Essays in International Economics

by

Laurien E. Gilbert

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics and Public Policy)
in The University of Michigan
2018

Doctoral Committee:

Professor Alan V. Deardorff, Co-Chair
Professor Andrei Levchenko, Co-Chair
Assitant Professor Dominick Bartleme
Assistant Professor Javier Cravino

Laurien E. Gilbert

legilb@umich.edu

ORCID iD: 0000-0001-8254-2433

DEDICATION

This dissertation is dedicated to my parents, Clifford and Patricia Gilbert, in gratitude for their constant encouragement and belief in me.

ACKNOWLEDGEMENTS

I am grateful to my advisers, Andrei Levchenko, Alan Deardorff, Javier Cravino and Dominick Bartelme for their indispensable guidance, support and encouragement throughout the doctoral program. I would like to thank the faculty and doctoral students of the Department of Economics at the University of Michigan for their feedback and encouragement as well.

I would like to acknowledge financial support from Rackham Graduate School, the Gerald R. Ford School of Public Policy, and the Board of Governors of the Federal Reserve System, as well as the Kilts Center at the University of Chicago Booth School of Business for data provision and support.

PREFACE

This dissertation makes use of data from Nielsen. I am required to note that some results are calculated (or derived) based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are mine and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF APPENDICES

ABSTRACT

The dissertation consists of two essays. The first considers product entry as a channel for business cycle propagation. Net product entry is procyclical, which amplifies fluctuations in consumer welfare over the cycle if consumers have love for variety. Using barcode-level data covering grocery expenditures in 26 major cities, I establish that differences in city-level product entry are largely uncorrelated with local economic conditions. I provide evidence that city-level changes in product variety over the business cycle are driven instead by multi-city retailers who introduce new products simultaneously in all cities in which they operate. This suggests that product introduction by multi-city retailers can propagate business cycle shocks. To quantify the impact of this mechanism, I develop a quantitative model of retailer product choice that relates the welfare gains from product entry in each city to demand growth in every other city. The model implies that the contribution of other cities' business cycle shocks to each city's price level is proportional to the share of other cities in retailer revenue. Since the share of outside cities in retailer revenue is 63 percent in the average city, the impact of other cities' shocks on product entry is substantial. The presence of multi-city retailers makes net product entry more correlated across cities than they would be if retailers operated in only one city. In a counterfactual in which retailers do operate in only one city, the variation in gains from new product entry across cities would be 47 percent higher than in the baseline.

The second essay focuses on the entry behavior of new exporting firms. New exporters take a few years to catch up to the average total sales of all exporters. Using Colombian export data from 2007-2012, I show that the slow catchup of new exporters is the result of low market participation and sequential market entry (the extensive margin), rather than low sales in individual foreign markets (the intensive margin). Motivated by my empirical findings, I develop a quantitative model in which firms enter markets sequentially, in order of decreasing profitability. I use this model to calculate the benefit from a policy designed to reduce bilateral entry costs and show that as the result of fixed costs that depend on export experience, trade promotion in one market may provide significant increases in exports to markets not targeted by the policy.

Chapter I: Gains from Product Variety and the Local Business Cycle

1 Introduction

The Great Recession in the United States demonstrated that aggregate business cycle measures can mask substantial differences in real household consumption and income across geographic regions. If households move in response to differences in local economic conditions, variation in productivity and employment across regions should be short-lived. Instead, output in some metropolitan areas remained depressed for years after the aggregate recovery, while other cities recovered quickly. Recent work has shown that shocks to local labor and housing markets can have persistent consequences for economic activity. However, relatively little is known about how markets for consumer products respond to local conditions over the business cycle. I focus on the introduction of new product varieties, which decreases the price level when consumers exhibit love for variety in their preferences.

This paper shows that multi-city retailers transmit the welfare consequences of business cycle shocks across cities through the synchronized introduction of new products across the cities in which they operate. I begin by using scanner data to establish three empirical findings that indicate that net product entry at the city level is primarily related to the decisions of retailers who syncronize their introduction of products across cities. Second, I build a quantitative model of retailer product choice to measure how much transmission of the welfare effects of city-level business cycle shocks occurs through common retailers.

I use barcode-store-level data from grocery retailers from Nielsen to study net product entry in 26 large US cities from 2006-2014. The data include revenue and average weekly prices at the barcode-store level. Because the scanner data set includes the universe of sales of barcoded products within participating retailer stores, the set of products in the data represents a complete picture of all varieties purchased by consumers at these stores.

Using the data, I document three findings that establish the role of retailers in explaining net product entry at the city level. First, I show that net product entry across cities varied from 9 to 15 percentage points during the period 2008-2010 and from 15 to 20 percent during the period 2011-2014, suggesting meaningful differences in the availability of new products to consumers across cities. Broda and Weinstein (2010) show that aggregate net product entry is procyclical. However, differences in city-level demand growth explain little of the cross-city variation. Alternative measures of the local business cycle such as local GDP growth, wage growth, or house price growth are also uncorrelated with city-level changes in local product variety.

Second, I find that retailers coordinate the introduction of new products across all cities in which they operate. Eighty-six percent of new products (by value) are introduced in all of a retailer's markets within a one-year period. Variation at the retailer level explains most of the variation in net product entry across city-retailers. I decompose net product entry in each retailer-city into retailer and city components and find that the retailer component can account for 88 percent of the variation in net product entry in an average year. Relationships with manufacturers do not appear to drive the coordination of product entry by retailers. On average, a retailer that already sells products from a particular manufacturer introduces only twelve percent of new barcoded products from that manufacturer.

Third, I find that the retailer's decision to supply new products is strongly related to changes in total sales by the retailer. I measure changes in demand two ways. First, I use growth in the retailer's sales across all markets as a direct measure of demand growth. I regress net product entry by the retailer in each city on growth in retailer sales and find a positive and significant relationship. However, the OLS estimate may be contaminated by simultaneity if the retailer experiences a productivity shock or by reverse causation if firms that introduce new products attract a significantly larger share of demand. In order to relate net product entry to a plausibly exogenous measure of changes in demand, I construct a measure of city-level demand growth weighted by the retailer's share in each market and use it as an instrument for retailer sales growth. The identifying assumption is that changes in MSA-level grocery expenditure are independent of the supply decisions of any particular retailer. I find that a 10 percent increase in a retailer's demand instrumented with the weighted demand growth measure is associated with a 3.6 percent change in new products by value introduced by the retailer.

In sum, the empirical findings show that there is significant variation in net product entry across cities, not driven by differences in local business cycles. Instead, I find that product entry occurs as the result of changes in the set of varieties offered by retailers in response to changes in their total sales. Product entry across a retailer's markets is highly coordinated, suggesting that retailers introduce the same set of products across cities in response to changes in their combined sales. As a result, retailer product entry can act as a mechanism through which the welfare impact of business cycle shocks is transmitted across cities. An increase in demand in one city will motivate retailers operating in that city to introduce new products. Because they also receive the newly introduced products, all other cities served by those retailers will inherit part of the welfare gain associated with the demand growth experienced by just one city.

In order to measure the extent to which retailer product entry decisions transmit shocks across cities, I construct a quantitative model that I calibrate using the microdata. Multi-city retailers choose a set of products to offer to consumers. Because consumer preferences exhibit love for variety, retailers can increase their revenue by increasing the set of products they offer. However, retailers face a convex fixed cost associated with increasing their set of products. Consistent with the empirical finding that retailers introduce the same products across all markets, all fixed costs associated with the introduction of new products are paid nationally. The optimal number of products offered by the retailer will increase as a function of its total revenue.

City-level productivity shocks affect the number of available varieties through a change in demand in each of each retailer's markets, prompting the retailer to change its product set. The same change in the number of products occurs within all the retailer's stores, redistributing the welfare impact of city-level shocks across all cities in which the retailer operates. I derive analytical expressions that relate the welfare gains associated with new products to changes in demand in the cities in which the retailer operates and two parameters: the elasticity of demand governing the welfare contribution of new products and a parameter governing the curvature of the fixed cost function.

I calibrate the two key parameters of the model using the microdata. First, I calibrate the parameter governing the curvature of the fixed cost function to the observed elasticity of net product entry to demand growth. Second, I calibrate the elasticity of substitution across goods, which will determine the magnitude of the response of the price level to net product entry. I begin by using the barcode-level microdata to estimate microelasticities as in Feenstra (1994) and Broda and Weinstein (2010). I consolidate the set of microelasticities into a single 'macro'-elasticity by matching the elasticity of city-level welfare gains to city-level net product entry. The 'macro'-elasticity between goods is essentially a weighted average of the microelasticities, a calibration strategy that avoids the bias associated with estimating a macroelasticity directly. I calculate the response of net product entry in the model to productivity shocks that I set equal to consumption growth over the periods of the recession and recovery.

To evaluate the model's fit, I compare the welfare gains from product entry experienced by consumers in the model with welfare gains calculated using barcode-level microdata. Because the productivity shock in the model affects demand for all products equally, net product entry in the model need not result in the same estimated welfare gains as I do using data on consumption of each barcoded product at each retailer in each city. Nevertheless, I find that the retailer model captures city-level welfare gains well.

I then decompose the model-generated gains from new products experienced in each city into a contribution coming from each other city. The model implies that the effect of outside demand shocks on each city's net product entry is proportional to the share of retailer revenues generated in outside cities. The average city generates less than half of total revenues for the retailers operating in that city, suggesting that outside demand shocks can significantly affect net product entry in the average city. The model predicts that 63 percent of the welfare gains in the average city result from shocks in other cities.

The model implies that the impact of city-level demand growth on consumer welfare in each city depends on the extent of product market integration between cities. In order to quantify the extent to which the empirical distribution of retailers amplifies city-level differences in business cycle outcomes, I compare the baseline model to two counterfactual retailer distributions. At one extreme, if all retailers operated in only one market, the entry of new products would respond only to local demand growth. On the other, if all retailers operated nationally, the entry of new products would respond to the aggregate business cycle. I compare the variation of gains from new products across cities in the calibrated baseline model to these two polar counterfactuals in which I generate predictions for gains from net product entry. In the local counterfactual, I generate welfare predictions under a distribution of shares such that each operates in only one city. In the national counterfactual, I set market shares such that all retailers operate nationally.

I find that the variation in welfare gains across cities in the local counterfactual is 47 percent higher than in the baseline model. This suggests that multi-city retailers smooth the welfare gains from local demand shocks across cities. However, the observed distribution of retailers still results in city-level welfare gains that amplify differences in business cycle outcomes across cities. Under the counterfactual distribution of retailers in which all retailers operate nationally, there is no variation in city-level welfare gains from product entry. The fact that there is still considerable idiosyncratic variation across cities under the observed distribution of retailer shares is consistent with the observation that most food retailers are regional rather than national: retailer demand responds to a subset of all idiosyncratic city-level demand shocks.

This paper relates to a recent literature that has connected product variety to productivity growth and the business cycle. Bernard, Jensen and Schott (2010) and Aghion et al (2017) study the implications of new product creation for measured productivity. Broda and Weinstein (2010) and Erickson and Pakes (2011) consider the impact of product entry and exit on the measured consumer price level. I contribute to the literature by considering cyclical net product entry at the local level, I am able to measure the welfare impact of new products available to consumers in their own product markets.

A second related literature studies changes in the consumer price level over the cycle. Coibion, Gorodnichenko and Hong (2012) and Jaimovich, Rebelo and Wong (2017) find that consumer substitution across stores and product qualities represents a significant source of bias affecting measured consumer prices and their business cycle implications. Argente and Lee (2016) focus on the distributional implications of cyclical substitution. Gagnon and Lopez-Salido (2014) present evidence that local prices do not respond to local shocks. I incorporate the findings of this literature by allowing for substitution across retailers and echoes the implication that prices are not determined at the local level. My paper is the first in this literature to explicitly consider retailers as a source of inter-city business cycle propagation.

Finally, this paper contributes to an extensive literature that studies regional business cycles. Moretti (2011) provides a recent review of the literature on local labor markets. Hanson (2005) and Bartelme (2015) emphasize the role of regional market access in explaining regional economic outcomes. Asdrubali et al (1996), Stumpner (2014) and Caliendo et al (2017) consider mechanisms through which regional shocks propagate within the U.S. This paper presents a novel channel by which local shocks may be transmitted through the price level.

The rest of this chapter is organized as follows. Section 2 describes the data. Section 3 presents empirical evidence that local product variety is determined by retailer product selection. Section 4 develops a model in which retailer product choice transmits demand shocks across cities. Section 5 describes how parameters of the model are calibrated to match moments in the data. Section 6 discusses the implications of the model for transmission of shocks across cities. Section 7 concludes.

2 Product Variety in Scanner Data

2.1 Data Description

Data on consumer products come from the Nielsen Retail Scanner data for the period 2006-2014.[1] The data contain average weekly prices and quantities for all UPC barcodes sold in participating stores across the United States along with descriptive information about the UPC created by Nielsen. Data are sent directly from the retailer or its parent company to Nielsen. Participating retailers include grocery, mass market, convenience and liquor stores, with the best coverage in the first two categories.[2] The data cover a significant proportion of consumer spending on food at home, but the fraction of total food-at-home expenditures captured in the data does vary across cities due to differences in the share of local expenditure within participating retailers. I restrict my analysis to 26 metropolitan areas in the sample with populations above 2 million and expenditure coverage above 25 percent in grocery and mass market stores as reported by Nielsen.[3] For convenience, I refer to metropolitan areas as 'cities.'

[1] These data are available through the Kilts Center for Marketing at the University of Chicago Booth School of Business.

[2] Mass market retailers carry the range of products commonly sold at grocery stores and non-grocery products, such as clothing or furniture.

[3] At least 40 percent of the US population lives within these metro areas according the 2010 Census.

This paper focuses on the entry and exit of final goods into retailers rather than entry and exit of retailers into city markets. Retailers freely select into providing data to Nielsen, making the data a noisy source of information about the entry and exit of individual grocery stores or retail chains. To avoid the issue of entry and exit of retailers themselves, I restrict the data to a sample that contains continuing retailers over each two-year horizon. For instance, retailers included in the sample for 2010 will include all retailers who also sold in 2008 and 2009 without change of ownership. These stores account for 60 percent of all expenditures in the Retail Scanner Data.

Because of the issues noted with interpreting entry and exit of retailers in the scanner data, I also consider the contribution of continuing retailers to total consumer expenditures using an alternate data source produced by Nielsen, the Homescan (HMS) data collected from household barcode scans. Because this data set records expenditures reported by the household at all retailers, not just those that share store-level data with Nielsen, it provides a more comprehensive look at the set of retailers available to the consumer in each city. In this data set, retailers that continue to sell over a two-year horizon account for 99 percent of expenditures in the average city-year, while retailers that continue to sell over the entire 2006-2014 period account for 96 percent. This suggests that retailer entry over the business cycle is not a quantitatively important phenomenon.

While the data contain both grocery and non-grocery UPCs, coverage is not uniform across categories. As an additional check on the data, I compare characteristics of the scanner data to the HMS data. I examine the revenue growth rates in each city-year for each product group as defined by Nielsen.[4] In the HMS grocery categories, average revenue growth is 1.9 percent on average, while it is 4.2 percent in the non-grocery data. Revenue growth in grocery categories is 0.9 percent on average in the retail scanner data and 0.8 percent in non-grocery. The very low growth in non-grocery sales is likely the result of

[4]Nielsen classifies UPCs into nested categories in order of increasing detail: departments, product groups and product modules. A 14 ounce can of Dole peaches would be found in the canned peach module, the canned fruit group, and the dry goods department.

less reliable coverage in these categories: most retailers in the sample specialize in grocery products. I restrict my analysis to grocery categories (other than fresh produce and meats, which are less likely to have UPCs). These data account for 62 percent of expenditures in the retail scanner data and 68 percent of expenditures reported by households in the HMS data. Last, to screen for potential problems in data reporting, I restrict the sample to product categories that are available in all years within a retailer. This step eliminates about 0.5 percent of all expenditures in the sample.

Table I.1 lists descriptive statistics by city in the final sample averaged over the years 2006-2014. Even after removing non-continuing stores, non-grocery UPCs and outlier product categories, the retail scanner data represent about twice as many products per city as the household scanner data collected by Nielsen over the same time period because products are included in the dataset if even one consumer purchases them from a participating retailer. These data include 18,629 stores corresponding to 71 retailers. Nielsen defines stores based on particular chain brands, rather than by parent company.[5] In the analysis, I focus on the retail parent. Total expenditures in the final dataset used for analysis represent 80.3 billion dollars of grocery expenditure on average per year, or about 12.5 percent of all expenditures on food at home in the United States.[6] The percentage of grocery expenditure in each city that is captured by the retail scanner data varies across cities, but expenditures are broadly increasing in city size. While the data only covers food expenditures, which make up about 10 percent of consumer expenditures each year, other categories of consumer expenditure likely exhibit gains from variety as well. Consumption of all non-durables accounts for about 22 percent of expenditures.

Most retailers in the data sell in only a fraction of all 26 cities of the data, but revenue per market increases with city size. Table I.2 describes how many one-city and multi-city

[5]The identities of retailers and parent companies in the sample are not released by Nielsen. One parent company ID may be associated with several retailer IDs. For example, XYZ Groceries, Inc. might have two types of retail location: XYZ Full and XYZ Express. These two chains would typically have different retailer IDs and the same parent ID. Some companies prefer to aggregate data to the parent-level before sending to Nielsen.

[6]Based on expenditure estimates from the Economic Research Service, USDA.

Table I.1: Descriptive Statistics

City	Number of Products	Number of Stores	Number of Retailers	Total Expenditures (millions)
Atlanta	197,893	1,651	13	2,746
Baltimore	208,079	1,559	15	4,556
Boston	183,186	2,110	15	4,854
Charlotte	175,999	1,263	13	2,077
Chicago	240,105	2,495	15	5,314
Cincinnati	156,133	664	10	1,717
Columbus	177,728	604	9	1,611
Dallas	193,581	1,596	12	2,967
Denver	200,491	1,215	16	3,317
Detroit	156,999	1,207	8	2,034
Houston	186,473	1,560	9	2,550
Los Angeles	206,582	3,014	14	8,277
Miami	142,462	963	9	1,245
New York	258,936	4,563	22	7,927
Orlando	155,209	1,174	13	986
Philadelphia	227,422	2,447	19	4,059
Phoenix	193,467	1,349	11	3,433
Pittsburgh	200,881	958	11	1,607
Portland	167,338	648	10	2,029
Raleigh-Durham	181,228	1,119	12	1,905
Sacramento	162,833	734	12	1,769
San Diego	171,210	516	10	1,790
San Francisco	166,459	1,271	10	3,930
Seattle	174,884	1,080	10	3,572
Tampa	148,770	1,305	13	1,187
Washington, D.C.	206,276	801	15	1,773

Notes: Values in each column represent the average over the years of the sample, 2006-2014.

Table I.2: Retailers by Number of Cities

Cities	Number of Retailers	Average City Revenue
1	23	475
2	14	777
3	9	1,170
4-6	8	1,090
7-10	7	1,540
11+	9	3,060

Notes: The table reports the maximum number of cities in which the retailer sells between 2006-2014 and the average per-market sales (in millions) per year for those retailers.

retailers there are in the data. While retailers operating in only one city make up almost one third of all retailers in the data, retailers that operate in multiple cities have significantly higher revenue per market.

This dataset captures expenditures at physical stores in each city in the sample, but consumers may have access to grocery products through online retailers. In principle, differences in the set of varieties available in local stores could be unimportant if consumers shop online, thereby access a common national set of varieties. During the period, expenditures at online retailers account for about 1 percent of all expenditures in the HMS data for any given year, suggesting that online retailers occupy a low share of grocery expenditure. Of course, it is possible both that online shopping is underreported compared to shopping at physical retailers and that the prevalence of online grocery shopping has grown since 2014, the last year of the sample. In section 6.2, I consider how the transmission of business cycle shocks and consumer welfare would differ if all retailers operated nationally. Were online retailers to represent a increasingly large fraction of expenditures, outcomes would approach this counterfactual, in which entry of variety responds to aggregate demand because retailers are active in all cities.

2.2 Defining the Entry and Exit Horizon

In the retail scanner data, entry and exit of products must be inferred from the time series of UPC sales, since there are no data on goods held in inventory. Over short horizons, inference about product entry and exit is likely to be complicated by inventory considerations, partial year effects, and clearance sales. To lessen this ambiguity, I choose to consider the extensive margin over periods of at least two years. In particular, a product is counted as new in year t if it is sold in year t, but not year $t-2$. Similarly, a product is said to exit the market in year t if it is sold in t but not in $t+2$. This timing convention allows direct comparison between the revenue of entering and exiting products in year t. Because this time horizon implicitly allows a product to be counted as entering or exiting for two years, I annualize the net product entry rates obtained using the two-year horizon.

Considering the cost of keeping unsuccessful products on the shelf and the fact that many grocery products are subject to spoilage, it is unlikely that a product would remain unsold but in inventory for a year or longer. Similarly, the wo-year horizon avoids the well-documented partial year problem: new products may enter partway through the year, making their sales hard to compare with those of existing products. The two year horizon also helps to avoid attributing a low expenditure weight to products that are initially less familiar to the consumer, resulting in low sales immediately after entry, or to those that go on clearance sale just before they are eliminated. These issues are likely to result in low revenue, but the resulting prices and expenditure will confound marketing or inventory considerations with the utility that the consumer derives from the product. After calculating entry and exit rates based on the two-year horizon, growth rates in this paper are annualized by dividing by two.

