs) can bias distance coefficients by upwards of 35%. Taken in conjunction with the results obtained from the mean-price difference specifications in Section 4.5.1, this suggests that arbitrage opportunities that may arise in any given period by the presence of a pre-tax price gap across locations are not merely limited by the geographical distance between locations. Rather, they are a function of the effective distance between locations, which takes into account both the geographical distance and the available supply network linking the two. The regression coefficients support the intuition that pipelines are potentially able to curtail arbitrage opportunities to a greater extent due to the speed and ease of coordination of shipping via pipeline, while larger price variations may be sustainable across locations connected by truck or rail, due to the larger time and monetary costs of coordinating the resources necessary to ship large quantities via these modes of transportation.

4.5.3 Robustness

To check the robustness of these results, I consider two alternative cases of the regression analysis: (1) omitting Quebec and Atlantic Canada cities from the data, and (2) omitting Whitehorse and Yellowknife from the data.

Case (1): While the Federal government does not regulate gasoline prices in Canada, several provinces do enforce some form of price regulation. Quebec

sets a *minimum* weekly price, dependent on estimated acquisition and transportation costs, while New Brunswick sets a *maximum* weekly price, indexed to crude prices and retail margins. Nova Scotia, Prince Edward Island and Newfoundland and Labrador all set the weekly price of gasoline in their province following similar standards based on spot crude prices and relative to estimated transportation costs.²³ This raises the possibility that the estimates in Tables C.1–C.2 are biased. I therefore perform the regression analysis while omitting all city-pairs that include at least one of the cities in the Quebec and Atlantic regions, effectively leaving only the Ontario and West regions.

Case (2): Due to their relatively extreme geographic remoteness, Yellowknife, NWT, and Whitehorse, YK may be outliers in the Canadian gasoline market. While theoretically possible, it may not be practical to consider that suppliers explore arbitrage opportunities between Yellowknife and St. John's, NFLD, in the same light as they would explore potential arbitrage opportunities between Edmonton and Calgary, AB, for example. I therefore repeat the regression analysis while omitting all city-pairs that include either Whitehorse or Yellowknife.

Omitting Quebec and Atlantic

In Table C.3, omitting Quebec and the Atlantic provinces produces relatively little change on the impact of the supply network on mean-price differences. The coefficient on pipeline remains significant and negative, while its value increases slightly — the existence of a pipeline connecting cities now results in a 4.1% reduction in price dispersion. The seaport coefficient is still statistically insignificant at the 90% confidence level, and the coefficient on terminal

²³A more detailed description can be found at the Consumer Council of Canada: <http://www.consumerscouncil.com/index.cfm?id=13904>.

distance becomes statistically insignificant as well once these cities have been omitted. The coefficient on distance remains significant, and rises slightly, with a 50% increase in distance between cities resulting in a 2.6% larger difference in mean-prices.

In Table C.4, the impact of the supply network on weekly price dispersion increases when omitting Quebec and Atlantic cities. A pipeline connection reduces weekly price dispersion by 3.0%, and a seaport connection reduces price dispersion by 4.3%. This change may be a function of the fact that for the remaining cities, in the West and Ontario regions, there are a larger variety of employed transportation methods, with pipelines in particular being more prominent. Contrarily, the omitted Quebec and Atlantic cities tend to be mainly connected by seaports and lack much pipeline infrastructure. The coefficient on distance increases, with a 50% increase in distance resulting in 2.7% less weekly price dispersion across locations.

Omitting Whitehorse and Yellowknife

In Table C.5, the supply network remains significant in impacting mean-price dispersion when omitting Whitehorse and Yellowknife. A pipeline connection between locations reduces mean-price dispersion by 1.6%, and a seaport connection between locations reduces mean-price dispersion by 1.0%, and the seaport variable now becomes significant at a 95% confidence level. Perhaps not surprisingly, removing the two cities that are the most extremely geographically located reduces the magnitude of the distance and region coefficients. A 50% increase in distance between cities results in only 0.2% less mean-price dispersion, and locations in the same region exhibit only 0.7% less price dispersion than city-pairs that span multiple regions.

In Table C.6, omitting Yellowknife and Whitehorse decreases the impact of the supply network on weekly price dispersion. The pipeline and seaport coefficients are still statistically significant, but smaller, with a pipeline connection reducing weekly relative price dispersion by only 0.4% while a maritime connection reduces weekly price dispersion by 1.3%. The distance coefficient, while still statistically significant, also decreases, with a 50% increase in distance resulting in a 0.6% increase in price dispersion between cities.

4.5.4 Summary: Impact of Supply Network on Price Dispersion

I find that mean price differences are functions of distance, as a proxy for transportation costs, and market-specific factors that limit arbitrage opportunities and allow price gaps to be sustained in mean prices over time. I also find that the supply network, both in regional supply effects and the existing pipeline infrastructure, significantly impact mean price differences. Weekly price differences are also significantly impacted by the structure of the supply network for both pipeline and marine transportation. Pipeline connections between locations limit weekly-price dispersion by the equivalent of a 53% reduction in distance, while seaport connections are akin to a 38% reduction in distance across locations. This suggests that the costs of arbitraging price gaps are significantly affected by the differences in the supply network's structure, with lower variable-cost methods (like pipelines) reducing the effective distance between locations more than alternative methods with higher per-unit shipping costs.

4.6 Case Studies of Supply Shocks

An alternative approach to identify how the supply network impacts retail prices in the Canadian gasoline market is how “supply shocks” affect retail price dispersion across locations. This paper specifically considers disruptions in production at the refinery level. This aspect of the supply chain is chosen for several reasons:

1. There are a relatively small number (16) of refineries that operate across Canada, and therefore occurrences of refinery shutdowns can be more easily identified via newspaper and petroleum industry reports;
2. Products are moved from refineries to terminals for distribution via all four modes of transportation, and therefore the variation in the transportation availability across geographic locations is larger than for the retail level, where all transportation is done by truck;
3. It allows for consideration of reactions in price to supply disruptions that may be felt by a group of cities within a geographic proximity to a particular refinery, but not necessarily at either the national level, or the extremely localized level.

Intuitively, one might expect that a shortage in supply caused by a temporary shutdown of a refinery in a particular region may impact relative prices for nearby cities, but not necessarily all cities nationwide. Further, one might expect that the ability of suppliers in a certain location to respond to supply shortages to depend on the available modes of transportation in their regions. Therefore, one way to investigate this conjecture is to examine cases of observed refinery shutdowns, to determine responses in production levels, net

imports and retail price changes that occur as a result. However, since refineries are such vital links in the supply chain, there are large incentives to keep them operating and resolve any disruptions as quickly as possible, regardless of repair costs. Thus it is difficult to find a large number of instances of refineries remaining closed for any prolonged period of time, in order to perform a thorough quantitative analysis.

The small number of shutdowns can be classified in one of two ways: planned shutdowns, defined as a scheduled, forecast shutdown of refinery operations, typically for maintenance purposes, performed during seasons when gasoline demand is lowest; and unplanned shutdowns, defined as an unexpected disruption caused by accidents or acts of nature. Examples of planned and unplanned shutdowns are listed in Table 4.3.

Table 4.3: Refinery Shutdowns

Refinery	Date	Reason	Capacity (bpd)	% of Reg. capacity
Planned shutdowns				
Calgary	Jun/Jul 2006	Maintenance	110,000	18.7
Saint John	Nov/Dec 2008	Maintenance	300,000	59.5
Edmonton	Aug 2009	Maintenance	135,000	23.0
Unplanned shutdowns				
Nanticoke (ON)	Feb 2007	Fire	112,000	29.2
Edmonton	Aug 2008	Cat. Conv.	135,000	23.0
Scotsford (AB)	Sept 2009	Unplanned maint.	100,000	17.0
Dartmouth (NS)	Sept 2010	Hurricane	89,000	17.6

In order to examine the effects of these shutdowns on refinery production volumes, data from Statistics Canada is employed for the period January 2001 to May 2017. This data is categorized by Statistics Canada into provincial aggregate levels, as well as regional levels, which can then be matched into the

regions suggested by the supply network infrastructure. Specifically, refinery data is grouped into the Atlantic provinces, Quebec and Ontario — the only notable exception is that the Western provinces are not aggregated together into one region, but rather by individual province. The one shortcoming of this data is that figures are withheld for refinery production in Saskatchewan and British Columbia, in accordance with confidentiality requirements.²⁴ However, as there are no refineries in Manitoba or the territories, and by capacity, the B.C. and Saskatchewan refineries make up relatively small portions of the total capacity of the region (less than 15%), the Alberta refinery data can serve as a reasonable approximation for the Western region. Summary statistics for the refinery data can be found in Table 4.4.

Table 4.4: Refinery Production

Summary Statistics (cubic metres per month)					
Region	Mean	Std. Dev.	Min	Max	Median
Canada	824670	116810	443136	1131445	829380
West	793820	107930	443136	970013	808240
Ontario	894240	128600	528807	1131445	872330
Quebec	796200	84740	513483	976323	807000
Atlantic	814410	112890	443624	995657	838090

4.6.1 Stylized Facts from the data

Production Planned shutdowns have minimal effect on refinery production volumes at the regional level. Planned shutdowns at the Saint John and Edmonton refineries coincided with 2% and 1% decreases in production levels


²⁴Refer to the Canadian Statistics Act, available for viewing at: <http://www.statcan.gc.ca/about-apercu/act-loi-eng.htm>.

from their seasonal (monthly) averages in their respective regions, while the Calgary shutdown coincides with a 4% increase in the regional production levels relative to the season average. The size of these deviations, in terms of standard deviations from the mean, are -0.26, -0.10 and +0.19 for Saint John, Edmonton and Calgary, respectively. Refinery production for the Alberta region can be seen in Figure 4.8 with the period 1 week prior to 6 weeks after the beginning of the shutdown highlighted. At the time of the planned shutdown, production in the Alberta region actually increases, and remains slightly above its seasonal average.



Conversely, unplanned shutdowns appear to have larger negative effects on regional production volumes. Unplanned shutdowns at the Nanticoke, Edmonton and Scotsford refineries coincide with 23.9%, 27.7% and 21.7% decreases in regional production relative to seasonal averages, which are analogous to -1.64, -2.80 and -1.85 standard deviations from their respective means. Refinery production for the Ontario region can be seen in Figure 4.9 with the period 1 week prior to 6 weeks after the beginning of the shutdown due to the Nanticoke fire highlighted. In this case, at the time of the fire, the region experience a notable decline in production and remains below its seasonal average in the following weeks.

Inventories Planned shutdowns typically correspond with increases in inventories in the month prior to the shutdown and little to no change in inventories during the shutdown period, relative to their seasonal averages. Conversely, unplanned shutdowns display no discernible pattern in inventories for the preceding month, with large decreases in inventories during the time of the shutdown, relative to seasonal averages. This suggests that suppliers may ramp up production near the end of the month preceding a scheduled closure

Figure 4.8: Calgary Planned shutdown

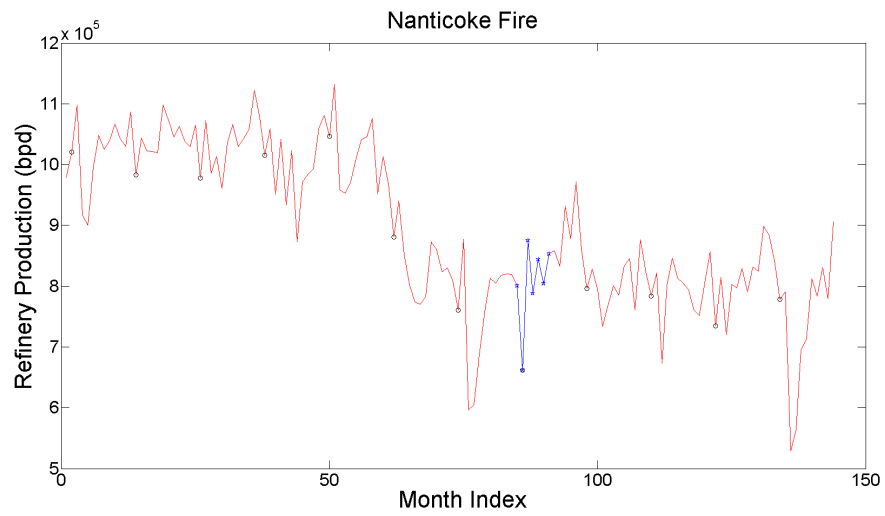


(a) Refinery Production





(b) Retail Prices

Figure 4.9: Nanticoke Unplanned shutdown



(a) Refinery Production



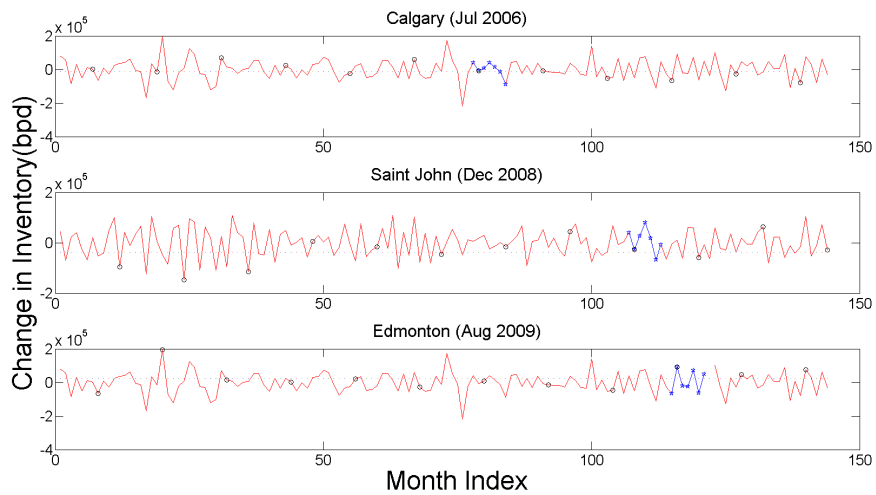
(b) Retail Prices

in order to utilize some of the additional inventories to cover production shortages, while this strategy is simply not available for suppliers who cannot predict the unplanned shutdowns, thus inventories must decrease to accommodate these unplanned decreases in production. Figure 4.10 shows the changes in net inventories during both planned and unplanned shutdowns.