Using Homescan data from Nielsen over an earlier period, Broda and Weinstein (2010) choose to consider product entry over a longer four-year horizon: a product is counted as 'entering' in 2010 if it is sold in 2010 and not in 2006, regardless of sales in 2007, 2008 or

2009. They argue that the long horizon captures the revenue share of new products better than a short horizon, in which new products may not have realized their full potential sales. However, the long horizon also excludes any products that are sold for less than four years. More than 55 percent of products that are observed to enter the national market between 2006-2012 sell for three years or less. Nearly 15 percent of these products sell for less than two years. Because the Retail Scanner data offer a larger sample of consumer expenditures than the Homescan data, it is possible to use a shorter horizon.

3 Three Findings about Net Product Entry

This section documents three findings about city-level net product entry and the associated reduction in the price faced by consumers. First, I show that despite the aggregate procyclicality of net product entry, local business cycles do not appear to explain variation in net product entry at the city level. Using an approach similar to that in Broda and Weinstein (2010), I calculate implied welfare gains from net product entry and show that they are both sizable and variable across cities. The magnitude of city-level welfare gains is weakly related to local demand growth. Second, I show that retailers, rather than city demand conditions, appear to drive local net product entry. Third, I show that net product entry at the retailer level is strongly related to the retailer's revenue growth stemming from demand shocks across all its markets.

3.1 City-level Demand Does Not Explain Local Net Product Entry

I use the Nielsen data to measure changes in the set of available products in each of 26 cities. Net product entry is pro-cyclical, but the extent of product variety changes across US cities varies considerably. Figure 3.1 displays the average net entry of new products by value each year over the period of the recession, 2008-2010, and the recovery, 2011-2014, against one-

Figure 3.1: Net Product Entry by Value and City-Level Expenditure Growth

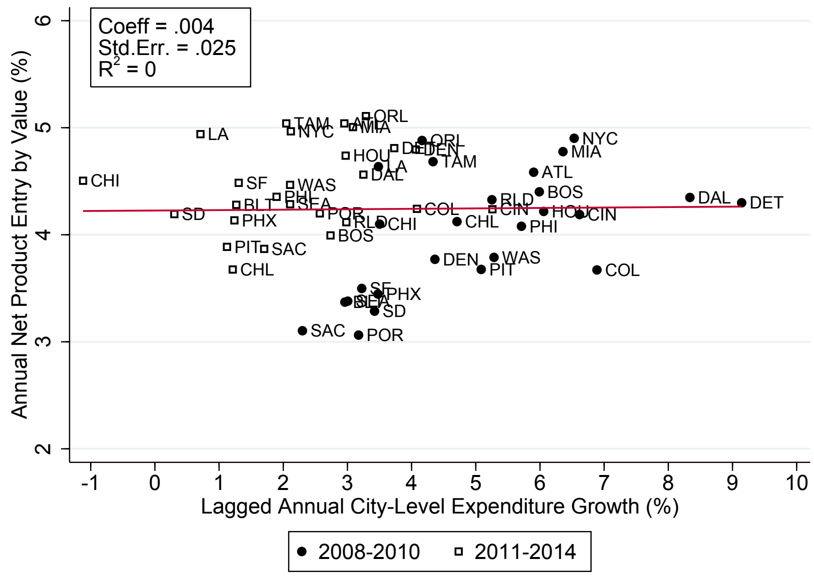

Notes: The figure plots average annual net entry of products for the periods 2008-2010 and 2011-2014 against average annual lagged city-level expenditure growth in the Nielsen Retail Scanner data in 26 US cities.

year-lagged city level expenditure growth over each period.[7] The figure also plots a linear regression of average annual net product entry on average lagged city-level expenditure growth for each period. The relationship is positive but statistically insignificant: city-level expenditure growth does not appear to explain variation in product entry across cities. While changes in expenditure growth in each city represent the most natural measure of business cycle movements in the Retail Scanner data, I also test the relationship between net product entry and other measures of the business cycle such as city-level GDP per capita growth, city-level wage growth, and changes in city-level house prices. Among these measures, lagged expenditure growth is most correlated with net product entry and the associated welfare gains. In what follows, I show that product entry is determined by retailers, rather than city-level shocks.

In order to understand how much new products matter for consumer welfare, it is necessary to take a stance on the consumer's utility function. I follow Feenstra (1994) and Broda and Weinstein (2010) in assuming that consumers have CES preferences. A key char-

[7]Net product entry exhibits a positive but insignificant relationship with contemporaneous consumption growth.

acteristic of a CES price index is that all else equal, an increase in the number of products reduces the price level because consumers have a love for variety. For smaller elasticities of substitution, the negative effect of new products on the price level is larger.

Following Broda and Weinstein (2010), I assume that consumers have a three-tier utility function over goods. However, I also allow grocery retailers to be imperfect substitutes, while retailers in their model of consumption are perfect substitutes. Evidence from the literature on shopping behavior over the business cycle suggests that consumer shopping behavior is consistent with imperfect substitutability across retailers. Griffith et al. (2009) find that most households shop for groceries once or twice a week. Nevo and Wong (2015) and Coibion, Gorodnichenko and Hong (2015) find evidence that households increase shopping intensity during recessions, including store-switching, implying that there is a cost associated with substitution across retailers.

The tiers of the price index consist of UPCs aggregated to a brand-product 'module', brand-product modules aggregated to a coarser product 'group', product groups aggregated within a retailer, and retailers aggregated to total grocery consumption.[8] The utility function first aggregates consumption over goods within a retailer, then aggregates across retailers.

Following Broda and Weinstein (2010), I estimate elasticities of substitution for each tier of the utility function following a GMM approach that relies on the identifying assumption that demand and supply shocks are uncorrelated. In Appendix 7, I discuss this approach in more detail and report statistics on the elasticities I calculate. I find a median cross-UPC elasticity of 6.23, consistent with magnitudes in the literature. In the appendix, I also discuss an alternate approach to estimating elasticities developed by Hausman (1996) that uses price changes in other markets as an instrument for supply shocks. As in other studies, I find a much smaller cross-UPC elasticity of 1.93 using the Hausman approach. To maintain comparability with Broda and Weinstein (2010), I report estimates using their approach in the main text. Note that using smaller elasticities would produce systematically

[8]Product modules represent fine product categories ('Canned Peaches' vs 'Canned Pineapple') while product groups represent broad product categories ('Canned Fruit' vs 'Coffee').

larger estimates of the welfare gains from new products, thereby amplifying the differences in outcomes across cities as well.

For brevity, I describe the price level of a generic tier x of the four-tier utility function, with elements v. In the absence of shocks to demand for x, the two-year growth in the exact price index, $\pi_{x,t}$ for each tier x can be written as

$$\pi_{x,t} = \prod_{v \in \Omega^*_{xt}} \left(\frac{p_{v,t}}{p_{v,t-2}} \right)^{\omega_{vt}} \left(\frac{\lambda_{x,t}}{\lambda_{x,t-2}} \right)^{\frac{1}{\sigma_x - 1}}. \tag{3.1}$$

Inflation is a composite of two components: the change in the prices of elements v of the set of common elements Ω^*_{xt}, which are sold in both t and $t-2$, and a correction reflecting the value of net product entry. The contribution of each common element is weighted by Sato (1976)-Vartia (1976) weights defined as:

$$\omega_{vt} = \frac{\frac{s_{vt} - s_{vt-1}}{ln\, s_{vt} - ln\, s_{vt-1}}}{\sum_{\Omega^*_{xt}} \frac{s_{vt} - s_{vt-1}}{ln\, s_{vt} - ln\, s_{vt-1}}}.$$

The term $\lambda_{x,t}$ is defined as the share of expenditures on common elements in the set Ω^*_{xt}, a subset of all products sold in the period Ω_{xt}:

$$\lambda_{x,t} = \frac{\sum_{v \in \Omega^*_{xt}} p_{v,t} c_{v,t}}{\sum_{v \in \Omega_{xt}} p_{v,t} c_{v,t}}.$$

Imagine that all else equal, a new product is introduced and consumed in period t that was not available in period $t-2$. In this case, $\lambda_{x,t} < 1$, and $\lambda_{x,t-2} = 1$. Because the share of common expenditures has fallen between the two periods, this lowers the price level. The degree to which the entry of new products lowers the price level is regulated by the elasticity of substitution σ_x between goods: when goods are more substitutable, entry has a smaller impact on the price level.

Because new products lower the price level, consumer welfare rises. Average annual implied welfare gains from new products under this model of consumer demand are displayed

Figure 3.2: Welfare Gains from New Products and City-Level Expenditure Growth

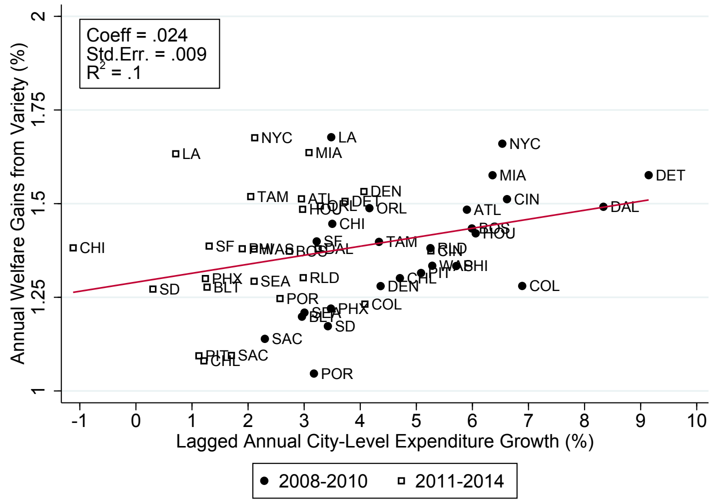

Notes: The figure plots welfare gains derived from a four-tier CES price index for the periods 2008-2010 and 2011-2014 against lagged city-level expenditure growth in the Nielsen Retail Scanner data in 26 US cities.

in Figure 3.2.[9] In total over the two periods, gains from 2008-2010 range from 3.2 to 5 percentage points across cities, while they range from 4.3-6.7 percent from 2011-2014. Once again, differences in welfare gains across cities within a given year are weakly explained by city-level expenditure growth.

3.2 Retailers Drive Product Entry

The weak relationship between the local business cycle and local product entry is surprising in light of the aggregate procyclicality of net product entry. This apparent contradiction is resolved if product entry decisions are driven by considerations at a more aggregate level. To understand what drives the entry of new products, I examine the entry pattern of products within a retailer. In principle, retailers may choose what products to offer at any level of disaggregation: from selecting an individual set of products for each physical store to choosing one set of products for every store across all markets. To explore this question, I study the behavior of products that enter or exit at least one market in a particular year.

[9] Appendix Figures C.1 and C.2 plots the relationship between welfare gains and GDP growth and housing price growth respectively.

Table I.3: Extensive Margin Entry and Exit Patterns by Number of Initial Markets

Markets 2009	Product Enters All Retailer's Markets (%)	Product Exits All Retailer's Markets (%)	Share of Value (%)
1	100	100	16
2-13	86	74	64
14-25	81	52	3
26	87	59	17
Overall	88	75	100

Notes: The 'Markets 2009' column shows the total number of cities in which the retailer sells within size bins. The entry and exit columns display the share of average quarterly revenue represented by products that enter all of the retailer's markets within a one-year period. The 'share of value' column displays the share of the overall value of entering or exiting products represented by retailers in each size bin.

To fix ideas, I focus on 2009, the year with the most striking extensive margin movement.[10]

In Table I.3, I report how much of new product entry by value within each retailer consists of products that are introduced in all of that retailer's markets within a one-year period. I divide retailers into size bins: while retailers that sell in just one city and national retailers each account for a significant fraction of new product entry, most new products are introduced in mid-sized retailers, often retailers with a regional presence. Because products may enter or exit in different quarters, shares of total entering or exiting value are calculated based on average quarterly revenue in the relevant time period. Across all retailer sizes, an average of 88 percent of new products are introduced simultaneously across all markets. While exiting products represent a smaller share of total value, I also look at product exit within retailers. Exit is less coordinated, but 70 percent of products that cease to be sold by multi-market retailers exit all cities within a one-year period.

The high percentage of total entering or exiting value that enters or exits all markets simultaneously in Table I.3 suggests that many retailers sell a similar set of products across all markets. This finding is particularly striking given the fact that the majority of retailers operate in a relatively small number of markets, suggesting that the cost of managing even a small number of markets individually is high. Another potential explanation might be that

[10]Results for all years 2008-2012 are similar. Entry/exit grids for the years 2009 and 2011 are displayed in Appendix Tables C.1 and C.2.

idiosyncratic regional business cycle shocks were simply insignificant in magnitude relative to the aggregate shock. However, local business cycles differed significantly over the period: the standard deviation of expenditure growth in the period 2008-2010 was 4.85 percentage points and 5.61 percentage points over the period 2011-2014. Later in the paper, I develop a model of retailer product selection that will help to distinguish these two cases.

It is possible that retailers introduce products simultaneously because the producer of the product itself has decided to sell its product in those markets. To address this possibility, I look at the fraction of retailers in each city in which new products are sold. I find that only about one third of retailers stock any given entering product, no matter how many markets the product is sold in. This remains true even among retailers that already stock another UPC from the same brand. While this is merely suggestive evidence, it indicates that retailers are able to select different product lines facing the same set of new products offered by producers.

Finally, to provide further support for the apparent finding that the retailer chooses its products globally rather than choosing a separate product line for each market, I decompose the growth in the entry of new products $\hat{n}_{ri,t}$ in a retailer-city at time t into common components coming from the retailer r in each year t, denoted $\hat{\alpha}_{rt}$, and from the city in each year, denoted $\hat{\beta}_{it}$ as follows

$$\hat{n}_{ri,t} = \hat{\alpha}_{rt} + \hat{\beta}_{it} + \epsilon_{rit}. \tag{3.2}$$

I report results year-by-year and pooled across all years 2008-2014 in Table I.4, as well as the F-statistic for the joint significance of each set of fixed effects and the associated p-value. The partial R^2 compares the share of the variation in new product entry that is explained by the full model reported in equation (3.2) with a model containing only one set of fixed effects. If the share of the variation explained by the model falls significantly when one set of fixed effects is omitted, that means that the partial R^2 associated with that set of fixed effects is high. Because there may be some covariance between retailer- and city-time fixed effects, the partial R^2 associated with both sets of fixed effects can be greater than one. Note

Table I.4: City-Retailer Level Product Entry Decomposition

	Retailer			City		
	Partial R^2	F-stat	P-val	Partial R^2	F-stat	P-val
2008	0.84	143.33	0.00	0.11	58.41	0.00
2009	0.88	206.16	0.00	0.12	77.73	0.00
2010	0.84	107.60	0.00	0.13	49.27	0.00
2011	0.89	207.36	0.00	0.15	74.48	0.00
2012	0.90	246.99	0.00	0.14	72.59	0.00
2013	0.90	249.11	0.00	0.17	66.13	0.00
2014	0.89	158.28	0.00	0.24	50.90	0.00
Mean	0.88	188.41	0.00	0.15	64.22	0.00
Pooled	0.75	16.11	0.00	0.13	5.90	0.00

Notes: The table reports statistics associated with regression equation (3.2) for each year, the mean across years, and a pooled regression. The pooled regression also includes a time-invariant city-retailer fixed effect.

that the pooled regression also includes time invariant city-year fixed effects. Both sets of fixed effects are highly statistically significant. Retailer fixed effects explain 88 percent of the variation in the entry of new products, while city fixed effects explain 15 percent. This finding suggests that coordinated introduction of new products by retailers explains most of the growth in product variety across cities.

3.3 Net Product Entry Co-moves with Retailer Revenue

The previous two findings suggest that gains from new products experienced in each city are better explained by retailer shocks than city-level shocks, but I have not yet explained what motivates a retailer to introduce new products. Now, I will show that retailers introduce new products in response to a positive shock to their total revenue. I begin by estimating a simple regression of new product entry at the retailer-city level on changes in the retailer's lagged total revenue across all markets. However, it is possible that this regression is contaminated by endogeneity: a retailer's revenue may increase because it has introduced new products. To address this concern, I construct a measure of plausibly exogenous demand growth at the retailer level by weighting city-level growth in total grocery expenditure, denoted $\hat{X}_{it}$, by

Table I.5: Net Product Entry and Retailer Revenue Growth

	OLS	IV
Revenue Growth$_{r,t-1}$	0.12	0.22
	(.01)	(.03)
Retailers	70	70
Observations	2,007	2,007
R^2	0.13	
F-stat		183
		(0.00)

Notes: The figure reports results of the regression equations (3.4) and (3.5) for each city-retailer-year.

the share of the retailer's revenue coming from that city, denoted ω_{rit}. The retailer demand shock $\hat{Z}_{rt}$ is given by

$$\hat{Z}_{rt} = \sum_i \omega_{rit} \hat{X}_{it}. \tag{3.3}$$

I estimate the relationship between net product entry at the city-retailer level and changes in retailer demand coming from shocks to total city-level demand via two-stage least squares as follows:

$$\hat{X}_{rt} = \alpha_0 + \alpha_1 \hat{Z}_{rt} + u_{rt} \tag{3.4}$$

$$\hat{n}_{rit} = \beta_0 + \beta_1 \hat{X}_{rt} + \epsilon_{rit}, \tag{3.5}$$

where $\hat{X}_{rt}$ denotes the growth of the retailer's total revenue across all cities. Table I.5 reports the results of the OLS regression of equation (I.5) and 2SLS regression of equations (3.4) and (3.5). The F-statistic of 183 suggests that the weighted city demand growth measure constructed in equation (3.3) satisfies the instrument relevance requirement. The IV estimate is highly statistically significant and suggests that a 10 percent increase in revenue growth due to a demand shock is associated with 2.2 percent net product entry. The average net product entry across retailer-city-years is 5.1 percent in the sample.

The finding that retailers drive differences in product entry across cities in response to changes in their demand motivates the model of retailer product choice that follows. Because net product entry in each retailer-city responds to total revenue, it suggests a mechanism through which demand shocks in one city can impact consumer welfare in another. Through the lens of the model I develop below, I am able to quantify the extent of inter-city demand shock transmission during the period covered by the data.

4 A Model of Retailer Product Choice

This section develops a model that can be used to demonstrate how retailer networks affect the propagation of local demand shocks across cities. The model environment consists of a set I of cities, indexed by i, with units of labor L_i supplied inelastically and immobile across cities. There are R retailers in the economy, with a set R_i active in market i, intended to capture the uneven distribution of retailers across cities in the data. The sets of active firms in each city are given, but the share of each retailer is determined by relative prices and city-specific retailer tastes.

Households choose consumption based on a three-tiered utility function. Grocery retailers sell a retailer-specific grocery bundle. They choose the length of their product lines subject to an increasing stocking cost and set a markup over the producer price of the grocery good. The model is static. Business cycle analysis compares steady states under different levels of city productivity a_i.

4.1 Preferences

Consumers in each city i derive utility from a three-tier utility function. The top tier expresses how consumers aggregate grocery consumption C_i and non-grocery consumption Y_i. The final consumption good is a Cobb-Douglas function of differentiated final grocery consumption and homogeneous non-grocery consumption:

$$U_i = C_i^{\beta} Y_i^{1-\beta}. \tag{4.1}$$

Final grocery consumption C_i is a two-tier CES aggregate that allows the elasticity of substitution to differ across retailers and across goods. The second tier describes preferences over the grocery consumption bundles available at each retailer r. The elasticity of substitution between retailers is given by σ_r. A taste shifter γ_{ri} allows for the possibility that retailers have characteristics beyond their pricing that affect consumer demand:

$$C_i = (\sum_{r \in R_i} \gamma_{ri}^{\frac{1}{\sigma_r}} C_{ri}^{\frac{\sigma_r - 1}{\sigma_r}})^{\frac{\sigma_r}{\sigma_r - 1}}. \tag{4.2}$$

Because the set of retailers R_i remains constant across the sample, there are no direct gains from variety stemming from the entry of new retailers. However, if prices across retailers vary as the result of differences in goods prices or the number of products available, σ_r determines how much consumers substitute toward the lower-cost store bundle.

Finally, the consumption aggregate available at each retailer, C_{ri}, is an aggregate over the set of varieties supplied by the retailer:

$$C_{ri} = (\int_{v \in \Omega_{ri}} C_{ri}(v)^{\frac{\sigma_g - 1}{\sigma_g}} dv)^{\frac{\sigma_g}{\sigma_g - 1}}. \tag{4.3}$$

Denote the measure of the set Ω_{ri} as n_{ri}. Because retailers choose how many products to offer, the within-retailer, across-good elasticity σ_g directly affects the number of varieties in a group. For the sake of tractability, I assume that the across-good elasticity is the same across all grocery goods.

4.2 Firms

4.2.1 Grocery Retailers

Grocery retailers r choose how many product varieties to sell and choose a markup μ_{ri} to set on all products in order to maximize their total profits subject to a stocking cost $F(n_r)$ with $F'(n_r) > 0$. For simplicity, I assume that the average marginal cost of goods sold at the retailer $\bar{mc}$ is unaffected by the number of products chosen and is common to all retailers. Retailers face total grocery expenditures X_i and a price level P_i in each city. The retailer's problem can be expressed as

$$\max_{\mu_{ri}, n_r} \sum_{i=1}^{I} \gamma_{ri} \left(1 - \frac{1}{\mu_{ri}}\right) \left(\frac{n_{ri}^{\frac{1}{1-\sigma_g}} \mu_{ri} \bar{mc}}{P_i}\right)^{1-\sigma_r} X_i - F(n_r) \tag{4.4}$$

The retailer's price level is given by

$$P_{ri} = n_{ri}^{\frac{1}{1-\sigma_g}} \mu_{ri} \bar{mc} \tag{4.5}$$

and P_i denotes the city grocery price level, defined as

$$P_i = \left(\sum_{r \in R_i} \gamma_{ri} P_{ri}^{1-\sigma_r}\right)^{\frac{1}{1-\sigma_r}}. \tag{4.6}$$

For all cities i such that $r \notin R_i$, the taste parameter γ_{ri} is equal to zero. Fixed costs are paid nationally based on the total product line across all markets $n_r = \max\{n_{ri}\}_{i \in I}$. Because the fixed cost is paid nationally, there is no motive for the retailer to choose an individual product line n_{ri} for each market. However, I do allow for a relationship of the form $n_{ri} = \tau_i n_r$ between the global product line and the set of varieties ultimately available in city i. The term τ_i can be interpreted as an additional cost associated with providing varieties to city i if $\tau_i < 1$.

Practically, this term allows the model to rationalize the fact that New York and Los Angeles experience slightly higher gains from variety than other cities even within retailers, though this is unrelated to changes in demand in those two cities. Setting $\tau_i = 1 \forall i$ will yield the same predictions quantitatively and qualitatively for every city but these two.

The first order condition with respect to the choice of μ_{ri} is standard:

$$d\mu_{ri} : \gamma_{ri} \left(\frac{1}{\mu_{ri}^2}\right) \left(\frac{n_{ri}^{\frac{1}{1-\sigma_g}} \mu_r \bar{mc}}{P_i}\right)^{1-\sigma_r} X_i + (1-\sigma_r)\left(1 - \frac{1}{\mu_{ri}}\right) s_{ri}(1-s_{ri}) \frac{X_i}{\mu_{ri}} = 0 \tag{4.7}$$

where $s_{ri} = \gamma_{ri} \left(\frac{n_{ri}^{\frac{1}{1-\sigma_g}} \mu_{ri} \bar{mc}}{P_i}\right)^{1-\sigma_r}$. The firm faces a trade-off between profit-per-unit and its share of the market: increasing the markup increases flow profits per unit, but tends to decrease the firm's overall share. I assume that firms are small enough that $s_{ri}^2 \approx 0$, i.e. that firms set the monopolistically competitive markup rather than a variable markup based on retailer share.

The first order condition with respect to the choice of n_r is

$$dn_r : \sum_{i=1}^{I} \left[\frac{1-\sigma_r}{1-\sigma_g}\left(1 - \frac{1}{\mu_{ri}}\right) s_{ri} \frac{X_i}{n_r} - \frac{1-\sigma_r}{1-\sigma_g}\left(1 - \frac{1}{\mu_{ri}}\right) s_{ri}^2 \frac{X_i}{n_r}\right] - F'(n_r) = 0. \tag{4.8}$$

As noted, increasing the length of the product line increases the retailer's share of demand because the consumer has love for variety: the same utility of consumption is less expensive at a retailer whose grocery bundle includes more varieties. I assume that $F'(0) = 0$ so that all retailers sell products.

Solving these two first order conditions yields the standard optimal gross markup

$$\mu_{ri} = \frac{1}{\sigma_r - 1} + 1 \tag{4.9}$$

The optimal product line is given by

$$n_r = \frac{1}{(\sigma_g - 1)\mu_{ri}} \frac{\sum_i s_{ri} X_i}{F'(n_r)}. \tag{4.10}$$

Higher marginal stocking costs and greater substitutability between goods both lead to shorter product lines. All else equal, a retailer with larger global sales will have a longer product line. However, if the degree of substitutability across retailers is low, μ_{ri} will be high and the product line will be shorter.

4.2.2 Goods Producers

Both grocery goods C and non-grocery goods Y are produced by perfectly competitive firms with technology $a_i L_i$. The price of the non-grocery consumption good is the numeraire, and the good is freely traded. Grocery goods are also freely traded and produced with a common productivity a_i in each market, giving a marginal cost $\bar{mc}_i = \frac{w_i}{a_i}$.

4.2.3 Retailer Profits

There are aggregate profits in the economy because free entry does not hold in the retail sector. They depend on the size of the retailer and on the stocking cost $F(n_r)$. Given equations (4.9) and (4.10), the general expression for the retailer's profits π_r is

$$\pi_r = \frac{X_r}{\sigma_r} - F(n_r) \tag{4.11}$$

Let the stocking cost be given by

$$F(n_r) = \left(\frac{n_r}{n_c}\right)^{\alpha}. \tag{4.12}$$

The term in the denominator n_c will allow for trend growth in the number of varieties and accommodate differences in trend growth across the four categories to which food retailers in the data belong: convenience stores, drug stores, grocery stores and mass retailers. The length of the product line under this stocking cost parametrization is given by

$$n_r = \left[\frac{\sum_i s_{ri} X_i}{\alpha(\sigma_g - 1)\mu_{ri}} \right]^{\frac{1}{\alpha}} n_c \tag{4.13}$$

If $\alpha > 1$, the stocking cost is strictly convex in the length of the product line n_r.

Profits can be expressed as a function of the firm's sales, its markup, and parameters:

$$\pi_r = \left[\frac{1}{\sigma_r} \left(1 - \frac{\sigma_r - 1}{\alpha(\sigma_g - 1)} \right) \right] \sum_{i=1}^{I} s_{ri} X_i \tag{4.14}$$

Note that weakly positive profits for all firms requires that $\alpha > \frac{\sigma_r - 1}{\sigma_g - 1}$. Since $\sigma_r \leq \sigma_g$, this condition only requires that the stocking cost not exhibit large increasing returns to scale in the size of the product line. It is always satisfied if the marginal cost of adding a product is constant or increasing.

4.3 Market Clearing

Goods and labor markets clear. Household expenditure X_i is divided across grocery and non-grocery goods according to their Cobb-Douglas utility shares β and $1 - \beta$. Goods market clearing in each city implies that

$$Y_i = (1 - \beta) X_i \tag{4.15}$$

and

$$P_i C_i = \mu X_i \tag{4.16}$$

Households supply labor inelastically. Labor demand associated with the production of non-grocery and grocery goods (the price of non-grocery goods is the numeraire) is given by

$$L_i^Y = \frac{Y_i}{a_i} \tag{4.17}$$

and

$$L_i^G = \frac{C_i}{a_i}. \tag{4.18}$$

Finally, there is labor demand associated with the stocking costs paid by retailers (not associated with one particular market):

$$L^F = \frac{\sum_r F(n_r)}{a_i} \tag{4.19}$$

Goods are freely tradable, so labor market clearing requires only that all labor be employed at a local wage $w_i = a_i$. Because the marginal value product of labor is equalized across cities, the location of production for each good type is indeterminate. Total labor market demand across all cities must equal total labor supply:

$$\sum_i \left(L_i^Y + L_i^G \right) + L^F = \sum_i L_i. \tag{4.20}$$

I assume that the share of total profits rebated to consumers in city i are both given by

s_i^π. For simplicity, I assume that s_i^π is proportional to the city's share of labor income:

$$s_i^\pi = \frac{w_i L_i}{\sum_j w_j L_j}. \tag{4.21}$$

As a result of this assumption, total expenditure in city i is a constant multiple of labor income across cities

$$X_i = \left(1 + \frac{\sum_r \pi_r}{\sum_j w_j L_j}\right) w_i L_i. \tag{4.22}$$

4.4 Equilibrium

An equilibrium in this economy is a set of retailer product lines and bundle prices $\{n_{ri}, P_{ri}\}_{i\in I, r\in R}$, city-level grocery and factor prices $\{P_i,\, w_i\}_{i\in I}$, local labor allocations $\{L_i^Y,\, L_i^G, L_i^F\}_{i\in I}$, local goods production $\{C_i, Y_i\}_{i\in I}$, and total profits per retailer $\{\pi_r\}_{r\in R}$. Consumers maximize utility over consumption of grocery and non-grocery goods and allocate grocery consumption across retailers as specified in equations (4.1) and (4.2). Retailers maximize profits choosing product lines and a markup according to equations (4.7) and (4.8).

The retailer-city price level and the overall grocery price for each city are defined in equations (4.5) and (4.6). Local output and labor markets satisfy equations (4.15), (4.18) and (4.20).

4.5 Business Cycle Interpretation

In order to understand how the decisions of retailers can transmit productivity shocks across cities, I compare equilibria under a set of city-level productivities $a_{i,t-1}$ and $a_{i,t}$. It is convenient to compare steady states in terms of log changes, where $\hat{x} = dlogx$. In this analysis, I assume that tastes are constant over time: $\hat{\gamma}_{ri} = 0$.

I focus on the impact of city-level shocks $\{\hat{a}_i\}_{i\in I}$ on the length of an arbitrary retailer's product line and markup and therefore on consumer welfare. First, log-linearizing equation

(4.13) gives an expression for the change in product line length:

$$\hat{n}_r = \frac{1}{\alpha}\hat{X}_r + \hat{n}_c, \tag{4.23}$$

where the second term $\hat{n}_c$ accommodates trend growth in the number of products for each retailer category.

Transforming equation (4.5), the change in the retailer's global price is given by

$$\hat{P}_r = \frac{1}{1-\sigma_g}\hat{n}_r + \hat{mc}. \tag{4.24}$$

The growth in the price set by each retailer in each city incorporates the cost term τ_i, which enters through $\hat{n}_{ri}$:

$$\hat{P}_r = \frac{1}{1-\sigma_g}\hat{n}_{ri} + \hat{mc}. \tag{4.25}$$

Note that because $w_i = a_i$ in every city, $\hat{mc} = 0$ in general. Combining these expressions, the change in the retailer's price can be expressed as a function of parameters and its own demand:

$$\hat{P}_{ri} = \frac{1}{1-\sigma_g}\frac{1}{\alpha}\hat{X}_r + \frac{1}{1-\sigma_g}(\hat{n}_c + \hat{\tau}_i) \tag{4.26}$$

Log-linearizing the city grocery price level P_i in equation (4.6) gives

$$\hat{P}_i = \sum_r s_{ri}\hat{P}_{ri}. \tag{4.27}$$

4.6 City-Level Contributions to Retailer Variety

In order to understand how shocks to a_i may be transmitted to other cities through changes in the set of available products, I decompose the change in the retailer price $\hat{P}_r$ into contributions coming from each city in which the retailer operates. I denote the impact of demand in city j on retailer r's price level by $\hat{T}_{rj}$. The full impact of demand in city j on city i's price level is a weighted sum of each contribution $\hat{T}_{rj}$, where the weights are the share of retailer r in city i's expenditure. I describe the derivation of the expression for the impact of city j on city i's demand, denoted $\hat{T}_{ij}$, in what follows.

Equations (4.23) and (4.24) describes the relationship between changes in total retailer revenue $\hat{X}_r$, changes in the length of the product line $\hat{n}_r$ and changes in the global price of the retailer $\hat{P}_r$. I begin by decomposing this relationship into a contribution coming from each city in which the retailer operates. The share ω_{ri} of retailer r's revenue derived from each city i is

$$\omega_{ri} = \frac{s_{ri} X_i}{\sum_i s_{ri} X_i}. \tag{4.28}$$

Retailer revenue growth can be expressed as the inner product of retailer-city revenue shares and city expenditure growth:

$$\hat{X}_r = \sum_i \omega_{ri} \hat{X}_i. \tag{4.29}$$

Combining equations (4.23) and (4.24), the retailer price level is:[11]

$$\hat{P}_r = \frac{1}{1-\sigma_g}\frac{1}{\alpha}\hat{X}_r + \frac{1}{1-\sigma_g}\hat{n}_c \tag{4.30}$$

I use equation (4.29) to decompose equation (4.30) into the contribution of each city j to the change in retailer r's price, denoting this contribution by $\hat{T}_{rj}$:

$$\hat{P}_r = \sum_j \omega_{ri}\hat{T}_{rj} \tag{4.31}$$

where the contribution of city j to the change in the price level of retailer r is given by

$$\hat{T}_{rj} = \frac{1}{1-\sigma_g}\frac{1}{\alpha}\hat{X}_j + \frac{1}{1-\sigma_g}\hat{n}_c. \tag{4.32}$$

The contribution of demand growth in city j to the price level in city i can be expressed as the sum of city j's contributions $\hat{T}_{rj}$ to retailers in set R_{ij}, weighted by the share s_{ri} of each retailer in city i's demand. I denote the total contribution of city j to city i's price level by $\hat{T}_{ij}$. Combining equation (4.27) with equations (4.31) and (4.32), the change in the price level in city i due to contributions to each retailer's price from city j is given by

$$\hat{T}_{ij} = \sum_{r\in R_{ij}} s_{ri}\omega_{rj}\hat{T}_{rj}. \tag{4.33}$$

Finally, combining equation (4.32) with equation (4.33), the connection between demand in city j and city i's price level is given by:[12]

[11] Note that this is almost equivalent to equation (4.26), omitting city-level differences in the retailer's price due to the term τ_i, which is unrelated to the retailer's choice of product line.

[12] Because $\sum_j \omega_{rj} = 1$, the contribution of each city to trend growth $\hat{n}_c$can be divided proportionally across cities.

$$\hat{T}_{ij} = \frac{1}{1-\sigma_g} \sum_{r \in R_{ij}} s_{ri} \omega_{rj} \left(\frac{1}{\alpha} \hat{X}_j + \hat{n}_c \right) \tag{4.34}$$

Equation (4.34) expresses an intuitive relationship between demand in city pairs. City j has a larger impact on price level changes in city i whenever common retailers R_{ij} represent a large fraction of consumption in city i (s_{ri} is large), these retailers derive a significant fraction of their revenue from city j (ω_{rj} is large), or demand shocks in city j are particularly significant ($\hat{X}_j$ is large).

5 Calibrated Model

The model links demand shocks in each city to consumer welfare gains from new products through the shares $\{s_{ri}, \omega_{ri}\}$, two key parameters α and σ_g, the cost parameter τ_i and the trend growth in varieties $\hat{n}_c$ specific to each category of food retailer. In this section, I describe how I choose each parameter.

5.1 Retailer-City Shares

The shares of each city-retailer in city-level expenditure, s_{ri}, and retailer revenue shares, ω_{ri}, are taken from data. I begin by characterising the relationship between cities implied by these shares. As equation (4.34) suggests, city j has a larger impact on city i's price level when retailers in the set R_{ij} make up a large share of city i's expenditure and when city j is a significant source of revenue for those retailers. The 'common retailer share' of city j in city i can be expressed as the sum of these shares, denoted $S_{ij} = \sum_{r \in R_{ij}} s_{ri} \omega_{rj}$, reflecting the share of common retailers in city i's demand. Denote $S_i = \sum_{j \neq i} S_{ij}$, the common share between city i and all other cities.

Table I.6 characterises both the bilateral common retailer shares S_{ij} as well as the overall

Table I.6: Common Retailer Shares: City Pairs and Total

	City Pairs S_{ij}		All Other Cities S_i
	Share		Share
Maximum		Maximum	
San Diego-LA	63	Columbus	91
Pittsburgh-Cleveland	37	Baltimore	88
Minimum		Minimum	
Boston-Louisville	<0.1	Chicago	25
Seattle-Louisville	<0.1	Boston	18
Mean	2	Mean	63

Notes: The table reports extremes of the distribution of city-level common retailer shares S_{ij} in the first column and the sum over all other cities j in the second. Note that $S_{ij} \neq S_{ji}$.

common retailer share between each city and all others, S_i. The first column characterises common retailer shares across city pairs. Not surprisingly, small cities tend to have large common retailer shares with large neighbors: L.A.'s share of revenue in retailers available to consumers in San Diego is 63 percent. Bilateral ties between small cities in distant regions are virtually non-existent. While the fact that the bilateral share between San Diego and Boston is small is intuitive, it is somewhat surprising that this bilateral connection is so close to zero given the presence of several national retailers in the data. It appears that expenditures at national retailers are not large enough to create strong bilateral connections: the average bilateral connection is 2 percent.

The second column characterizes the common retailer share between city i and all other cities. Consumers in small cities, particularly those geographically close to major metros, conduct most grocery purchases at retailers they share with other cities. At the highest extreme, the total common retailer share between Columbus and all other cities is 91 percent. Similarly large linkages can be seen in Baltimore. Chicago and Boston have strikingly low common retailer shares: these markets are large enough and isolated enough that outside cities contribute little to the revenue of retailers operating in these two cities. Boston is the most extreme, with a total common retailer share of just 18 percent.

Looking only at the extremes of the bilateral share distribution may give the impression that geography explains the degree to which two cities are connected. However, while some logical geographic patterns exist, strong common retailer shares are not restricted to cities in the same geographic region. I define regions based on Census divisions.[13] Table I.7 describes the share of all linkages accounted for by cities within and outside city i's region according the expression

$$s_{i,D_x} = \sum_{j \in D_x} \sum_{r} s_{ri} \omega_{rj}.$$

The common share of each city with other cities in the same region is generally high. However, the table also reports the region other than city i's own that accounts for the largest share of shock transmission. For example, cities in the West Coast region have their highest common share with other cities in the region, at 80 percent. Their highest outside common retailer share with another region is 7 percent with the Mountain region. Similarly, the Northeast has a high within-region common retailer share at 73 percent. Its largest common retailer share with another region is 14 percent with the Midwest. Other regions feature stronger links with cities in other regions. Texas accounts for just 54 percent of its own retailers' revenue with strong connections to the Midwest, and the common retailer share of Texas cities in the Midwest is similarly large.

5.2 Parameters

Next, I describe how the values of the fixed cost parameter α from equation (4.12), elasticity of substitution across goods σ_g, cost parameter τ_i and trend growth rate of product variety$\hat{n}_c$ are chosen.

[13] I combine the New England and Mid Atlantic divisions and separate Tampa, Miami and Orlando from the rest of the South Atlantic division. For clarity, I refer to the South Central division as Texas because the cities representent in the data are Dallas and Houston.

Table I.7: Common Retailer Shares by Region

Region i	s_{i,D_i}	Region j	s_{i,D_j}
West Coast	80	Mountain	7
Northeast	73	Midwest	14
Mountain	71	West Coast	18
Florida	64	Northeast	13
South	62	Northeast	9
Texas	54	Midwest	22
Midwest	47	Texas	25

Notes: The table reports the share of all common retailer links accounted for by other cities in the region (first two columns) and by cities in the outside region with the largest common retailer share (last two columns).

Equation (4.23) demonstrates that the shape parameter of the cost function, α, also governs the elasticity of growth in the product line to growth in the retailer's demand. In practice, changes in the value of new products offered by each retailer are better predicted by one year lagged than contemporaneous expenditure growth. As in section 3.3, I use the weighted average growth in a retailer's markets as an instrument for exogenous growth in the retailer's revenue:

$$\hat{Z}_{rt} = \sum_i \omega_{rit} \hat{X}_{it} \tag{5.1}$$

$$\hat{X}_{r,t-1} = \gamma_0 + \gamma_1 \hat{Z}_{r,t-1} + u_{ri,t} \tag{5.2}$$

$$\hat{n}_{ri,t} = \beta_1 \hat{X}_{r,t-1} + \beta_2 (1 - LANYC_i) + \Gamma_c + \epsilon_{ri,t} \tag{5.3}$$

The coefficients β_1 can be interpreted as $\beta_1 = \frac{1}{\alpha}$. Note that I allow for a set of category-specific trend growth rates of product variety Γ_c corresponding to the parameter $\hat{n}_c$. I estimate an elasticity of new products to retailer revenue of 0.16, implying a stocking cost

Table I.8: Parameter Calibration Regression

	$\hat{n}_{ri,t}$	Associated Parameter	Value
β_1	.155	α	6.46
	(.011)		(.45)
β_2	-.007	$\hat{\tau}$	-.007
	(.001)		(.001)
Number of Obs	2,007		
Retailers	70		
R^2	0.39		

Note: The table reports the results of regression equation (5.3), which is used to calibrate parameters α and $\hat{\tau}$.

shape parameter $\alpha = 6.46$. Because this elasticity is less than one, the retailer's marginal cost of supplying additional products $F(n_r)$ increases with the number of products offered.

I allow for a difference in the trend growth of new products between the two largest cities in the sample, New York and Los Angeles, and all other cities, given by the indicator variable $LANYC_i$. This trend difference $\hat{\tau}$ implies that all other cities experience new product entry that is 0.7 percentage points lower. Importantly, this is not driven by particularly high demand in New York and Los Angeles, but rather by the fact that some retailers appear to offer a slightly longer product line in large cities. Note that this is one parameter, not a trend adjustment specific to each city.

To avoid the bias associated with macroelasticities highlighted by Imbs and Mejean (2015), I choose an elasticity to match the average welfare gains calculated using the microelasticities in Appendix Table A.1 rather than calculating a macro elasticity pooling all goods. To do so, I regress the model-predicted change in product variety in each city, $\hat{n}_{i,t}$, on the microestimates reported in Section 3.1. I estimate a value of σ_g of 4.09. The elasticity of substitution σ_g governs the welfare impact of new products. Parameter values are summarized in Table I.9.