Net Imports Planned shutdowns do not noticeably impact net imports in a region, relative to the seasonal average, while unplanned shutdowns tend to coincide with increases in regional net imports during periods of decreased production. This suggests that suppliers may be able to compensate for regional production shortages by increasing net imports into the region when unplanned shutdowns occur, while net imports may not be affected during periods of planned shutdowns, as prior planning may be able to adequately compensate for these production shortages and not necessitate additional net imports over seasonal averages. Figure 4.11 shows changes in net imports during both planned and unplanned shutdowns.

Taken together, these facts from the data suggest a narrative in which the scenarios of planned vs. unplanned shutdowns induce different responses from suppliers. Since planned shutdowns are generally scheduled for periods when demand is at its lowest, suppliers may be able to mitigate any decreases in current production by using previously increased inventories to account for any shortages. Further, as these planned shutdowns are scheduled and known, net imports can be pre-arranged to compensate for any further shortages in net production, so that net imports during planned shutdowns do not vary from seasonal averages. However, when unplanned disruptions occur, suppliers may resort to covering supply shortages by increasing net imports above seasonal averages. Suppliers' ability to adjust quickly to these unplanned shortages

Figure 4.10: Refinery Inventories

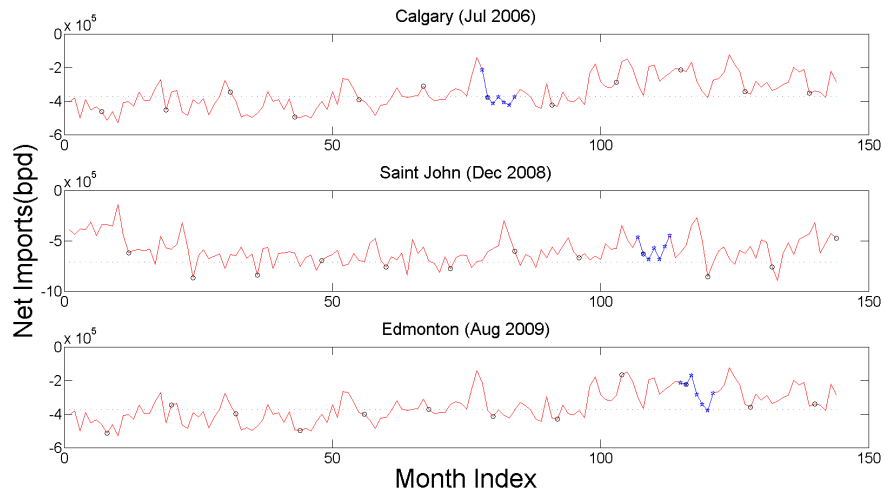


(a) Planned Shutdowns




(b) Unplanned Shutdowns

Figure 4.11: Refinery Net Imports



(a) Planned Shutdowns



(b) Unplanned Shutdowns

will therefore be constrained by the existing supply network available to move product across and within each region. Therefore the structure of the supply network may have a large influence on the capability and speed at which supply shocks can be dissipated in the retail market.

4.6.2 Retail prices

As there are so few instances of prolonged refinery-level shutdowns, this paper focuses on two specific examples: a planned maintenance shutdown in Calgary in late June/early July 2006 and an unplanned shutdown due to a fire at the Nanticoke, ON refinery in February 2007. These two cases studies were chosen for closer examination due to each refinery's relatively similar share of regional production capabilities, ranging from 20–30% of their respective region's typical gasoline production and the relatively central location of each refinery within their regional supply hub.

A preliminary look at retail prices following these supply shocks indicates that retail prices spike slightly more in locations in closer proximity to the refinery that is shutdown than in other regions following an unplanned shutdown. However, following unplanned shutdowns, prices generally move in similar fashions across all regions, and these supply shocks do not appear isolated to local prices. Further, locations close to the shutdown refinery that share pipeline connections tend to exhibit less variation in price changes than those locations with no pipeline access. Once again, this suggests that the supply infrastructure impacts retail price dispersion, as locations connected via pipeline are separated by a smaller effective distance, *ceteris paribus*, and can therefore disseminate supply shocks more quickly, resulting in less variation in retail price.

In the case of the Nanticoke fire, over the three week period following the refinery shutdown, the Ontario average retail price spiked by 15.3% while the Canadian average price increased by 12.6% and Western and Atlantic prices grew by only 6.4% and 6.3% respectively.²⁵ Additionally, the standard deviation of price spikes for Ontario cities that share a common pipeline connection is 0.0263, whereas that of cities without pipeline access was 0.0384. Conversely, in the case of a planned shutdown in Calgary, in the three week period following the shutdown, the Western average retail price decreased by 0.01% while the Canadian average price decreased by 0.03% and the Quebec and Ontario averages changed by -0.02% and +0.03% respectively, all of which reflect almost no change in prices and minimal differences across regions.²⁶ Contrasting these two case studies provides preliminary evidence to reinforce the significance of the supply network in impacting the way in which supply shocks are transmitted into retail price changes across locations.

4.7 Conclusions

The supply network has a significant impact on relative price dispersion in the Canadian gasoline market, limiting arbitrage opportunities that arise from price gaps across locations. Like previous studies on price dispersion, geographic distance and market size are found to be significant factors in accounting for price gaps and violations of the law of one price for a homogeneous good like gasoline. However, the contribution of this chapter is to determine that variation in the available methods of transportation of gasoline products, via pipeline or seaport connections, also significantly impacts price differences

²⁵These results can be seen in the bottom panel of Figure 4.10(b).

²⁶These results can be seen in the bottom panel of Figure 4.9(b).

across locations. Pipelines decrease weekly price dispersion by the equivalent of a 53% reduction in distance, while seaport connections reduce price dispersion by the equivalent of a 38% reduction in geographical distance.

This result is reinforced by a case study of supply shocks in the supply chain for gasoline during refinery shutdowns. Examining these incidents suggests that unplanned refinery shutdowns coincide with decreased regional production levels that are accompanied by spikes in retail prices. These price spikes are larger (in percentage terms) in areas closer to the shutdown, and also exhibit less variation across locations that are connected by pipelines than across those that are not. Planned shutdowns, however, exhibit minimal changes to production levels and no clear price spikes, regardless of geographic location or proximity to the refinery shutdown.

While this chapter offers some initial insight into the role of the supply network on price dispersion in the Canadian gasoline market, more extensive data on pre-tax and retail price dispersion across a larger number of locations, of varying market size, would allow a more precise examination. As more cities are included in the Kent dataset, further work may be able to exploit variation in more remote and smaller cities, as opposed to the larger and more centrally located cities that are currently available. Future potential work may also focus on the role of the supply network on an international scale, by incorporating multiple countries and investigating retail price dispersion on a larger geographic scale, including the impact of border effects. Data from American retail gasoline markets at a level of detail comparable to the Canadian data set could offer a first step in understanding what price differences arise when policy and trade barriers potentially affect supply distribution and price dispersion across countries.

4.8 References

- ANDERSON, J.E. & YOTO V. YOTOV (2010). “The Changing Incidence of Geography,” *American Economic Review*, *American Economic Association*, Vol. 100(5), pages 2157–86
- ADAMS, A.F. (1997). “Search Costs and Price Dispersion in a Localized, Homogeneous Product Market: Some Empirical Evidence”, *Review of Industrial Organization*, Volume 12, Numbers 5–6, pp. 801–808
- BURDETT, K. AND K.L. JUDD (1983). “Equilibrium Price Dispersion”, *Econometrica*, Vol. 51, No. 4, pp. 955–969
- CANADIAN ENERGY PIPELINE ASSOCIATION, Available at: <http://www.cepa.com/>
- CEGLOWSKI, J. (2003). “The Law of One Price: Intranational Evidence for Canada”, *The Canadian Journal of Economics*, Vol. 36, No. 2, pp. 373–400
- CHAKRAVORTY, U., C. NAUGES AND A. THOMAS (2008). “Clean Air Regulation and Heterogeneity in US gasoline Prices”, *Journal of Environmental Economics and Management*, Vol. 55 pp. 106–122
- CHANDRA, A., AND M. TAPPATA (2011). “Consumer Search and Dynamic Price Dispersion: An Application to Gasoline Markets.” *The RAND Journal of Economics* 42.4: pp.681–704.
- CRUCINI, M., M. SHINTANI AND T. TSURUGU (2015). “Noisy Information, Distance and the Law of One Price Dynamics Across US Cities”, *Journal of Monetary Economics*, Vol 74(C), pp. 52–66
- CRUCINI, M., C. TELMER AND M. ZACHARIADIS (2003). “Price Dispersion: The Role of Borders, Distance and Location”, *Tepper School of Business*, Paper 490
- CRUCINI, M., AND H. YILMAZKUDAY (2014). “Understanding Long-Run Price Dispersion”, *Journal of Monetary Economics*, pp. 226–240
- ECKERT, A. (2013). “Empirical Studies of Gasoline Retailing: A Guide to the Literature”, *Journal of Economic Surveys*, 27(1), 140–166.
- ENGEL, CHARLES M. AND J. ROGERS (1996). “How Wide Is the Border?”, *American Economics Review*, Vol 86 No 5, pp. 1112–1125
- FEYRER, J., E. MANSUR AND B. SACERDOTE (2017). “Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution”, *The American Economic Review*, 107.4, pp. 1313–1334

- GORODNICHENKO, Y. AND L. TESAR, (2009). "Border Effect or Country Effect? Seattle May Not Be so Far from Vancouver After All", *American Economic Journal*, 1:1, pp. 219–241
- KANO, K., T. KANO AND K. TAKECHI (2013). "Exagerrated Death of Distance: Revisiting Distance Effects on Regional Price Dispersions", *Journal of International Economics*, Vol 90(2), pp. 403–413
- KENT MARKETING SERVICES, THE KENT GROUP, "Petroleum Price Data", Available at: <http://www.kentmarketingservices.com/dnn/PetroleumPriceData.aspx>
- LEWIS, M. (2011). "Asymmetric Price Adjustment and Consumer Search: An Examination of the Retail Gasoline Market." *Journal of Economics & Management Strategy* 20.2, pp. 409–449.
- MARVEL, H. (1976). "The Economics of Information and Retail Gasoline Price Behavior: An Empirical Analysis", *Journal of Political Economy*, Vol. 84, No. 5, pp. 1033–1060
- MILLS, J., J. SCHMITZ, AND G. FRIZELLE (2004). "A Strategic Review of Supply Networks", *International Journal of Operations & Production Management*, Vol. 24 Iss 10, pp.1012–1036
- MUEHLEGGGER, E.J. (2004). "Gasoline Price Spikes and Regional Gasoline Content Regulations: A Structural Model", Center for Energy & Environmental Policy Research, Nov. 2004
- NAGURNEY, A., J. DONG AND D. ZHANG (2002). "A Supply Chain Network Equilibrium Model", *Transportation Research Part E: Logistics and Transportation Review*, Vol. 38, Iss. 5, pp. 281–303
- NATURAL RESOURCES CANADA, "Infrastructure for Crude Oil and Petroleum Products", Available at: <http://www.nrcan.gc.ca/energy/sources/infrastructure/1390>
- NOEL, M. (2007). "Edgeworth Price Cycles, Cost-Based Pricing and Sticky Pricing in Retail Gasoline Markets", *The Review of Economics and Statistics*, Vol 89, No. 2, pp. 324–334
- PENNERSTORFER, D., P. SCHMIDT-DENGLER, N. SCHUTZ, C. WEISS AND B. YONTCHEVA, "Information and Price Dispersion: Theory and Evidence", CEPR Discussion Paper No. DP10771
- SHEN, Z.J.M. (2006). "A Profit-Maximizing Supply Chain Network Design Model with Demand Choice Flexibility", *Operations Research Letters*, Vol. 34, Iss. 6, pp. 673–682
- STATISTICS CANADA, "Refinery Production, Supply and Disposition of Refined Petroleum Products", Table 134-0004

STATISTICS CANADA, “Opening/Closing Inventories, Supply and Disposition of Refined Petroleum Products”, Table 134-0004

STATISTICS CANADA, “Inter-provincial transfers, Imports, Exports, Supply and Disposition of Refined Petroleum Products”, Table 134-0004

STIGLER, G. (1961). “The Economics of Information”, *Journal of Political Economy* , Vol. 69, No. 3, pp. 213–225

YILMAZKUDAY, D. AND H. YILMAZKUDAY (2016). “Understanding Gasoline Price Dispersion”. *The Annals of Regional Science*, 57(1), 223–252.

Chapter 5

Conclusion

My thesis consists of three essays in international economics. The first two chapters provide a quantitative analysis of bilateral trade growth during rapid growth episodes over the past 60 years. I use bilateral trade data to determine stylized facts of trade growth across goods classifications. I build a standard trade model to assess the role of various sources of heterogeneity in matching these stylized facts — in Chapter 2, I use heterogeneous productivity and tariff changes imputed from trade data, while in Chapter 3, I include a choice for exporters among multiple distribution technologies. Chapter 4 investigates the role of variation in transportation options for gasoline in accounting for observed price dispersion across Canadian cities.

In Chapter 2, I use bilateral trade data to decompose trade growth across goods classifications during episodes of rapid growth in trade. I find that trade growth is granular — a small number of goods account for a the majority of total trade between country-pairs, while most goods exhibit little to no growth in trade. I calibrate a standard Melitz-style trade model, and find that the model predicts less granularity in trade growth — only 10% of that found in the data,

as measured by the share of total trade growth accounted for by various quantiles of goods classifications. To quantitatively assess the model's ability to match the observed granularity in the data, I include heterogeneous productivity and tariff changes imputed from production and trade data and find the model predicts roughly 70% of the observed granularity in the data.