To evaluate the model's predictions, I use the calibration of the parameters in equations (4.26) and (4.27) to predict city-level welfare gains based on demand growth in each city. I compare the model's predictions to the welfare gains calculated in section 3.1 using microdata

Table I.9: Parameter Values in Baseline Calibration

		Value
Expenditure and Revenue Shares	$\{s_{ri}, \omega_{ri}\}$	Table I.6
City Trend Growth Adjustment	$\hat{\tau}$	-.007 (.001)
Goods Elasticity of Substitution	σ_g	4.09 (.03)
Stocking Cost Shape Parameter	α	6.46(.45)

Notes: The table summarizes the parameter values used in the calibration. The distribution of shares, which implicitly defines γ_{ri}, was discussed previously.

Figure 5.1: Baseline Model Predicted Welfare Gains vs. Microdata-Based Welfare Gains

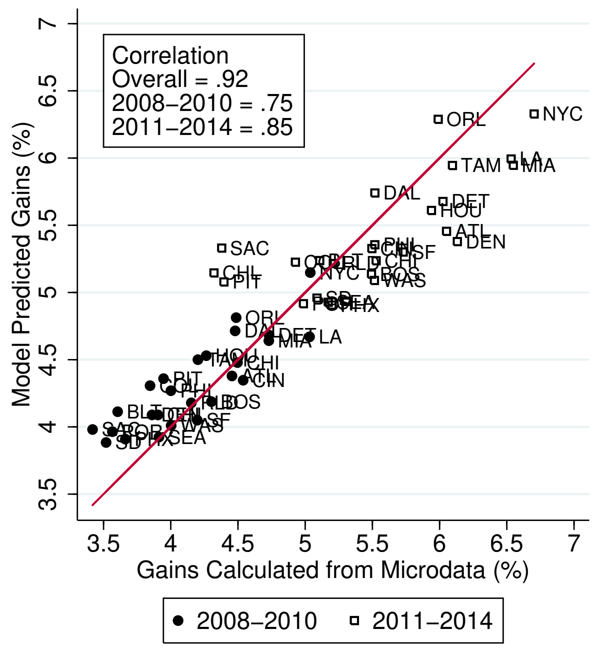

Notes: The figure compares the prediction of the calibrated model to the welfare gains estimated using UPC-level microdata for each retailer.

at the UPC level. The predictions of the model are plotted against the gains based on micro data in Figure 5.1. Within the period 2008-2010, the correlation between the baseline model and the microdata estimate is 0.75, while within the period 2011-2014 it is 0.85. Overall, the correlation between the baseline model and the microdata estimates is 0.92.

6 Decomposing Welfare Gains into City-level Contributions

6.1 City-level Shocks in the Baseline Model

The model implies that demand shocks in other cities impact consumer welfare through the product selection decisions of common retailers. I return to the decomposition in equations (4.29) and (4.32) and calculate the implied contribution of each city to consumer welfare in other cities in the baseline model. The contribution of city j to city i 's price level can be expressed as

$$\hat{T}_{ij} = \frac{1}{1-\sigma_g} \sum_{r \in R_{ij}} s_{ri} \omega_{rj} \left(\frac{1}{\alpha} \hat{X}_j + \hat{n}_c \right) \tag{6.1}$$

Table I.10 reports the results for the average city as well as two particular examples: other cities account for 63 percent of welfare gains in the average city (this share exceeds 50 percent in part because the largest cities in the sample are much larger than the average). The average city receives welfare gains of 1.5 percentage points from its own demand shocks between 2008-2010 and 1.8 percentage points from 2011-2014. Welfare gains derived from outside shocks are larger at 2.8 and 3.6 percentage points respectively. The table also reports the contribution of own and outside demand shocks for two cities on opposite extremes: Boston and Columbus. Not surprisingly, Boston, which has very low connections to other cities through retailers, inherits little of its welfare gains from other cities. Columbus, which is highly integrated with other Midwestern cities, derives almost all its welfare gains from other cities.

Next, I ask whether welfare shock transmission results from intra-regional shocks: in other words, whether understanding regional outcomes is sufficient to predict city-level welfare. Table I.11 reports the share of welfare gains coming from cities in the same region and the

Table I.10: Contribution of Own and Outside Shocks to Welfare Gains

		2008-2010		2011-2014	
	Outside Share	Own Shocks	Outside Shocks	Own Shocks	Outside Shocks
Average City	0.63	1.5	2.8	1.8	3.6
Boston	0.09	3.3	0.9	3.9	1.2
Columbus	0.91	0.4	4.0	0.5	4.7

Table I.11: Contribution of Intra- and Cross-Regional Shocks to Welfare Gains

Region	Gains	Largest Outside Region	Gains
West Coast	75	Mountain	7
Northeast	70	Midwest	14
Mountain	66	West Coast	18
South	58	Northeast	10
Florida	54	Northeast	15
Texas	49	Midwest	22
Midwest	44	Texas	24

Notes: The table records the average across cities in each region of the share of welfare gains attributable to cities within and outside its own region.

outside region that makes the largest contribution. While regional shocks account for a substantial share of welfare gains, consumers in Florida, the South, Texas and the Midwest derive a substantial fraction of their gains from new products from demand shocks in other regions.

Large welfare gains from outside shocks can be generated two ways: if one city has an unusually large shock, it transmits some welfare gains (or losses) even to cities with which it has few common retailers. On the other hand, shock transmission may be the result not of a significantly large shock in some locations, but of large connections between cities through common retailers that account for a significant share of demand. I decompose the relationship further to explore whether the large contribution of outside cities is attributable to large shocks or to large retailer connections.

Define the total common retailer share between city i and city j as $S_{ij} = \sum_{R_{ij}} s_{ri}\omega_{ri}$. The contribution of city j to city i's price level can be expressed as the impact of variation in

Table I.12: Contribution of Shocks and Shares to Demand Shock Transmission

	T_{ij}	Shocks	Shares	Covariance
Partial R^2		0.04	0.86	0.11
Standard Deviation (pp)	0.67	0.17	0.64	0.23

Notes: The table records the share of variation orthogonal to other variables that can be explained by each component of equation (I.12).

demand shocks X_j across locations, the impact of variation in shares S_{ij}, and an expression closely related to the covariance between these two:

$$T_{ij} = \frac{1}{1-\sigma_g}\frac{1}{\alpha}\left(\underbrace{\bar{S}_i X_j}_{Shocks} + \underbrace{S_{ij}\bar{X}}_{Shares} - \underbrace{\bar{S}_{ij}\bar{X}}_{Avg} + \underbrace{(S_{ij}-\bar{S}_i)(X_j-\bar{X})}_{Covariance}\right). \tag{6.2}$$

The average $\bar{S}_i$ is defined at the city level to accommodate the fact that cities vary significantly in their degree of integration:

$$\bar{S}_i = \frac{1}{J}\sum_{j\neq i} S_{ij}. \tag{6.3}$$

The average $\bar{X}$ is the average demand shock across cities

$$\bar{X} = \frac{1}{J}\sum_{j} \hat{X}_j. \tag{6.4}$$

The results of the decomposition are reported in Table I.12. Differences in shares orthogonal to other factors account for 86 percent of the variation in total welfare gains from outside cities, while the covariance term accounts for 11 percent of the variation. Variation in shocks accounts for just 3 percent of the variation. The intuition is clear after looking at the amount of variation in the three terms: the standard deviations of the share and covariance terms are significantly greater than the standard deviation of the shock term.

The results indicate the importance of inter-city common retailer shares in explaining transmission of shocks. Though demand shocks range from -2.4 to 27 percent, the average inter-city share is just 2 percent. With such a low common retailer share between i and j, there will be relatively low transmission even of a large demand shock. Shock transmission occurs primarily between the city pairs that have larger common retailer shares.

6.2 Counterfactuals

For a given business cycle, a city may benefit or suffer from its connection with other cities. If a city receives a better shock than its neighbor, their common retailers offer fewer products than had both cities experienced the better shock. The converse is also true. However, the transmission of shocks across cities unambiguously decreases the variation in welfare gains if idiosyncratic business cycle shocks are random across cities.

To understand how much smoothing of welfare gains from new variety is generated by existing multi-city retailers, I compare the predictions of the baseline model to the predicted welfare gains under three counterfactual retailer distributions. At one extreme, I conduct a 'local' counterfactual in which all retailers are local: i.e., setting $\omega_{ri} = 1$. In this counterfactual, new product entry will only respond to local demand shocks. At the other extreme, I consider a counterfactual in which retailers are all national, i.e., $\omega_{ri} = \frac{X_i}{\sum_{j \in I} X_j} \ \forall r$. In this case, even though local shocks vary, this city-level variation will be smoothed out within each retailer, and new product entry will respond only to aggregate demand. Finally, I consider a counterfactual in which retailers operate in only in one region, i.e. $\omega_{ri} = \frac{X_i}{\sum_{j \in D_i} X_j}$ if $r \in R_{D_i}$. The results reported in Table I.11 suggest that the welfare gains experienced by cities in most regions are primarily driven by demand growth within the same region, although there are significant retailer connections between some cities that are geographically distant.

Table I.13 reports the standard deviation of demand growth and welfare gains across cities in the periods 2008-2010 and 2011-2014. Demand growth was more variable in the recovery period, so the standard deviation of welfare is higher in this period. Because retailers in all cities are exposed to the same demand shock under the national counterfactual, the

Table I.13: Baseline and Counterfactual Models: Standard Deviation of Welfare Gains

	(1) $\hat{X}_i$	(2) Baseline	(3) Local Variety	(4) Regional Variety	(5) National Variety
2008-2010	4.85	0.17	0.24	0.18	0.00
2011-2014	5.61	0.19	0.28	0.16	0.00

Notes: The table reports the standard deviation of demand shocks and the standard deviation of business cycle product entry in the baseline model, local, regional and national counterfactuals.

standard deviation of business cycle welfare gains is zero in this scenario. In both periods, the standard deviation of welfare gains from demand shocks across cities under the baseline distribution of retailers is about 68 percent of the standard deviations when retailers are local. While the existence of multi-city retailers appears to result in smoother welfare gains across cities, the baseline model suggests that the fact that most retailers are regional rather than national results in much more variable welfare gains from product variety than would result from national retailers. Indeed, the standard deviation of welfare gains in the regional counterfactual is quite similar to that of the baseline model.

7 Conclusion

This paper documents a novel channel through which business cycle shocks are transferred across locations. While previous work has shown that booms are associated with entry of new products at the national level, entry of new products across cities is not well-explained by local business cycles. If consumers value variety, entry of new products effectively lowers the consumer price level, raising the real wage.

The welfare implications of local changes in product variety are relatively large: from 2008-2010, the average city experienced welfare gains ranging from 3.9 to 5.7 percentage points, and they experienced gains of 4.7 to 6.7 percentage points in 2011-2014. I present evidence that the welfare gains associate with net product entry are driven by retailers who operate in multiple markets and coordinate product entry across markets. The role of retailers in explaining product entry motivates a model in which multi-city retailers choose how many products to sell to maximize their profits subject to an increasing cost of extending the product line.

I use the model to decompose the welfare gains in each city into a contribution coming from each other city in the data. The model implies that consumers in the average city derive 63 percent of welfare gains from demand in other cities. The transmission of shocks acts as a mechanism through which the idiosyncratic component of business cycle shocks may be 'smoothed' across cities. However, comparison of the level of inter-city variation in welfare gains under the basline and counterfactuals in which variety is determined purely locally or at the aggregate level suggest that the predominance of regional rather than national retailers results in substantially higher variation in consumer welfare than under full product market integration.

Chapter II: New Exporter Growth and Sequential Entry

1 Introduction

New exporters start out small and grow to meet average total export sales over a multi-year period. The period of growth suggests that firms face frictions that prevent them from reaching their full potential export sales immediately. Previous work such as Ruhl and Willis (2017) has represented the growth of new exporters as a falling marginal cost of serving a foreign market, either the result of a random process or investment decisions by the exporting firm. Because the friction affects the intensive margin of trade in these models, it depresses export sales for all new exporters, regardless of their productivity.

I show that much of new exporter growth is actually explained by sequential entry into new foreign markets, and that as the result of this market entry strategy, policies which aim to help new exporters grow should focus on helping firms overcome entry costs. Moreover, I show that the immediate bilateral benefit of reducing entry costs is significantly lower than the full benefit, which takes into account long-run bilateral effects as well as potential effects on other export markets. I begin using Colombian customs data from 2007 to 2012 to document novel patterns in the behavior of new exporters that suggest a pattern of sequential market entry by exporting firms in the face of extensive margin friction. I then build a quantitative model of the exporter's market entry decision that allows me to evaluate the consequences of trade policy aimed at reducing these frictions. The model illustrates that empirical evaluations of the effect of trade policy ignore a substantial fraction of the total effect of such policy.

These data allow me to document four findings. First, although new exporters take a few years to reach average total exports, the value of their exports to an individual foreign market is surprisingly similar to that of older exporters: new exporters actually exceed average sales per market within the first two years. Second, new exporters sell to fewer-than-average markets, even after five years of continuous exporting. Third, conditional on entry, newer exporters are more likely to enter one of the largest export markets for Colombia, while older exporters enter less popular markets, suggesting that firms begin with the largest markets and enter successively more difficult markets over the life cycle. The choice of export location and export product conform to the pattern demonstrated in Eaton, Kortum and Kramarz (2011): new exporters typically enter the top-ranked markets based on sales of all Colombian exporters. New exporters who survive export to new destinations over time, resulting in growth in total sales over the first five years of export. In contrast, experienced exporters who enter a new destination have lower sales, and are less likely to enter a top export destination. Fourth, I show that the entry hazard rises with the number of countries served by an exporter, suggesting that export experience creates efficiencies in other export markets.

It is well known that trade policies targeting the intensive and extensive margin of export will have different effects on the value of trade and the characteristics of domestic firms who are able to export after the policy. In order to understand how a trade policy aimed at encouraging the growth of new exporters might affect the behavior of firms who engage in sequential exporting, I build a model that captures my findings about exporter behavior. I assume that firms are constrained to enter markets sequentially and that the entry cost for any market falls in the number of foreign markets already served by the exporter. In this framework, exporters will choose to enter markets in order of decreasing profitability. Firms will enter markets that were not profitable to enter in the first year of export during later periods as their export experience reduces the cost of entry facing some constant probability that their export spell ends before they have entered all profitable markets.

The model is parsimonious in the number of parameters. Most important are the estimated size and fixed cost of each market and the shape of the distribution of firm productivity. In bringing the model to the data, I construct measures of relative market size based on average exports in each market and distributional assumptions on firm productivity. The steady state of the model offers clean relationships between moments of the distribution of firm sales and parameters of the model. Most parameter values come directly from moments of the export data. However, the elasticity of substitution across goods, which determines the correspondance between the distribution of firm sales and firm productivity, is difficult to calculate with this data. I adopt a commonly used value from Broda and Weinstein (2006). I compare the predictions of the model to key findings such as the patterns of new exporter growth and the correlation between entry rank in the data and as predicted by the model.

I use the calibrated model to predict the effect of reductions in the cost of entry for each foreign market one-by-one. I demonstrate that the potential benefits are large and different from those assumed by many empirical studies. It is common when evaluating empirical instances of export promotion to consider a bilateral effect over a short time frame. The model suggests that in addition to the bilateral change in market sales coming from a change in entry costs, there may be significant 'multilateral' effects on entry into other markets. Because firms adopt a sequential market entry strategy, the multilateral effect of a reduction in entry costs may not be fully realized in the first year after the change.

The multilateral effect can be positive or negative depending on whether the entry cost reduction increases or decreases the probability that an exporter sells to other markets. For example, consider an initial scenario in which an exporter will only export to Chile if it has first exported to Mexico. A reduction in entry costs to Chile could reverse the order in which the exporter enters these markets. It will now be less costly to enter Mexico because exporters have more export experience before attempting entry, meaning that lower productivity exporters will be able to enter Mexico. This will boost total exports to Mexico. However, deferring entry into an already profitable market means that some firms will be forced to terminate their export spell before they have the opportunity to export to Mexico. This will decrease total exports to Mexico. The total effect will depend on parameter values and market sizes. Estimates based on the model suggest that the net benefit of reducing entry costs may be as high as $336.6 million dollars per year for Venezuela (the market served by the most exporters with $2.9 billion in total exports per year) or as low as $4.6 million per year for Guatemala (the ninth largest export market with $311 million in total exports per year). The benefit to reducing entry costs to Guatemala is so low in part because the multilateral effect, which delays entry into Mexico and the Dominican Republic, is a loss of $9 million. For Mexico, the multilateral effect is positive: the increase in entry into Costa Rica and Peru conveys a benefit of $8.3 million.

The insights of the model contribute to a literature evaluating the impact of state export promotion activities on outcomes for exporting firms. Bernard and Jensen (2004) find no effect from state subsidies to exporting in the US, while evidence from Spain, Peru, and Uruguay suggests that increase their activity on the extensive margin in response to receiving export assistance (Gil-Pareja et al. 2008, Volpe Martincus and Caballo 2008, 2010). Other work has evaluated the pro-trade impact of less direct export promotion activities: permanent diplomatic missions (Rose 2011), visits by heads of state (Nitsch 2007), and formal trade missions (Head and Ries 2010). The modal result is that under plausible identification strategies, bilateral trade policies have modest or no effects on trade in the year immediately following implementation of the trade policy. Broad export promotion organizations of the type studied by Volpe Martincus and Carballo (2010) provide modest increases in the number of countries and products served by the exporter over the short term. These studies generally do not ask whether there might be persistent multilateral effects of trade policy. My findings suggest that when exporters enter markets sequentially, the long-term benefit is likely to exceed immediate gains from trade policy. Efficiencies in overcoming entry costs may lead exporters to add more than one new market as the result of purely bilateral trade promotion activities, but this multilateral effect will likely not appear immediately after a trade promotion policy is implemented.

My approach is most similar to the 'testing the waters' model of Albornoz et al. (2012) and the extended gravity model in Morales, Sheu and Zahler (2017). They show that under linear demand and uncertainty about profitability, firms may choose to start out by exporting a small quantity to the destination(s) with highest expected demand, entering additional markets after the uncertainty is resolved. The strongest evidence for the 'testing the waters' hypothesis comes from the particularly high growth rate of sales between year 1 and year 2, but Bernard et al. (2017) have shown that correcting for the partial year problem drastically reduces the growth rate of export sales. I build on the intuition of their paper by showing that in the presence of dynamic entry costs, exporters may have a motive to export sequentially even without reducing quantities exported in the first year. The approach to rationalizing sequential entry through falling entry costs echoes the intuition of Morales et al. (2017), who argue that there are entry cost efficiencies in exporting to markets with similar characteristics. While my findings and modelling assumptions are broadly consistent with the existence of extended gravity, I present evidence that the tendency of firms to export to markets in the same region can be rationalized through the concentration of demand for specific goods types in geographic regions.

The rest of the paper is organized as follows. Section 2 describes the Colombian export data used in the paper. Section 3 presents new facts about new export growth based on these data. Section 4 discusses the relationship between the state-dependent entry costs in Morales et. al (2017) and my findings. Section 5 develops the theoretical framework, in which entry costs decline in exporter experience, and shows that this simple model matches key findings in the data. Section 6 calculates the effect of trade policy aimed at reducing market-specific entry costs. Section 7 concludes.

Table II.1: Annual Exporter Sales, Millions

	Total	Firm-Level	Single Market
2007	13,680	1.74	0.29
2008	13.,700	1.76	0.30
2009	10,910	1.46	0.26
2010	10,930	1.66	0.27
2011	12,850	1.92	0.31
2012	13,100	1.91	0.31

2 Data

Colombian customs data come from the National Administrative Department of Statistics of Colombia. They include the f.o.b. value in US dollars of all Colombian manufacturing exports at the shipment level between 2007 and 2012 and reported quantities and weights in kilograms. Data on shipping and insurance costs are sporadically reported. Shipments are identified with the firm's unique national identification number (NIT), an HS10 product code, and the destination country and port.

In this paper, I focus on private manufacturing firms, which account for about 40 percent of total Colombian exports. Sales of crude oil and coal account for almost half of Colombian exports, but their sales are heavily influenced by world commodity pricing, which was particularly volatile during the period covered by the data. Furthermore, the state-controlled firm Ecopetrol accounts for most of Colombian crude oil exports.[1]

Table II.1 describes total manufacturing exports, average total firm sales, and average firm sales within individual markets (defined as an HS4 product-country pair) in millions of US dollars. Total exports range from \$10.9 to \$13.7 billion per year. The average firm exports about \$1.75 million per year, with an average of \$290 thousand per market. The pattern of growth in Colombian exports in my sub-sample of the data is consistent with Colombian national statistics: the trough of exports occurs in 2009. The trade collapse

[1] The Ecopetrol petroleum company, one of the largest Colombian exporters, launched an IPO in late 2007. The bulk of the company's stock continues to be held by the Colombian government and government officials represent about half of its board of directors.

Table II.2: Number of Firms by Type, 2008-2012

	Type of Firm			
	All	Exp. Non-Entrant	Exp. Entrant	New Exporter
2008	46013	41353	1283	241
2009	42674	36447	1283	241
2010	39926	33617	1283	241
2011	4612	33664	1283	241
2012	42712	33139	1283	241

looks relatively mild in Colombia: total exports fall by 14 percent.