To account for the higher level of granularity in the data, Chapter 3 includes a choice for exporters among multiple distribution technologies. When firms export, they choose among multiple distribution options, which can be broadly grouped into two categories — those with high-fixed and low-variable costs, and those with low-fixed and high-variable costs. Characterizing equilibrium, I find a new channel that generates added granularity in trade growth. Following productivity increases or tariff reductions during periods of trade liberalization, some firms may find it profitable to switch distribution methods — either from not-traded to the low-fixed cost method, or from the low-fixed to high-fixed cost method. These “switchers” exhibit disproportionately larger growth than firms that retain their prior distribution methods, as they experience a compound effect — a direct increase in exports due to the productivity increase or tariff reduction, and the indirect effect of switching to a lower variable cost method. Calibrating the model, I find it now generates roughly 90% of the observed granularity in bilateral trade data, compared to the 70% generated by the model with a single distribution technology.

Chapter 4 examines how variation in transportation methods impacts retail price dispersion across locations. Economic theory suggests that the transportation costs of shipping goods between locations bounds the arbitrage condition that allows for sustained price gaps over time. Many studies use distance

as a proxy for these transportation costs. In Chapter 4, I also consider how variation in transportation methods impacts observed price dispersion for a specific good, gasoline, across Canadian cities. I use a unique data set on weekly average prices in 44 Canadian cities between 2001 and 2017 to quantify the impact of variation in the availability of the four main modes of transportation for gasoline — pipeline, marine tanker, rail and truck — on mean-price and weekly-price differences between locations. Regression analysis finds that the supply network has a significant effect on price dispersion in the Canadian gasoline market. City-pairs connected via pipeline exhibit 3.5% less mean-price dispersion than those connected by rail or truck. Further, the existence of pipelines connecting cities has the effect of reducing weekly city-pair price dispersion by the equivalent of a 53% reduction in geographical distance, while a seaport connection between cities reduces the effective distance by 38%, in terms of weekly price differences, compared to land-route alternatives. These quantitative results suggest that the structure of the supply network is significant in accounting for observed price dispersion across locations.

5.1 References

BURDETT, K. AND K.L. JUDD (1983). “Equilibrium Price Dispersion”, *Econometrica*, Vol. 51, No. 4, pp. 955–969

CHANEY, T. (2008). “Distorted Gravity: The Intensive and Extensive Margins of International Trade”. *American Economic Review*, 98:4, 1707–1721.

CRUCINI, M., C. TELMER AND M. ZACHARIADIS (2003). “Price Dispersion: The Role of Borders, Distance and Location”, *Tepper School of Business*, Paper 490

ENGEL, C. AND J. ROGERS (1996). “How Wide Is the Border?”, *American Economics Review*, Vol 86 No 5, pp. 1112–1125

MELITZ, M.J. (2003). “The Impact of Trade on Intraindustry Re-allocations

and Aggregate Industry Productivity.” *Econometrica*, Vol. 71.6 (2003): 1695–1725.

Appendix A

Chapter 2 Appendix

Figure A.1: Worldwide Trade: 1960-2015

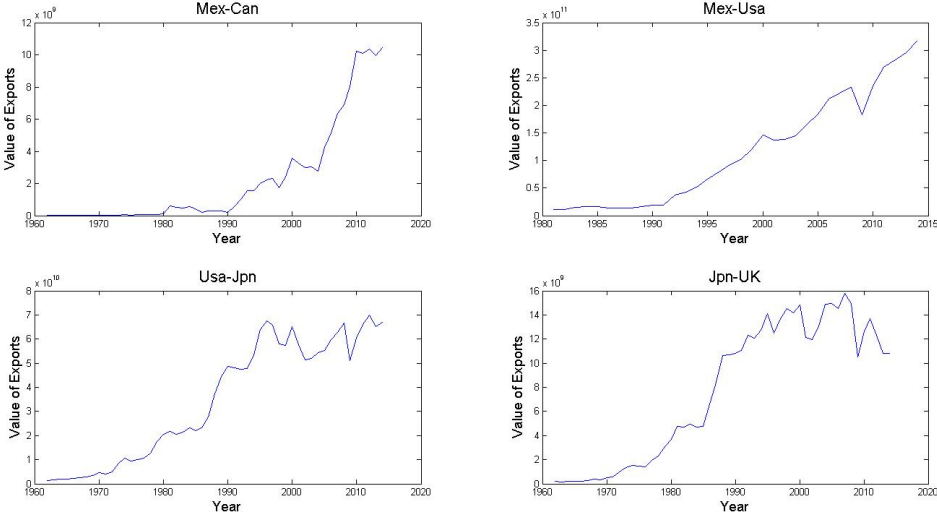


Figure A.2: Growth by Good: Least-Traded Goods (Canada-Mexico)
Canadian Exports to Mexico

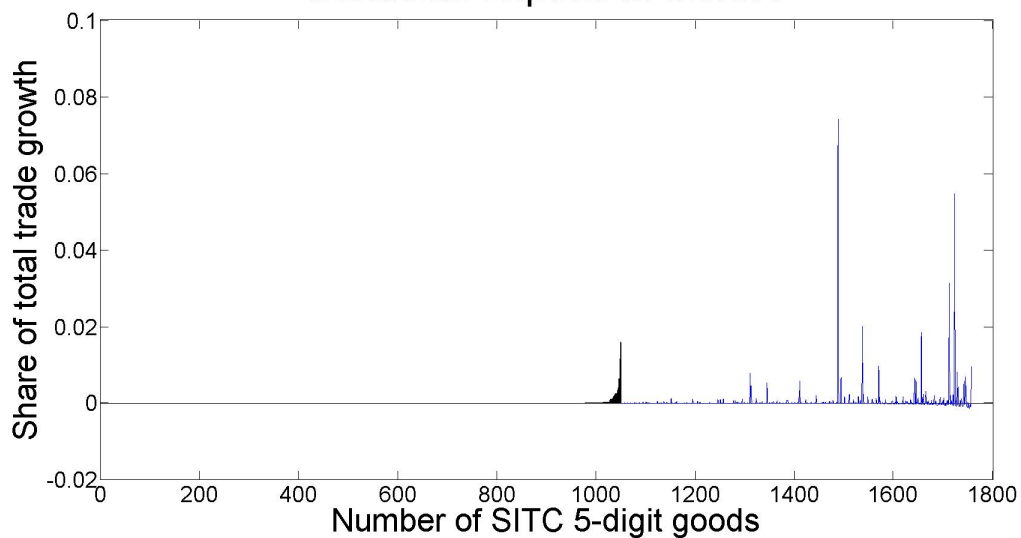


Figure A.3: Growth by Good: Mid-Traded Goods (Canada-Mexico)
Canadian Exports to Mexico




Figure A.4: Growth by Good: Most-Traded Goods (Canada-Mexico)
Canadian Exports to Mexico