In order to study the behavior of new exporters in the data, I establish the following definitions. I refer to an exporter that does not export in 2007 as a new exporter. For consistency with the literature, I focus particularly on new exporters that survive for five years, i.e., that export every year from 2008 to 2012. Of course, some firms that I define as new exporters could be firms that exported in the years before 2007. I evaluate the likelihood that this occurs by applying the same definitions shifted forward one year: new exporters are those that export in 2009 but not 2008. Since I have the data from two years earlier for these firms, I am able to calculate how many of these firms actually exported in 2007. I find that 7 percent of firms and 12 percent of firm-product-destination observations are misclassified as new exporters. Given that 2008-2009 represents the trough of the US recession, the amount of misclassification between 2007-2008 is probably lower since there is likely to have been less firm exit.

Among market entrants, I distinguish between new exporters and firms that already export another product or ship to another country. I refer to totally new exporters as 'new entrants' and to firms that have export experience and are adding a new market as 'experienced entrants'. I refer to existing exporters who do not add a new market in 2008 as 'experienced non-entrants'.

Based on these definitions, I am able to observe new firms for their first five years of export and incumbent firms for six. Ruhl and Willis (2017) find that about 90 percent of firms that have exported for five years survive to their sixth, suggesting that the panel is long enough to distinguish between successful and unsuccessful entrants. Table II.2 shows the number of firms of each type, as well as the total number of firms.

In the main analysis, I collapse export sales to the total in each year for each firm-destination country-HS4 and drop observations with total sales less than $1,000. The advantage of aggregating sales to the HS4 level is that products in different HS4 categories are likely to be clearly distinct goods in consumption (e.g. 6101, men's and boys' coats vs. 6103, men's and boys' suits). I assume that entry decisions for each product are independent even in multiproduct firms. The validity of this assumption is likely to vary across broad product types, but defining products at the HS2 or HS6 level generates similar results.

3 Empirical Findings

In this section, I start by showing that the data reproduce the growth pattern in Ruhl and Willis (2017): average total sales for new exporters take about five years to reach the average for all firms. I then document four stylized facts about new exporter growth: (1) new exporter sales are similar to the average in markets in which they are active, but they (2) enter different countries and (3) participate in fewer markets than more experienced exporters. Finally, (4) the probability that an exporter enters a new market increases in the number of markets served, even after accounting for differences in firm productivity. Taken together, these findings suggests that the extensive margin, not the intensive, must drive growth over the first five years of exporting.

I begin by demonstrating that new exporters in the data exhibit low total sales over a period of five years. The analysis concentrates on firms that enter a new market in 2008 and survive in that market for the five years over which I can observe them. New entering firms in 2008 reach the average total export sales of all exporters slightly less than five years after entry. Figure 3.1 shows the growth in average total sales and the average total export sales of all exporters between 2007 and 2012. The data produces a pattern similar to the

Figure 3.1: Average Total Sales by Export Experience

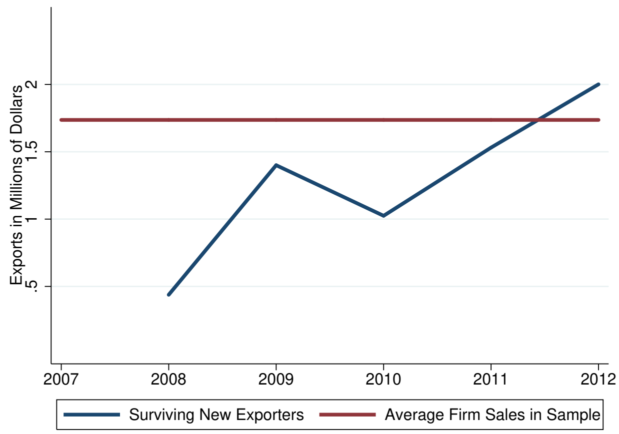

well-known stylized fact in Ruhl and Willis (2017), which used an earlier panel of Colombian export data. The authors express the export intensity as the ratio of export to domestic sales. Because I am looking at export data only, I compare new exporter sales to the average exporter sales over the sample.

3.1 High Single-Market Sales

Next, I turn to export sales in individual markets. Figure 3.2 compares sales for incumbents, surviving new exporters, and the average firm. Both surviving incumbent exporters and surviving new exporters exhibit growth during this period, but new exporters grow faster. However, average sales of new exporters start out as high as sales of other firms, including surviving incumbents, and grow to be about twice as large over the period. The high starting sales of new exporters does not immediately suggest that these firms are more constrained that the average exporter.

High sales by new entrants can be explained in two ways: it may be that surviving entrants have significantly higher productivity than incumbents, or they may be entering markets that are larger or growing relative to the average market. I regress log sales in each market on indicators for whether the firm is considered a new exporter or is an experienced exporter adding an additional market, interacted with the year of the observation. The

Figure 3.2: Average Market Sales by Export Experience

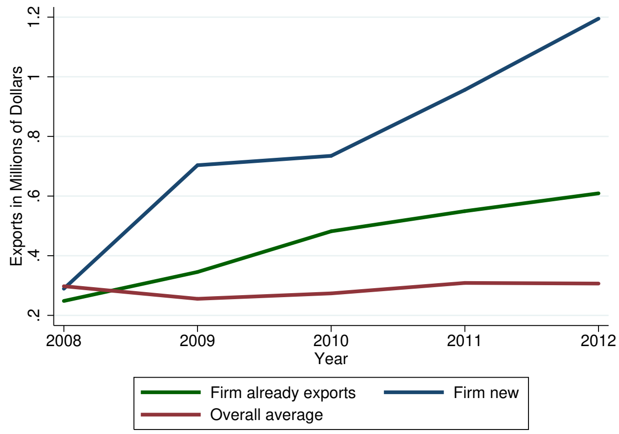

regression equation is as follows

$$logsales_{icp,t} = \alpha_0 + \sum_{i=1}^{5} \alpha_{it} New_i * Year_t + \sum_{i=1}^{5} \beta_{it} Experienced_i * Year_t + \Gamma X_{cp,t} + \varepsilon_{icp,t}$$

where $\Gamma X_{cp,t}$ represents an optional set of year t, country c and/or country-industry$p$ dummies (HS2 product codes correspond roughly to four-digit NAICS industries). These additional covariates may be necessary to account for time trends in export sales and differences in market size across industries or countries. If high sales by new exporters stem from high productivity, controlling for export destination or industry should have no effect on the result. On the other hand, if high sales stem from market selection, comparing exporters in the same market should clarify the result.

Table II.3 reports the results. Column (1) shows the results of the regression with year fixed effects only. New exporter sales in the first year are similar to those of other firms, growing to be almost double in five years. Experienced exporters who add a market actually have significantly lower sales than other firms in the first year according to this specification, eventually growing to about 72 percent larger. Columns (2) and (3) add destination and HS4 product dummies, accounting for differences coming from market selection: new entrants

Table II.3: Log Sales by Exporter Experience, 2008-2012

	(1)	(2)	(3)
New Exporters			
Year 1	0.286**	0.378***	-0.129
	(0.128)	(0.124)	(0.109)
Year 2	0.931***	0.991***	0.523***
	(0.126)	(0.125)	(0.111)
Year 3	1.158***	1.163***	0.658***
	(0.125)	(0.125)	(0.111)
Year 4	1.211***	1.222***	0.704***
	(0.129)	(0.128)	(0.116)
Year 5	1.111***	1.131***	0.588***
	(0.131)	(0.129)	(0.116)
Experienced Exporters			
Year 1	-0.198***	-0.113**	-0.271***
	(0.0526)	(0.0514)	(0.0501)
Year 2	0.300***	0.352***	0.233***
	(0.0540)	(0.0526)	(0.0495)
Year 3	0.611***	0.610***	0.452***
	(0.0545)	(0.0533)	(0.0493)
Year 4	0.767***	0.771***	0.601***
	(0.0557)	(0.0543)	(0.0506)
Year 5	0.704***	0.716***	0.521***
	(0.0556)	(0.0544)	(0.0517)
Fixed Effects	Year	Year, Dest	Year, Dest, HS4
Observations	212,937	212,937	212,937
R-squared	0.005	0.052	0.368

Standard errors are clustered at the firm (NIT) level. ** signifies p<0.01, * signifies p<0.05

may be over-represented in particular export markets. Allowing for differences in destination produces very similar results, except that new entrants actually start out larger than the average firm. The coefficients in column (4) do change significantly: new exporters still start out close in size to other firms in the same market, but perhaps a little smaller (the coefficient is significant at 90 percent), and grow to be 50 percent larger than the average. Experienced entrants start out significantly smaller and growth to be similar in size to new exporters. These results do not support the view that new exporters are significantly more constrained on the intensive margin than incumbent exporters. Both new and incumbent exporters grow over the sample period, but they do so at a similar rate. The fact that sales are higher than average for these two groups of firms likely reflects the fact that most exporters do not export continuously over a five or six year period; these firms are likely to be more productive.

Appendix Tables C.3 and C.4 show the results of the same regressions with products defined at the HS2 or HS6 level respectively. The results are similar using more aggregated products or more disaggregated products. In sum, the results of Table II.3 supports the conclusion that successful new exporters have sales revenue comparable to that of the average firm within their first year of export, generating the second stylized fact.

Finding 1: New entrants have similar sales to more experienced firms in individual markets in which they participate.

If low total sales are to be explained by frictions on the intensive margin, Finding 1 presents a puzzle. Low total sales combined with relatively high sales in individual markets suggests that market selection may play a role. The second and third empirical findings provide support for this hypothesis.

3.2 Low Market Participation

I now study market participation by new and existing exporters. Despite similar performance in individual markets, new exporters start out with fewer markets than incumbent

Table II.4: Exporter Markets in 2008

	Countries	Products	Total Markets
All	2.8	3	5.9
	(4.1)	(4.5)	(11.9)
Incumbent, No Entry	3.6	3.7	7.7
	(4.7)	(5.2)	(13.9)
New Entrant	1.2	1.7	1.9
	(0.9)	(1.8)	(2.3)
Incumbent, Entry	3.4	3.3	7.1
	(4.6)	(4.8)	(13.3)

firms. Table II.4 displays the average number of destinations, products and total markets (product-destination combinations) in which exporters participate in 2008, the year of entry for the new exporters I follow. New, successful exporters export to fewer markets than the average exporter. Experienced exporters appear to export to nearly twice as many product-destination combinations as new entrants, though the associated standard errors are large.

Because market participation is lower, total export sales for the new exporter are lower than average despite above-average sales in their active markets. This reconciles the apparently conflicting pictures of new exporter performance in Figures 3.1 and : new exporters select their best markets to enter first, explaining their strong single-market sales, but enter fewer markets, resulting in lower total sales across all export markets.

Next, I evaluate how the number of active market this interpretation further using a negative binomial regression of the number of markets in which exporters sell, with an indicator for $New_i * Year_t$ as in Table II.3. Column (1) of Table II.5 reports the results for the number of countries: new exporters sell in 81 percent fewer countries than other exporters in the first year, and 54 percent fewer in the fifth year after entry. Column (2) reports that new exporters sell 74 percent fewer products in the first year, and 46 percent fewer after five years. Last, Column (3) reports that in terms of total markets (destination-product pairs), new exporters sell in 122 percent fewer in the first year, and 84 percent fewer in the fifth year. While the increase in the number of markets served between the first and fifth years is statistically significant, the results of Table II.5suggest that new exporters are slow

Table II.5: Number of Markets by Years after Entry

	(1) Countries	(2) Products	(3) Total Markets
New Exporters			
Year 1	-0.813***	-0.738***	-1.221***
	(0.0810)	(0.0922)	(0.0908)
Year 2	-0.511***	-0.362***	-0.678***
	(0.0906)	(0.104)	(0.117)
Year 3	-0.585***	-0.314***	-0.720***
	(0.0855)	(0.114)	(0.120)
Year 4	-0.511***	-0.396***	-0.732***
	(0.0960)	(0.0945)	(0.121)
Year 5	-0.541***	-0.460***	-0.837***
	(0.0886)	(0.0848)	(0.0917)
Fixed Effects	Year	Year	Year
Observations	21,222	21,222	21,222

to add new countries and products, even after five years of export. Overall, new exporters participate in fewer markets even five years after successful entry into exporting. The delay is particularly remarkable since it has been shown that within active markets, surviving new entrants have high sales.

Finding 2: New entrants participate in fewer markets than the average firm, even after exporting for five years.

3.3 Market Selection

Finding 1 showed that new entrants have comparable sales to those of the average firm in individual markets. At first look, experienced exporters who enter a new market actually have smaller sales than completely new firms. At least part of this finding appeared to be explained by market selection: new exporters and experienced entrants looked more similar once destination-product differences in average sales were taken into account. This strongly suggests that market selection plays a role in explaining the differences between new and other exporters: if new exporters enter fewer markets, as Finding 2 showed, but these markets

Table II.6: Market Selection

	(1) Top Destinations	(2) Top HS2
New	0.176*	-0.0689
	(0.0822)	(0.0827)
Experienced	-0.281**	-0.0803*
	(0.0357)	(0.0363)
Observations	46,013	46,013

tend to be the largest export markets for Colombia, it would explain how total export sales could be low while individual market sales are high.

To demonstrate the role of market selection, I compare the probability that a new versus and experienced exporter begins exporting to one of the most popular destinations in 2008.[2] Overall, 80 percent of firms shipped to at least one of the top five export destinations in 2008. I also consider selection into the top product groups, focusing on the top five HS2 categories because the number of possible HS4 categories is much larger than the number of possible export destinations.

Whether focused on top destinations or top products, the regression equation is given by

$$Pr(shiptotopfive)_{icp,year} = \beta_0 + \beta_1 NewEntrant_{icp} + \beta_2 ExperiencedEntrant_{icp} + \varepsilon_{icp,y} \quad (3.1)$$

Column (1) of Table II.6 shows results of a probit regression of equation 3.1: the probability that a firm is active in a top five market on exporter experience in 2008. Because the new and experienced entrant variables are 0/1, their coefficients reflect the probability that the firm-product-destination is operating in a top five market conditional on being a new or experienced entrant. New exporters are 17.6 percent more likely to participate in a top destination. In other words, firms that begin their first export spell are particularly likely to be selling to a top 5 destination country. Experienced entrants, who already participate in

[2] In 2008, Ecuador, the United States, Panama, Venezuela and Costa Rica were the most popular manufacturing export destinations in terms of number of exporters.

some foreign markets in 2007, are 28 percent less likely to enter one of the top five destinations in 2008, implying that they are often already active in these markets if it is profitable. In fact, 92 percent of experienced entrants sell the same product to a top five destination.

Product market selection appears to be less relevant. Column (2) shows the result of a regression on the probability of participating in a top 5 HS2 product market. New exporters are as likely to do so as the average firm, while experienced entrants are 8 percent less likely to be entering a top product market. The type of products a firm produces may be a fundamental attribute of the firm, making the difference in the result for destination selection and product selection somewhat expected. It is realistic to assume that a firm may more easily be able to ship the same product line to a new country than to begin producing a completely different type of good.

Finding 3: Experienced entrants are less likely to enter the most popular export destinations than are completely new exporters, suggesting that firms enter markets sequentially, in order of decreasing profitability.

3.4 Sequential Exporting: Increasing Entry Hazard

Last, I document the probability that an exporter enters a new market as a function of the number of foreign markets they exported to in the previous year. For this analysis, I consider all exporters in the data who export for at least two consecutive years. I start by documenting the basic relationship between the number of foreign markets served and the entry hazard that appears in the data. I group exporters into eight bins based on the number of markets served last period to ensure a large enough number of observations in each bin, thereby smoothing out some noise in the entry hazard. Figure 3.3 plots the percentage of exporters entering a new market in a given year by their years of export experience, both unweighted and weighted by total sales in that year. In either case, the probability that a firm enters a new destination in a given year rises with the number of markets already served by the exporter. The weighted and unweighted entry hazards are relatively similar because

Figure 3.3: Fraction of Exporters Entering New Markets by Number of Markets Served

Note: I group the data into eight bins: 1 market, 2-3 markets, 4-5 markets, 6-8 markets, 9-11 markets, 12-15 markets, 16-20 markets, and greater than 20 markets. The smallest bin contains 187 observations.

conditioning on the number of markets already selects toward the largest firms. It may be somewhat surprising to note that the entry hazard continues to rise even for exporters who already export to more than ten markets: the entry hazard for firms that already export to about 30 markets is about 95 percent.

Striking though this empirical pattern may be, it is hard to draw conclusions about whether entering one additional market would make it easier for an exporter to enter new markets in the future. Underlying firm productivity could produce a correlation between the number of export destinations served by a firm and that firm's probability of entering new markets. I address this question by comparing the firm's total exports to the average total exports for the markets it serves, given by

$$MktAvgExp_{it} = \frac{1}{I} \sum_{k \in I} \bar{S_{kt}}$$

where I represents the set of export markets served by firm i and $\bar{S}_{kt}$ denotes average exports to market k in year t. I include this measure in a regression of a variable indicating whether firm i enters a new market at time t on the log of the number of markets:

$$Enters_{it} = \beta_0 + \beta_1 log(Markets)_{it-1} + \beta_2 log(AvgExp) + log(MktAvgExp)_{it} + \delta_t + \epsilon_{it} \quad (3.2)$$

I assume the effect of an additional year of exporting experience is log-linear, an assumption that appears to be consistent with the shape of the relationship in Figure 3.3. Table II.7 records the results of estimating regression model equation (3.2). The unweighted regression is reported in column (1): an increase in the number of markets served of one log point is associated with a 32 percent increase in the probability that the firm enters a new market that year. Weighting observations by the firm's lagged total export sales in column (2) decreases the strength of that relationship slightly: an increase in the number of markets served of one log point increases the probability of entering a new market by 27 percent. Because firms that exported to more foreign destinations are likely to have higher average total sales, the weighted regression likely weights the regression toward capturing the behavior of firms that already export to multiple markets. However, the point estimate is only slightly smaller. Adding year fixed effects in column (3) has a negligible effect on the estimated relationship: the estimate increases to 28 percent.

In principle, autocorrelation in the probability of market entry within firms might be a concern. In column (4), I run the full regression, accounting for firm sales relative to the average firm in the markets it serves. As expected, contribution of markets served last year to the probability of entry is smaller after accounting for a measure of productivity, but the effect is slight. An increase in the number of markets served of one log point is associated with a 27 percent increase in the probability of entry. While equation (3.2) may not succeed in fully controlling for firm productivity, the results reported in Table II.7 provide some evidence that there is a significant direct effect of market participation on the entry hazard.

The strong relationship between the number of markets to which the firm exports and market entry suggests that entering new markets over time is part of the typical exporter

Table II.7: Regression Results: Probability of Entering a New Market

	(1)	(2)	(3)	(4)	(5)
Number of Markets (log)	0.321	0.270	0.279	0.266	0.295
	(0.006)	(0.003)	(0.003)	(0.003)	(0.005)
Avg Firm Exports per market (log)				-0.010	-0.020
				(0.002)	(0.003)
Avg. Exports, Same markets (log)				-0.095	-0.064
				(0.008)	(0.010)
Total Sales Weighted		Yes	Yes	Yes	Yes
Year Fixed Effects			Yes	Yes	Yes
HS4 Fixed Effects					
Number of Obs	17,688	17,688	17,688	17,688	
R^2	0.16	0.34	0.38	0.38	

life cycle. This fact may be surprising because standard models of export behavior are static and assume that exporters pay an entry cost to sell to all markets in which they expect to be profitable at the same time. A decreasing cost of exporting over the life cycle of the firm suggests that the firm's decision to begin exporting should incorporate not only the expected flow profits from the markets it enters initially, but also the flow profits from markets it will expect to enter in the future.

4 Gravity and Sequential Exporting

In this section, I briefly discuss the relationship between my empirical findings and the widely-cited findings of Morales, Sheu, and Zahler (2017) that also suggest a pattern of sequential exporting. The authors find evidence of what they term 'extended gravity' in the behavior of exporting firms within the Chilean chemicals industry: firms that have recently begun exporting are particularly likely to enter markets that share characteristics such as a common border or language with the first set of export destinations the firm chose. Within a structural model of the export decision, this pattern suggests significantly smaller fixed costs to these destinations conditional on prior exporting. In principle, this mechanism could underlie and complement Findings 3 and 4, suggesting a specific reason why new exporters might enter markets sequentially and selectively.

First, I note that export industries (thought of here as HS2 codes) concentrate in distinct subsets of foreign markets. For example, Animal and Vegetable Oils and Fats (HS Chapter 15) account for 51 percent of the value of manufacturing exports outside of Spanish-speaking Latin America and 7 percent within. Within Latin America, trade of Vehicles (HS Chapter 87) account for 33 percent of manufacturing trade. Outside Latin America, this sector accounts for less than 1 percent of trade. As a result of this geographic concentration in exports, firms are almost guaranteed to export in a way that looks consistent with extended gravity, but could also be the result of geographic correlation in demand for the exporter's good.

Without data on domestic sales, as in Morales et al. (2017), I cannot implement their structural model on my data and derive specific estimates of the entry cost reduction from extended gravity with and without controlling for industry. However, in order to test whether the Colombian export data replicates their finding at least in reduced form, I calculate the conditional probability that an exporter enters a new foreign market in the same geographic region as a market to which it already exports. I calculate this conditional probabilty first without controlling for industry, then with HS2 products as an industry control. Appendix Table 1 gives the results of regressions giving the probability of adding a new market in a destination group conditional on which groups were added in year 1. For simplicity, I concentrate on Latin America, US-Canada, and the European Union and exclude all other destinations. Without conditioning on industry, firms that already export to Latin America or the EU are significantly more likely to enter a country in a region to which they already export, conditional on adding at least one market. After conditioning on industry, the sign reverses: exporters are if anything less likely to export to a new country in a region to which they already sell.