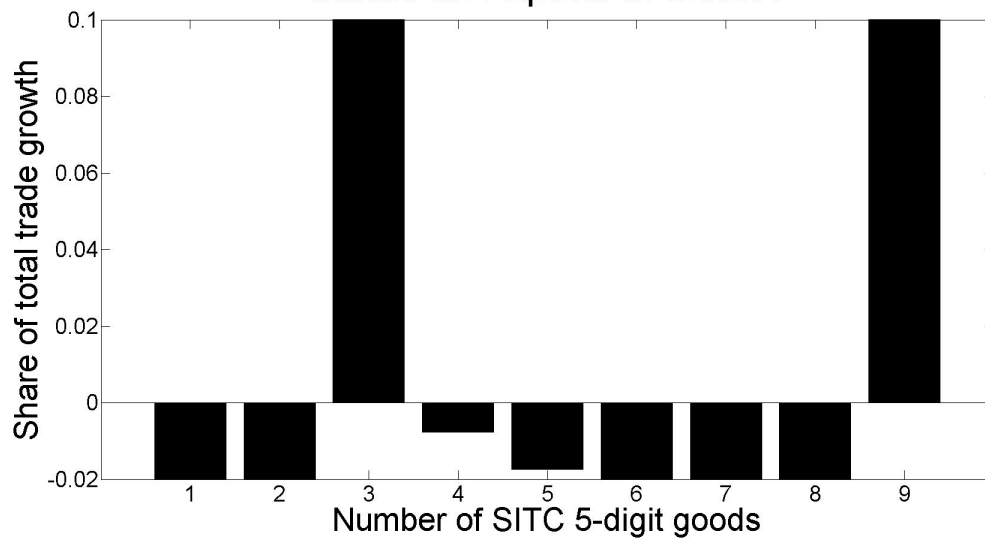


Figure A.5: Correlation: Imports vs. Tariff changes (U.S.-Canada)




Figure A.6: Correlation: Imports vs. Tariff changes (U.S.-Mexico)

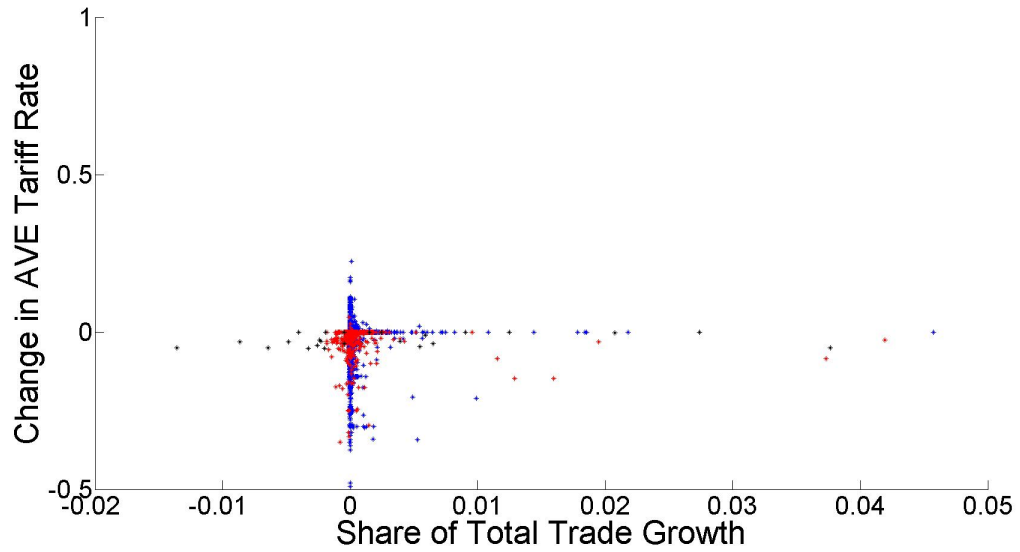


Figure A.7: Correlation: Imports vs. Tariff changes (U.S.-U.K.)


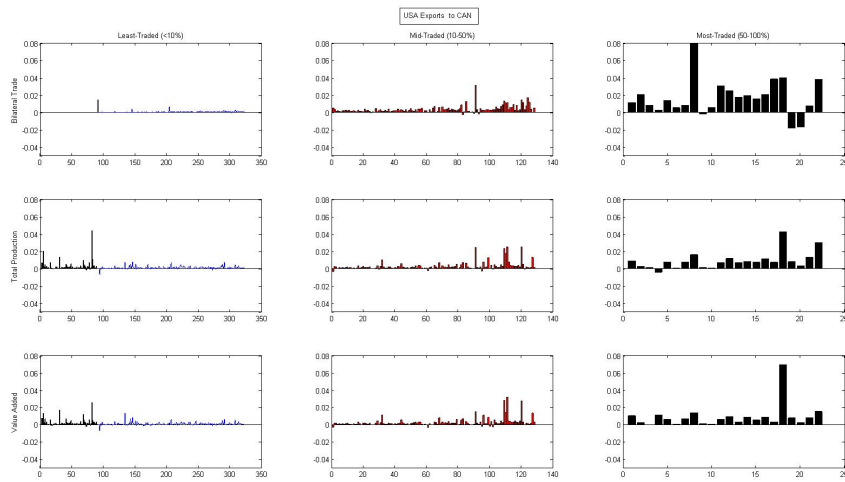
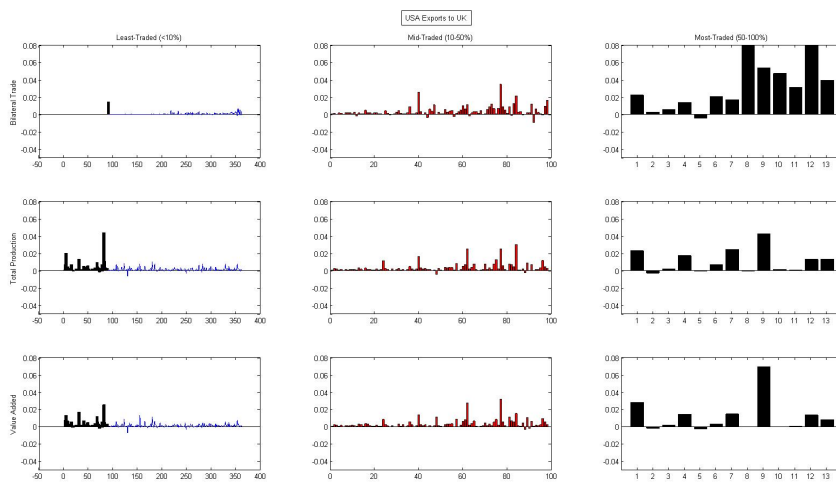


Figure A.8: Growth by Good: 6-digit NAICS bilateral trade



(a) U.S.-Canada



(b) U.S.-U.K.

Appendix B

Chapter 3 Appendix

Figure B.1: Growth by Goods Category: US Exports to Mexico




Figure B.2: Growth by Goods Category: US Exports to Japan




Figure B.3: Productivity Distributions: Heterogeneous Productivity Changes




Figure B.4: Heterogeneous Tariff Changes: Mexico

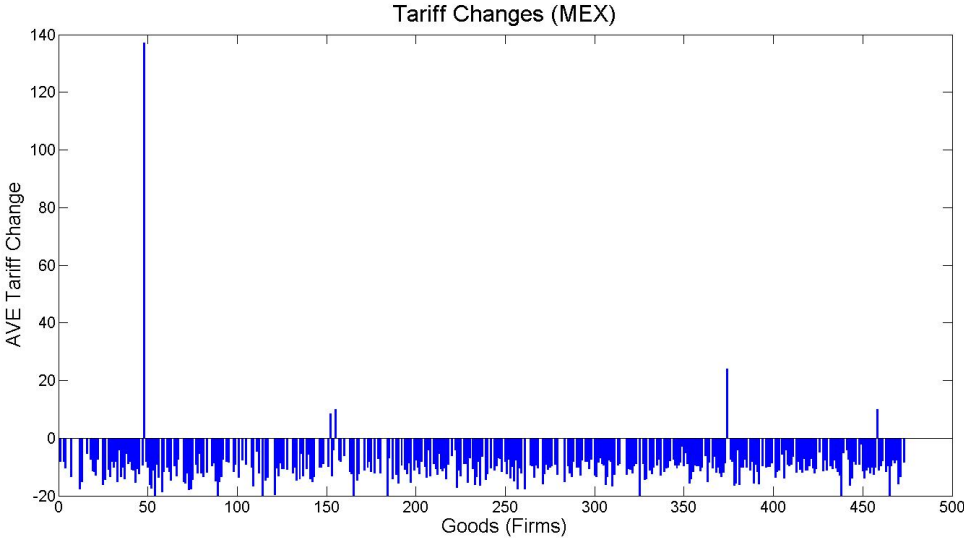


Figure B.5: Heterogeneous Tariff Changes: Canada


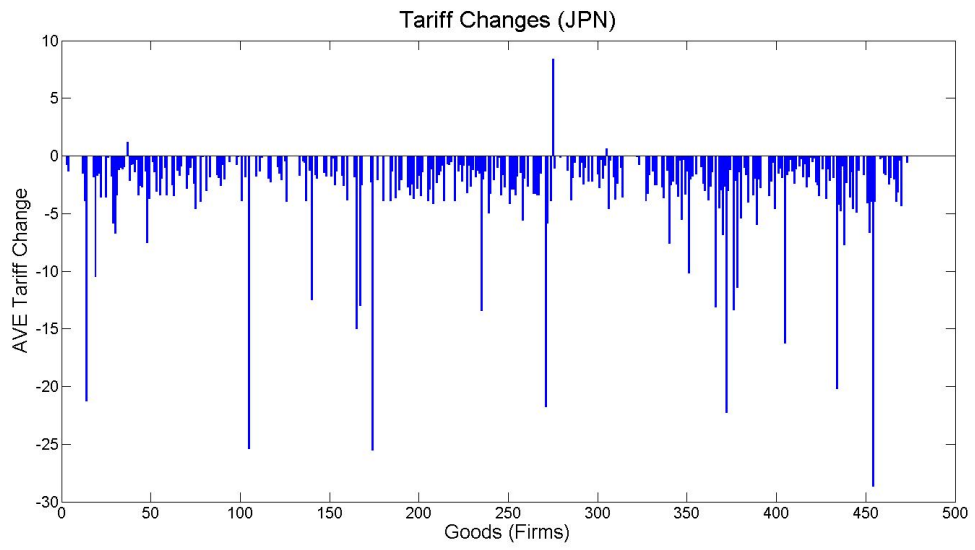


Figure B.6: Heterogeneous Tariff Changes: Japan



Appendix C

Chapter 4 Appendix

C.1 Data Sources and Implementation

The gasoline price data comes from the Kent Group Ltd. website, publicly available at <http://charting.kentgroupltd.com/>. I use weekly data on retail prices, excluding taxes, for regular gasoline, from the 44 cities listed, and compile it over the years 2001 to 2017. These prices represent a city-wide weekly average of gasoline prices in each city, as sampled by the Kent Group from a wide selection of branded and independent gasoline retailers every Tuesday morning at 10:00AM local time.

The distance between cities is calculated as the shortest driving distance as suggested by online navigation system Mapquest.¹ While many papers use variations of the Great-circle or Euclidean distances between locations to measure distance, this paper focuses on the arbitrage condition that governs price dispersion, which is a function of the costs associated with physically transporting gasoline products between locations. The default alternative for transporting gasoline is by truck, as it is the only method accessible to all locations.

¹available at <https://www.mapquest.ca/>.

I therefore use the shortest highway route provided by Mapquest, measured in kilometers, as the measure of distance between any given cities.

The region dummies are calculated to correspond with the natural supply orbits suggested by Natural Resources Canada- cities in British Columbia, Alberta, Saskatchewan and Manitoba, as well as the Yukon and Northwest territories fall into the West region; cities in Ontario and Quebec fall into their respectively named regions; and cities in New Brunswick, Nova Scotia, Prince Edward Island and Newfoundland and Labrador fall into the Atlantic region.

The pipeline dummies are calculated according to the current pipeline infrastructure in Canada, available from the Canadian Association of Petroleum Producers (CAPP). Cities that considered to be connected by pipeline if they are directly connected by an existing pipeline, or if they are both connected to a common third city by pipeline. For example, while Edmonton and Regina may not share a direct pipeline link, they are both connected to Calgary, and are thus considered to be linked via pipeline.

The seaport dummies are set to one if there exists a plausible maritime connection between the two cities, whether or not gasoline products are currently shipped via marine tanker between the two cities. This reflects the fact that the price dispersion between cities is defined by the arbitrage condition governed by the available transportation options, whether they have been employed in the past or not. For example, Saint John, N.B. and St. John's, NFLD are considered to be connected via seaport, where existing shipping routes exist; Thunder Bay, ON and Sault Ste. Marie, ON are considered to be connected by seaport, even though they are not currently serviced by marine tanker; however Thunder Bay and St. John's are not considered to be linked via seaport, since it is not feasible for marine tankers to navigate the sea route between the

cities, due to their size and scale, even though both cities have accessible ports.

The distances to the nearest terminals are calculated by the shortest highway distance provided by Mapquest, as well, using the terminal locations provided by Petroleum industry suppliers, like Esso. ²

C.2 Regression Tables

Table C.1: Mean price differences regressions

Regression estimation: Dependent variable = $ \log(\bar{P}_i) - \log(\bar{P}_j) $									
Variable (Std. Dev)	Reg. 1	Reg. 2	Reg. 3	Reg. 4	Reg. 5	Reg. 6	Reg. 7	Reg. 8	Reg. 9
α (Constant)	0.0013 (0.0163)	0.0059 (0.0128)	0.0110 (0.0112)	0.0136*** (0.0116)	0.0231*** (0.0074)	0.0214*** (0.0074)	0.0199*** (0.0081)	0.0466*** (0.0046)	0.0446*** (0.0047)
β (distance)	0.0374*** (0.0171)	0.0374*** (0.0173)	0.0379*** (0.0175)	0.0370*** (0.0172)	0.0374*** (0.0175)	0.0387*** (0.0175)	0.0398*** (0.0182)	0.0219*** (0.0116)	0.0233*** (0.0114)
γ (Region)	0.0627*** (0.0226)	0.0634*** (0.0233)	0.0642*** (0.0240)	0.0621*** (0.0232)	0.0628*** (0.0241)	0.0621*** (0.0243)	0.0523*** (0.0235)		
δ (Pipeline)	-0.0353*** (0.0115)	-0.0363*** (0.0119)	-0.0385*** (0.0118)	-0.0374*** (0.0122)	-0.0408*** (0.0122)	-0.0385*** (0.0120)		-0.0125** (0.0055)	
λ (Seaport)	-0.0126 (0.0086)	-0.0127 (0.0087)	-0.0129 (0.0095)	-0.0131 (0.0095)	-0.0135 (0.0095)				
θ (Pop diff)	0.0052 (0.0037)	0.0055 (0.0035)	0.0067*** (0.0032)						
θ (Pop dens)	0.0045 (0.0044)	0.0050 (0.0049)		0.0065 (0.0047)					
ψ (Term. distance)	0.0029 (0.0033)								
R^2	0.1974	0.1939	0.1875	0.1861	0.1747	0.1727	0.1544	0.0947	0.0925
Adj. R^2	0.1906	0.1879	0.1824	0.1809	0.1704	0.1692	0.1517	0.0918	0.0906
N (obs.)	946								