In this data, extended gravity, or the idea that entry costs are lower when exporters already serve similar markets, does not appear to explain more of market selection than what is captured by correlation between the top export industries in a region. These results are not contradict the hypothesis that entry cost efficiencies of the type found in Morales et al. (2017) exist for Colombian exporters, they merely indicate that the data does not afford enough variation to disambiguate these two potential explanations. In the model that follows, I incorporate the intuition that firms may face different costs of entry based on their export experience, but I abstract from the idea that these efficiencies are particularly large for markets with similar geographic characteristics.

5 A Model of Sequential Export Entry

I develop a model of the life cycle of exporters in order to understand the consequences of sequential exporting for the behavior of exporters and the consequences of export promotion policy. The model environment consists of a mass m of firms from the home country of heterogeneous productivity, a fraction δ of which enter and exit exporting status each period, that choose whether or not to export to a series of I foreign export markets trading off expected profits against the cost of entry. Entry costs decline in export experience.

For simplicity, I assume that firms can only enter one market per period. As a result of entry in prior periods, an exporter may eventually choose to enter a market that would not have been profitable earlier in the firm's export history. Firms face demand in each market that depends on their productivity. I analyze the model in partial equilibrium: it is assumed that the share of foreign consumption and domestic factor employment attributable to the activity of exporters from the home country is small enough that general equilibrium effects are negligible. I rule out a circumstance in which an exporter might enter a market in which its total profits will be negative in order to reduce its entry costs to another market by assuming that firms are myopic with respect to the source of the efficiency in entry costs.

5.1 Exporters

Firms from the home country differ in productivity ν , distributed $g(\nu)$. In each year of operation t, exporters can choose to enter (at most) one of I total export markets. Their flow profits are $\pi_i(\nu)$. In order to begin selling to market i, firms must pay a sunk entry cost of $F_i(N_{t-1})$. Expected net profits for each destination are given by

$$E[\pi_i(\nu)] = \frac{\pi_i(\nu)}{\delta} - F_i(N_{t-1}), \tag{5.1}$$

where $\pi_i(\nu) = \frac{1}{\epsilon} A_i \nu^{\epsilon-1}$. A_i represents the size of market i, which is known before entry.

Entry costs are a function of the number of markets to which the firm exported last period, N_{t-1},

$$F_i(N_{t-1}) = \frac{f_i}{N_{t-1}}. \tag{5.2}$$

Firms that have not yet begun to export or choose not to sell to foreign markets operate in the domestic market only, i.e. $N_{t-1} = 1$. For each market, there is a productivity threshold as a function of N_{t-1} above which all firms will make positive profits in expectation:

$$\nu_{it}(N_{t-1}) = \left[\frac{\delta \epsilon f_i}{A_i N_{t-1}}\right]^{\frac{1}{\epsilon-1}}$$

Given the entry constraint, firms will enter markets in order of highest profit net of fixed cost. Firms will enter market i before market j whenever $\nu_{it}(N_{t-1}) > \nu_{jt}(N_{t-1})$. Markets can be ranked and indexed in order of year of entry: $\frac{A_1}{f_1} > \frac{A_2}{f_2} > ... > \frac{A_I}{f_I}$. The productivity cutoff for market i can be simplified to

$$\nu_i \geq \left(\frac{\delta \epsilon f_i}{A_i i}\right)^{\frac{1}{\epsilon-1}}. \tag{5.3}$$

All exporters that survive to year i and meet the productivity threshold will enter market i. Note that because entry is sequential, it may be that $\nu_i < \nu_{i-1}$, in which case all surviving firms that export to market $i-1$ will enter market i.[3] The binding export cutoff is the maximum of the cutoff described in equation 5.3 for markets $i-1$ and i:

$$\nu_i^* = \max\{\nu_i, \nu_{i-1}^*\}. \tag{5.4}$$

5.2 Exports in Steady State

Because the same fraction δ of firms enter and exit export status every period, the model has a steady state in which, absent any shocks, the number of exporters and the total volume of sales to each market are constant. Given the decision rule for market entry in equations (5.3) and (5.4), total export sales to market i are given by

$$TS_i = m(1-\delta)^{i-1} \int_{\nu_i^*}^{\infty} A_i \nu^{\epsilon-1} g(v) d\nu. \tag{5.5}$$

The number of exporters in each market is given by

$$n_i = m(1-\delta)^{i-1}(1 - G(\nu_i^*)), \tag{5.6}$$

where $G(\nu)$ is the cumulative distribution function of firm productivity. Notice that a fraction $(1-\delta)^{i-1}$ of exporters will be exogenously shocked out of exporting status before reaching the period in which they may enter market i.

Average export sales in market i are then given by

$$\bar{S}_i = \frac{1}{1 - G(\nu_i^*)} \int_{\nu_i^*}^{\infty} A_i \nu^{\epsilon-1} g(v) d\nu. \tag{5.7}$$

[3] When the cutoff for market $i-1$ binds, I am implicitly assuming that firms are myopic about the decrease in fixed costs. Firms with productivity $\nu_{i-1} > \nu \geq \left(\delta\epsilon \frac{\frac{f_{i-1}}{i-1} + (1-\delta)\frac{f_i}{i}}{A_{i-1} + (1-\delta)A_i} \right)^{\frac{1}{\epsilon-1}}$ would have an incentive to incur losses in market i in order to enter market j. In the calibrated model, less than 0.1 percent of firms behave differently if I relax the assumption of myopia.

5.3 Change in f_i

Next, I consider the effect of a decrease in the fixed costs of entry f_i.[4] A change in bilateral entry costs can have multiple effects because both the productivity cutoff and the market entry order can potentially change. Some firms that would not enter market i facing fixed costs f_i will do so immediately facing $f_i^{'}$. In the long term, other firms for whom entry into market i is eventually profitable will enter as well. The decrease in fixed costs may affect other markets as well (I refer to these as 'multilateral' effects). If the cutoff for market i is binding on market $i+1$ ($\nu_i = \nu_{i+1}^*$), new exporters will enter market $i+1$ after the decrease in f_i. If the decrease in f_i is large enough to affect the entry order of markets, exports into markets $i-k$ may be affected. On the one hand, deferring entry into market $i-k$ means that some exporters will exit before they have an opportunity to sell to market $i-k$. However, exporters will also face lower fixed costs in market $i-k$: $F_{i-k}^{'} = \frac{f_i}{i-k+1} > F_{i-k}$.

I start with the long-term (steady state) effect of the change in f_i, assuming no change in entry order:

$$\frac{dTS_i}{df_i} = -m(1-\delta)^{i-1} A_i \left(\nu_i^*\right)^{\epsilon-1} g(\nu_i^*)(\epsilon-1)\frac{d\nu_i^*}{df_i}.$$

Note that if $\nu_i^* = \nu_{i-1}$, $\frac{d\nu_i^*}{df_i} = 0$. Otherwise,

$$\frac{d\nu_i^*}{df_i} = \frac{1}{\epsilon-1}\left(\frac{\delta\epsilon f_i}{A_i i}\right)^{\frac{1}{\epsilon-1}-1}.$$

If $\nu_{i+1}^* = \nu_i$, the change in the entry threshold for market $i+1$ may be affected. The change in the entry threshold for market $i+1$ is given by

$$\frac{dv_{i+1}^*}{df_i} = \max\{\nu_{i+1} - \nu_i, \frac{d\nu_i^*}{df_i}\}.$$

Large changes in the entry cost from f_i to $f_i^{'}$ will affect the entry ordering of markets.

[4] The intuition behind the effects of an increase in entry costs will be similar, but because increasing entry costs faced by domestic firms is rarely a goal of policy, I focus on a decrease.

Denote the productivity cutoff for market with previous order i and new order $i^{'}$by

$$v_{i,i'}^{*}=\max\{\nu_{i'},\nu_{i^{'}-1}^{*}\}.$$

A change of market order produces a discrete change in steady-state sales in market i, given by

$$TS_{i'}-TS_{i}=m\left[(1-\delta)^{i^{'}-1}\int_{\nu_{i'}^{*}}^{\infty}A_{i'}\nu^{\epsilon-1}g(v)d\nu-(1-\delta)^{i-1}\int_{\nu_{i}^{*}}^{\infty}A_{i}\nu^{\epsilon-1}g(v)d\nu\right].$$

The direction of the change in productivity is monotonic. In response to a decrease in f_i, $v_{i,i'}^{*}\leq v_{i}^{*}$. For simplicity, consider the case where a decrease in f_i causes only only two markets to switch order. Market $i^{'}-1$ (previous market i) has a lower cutoff than market $i-1$:

$$v_{i,i^{'}-1}^{*}=\max\{\nu_{i'-1},\nu_{i^{'}-2}^{*}\}.$$

The cutoff for market $i^{'}-2$ is unaffected because fixed costs and the order are both unchanged:

$$\nu_{i'-2}^{*}=\nu_{i^{'}-2}^{*}$$

The cutoff to enter the $i^{i}-1$th market falls ($\nu_{i^{'}-1}<\nu_{i-1}$) because

$$\frac{A_i}{f_i}<\frac{A_{i^{'}-1}}{f_{i^{'}-1}}.$$

In other words the fixed costs associated with the new market $i^{'}-1$ are low enough that the size of market i relative to fixed costs is now large enough to precipitate a change in market ordering. Finally, the cutoff to enter the $i^{'}$th market falls:

$$\nu_{i'} < \nu_i$$

because exporters already experienced positive profits from entering this market one period earlier:

$$\frac{A_i}{f_i} < \frac{A_{i'}}{f_{i'}}.$$

In the period immediately following the change, all exporters of age k, where $i'-1 \leq k \leq i$, that have productivity above the new cutoff will enter:

$$\Delta TS_{i'}(0) = m\left[\sum_{k=0}^{i-i'}(1-\delta)^i \int_{\nu^*_{i',i-k}}^{\nu^*_i} A_{i'}\nu^{\epsilon-1}g(v)d\nu\right]$$

Note that on the transition path to the new steady state, firms of age k have actually entered more markets than prescribed by the new entry order. This implies that in the short term, lower productivity firms will enter market i' than in steady state. As these firms that are below the steady-state cutoff exit exporting status, average firm exports will increase slightly.

5.4 Calibration

In order to quantify how high the spillovers from market entry may be, I calibrate the model to match key features of the data. I assume that export markets in the data are in steady state, which allows me to equate the expressions in equations (5.5), (5.6), and (5.7) to their analogs in the data. The values of each parameter of the model are reported in Table II.8. Most values are estimated using my data, but I take values for certain parameters that are difficult or impossible to estimate with customs data from the literature. I set the elasticity of substitution ϵ equal to the widely-used value from Broda and Weinstein (2006). In order to pin down the probability of entering the first export market, which depends on the ratio $\frac{\nu_m}{\nu_i^*}$ (only the ratio matters here, so I assume $\nu_m = 1$ for simplicity), I use the finding from

Table II.8: Calibration

Parameter	Target Moment	Value
δ	Average Dropout Rate	.039
ϵ	Broda and Weinstein (2006)	6
β	Shape of distribution function	6.6
ν_1	Share of firms exporting, Ruhl and Willis (2017)	0.25
m	Number of exporters	11,700
$\{A_i\}$	Fraction of exporters in market	-
$\{f_i\}$	Average sales per market	-

Ruhl and Willis (2017) that an average of 25 percent of Colombian plants export.

I choose δ to match the weighted average one-year survival rate, finding an average survival rate of 96 percent (high-revenue exporters survive at a far higher rate than the unweighted average).

In order to parametrize the distribution of firm productivity parsimoniously, I assume that firm productivity is Pareto distributed with cdf

$$G(\nu) = 1 - \left(\frac{\nu_m}{\nu}\right)^{\beta}. \tag{5.8}$$

I calibrate the shape parameter of the Pareto distribution by exploiting its relationship with the distribution of firm sales as is common in the literature. The probability that an exporting firm's sales to market i exceed S_i is given by

$$1 - Pr(S_i) = \left(\frac{S_i}{\delta \epsilon F_i}\right)^{-\frac{\beta}{\epsilon-1}}. \tag{5.9}$$

I estimate $\frac{\beta}{\epsilon-1}$ using the relationship

$$logPr(S_i) = -\frac{\beta}{\epsilon - 1} logS_i + \delta_i + \epsilon_i. \tag{5.10}$$

I focus on a sample of the fifteen largest export markets in terms of total annual exports. I then divide the sales in each foreign market into centiles to implement the regression, for which I use the top fifty percent of exporters.[5] Given the assumed value of ϵ, I estimate $\hat{\beta} = 2.2$ (0.02). Using this regression, it is possible to predict a minimum value of export sales in each market. I use the average value for the top thirty percent of firms to mitigate problems associated with local deviations between the empirical distribution of firm sales and the assumed Pareto distribution.

I use the mean sales in each market $\bar{S}_i$ given in equation (5.7) under the distributional assumption in equation (5.8) to solve for the set of fixed costs paid for each market $\{F_i\}$:

$$\bar{S}_i = \frac{\beta}{\beta - (\epsilon - 1)} \delta \epsilon F_i. \tag{5.11}$$

Notice that while the minimum value of sales can pin down F_i, this relationship cannot inform about the fundamental parameter values $\{f_i\}$ without knowing the market ranking. The market ranking will also be necessary in order to estimate values for the market-specific terms $\{A_i\}$. I use the number of exporters that export to each market according to equation (5.6) to calculate the probability that an exporter serves market i, conditional on surviving, which is given by

$$\frac{n_i}{n_{i-1}} = (1-d)\frac{1 - G(\nu_i^*)}{1 - G(\nu_{i-1}^*)}$$

This simplifies to

$$\frac{n_i}{n_{i-1}} = (1-d)\left(\frac{F_{i-1}}{F_i}\frac{A_i}{A_{i-1}}\right)^{\frac{\beta}{\epsilon-1}}$$

Bearing in mind that the fraction of firms who export to the first market is assumed to be 0.25, I use this expression to pin down the values of $\{A_i\}$.

Finally, the fraction of firms operating in each market implies the order of entry ($n_1 > n_2 > .. > n_I$). I solve for the parameter f_i to satisfy equation (5.2). I find that the markets into which exporters enter in early periods have significantly lower estimated entry costs f_i, while market size A_i is largely unrelated to entry order.

[5]The fact that a Pareto distribution fits best for the right tail of sales has been extensively documented.

Table II.9: Export Market Sales and Entry Costs

Rank	Country	Average Sales	Entry Costs F_i	Percent of All Exporters
1	Venezuela	2,306	2,883	100
2	Ecuador	1,296	1,621	80
3	USA	1,707	2,134	66
4	Panama	406	507	54
5	Peru	1,469	1,836	39
6	Costa Rica	490	613	38
7	Mexico	1,093	1,367	34
8	Dominican Rep.	550	687	25
9	Guatemala	436	545	24
10	Chile	1,290	1,613	19
11	Spain	1,414	1,768	15
12	Brazil	3,897	4,872	9
13	Italy	3,554	4,444	6
14	China	10,818	13,526	5
15	Netherlands	4,529	5,663	4

Note: The table lists average sales per market and sunk entry costs paid in thousands of USD in the calibrated model.

Finally, the mass of firms can be calculated by using the relationship

$$m_1 = mG(\nu_1^*).$$

5.5 Model and the Data

In order to understand how well this simple model can capture salient features of the data, I compare the predictions of the model to the closest empirical equivalent of the entry order the model describes and to my empirical findings.

For the top fifteen markets in terms of export sales, I assign a rank based on the share of exporters, weighted by sales, that export to the market. I compare this market rank with the prediction of the model listed in Table II.9. I find a correlation of 0.95 between the rank predicted by the model and the empirical ranking of markets. Given that the ordering in the model relies essentially on the unweighted share of exporters that export to each market,

Figure 5.1: Correlation Between the Empirical and Model Market Rank

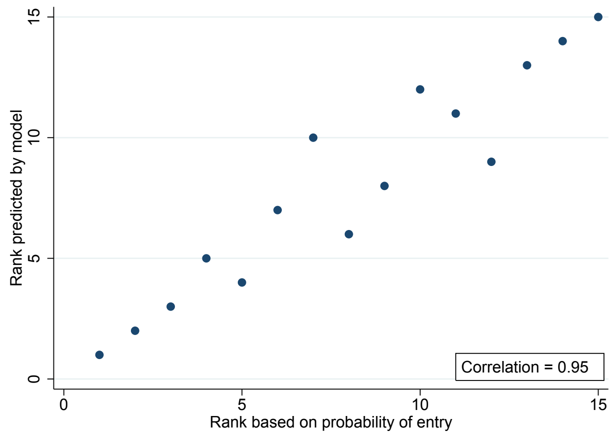

the correlation between the two market ranks should be high. Naturally, the prediction of the data is less stark than that of the model: there is a positive probability that a firm with only one export market will be selling to any of the top fifteen export markets.

Next, I compare the predictions of the model with my empirical findings. Two of the findings are true in the model by assumption. First, new exporters have average export sales in each individual foreign market because there are no intensive margin frictions in the model. Second, exporters are constrained to enter markets sequentially. This result is more extreme in the model for tractability but is intended to reflect the fact that new exporters enter new markets over time.

There are three empirical findings that the model may or may not replicate depending on how the parameterization affects average exporting behavior. The first finding to verify is that new exporters have below average total sales for a period of about five years. The average total sales of a new exporter that survived to year t of an export spell are simply the sum of average exports to markets 1 to t. Average total exports for all exporters are given by

$$AvgTtlSales = \sum_i (1-\delta)^{i-1} \left(\frac{\nu_1^*}{\nu_i^*}\right)^{\beta} \bar{S}_i$$

where $\bar{S}_i$ is defined by equation (5.11).

Figure 5.2: Total Sales of New and Average Exporters, Model

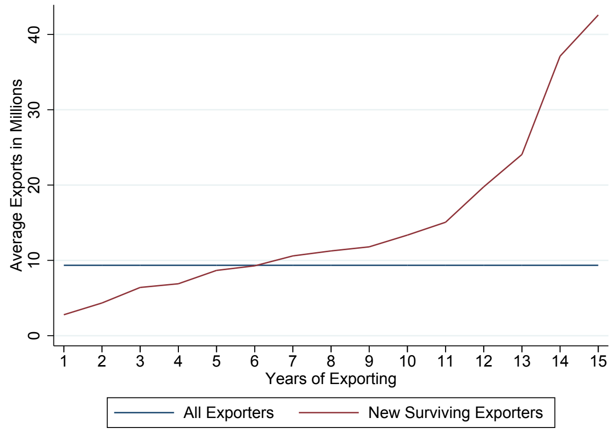

Figure 5.2 shows what the relationship between new and average exports looks like in the model. The pattern looks similar to the empirical pattern shown in Figure 3.1. It takes just over six years for new exporters to reach average size. This is because I target the death rate of exporters δ to match the weighted average across exporters, thereby overstating the lifespan of the average firm. Note that average exports of both new and average firms are also higher as a result of the decision to fit the model to the top 50 percent of exporters in terms of sales.

The second finding to verify is the fact that new exporters have below average market participation and reach the average slowly over a period of five years or more. A new exporter in the model that survives for t years will sell to t markets provided that its productivity exceeds the cutoff in equation (5.4):

$$AvgMarkets_t = \sum_{i=1}^{t} \left(\frac{\nu_1^*}{\nu_i^*} \right)^{\beta}.$$

The number of markets served by the average exporter is given by

$$AvgMarkets = \sum_{i} (1-\delta)^{i-1} \left(\frac{\nu_1^*}{\nu_i^*} \right)^{\beta}.$$

Figure 5.3 shows that the average exporter sells to about 6 markets. It takes 9 years for

Figure 5.3: Number of Markets, New and Average Exporters, Model

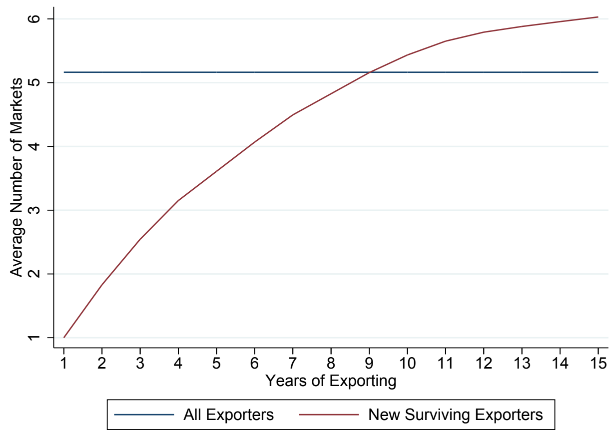

Figure 5.4: Number of Markets, New and Average Exporters, Model

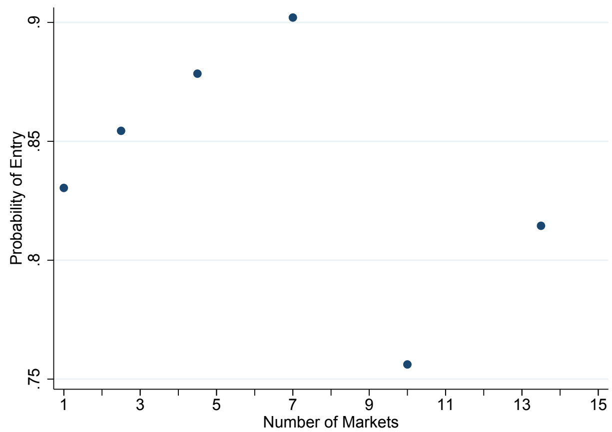

new exporters to catch up. Just as in the empirics, I find that even after total exports reach the average, the firm still participates in fewer markets.

Finally, I find empirically that the entry hazard increases in the number of markets served by an exporter and appears to be concave. I compare this empirical regularity with the model-generated hazard $G(\nu_i^*|\nu > \nu_{i-1}^*)$. Figure 5.4 shows the result. For exporters in less than 10 foreign markets, there increasing entry hazard as a function of the number of markets. As the number of markets increases beyond 10, market size relative to fixed costs falls, resulting in a decreasing probability of entry.

6 Reducing Entry Costs

I evaluate the bilateral and multilateral effect of decreasing entry costs f_i by 50 percent one-by-one for each of the top fifteen export markets. A decrease in market-specific fixed costs affects the number of exporters and the volume of exports to market i whenever the new cutoff ν_i' is binding. Table II.10 reports the effects of reducing fixed costs f_i market-by-market on the number of exports and total exports. For the first market entered by exporters, Venezuela, a 50 percent decrease in fixed costs results in a 92 percent increase in the number of exporters serving the market. Average sales per exporter fall by 69 percent because lower productivity exporters are able to sell to Venezuela, resulting in a 22 percent increase in total annual exports. The largest bilateral benefit actually comes from decreasing entry costs to China, which increases total exports by 25 percent.

Reducing entry costs to some markets results in effects that are not captured by the bilateral benefits of the decrease in export costs. The 'multilateral' impact of reducing f_i can have a positive or negative impact on destinations $j \neq i$. More firms may export to market j either because the decrease in f_i decreases ν_j^* (this happens when $\nu_i^* = \nu_j^*$) or because firms used to enter market i after market j, but now the entry order flips. This implies that entry costs F_j fall as implied by equation (5.2). However, the multilateral effect may also be negative when market i and market j switch entry order because a fraction of exporters that used to export to market j will now drop out of exporter status before entering j. Multilateral effects are generally modest. The largest multilateral gain comes from reducing entry costs to Mexico, which results in an increase in exporters of 20 percent to both Costa Rica and Peru and an increase in total export sales of 2 percent in each of those markets. The largest negative effect comes from reducing entry costs to Chile, where the reduction of entry costs results in an average 3 percent increase in the number of exporters to Mexico, the Dominican Republic and Guatemala but a 2 percent decrease in total export sales.

Table II.10: Effect of Fixed Cost Reduction

		Effect on Market i			Effect on Affected Markets $j \neq i$		
Rank	Country	$\%\Delta n_i$	$\%\Delta \bar{S}_i$ (\$)	$\%\Delta$Total Sales (\$)	$\%\Delta n_j$	$\%\Delta \bar{S}_j$ (\$)	$\%\Delta$Total Sales (\$)
1	Venezuela	92	-69	22	0	0	0
2	Ecuador	19	-14	5	0	0	0
3	USA	38	-26	12	15	-14	1
4	Panama	36	-24	12	11	-11	0
5	Peru	48	-34	15	8	-8	0
6	Costa Rica	49	-31	18	17	-16	1
7	Mexico	43	-27	17	20	-18	2
8	Dominican Rep.	41	-25	16	1	-5	-2
9	Guatemala	45	-25	20	3	-7	-2
10	Chile	56	-34	23	3	-5	-2
11	Spain	61	-38	24	6	-8	-2
12	Brazil	72	-51	20	8	-9	-1
13	Italy	85	-61	24	7	-8	-1
14	China	79	-54	25	6	-7	-2
15	Netherlands	74	-50	24	5	-7	-2

Table II.11: Net Benefit of Fixed Cost Reduction, ($,1000)

		Effect on Market i		Total Effect on Markets $j \neq i$	Total Benefit
Rank	Country	Δf_i	ΔTotal Profits	ΔTotal Profits	ΔTotal Profits
1	Venezuela	-1,441	336,601	0	336,601
2	Ecuador	-1,621	28,027	0	28,027
3	USA	-3,200	85,398	3,215	88,613
4	Panama	-1,014	15,789	-2,133	13,656
5	Peru	-4,590	53,515	890	54,405
6	Costa Rica	-1,838	21,857	5,404	27,261
7	Mexico	-4,783	38,663	8,297	46,960
8	Dominican Rep.	-2,748	13,916	-8,505	5,412
9	Guatemala	-2,451	13,909	-9,224	4,685
10	Chile	-8,065	37,367	-8,181	29,186
11	Spain	-9,725	32,754	-4,553	28,201
12	Brazil	-29,230	46,972	-1,438	45,534
13	Italy	-28,882	30,884	-2,929	27,955
14	China	-94,674	82,537	-5,352	77,186
15	Netherlands	-42,469	30,046	-6,902	23,144

While the effects in percentage terms are easier to digest, the variation in the size of fixed costs and market size make it hard to evaluate the magnitude of these effects. Table II.11 displays the net effect of the decrease in bilateral fixed costs, assuming that the decrease in fixed effects faced by exporters takes the form of a subsidy for exporting, such that the policy maker pays 50 percent of the fixed costs f_i. The biggest benefit comes from reducing fixed costs to Venezuela, resulting in a total benefit of $336.6 million. This benefit is about 72 times larger that the benefit from reducing entry costs to Guatemala, despite the fact that entry costs are relatively similar. Reducing entry costs to Guatemala results in a benefit of only $4.7 million. The benefit is so small because delaying entry into Mexico and the Dominican Republic results in loss in levels of $9.2 million.

In Mexico, where the multilateral benefits are highest in percentage terms, the net effect is an increase in total profits from exporting of $42.2 million, of which $8.3 million come from the multilateral effect on entry in Costa Rica and Peru. The multilateral increase in profits plays a similarly large role in the benefits of reducing entry costs to Costa Rica.

7 Conclusion

This paper demonstrates that new exporter sales patterns are consistent with the existence of extensive margin constraints resulting in a pattern of sequential market entry. I show that intensive margin sales for new Colombian exporters are similar in magnitude to the sales of other exporters after the first year, despite the fact that total sales reach the average only in the fifth year, and that new exporters are disproportionately present in only the top export markets despite the fact that they are likely to have higher productivity than the average exporter. Furthermore, the probability that an exporter enters a new market rises sharply in the number of markets served. These findings suggests that new exporters may face higher costs of entry than other firms. A model in which exporters enter markets sequentially and fixed costs fall with export experience fits these findings well. I show that as the result of lowering bilateral entry barriers, trade volumes can increase by up to \$336.6 million bilaterally but also by up to \$8.3 million in other markets as the result of the decreased entry cost to these markets coming from the firm's accumulation of export experience.

APPENDICES

Appendix A: A Variety-Adjusted Price Index

In this section, I use shares and prices of grocery goods for the 26 cities in the data to calculate an adjusted price index consistent with the CES utility structure in section 4.1. The measured CPI relies upon a constant basket of goods, but under CES preferences, the price level decreases in the number of varieties. I follow Broda and Weinstein (2010) in defining a variety-adjusted price index consistent with a three-tier CES preference structure within each retailer, but add an additional, uppermost tier describing consumer's substitution across retailers.[6] I discuss this modeling decision and evidence from previous studies that retailers are imperfect substitutes in section 3.1. Here, I begin by outlining the structure of the price index and the assumptions needed to estimate elasticities of substitution. In sections 7 and 7, I will consider two well-known approaches to estimating elasticities of substitution, one based on Feenstra (1994) and one on Hausman (1996).

In order to identify elasticities of substitution for each tier of the price level, I assume that each coarse product group has a constant elasticity of substitution within each tier. For brand-modules in the same group (e.g. Dole Canned Peaches in the Canned Fruit group), the elasticity of substitution between UPCs in the brand-module will be the same. Similarly, the elasticity of substitution across brand-modules in the group is constant. For brevity, I

[6] I abstract from the quality adjustment considerations discussed in Broda and Weinstein (2010). Differences in quality should be captured by products' in-category market share.

list only the expression for the most disaggregated tier: the growth in the price level for brand-module b in group g:

$$\frac{P_{b,rt}}{P_{b,rt-2}} = \prod_{x \in \Omega^*_{b,t}} \left(\frac{p_{xb,rt}}{p_{xb,rt-2}} \right)^{w^g_{xb,rt}} \left(\frac{\lambda^g_{b,rt}}{\lambda^g_{b,rt-2}} \right)^{\frac{1}{\sigma^b_g - 1}} \tag{.1}$$

where $\Omega^*_{b,t}$ denotes the set of UPCs common to periods t and $t-2$ in group g and $\mathbf{w^g_{bt}} = \{w_{xb,rt}\}_{\Omega^*_{b,rt}}$ are consumption weights to be specified below. The set of common UPCs $\Omega^*_{b,t}$ is a subset of all UPCs sold in a given time period, denoted $\Omega_{b,t}$. $\lambda^g_{b,rt}$ is defined as

$$\lambda^g_{b,rt} = \frac{\sum_{\Omega^*_{b,t}} p_{xb,rt} q_{xb,rt}}{\sum_{\Omega_{b,t}} p_{xb,rt} q_{xb,rt}}, \tag{.2}$$

In words, $\lambda^g_{b,rt}$ represents the share of all revenue in the brand-module attributable to products sold in both $t-2$ and t. Notice that, all else equal, if the brand-module experiences sales of new varieties in year t that exceed those in year $t-2$, $\lambda^g_{b,rt}$ falls relative to $\lambda^g_{b,rt-2}$, and the relative price level in year t also falls. Conversely, if new products represent a smaller share of expenditures in the category in year t than in year $t-2$, the price level rises. The strength of the impact of the share ratio $\frac{\lambda^g_{b,rt}}{\lambda^g_{b,rt-2}}$ is governed by the elasticity of substitution σ^b_g. Again following the literature, I define the weight on price inflation for common products as the Sato-Vartia (1976) weights based on the share of variety x in total expenditures on common products, $s_{xt} = \frac{p_{xb,rt} q_{xb,rt}}{\sum_{\Omega^*_{b,t}} p_{xb,rt} q_{xb,rt}}$,

$$w^g_{xb,rt} = \frac{\frac{s^g_{xb,rt} - s^g_{xb,rt-2}}{lns^g_{xb,rt} - lns^g_{xb,rt-2}}}{\sum_{\Omega^*_{b,t}} \frac{s^g_{xb,rt} - s^g_{xb,rt-2}}{lns^g_{xb,rt} - lns^g_{xb,rt-2}}}. \tag{.3}$$

Sato-Vartia weights represent a method of chain-weighting. While it is possible to compute an adjusted price index using either Laspeyres or Paasche weights, thereby imitating the conventional measured CPI, each of these expenditure weight schemes introduces a set of mismeasurement concerns due to substitution, precisely the issue this adjusted price index is attempting to resolve. The Sato-Vartia is an ideal weight in the sense that

$c(p_t, \Omega_{b,t})/c(p_{t-1}, \Omega_{b,t-2}) = P_t/P_{t-2}$ in the absence of demand shocks to common goods. Redding and Weinstein (2016) note that these weights are different from those implied by a ideal CES price index with time-varying taste shocks and propose an alternative weighting scheme. However, the common goods price growth and the variety adjustment remain log-separable in their ideal CES price index. Because my primary interest is in producing a comparison between a typical measured price index (here, chain-weighted) and a variety-adjusted price level, I retain this weighting scheme.[7] If there are additional sources of mismeasurement in the conventional CPI, the price index defined in (.1) is no longer exact, but as long as the sources of mismeasurement are multiplicatively separable, it remains an exact measure of the bias created by considering only a common basket of varieties.

The price indices for the middle and top tier are constructed similarly. The data includes a product group (coarse category), product module (fine product category) and brand for each UPC. In calculating the empirical price index, I allow for two different elasticities within each group: the elasticity of substitution across UPCs within a brand-module σ_b^g, and across brand-modules in the same group σ_a^g. There is one elasticity across product groups within each retailer, σ^g , and across retailers, σ_r. The four-tiered utility function allows for foure different variety adjustment terms in principle. The addition of a new brand within a module or a new module within a group could imply greater welfare gains, since fewer close substitutes for the new variety may exist. In practice, I restrict the sample to a constant sample of product groups and consider a set of retailers that are common to both periods. There are only two relevant variety adjustment terms: the within brand-module adjustment defined in equation (.2) and the across brand-module adjustment within each retailer that is calculated similarly.

The price index for final consumption is given by

[7]A requirement of this or any other measure of the price level using data on prices and expenditures is that demand shocks 'net out' across the set of products common to the two periods. Otherwise, the consumer experiences a wholesale increase in utils per dollar which cannot be measured.

$$\frac{P_t}{P_{t-2}} = \prod_{r\in R}\left[\prod_{g\in\Omega}\left(\prod_{b\in\Omega^*_{g,t}}\left(\prod_{x\in\Omega^*_{b,t}}\left(\frac{p_{xb,rt}}{p_{xb,rt-2}}\right)^{w^g_{xb,rt}}\left(\frac{\lambda^g_{b,rt}}{\lambda^g_{b,rt-2}}\right)^{\frac{1}{\sigma^b_g-1}}\right)^{w^g_{b,rt}}\left(\frac{\lambda^g_{rt}}{\lambda r^g_{t-2}}\right)^{\frac{1}{\sigma^g_a-1}}\right)^{w^g_{rt}}\right]^{w_{rt}} \tag{.4}$$

where $\Omega^*_{g,t}$ represents the set of common brand-modules within the group and the shares of common brand-modules within the group, $w^g_{b,rt}$, shares of each group within the retailer w^g_{rt}, and shares of each retailer w_{rt} in total consumption are defined analogously to $w^g_{xb,rt}$ in equation (.3).

The variety-adjusted price index in equation (.4) allows for a flexible estimate of gains from the entry of new product varieties. The contribution to gains from variety in each product category is affected by three elements of the price level: the relative share of new products, the elasticity of substitution, and the category consumption weight.

First, the value of new products to consumers is measured as their share of total revenue. Varieties within a product category that are most valued by consumers are likely to have a high revenue share, either because they are purchased in large quantities or because they are expensive as a result of high demand. Therefore, the share of common products λ is likely to fall significantly between $t-2$ and t when a new variety that is highly valued by consumers enters the market in period t, while it will rise when a product with a high revenue share leaves the market between $t-2$ and t.

Second, the welfare implications of the growth or fall in the share of common products are adjusted by the elasticity of substitution in each category. Because the substitutability of products may vary widely across categories, it is essential to adjust the contribution of the extensive margin accordingly. Suppose that products in a brand-module are near-perfect substitutes, implying that if a new and relatively inexpensive product enters the market, it will represent a very large share of revenues in the brand-module. This will occur even if the difference in prices between the two goods is small, resulting in a very large λ ratio. Clearly, the welfare implications of a slight decrease in the price should also be small. Raising the expression to the power $\frac{1}{\sigma_b^g - 1}$, the same expression found in the CES price index, adjusts the price level by scaling down the contribution of the extensive margin in highly substitutable product categories relative to less elastic categories. I restrict $\sigma > 1$ for the elasticity of substitution within brand-modules and across brand-modules within the same product group.[8] Because the extensive margin at the group level is inactive, the upper tier elasticities do not need to be estimated: their values are implicit in the relative group weights w_{rt}^g and w_{rt}. The elasticities of substitution σ_b^g and σ_a^g play a large role in capturing the welfare implications of the extensive margin. The procedure for estimating them is addressed in detail in the next section.

The third and final aspect of the price level that affects the variety-adjustment to the price level is the set of weights $\mathbf{w}_t$. The level of activity on the extensive margin may vary widely across product types, which could be driven by demand- or supply-side differences between types of producers. The relative weights between categories, which are applied to the common goods price index and the variety adjustment, re-weight the contribution of each category's extensive margin by its share in consumption. This again captures the intuition that the consumer does not value variety *per se*, but rather a larger consumption set in important expenditure categories.

The data contain prices and expenditures for each product, which provides all the information necessary to calculate an adjusted price index except the within and across elasticities σ_b^g and σ_a^g. I now describe how I estimate these elasticities following approaches based on Feenstra (1994) and Hausman (1996).

Feenstra (1994) Elasticities

The difficulty of estimating the elasticity of demand based on price and quantity data is well-known. Without more information or assumptions, it is impossible to disentangle demand and supply shocks. Here, the large number of observations for a panel of many related

[8] With inelastic demand for varieties, the price level is infinite whenever one variety is unavailable. This is true, for example, in Cobb-Douglas or Leontief preference structures.

products offers the opportunity to make some reasonable and relatively non-restrictive assumptions. Before writing down the estimation structure formally, I will sketch out the approach.

Following Feenstra (1994) and Broda and Weinstein (2010), I assume that the elasticities of demand and supply are constant within the product group tier. For example, take the canned fruit product group, containing the canned peach and canned pineapple product modules. This structure allows the within brand-module (different types of canned peaches sold by the same brand) and across brand-module (different types of canned fruit sold by any brand) elasticities to differ, but their values are the same for all brand-modules in the group. Without this type of assumption, each product pair could presumably have a different elasticity of substitution. Holding the elasticities constant within the product group allows for differences in substitutability in clearly distinct groups (e.g. Canned Fruit vs. Tea) while ensuring that there are enough observations per elasticity to make estimation possible.

One challenge to identification concerns the presence of common shocks to all products in a category. Seasonality, changes in taste, changes in cost from common suppliers, or changes in the competitiveness of the industry could all contribute to shocks that change not only the prices of individual products, but also the price level of goods in the category, potentially biasing elasticity estimates. I identify a benchmark product and calculate a log difference-in-difference of all other products' sales shares and prices versus the change in these variables for the benchmark product in order to net out common time trends. Finally, I assume that idiosyncratic demand and supply shocks are uncorrelated.

I now will describe the procedure for estimating the brand-module elasticities, which I will call the 'within' elasticities, more concretely. Allowing s^b_{vt} to denote the share of variety v within brand-module b in group g, I note that

$$s^b_{vt} = \epsilon^b_{vt} \left(\frac{p_{vt}}{P^b_t} \right)^{1-\sigma^g_b} \qquad (.5)$$

where ϵ^b_{vt} is an idiosyncratic demand shifter for variety v at time t. Taking logs and first-differencing,

$$\Delta lns^b_{vt} = \xi^b_t + (1 - \sigma^g_b)\Delta lnp_{vt} + \Delta ln\epsilon^b_{vt} \tag{.6}$$

where $\xi^b_t = (1 - \sigma^g_b)\Delta lnP^b_t$ and $E[ln\epsilon^b_{vt}] = 0$.

The problem of the firm will be developed further below. For the purpose of estimation, I assume that the inverse elasticity of supply in each product group is given by $\omega^b_g \geq 0$, allowing for the possibility that marginal costs are increasing in production.[9] In logged first-differences,

$$\Delta lnp_{vt} = \zeta^b_t + \frac{\omega^b_g}{1 + \omega^b_g}\Delta lns^b_{vt} + \Delta ln\delta^b_{vt} \tag{.7}$$

where $\zeta^b_t = \frac{\omega^b_g}{1+\omega^b_g}\Delta lnExp^b_t$, δ^b_{vt} represents an idiosyncratic shock to marginal cost of variety v at time t, and $E[ln\delta^b_{vt}] = 0$. I restrict $E[\epsilon^b_{vt}\delta^b_{vt}] = 0$: demand and supply shocks are assumed to be uncorrelated. Last, to remove the brand-time fixed effects, I take the first difference of the share and price data of each variety v versus a reference variety v' in the same brand-module. Define $\Delta' lnx_v = \Delta lnx_v - \Delta lnx_{v'}$. Equation (.6) becomes:

$$\Delta' lns^b_{vt} = (1 - \sigma^g_b)\Delta' lnp_{vt} + \Delta' ln\epsilon^b_{vt}, \tag{.8}$$

while Equation (.7) becomes

$$\Delta' lnp_{vt} = \omega^b_g \Delta' lns^b_{vt} + \Delta' ln\delta^b_{vt}. \tag{.9}$$

The intuition of the procedure to follow, as set out in Leamer (1981), is that given a pair of price and quantity observations for one variety, the combinations of demand and supply elasticities that could possibly give rise to these observations are described by a hyperbola.

[9] This approach was developed in Feenstra (1994).