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

²available at <http://www.essomaps.ca/terminals>.

Table C.2: Weekly relative price regressions

Regression estimation: Dependent variable = $|\log(P_{it}) - \log(P_{jt})|$

Variable (Std. Dev)	Reg. 1	Reg. 2	Reg. 3	Reg. 4	Reg. 5	Reg. 6	Reg. 7	Reg. 8	Reg. 9
α (Constant)	0.0288*** (0.0042)	0.0331*** (0.0036)	0.0384*** (0.0033)	0.0401*** (0.0029)	0.0499*** (0.0021)	0.0478*** (0.0020)	0.0468*** (0.0020)	0.0703*** (0.0009)	0.0701*** (0.0006)
β (distance)	0.0416*** (0.0017)	0.0416*** (0.0018)	0.0421*** (0.0019)	0.0412*** (0.0018)	0.0417*** (0.0019)	0.0432*** (0.0018)	0.0439*** (0.0018)	0.0282*** (0.0011)	0.0283*** (0.0010)
γ (Region)	0.0563*** (0.0027)	0.0569*** (0.0028)	0.0578*** (0.0030)	0.0558*** (0.0028)	0.0564*** (0.0031)	0.0556*** (0.0031)	0.0493*** (0.0031)		
δ (Pipeline)	-0.0221*** (0.0010)	-0.0231*** (0.0009)	-0.0254*** (0.0008)	-0.0241*** (0.0009)	-0.0276*** (0.0008)	-0.0247*** (0.0007)		-0.0015*** (0.0018)	
λ (Seaport)	-0.0160*** (0.0013)	-0.0161*** (0.0013)	-0.0164*** (0.0014)	-0.0165*** (0.0013)	-0.0170*** (0.0014)				
θ (Pop diff)	0.0048*** (0.0006)	0.0050*** (0.0006)	0.0064*** (0.0007)						
θ (Pop dens)	0.0048*** (0.0004)	0.0054*** (0.0005)		0.0067*** (0.0006)					
ψ (Term. distance)	0.0027*** (0.0004)								
R^2	0.1349	0.1331	0.1288	0.1291	0.1218	0.1199	0.1154	0.0823	0.0823
Adj. R^2	0.1349	0.1331	0.1288	0.1291	0.1218	0.1199	0.1154	0.0823	0.0823
N (obs.)	856130								

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

Table C.3: Mean price differences regression: Excluding Que/ATL
 Regression estimation: Dependent variable = $|\log(\bar{P}_i) - \log(\bar{P}_j)|$

Variable (Std. Dev)	Reg. 1
α (Constant)	-0.0188 (0.0197)
β (distance)	0.0520*** (0.0176)
γ (Region)	0.0751*** (0.0200)
δ (Pipeline)	-0.0405*** (0.0146)
λ (Seaport)	0.0358 (0.0299)
θ (Pop diff)	0.0240*** (0.0069)
θ (Pop dens)	0.0042 (0.0049)
ψ (Term. distance)	-0.0003 (0.0026)
R^2	0.3966
Adj. R^2	0.3835
N (obs.)	378

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

Table C.4: Weekly relative price regression: Excluding Que/ATL
 Regression estimation: Dependent variable = $|\log(P_{it}) - \log(P_{jt})|$

Variable (Std. Dev)	Reg. 1
α (Constant)	0.0148*** (0.0054)
β (distance)	0.0544*** (0.0022)
γ (Region)	0.0660*** (0.0029)
δ (Pipeline)	-0.0298*** (0.0017)
λ (Seaport)	0.0433*** (0.0033)
θ (Pop diff)	0.0209*** (0.0008)
θ (Pop dens)	0.0054*** (0.0005)
ψ (Term. distance)	-0.0004 (0.0004)
R^2	0.2288
Adj. R^2	0.2288
N (obs.)	342090

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

Table C.5: Mean price differences regression: Excluding YK/NWT
 Regression estimation: Dependent variable = $|\log(\bar{P}_i) - \log(\bar{P}_j)|$

Variable (Std. Dev)	Reg. 1
α (Constant)	-0.0299*** (0.0037)
β (distance)	0.0044** (0.0023)
γ (Region)	0.0076* (0.0045)
δ (Pipeline)	-0.0157*** (0.0038)
λ (Seaport)	-0.0095** (0.0044)
θ (Pop diff)	0.0019 (0.0014)
θ (Pop dens)	0.0030* (0.0018)
ψ (Term. distance)	-0.0000 (0.0007)
R^2	0.0791
Adj. R^2	0.0704
N (obs.)	861

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

Table C.6: Weekly relative price regression: Excluding YK/NWT
 Regression estimation: Dependent variable = $|\log(P_{it}) - \log(P_{jt})|$

Variable (Std. Dev)	Reg. 1
α (Constant)	0.0556*** (0.0024)
β (distance)	0.0114*** (0.0008)
γ (Region)	0.0054*** (0.0014)
δ (Pipeline)	-0.0040*** (0.0007)
λ (Seaport)	-0.0130*** (0.0009)
θ (Pop diff)	0.0014*** (0.0004)
θ (Pop dens)	0.0035*** (0.0003)
ψ (Term. distance)	-0.0001* (0.0003)
R^2	0.0381
Adj. R^2	0.0381
N (obs.)	779205

***-significant at 99% CL, **-sig. at 95% CL, *-sig. at 90% CL

CURRICULUM VITAE

Name: Brandon Malloy

Place of Birth: New Brunswick, Canada

Year of Birth: 1984

Post-Secondary Education and Degrees: 2003–2007 B.A. Mathematics & Economics
Bowdoin College. Brunswick, Maine

2010–2017 Ph.D.
The University Of Western Ontario. London, Ontario

Honors and Awards: Ontario Graduate Scholarship (OGS)
2014–20015

Tutorial Leader of the Year Award, UWO
2010-11, 2011-12

Related Work Experience: Research and Teaching Assistant
The University of Western Ontario
20010-2016

Instructor
The University of Western Ontario
2012–2016

Lecturer
King’s University College, UWO
2013–2016

Assistant Professor
St. Francis Xavier University
2016–present