The assumption that these elasticities are constant within the product group allows us to choose the (σ, ω) point where the hyperbolas for each variety intersect. However, equations (.6) and (.7) suffer from bias due to the fact that the errors ϵ_{vt}^{b} and δ_{vt}^{b} are correlated with p_{vt}^{b} and s_{vt}^{b}. The assumption that the parameters to be estimated are constant across varieties in the same brand-module allows for consistent parameter estimates following the approach in Leamer (1981). Exploiting the uncorrelated errors, I take the product of equations (.8) and (.9):

$$\left(\Delta' lnp_{vt}\right)^2 = \theta_1 \left(\Delta' lns_{vt}\right)^2 + \theta_2 \left(\Delta' lnp_{vt}\right)\left(\Delta' lns_{vt}\right) + u_{vt}$$

where $\theta_1 = \frac{\omega_g}{(1+\omega_g)(\sigma_b^a - 1)}$, $\theta_2 = \frac{1-\omega_g(\sigma_b^a - 2)}{(1+\omega_g)(\sigma_b^a - 1)}$, and $u_{vt} = \Delta' ln\epsilon_{vt}^{b} \Delta' ln\delta_{vt}^{b}$, Note that the error u_{vt} is still correlated with the transformed price and share variables. Following Feenstra (1994) by averaging across t within a variety v, I obtain the following equation, the bar denoting the average:

$$\overline{\left(\Delta' lnp_{v}\right)^2} = \theta_1 \overline{\left(\Delta' lns_{v}\right)^2} + \theta_2 \overline{\left(\Delta' lnp_{v}\right)\left(\Delta' lns_{v}\right)} + \overline{u_v} \tag{.10}$$

and estimate by GMM with Hansen (1982) weights, with the number of moment conditions equal to the number of varieties in set Ω_{vb}. A weak condition must hold for any pair of varieties v and v', namely

$$\frac{var(\Delta' ln\epsilon_{vt}^{b})}{var(\Delta' ln\epsilon_{v't}^{b})} \neq \frac{var(\Delta' ln\delta_{vt}^{b})}{var(\Delta' ln\delta_{v't}^{b})}.$$

This restriction rules out the case in which products in the same category have exactly proportional demand and supply shock distributions, which would imply that their $\sigma$and ω estimates were co-linear.

Broda and Weinstein (2006) suggest a modification of Feenstra's procedure in the case where the estimates θ_1 and θ_2 from equation (.10) yield imaginary or non-feasible estimates of the elasticities (i.e. $\sigma < 1$, $\omega < 0$). In these cases, I conduct a grid search in the interval

Table A.1: Demand and Supply Elasticity Estimates

	(1)	(2)	(3)	(4)
	σ_b	ω_b	σ_a	ω_a
Number of Product Groups	61	61	61	61
Number of Significant Estimates	60	59	59	56
Percentiles of Elasticity Dist.				
10th	3.94	0.14	3.31	0.07
25th	4.79	0.28	4.50	0.25
50th	6.08	0.49	6.28	0.87
75th	8.13	0.70	35.48	1.95
90th	17.14	1.79	50.50	4.28

Note: The table describes the within and across elasticity estimates for the categories in which they could be estimated.

$\sigma^g \in [1.05, 131.5]$ at intervals of 0.05 to identify the value that minimizes the GMM objective. In cases in which I am unable to calculate an elasticity that is statistically different from zero, I impute the median elasticity estimate from the rest of the sample. There are 61 product groups; I calculate 60 significant 'within' elasticities and 59 significant 'across' elasticities.[10] The process to estimate the 'across' elasticities is similar, using the weighted average price for items in each brand within the module and the share of each brand instead of individual goods' prices and shares.

Table A.1 reports descriptive statistics on these estimates in columns (1) to (4), as well as the distribution of (5) within brand-module common expenditure ratios $\frac{\lambda_t}{\lambda_{t-2}}_b$ and (6) across brand-module common expenditure ratios $\frac{\lambda_t}{\lambda_{t-2}}_b$ for each group. My elasticity estimates are a little smaller than those in Broda and Weinstein (2010), who use a different set of product groups and report median within and across elasticities of 11.5 and 7.5. In the adjusted price index calculations that follow, I use the baseline estimates in (1) and (3), but also compare them with median elasticity estimates from Broda and Weinstein (2010). Perhaps surprisingly, the 'across' elasticities are slightly larger than the 'within' elasticities. Large elasticities push the gains from variety toward zero.

[10] One of the 'within' elasticities and three of the 'across' are generated using grid search.

However, as columns (5) and (6) demonstrate, most of the gains from variety occur within brand-module, making the 'across' elasticities relatively less important. The median 'within' common expenditure ratio is 0.93, indicating that the share of products that existed in both t and $t-2$ in total annual sales fell by 7 percent in t relative to $t-2$. The median 'across' expenditure ratio is 0.99, indicating that the share of existing brands within the group fell by 1 percent in the two year period. Not surprisingly, new brands arrive less frequently than new products within an existing brand. 90 percent of groups experience increases in variety within brand-modules over a two year period, suggesting that gains will not be driven by one particular type of product.

Calculating the CPI Bias

Next, I apply these estimates to the calculation of city-level price indices according to Equation (.4). Table A.2 reports the distribution of the average annual CPI bias across the 26 cities in the sample. Each city experiences an average annual welfare gain of at least 1.5 percent. The distribution of city-year annual bias estimates is reported in column (1). The variability reflected in the average annual bias is amplified in the cumulative gains for the period 2008-2014 reported in column (2). The median city experiences welfare gains equivalent to a 9.77 percent drop in the price level over the period. The city with the highest welfare gains experiences a cumulative variety adjustment that is 3.9 percentage points larger than the adjustment in the city with the lowest welfare gains.

Hausman (1996) Elasticities

While the Feenstra (1994) technique is frequently used in settings that require the estimation of a large number of elasticities, there are also alternative methods commonly used to estimate elasticities. In this section, I describe a technique based on Hausman (1996). The identifying assumption in this approach is that the price of a UPC does not vary based on

Table A.2: CPI Bias Estimates

	Baseline Elasticity Measures	
	(1)	(2)
	Annual Bias	Cumul. Bias 2007-2014
Percentile		
5th	0.92	8.13
10th	0.96	8.22
25th	1.16	8.77
50th	1.37	9.77
75th	1.60	10.27
90th	1.77	11.27
95th	1.85	11.56

Note: The table reports annualized and cumulative percentage estimates of the bias of the CPI due to changes in product variety.

city-specific cost as well as that there are no national UPC-level demand shocks. As a result, changes in price in other cities can be used as an instrument for shifts in the supply curve. Specifically, the inverse supply function for UPC v belonging to brand b sold in city i at time t is given by

$$ln\, p^b_{v,it} = \delta_b ln\, c^b_{v,t} + \alpha^b_{v,i} + \beta^b_{it} + w^b_{v,it}$$

Notice that cost $c_{vb,t}$ does not vary at the city level. The estimation allows for city-specific, time-invariant differences in cost across cities for the UPC and for city-specific cost shocks at the brand level. The error $w^b_{v,it}$ reflects mean-zero idiosyncratic variation in the price of the variety stemming from local disturbances such as variation in the timing of sales by store managers across cities. The fixed cost structure invites the same double-differencing approach taken above. The change in the price of UPC v in city i is given by

$$\Delta ln\, p^b_{v,it} = \delta_b \Delta ln\, c^b_{v,t} + \Delta\beta^b_{it} + \Delta w^b_{v,it}$$

To remove the brand-time fixed effects, I take the first difference of the share and price data of each variety v versus a reference variety v' in the same brand-module. Define $\Delta' lnx_v = \Delta lnx_v - \Delta lnx_{v'}$.

$$\Delta' ln\, p^b_{v,it} = \delta_b \Delta' ln\, c_{vb,t} + \Delta' w_{vb,it} \tag{.11}$$

No new assumptions are required on consumer demand. As in the Feenstra technique, demand in logged first-differences is given by

$$\Delta' ln s^b_{vt} = (1 - \sigma^g_b)\Delta' ln p_{vt} + \Delta' ln \epsilon^b_{vt}, \tag{.12}$$

The structure of the inverse supply equation (.11) suggests that supply-side changes in the differenced price $\Delta' ln\, p_{vb,it}$ can be identified using the price in other markets as an instrument. Any cmobination of the price in other markets could be a valid vector of instruments, but the computational burden associated with using the price in indiudual cities as an instrument is large in this setting. I adopt the commonly used strategy of computing, for each market, the average price in all other markets. The first stage is given as follows:

$$\Delta' ln\, p^b_{v,it} = \gamma_{b0} + \gamma_{b1} \sum_{j \neq i} \frac{\Delta' ln\, p^b_{v,jt}}{I - 1} + u^b_{v,it}$$

Equation (.12) is the second stage.

The process to estimate the 'across' elasticities is similar, using the weighted average price for items in each brand within the module and the share of each brand instead of individual goods' prices and shares. I report the elasticities calculated using this approach in the table below. In cases in which I am unable to calculate an elasticity that is statistically different from zero, I impute the median elasticity estimate from the rest of the sample. There are 61 product groups; I calculate 53 significant 'within' elasticities and 33 significant 'across' elasticities.

Table A.3 reports descriptive statistics on these estimates in columns (1) to (4). The median 'within' elasticity is 1.99, significantly smaller than the median elasticity calculated using the Feenstra technique. The median 'across' elasticity is 1.35. Because an elasticity closer to one will yield a larger implied welfare gain given the same ratio $\frac{\lambda_t}{\lambda_{t-2}}$, the set of

Table A.3: Demand and Supply Elasticity Estimates

	(1)	(2)
	σ_b	σ_a
Number of Product Groups	60	60
Number of Significant Estimates	53	33
Percentiles of Elasticity Dist.		
10th	1.22	1.07
25th	1.77	1.19
50th	1.99	1.35
75th	2.37	1.51
90th	2.64	1.73

Note: The table describes the within and across elasticity estimates for the categories in which they could be estimated.

elasticities estimated using the Hausman approach will yield larger estimates for each city.

Calculating the CPI Bias

Next, I apply these estimates to the calculation of city-level price indices according to Equation (.4). Table A.4 reports the distribution of the average annual CPI bias across the 26 cities in the sample. Each city experiences an average annual welfare gain of at least 6.6 percent. The distribution of city-year annual bias estimates is reported in column (1). The variability reflected in the average annual bias is amplified in the cumulative gains for the period 2008-2014 reported in column (2). The median city experiences welfare gains equivalent to an 86 percent drop in the price level over the full sample period.

The welfare gains calculated using Hausman elasticities are almost ten times as large as those calculated using the Feenstra method. I choose to use the Feenstra elasticities in the main text because it leads to more conservative estimates of the welfare impact of new varieties and the magnitude of demand shock transmission across cities. Results using the Hausman elasticities would be qualitatively similar though the magnitude of welfare gains transmitted across cities and their business cycle variation would be larger.

Table A.4: CPI Bias Estimates

	Baseline Elasticity Measures	
	(1)	(2)
	Annual Bias	Cumul. Bias 2008-2014
Percentile		
5th	6.6	73.8
10th	7.5	73.9
25th	9.7	78.5
50th	12.1	86.0
75th	15.2	92.0
90th	18.0	106.7
95th	19.4	109.0

Note: The table reports annualized and cumulative percentage estimates of the bias of the CPI due to changes in product variety.

Appendix B: Model Extension, Variable Markups

The model of product choice with multi-city retailers developed in section 4 made the simplifying assumption that firms are small enough that they do not internalize their effect on the price level in making product decisions. However, with a finite number of retailers in each market, it is possible that retailers could be large enough to have significant market power. This section develops a variant of the model that allows for changes in market share to affect the product entry decision made by each retailer. Household preferences remain as described in section 4.1 and the production of grocery and non-grocery goods is as described in section 4.2.2.

The only change will be to the problem of the grocery retailer. I develop the retailer's problem below, allowing the retailer to choose a city-specific markup. I make this choice based on evidence from the literature. Handbury and Weinstein (2015) find that there is relatively little within-retailer variation in the price of each product, despite the fact that costs are presumably higher in large cities, implying that the markup on goods is smaller in large cities. Hottman (2015) rationalizes the decline in markup with city size as the result of greater competition between retailers in larger locations. In my model, costs do not vary at the city level, so allowing the markup to depend on the price at the city level will result in different prices across locations.

Firms

Grocery Retailers

Grocery retailers r choose how many product varieties to sell and choose a markup μ_{ri} to set on all products in each city i order to maximize their total profits subject to a stocking cost $F(n_r)$ with $F'(n_r) > 0$ as in the baseline model. For simplicity, I again assume that the average marginal cost of goods sold at the retailer $\bar{mc}$ is unaffected by the number of products chosen and is common to all retailers. Retailers face total grocery expenditures X_i and a price level P_i in each city. The retailer's problem can be expressed as

$$\max_{\mu_{ri},n_r} \sum_{i=1}^{I} \gamma_{ri} \left(1 - \frac{1}{\mu_{ri}}\right) \left(\frac{n_{ri}^{\frac{1}{1-\sigma_g}} \mu_{ri} \bar{mc}}{P_i}\right)^{1-\sigma_r} X_i - F(n_r) \tag{.13}$$

The retailer's price level is given by

$$P_{ri} = n_{ri}^{\frac{1}{1-\sigma_g}} \mu_{ri} \bar{mc} \tag{.14}$$

and P_i denotes the city grocery price level, defined as

$$P_i = \left(\sum_{r \in R_i} \gamma_{ri} P_{ri}^{1-\sigma_r}\right)^{\frac{1}{1-\sigma_r}}. \tag{.15}$$

For all cities i such that $r \notin R_i$, the taste parameter γ_{ri} is equal to zero. As before, fixed costs are paid nationally, so there is no motive for the retailer to choose an individual product line n_{ri} for each market. However, I allow for a relationship of the form $n_{ri} = \tau_i n_r$ between the global product line and the set of varieties ultimately available in city i. The term τ_i can be interpreted as an additional cost associated with providing varieties to city i if $\tau_i < 1$.

Practically, this term allows the model to rationalize the fact that New York and Los

Angeles experience slightly higher gains from variety than other cities even within retailers, though this is unrelated to changes in demand in those two cities. Setting $\tau_i = 1 \forall i$ will yield the same predictions quantitatively and qualitatively for every city but these two.

The first order condition with respect to the choice of μ_{ri} is standard:

$$d\mu_{ri}: \left(\frac{1}{\mu_{ri}^2}\right) s_{ri} X_i + \left[(1-\sigma_r)\left(1-\frac{1}{\mu_{ri}}\right) s_{ri}(1-s_{ri})\frac{X_i}{\mu_{ri}}\right] = 0 \tag{.16}$$

where $s_{ri} = \gamma_{ri}\left(\frac{n_{ri}^{\frac{1}{1-\sigma_g}} \mu_r \bar{mc}}{P_i}\right)^{1-\sigma_r}$. The firm faces a trade-off between profit-per-unit and its share of the market: increasing the markup increases flow profits per unit, but tends to decrease the firm's overall share. I depart from the previous section in that I assume that firms are small enough that $s_{ri}^2 \approx 0$, i.e. that firms set the monopolistically competitive markup rather than a variable markup based on retailer share.

The first order condition with respect to the choice of n_r is

$$dn_r: \sum_{i=1}^{I}\left[\frac{1-\sigma_r}{1-\sigma_g}\left(1-\frac{1}{\mu_{ri}}\right) s_{ri}\frac{X_i}{n_r} - \frac{1-\sigma_r}{1-\sigma_g}\left(1-\frac{1}{\mu_{ri}}\right) s_{ri}^2\frac{X_i}{n_r}\right] - F'(n_r) = 0. \tag{.17}$$

As noted, increasing the length of the product line increases the retailer's share of demand because the consumer has love for variety: the same utility of consumption is less expensive at a retailer whose grocery bundle includes more varieties. I assume that $F'(0) = 0$ so that all retailers sell products.

Solving these two first order conditions yields a gross markup that varies with the retailer's share of market i

$$\mu_{ri} = \frac{1}{(\sigma_r - 1)(1 - s_{ri})} + 1 \tag{.18}$$

The optimal product line is given by

$$n_r = \frac{1}{(\sigma_g - 1)} \frac{\sum_i \frac{s_{ri} X_i}{\mu_{ri}}}{F'(n_r)}. \tag{.19}$$

Higher marginal stocking costs and greater substitutability between goods both lead to shorter product lines. All else equal, a retailer with larger global sales will have a longer product line. However, if the gross markup in a particular city is high, the retailer will 'discount' the revenues that it receives from that city. From equation (.18) , the gross markup will be larger if retailer substitutability σ_r is low or the share of the retailer s_{ri} is high. There are two reasons why a retailer might have high sales in city i: it may be that total expenditure X_i is large, in which case the incentive to increase the product line to capture a larger share is high, or it may be that retailer r already has a high market share s_{ri} in that city, in which case it faces a reduced incentive to compete.

Retailer Profits

There are aggregate profits in the economy because free entry does not hold in the retail sector. They depend on the size of the retailer and on the stocking cost $F(n_r)$. Given equations (.18) and (.19), the general expression for the retailer's profits π_r is

$$\pi_r = \sum_i \frac{s_{ri} X_i}{\sigma_r (1 - s_{ri}) + s_{ri}} - F(n_r) \tag{.20}$$

Let the stocking cost be given by

$$F(n_r) = \left(\frac{n_r}{n_c} \right)^{\alpha}. \tag{.21}$$

The term in the denominator n_c will allow for trend growth in the number of varieties and accommodate differences in trend growth across the four categories to which food retailers in the data belong: convenience stores, drug stores, grocery stores and mass retailers. The length of the product line under this stocking cost parametrization is given by

$$n_r = \left[\frac{\sum_i \frac{s_{ri} X_i}{\mu_{ri}}}{\alpha(\sigma_g - 1)} \right]^{\frac{1}{\alpha}} n_c \tag{.22}$$

If $\alpha > 1$, the stocking cost is strictly convex in the length of the product line n_r.

Profits can be expressed as a function of the firm's sales, its markup, and parameters. For convenience, define $\sigma_{ri} = \sigma_r(1 - s_{ri}) + s_{ri}$:

$$\pi_r = \sum_{i=1}^{I} \frac{1}{\sigma_{ri}} \left(1 - \frac{\sigma_{ri} - 1}{\alpha(\sigma_g - 1)} \right) s_{ri} X_i \tag{.23}$$

Weakly positive profits for all firms requires that $\alpha > \frac{\sigma_{ri}-1}{\sigma_g - 1}$. Since $\sigma_{ri} \leq \sigma_r \leq \sigma_g$, this condition only requires that the stocking cost not exhibit large increasing returns to scale in the size of the product line. It is always satisfied if the marginal cost of adding a product is constant or increasing.

Market Clearing and Equilibrium

Markets clear as described in section 4.3 and the equilibrium is as described in section 4.4.

Business Cycle Interpretation

In order to understand how the decisions of retailers can transmit productivity shocks across cities, I compare equilibria under a set of city-level productivities $a_{i,t-1}$ and $a_{i,t}$. It is convenient to compare steady states in terms of log changes, where $\hat{x} = dlogx$. In this analysis, I assume that tastes are constant over time: $\hat{\gamma}_{ri} = 0$.

I focus on the impact of city-level shocks $\{\hat{a}_i\}_{i \in I}$ on the length of an arbitrary retailer's product line and markup and therefore on consumer welfare. First, log-linearizing equation (.22) gives an expression for the change in product line length:

$$\hat{n}_r = \frac{1}{\alpha} \sum_i \omega_{ri} (\hat{s}_{ri} + \hat{X}_i - \hat{\mu}_{ri}) + \hat{n}_c. \tag{.24}$$

The change in the length of the retailer's product line not only depends on exogenous changes in demand, but also on the change in its share and therefore markup.

The growth rate of the markup of the retailer in city i is given by

$$\hat{\mu}_{ri} = \frac{s_{ri}}{(1 - s_{ri})\, \sigma_{ri}} \hat{s}_{ri}. \tag{.25}$$

The change in the share is given by

$$\hat{s}_{ri} = (1 - \sigma_r)(\hat{P}_{ri} - \sum_{r' \in R_i} s_{ri} \hat{P}_{r'i}) \tag{.26}$$

Transforming equation (.14), the change in the retailer's global price is given by

$$\hat{P}_{ri} = \frac{1}{1 - \sigma_g} \hat{n}_{ri} + \hat{\mu}_{ri}. \tag{.27}$$

The growth in the price set by each retailer in each city incorporates the cost term τ_i, which enters through $\hat{n}_{ri}$. Note that because $w_i = a_i$ in every city, $\hat{mc} = 0$ in general.

Combining these expressions, the change in the retailer's price can be expressed as a function of parameters and its own demand:

$$\hat{P}_{ri} = \underbrace{\frac{1}{1-\sigma_g}\frac{1}{\alpha}\sum_i \omega_{ri}\left[\left(1-\frac{s_{ri}}{(1-s_{ri})\,\sigma_{ri}}\right)\hat{s}_{ri}+\hat{X}_i\right]+\frac{1}{1-\sigma_g}(\hat{n}_c+\hat{\tau}_i)}_{\text{Change in Product Variety}} \\ + \underbrace{\frac{s_{ri}}{(1-s_{ri})\,\sigma_{ri}}\hat{s}_{ri}}_{\text{Change in Markup}} \tag{.28}$$

Log-linearizing the city grocery price level P_i in equation (.15) gives

$$\hat{P}_i = \sum_r s_{ri}\hat{P}_{ri}. \tag{.29}$$

Notice that in the model with variable markups, demand shocks may lead to changes on the intensive margin of the price level as the markup on all goods rises or falls, along with the impact on the price level coming from new products, a mechanism common to the baseline model.

City-Level Contributions to Retailer Variety

In order to understand how shocks to a_i may be transmitted to other cities through changes in the set of available products, I decompose the change in the retailer price $\hat{P}_r$ into contributions coming from each city in which the retailer operates. I denote the impact of demand in city j on retailer r's price level by $\hat{T}_{rj}$. The full impact of demand in city j on city i's price level is a weighted sum of each contribution $\hat{T}_{rj}$, where the weights are the share of retailer r in city i's expenditure. I describe the derivation of the expression for the impact of city j on city i's demand, denoted $\hat{T}_{ij}$, in what follows.

I use equation (4.29) to decompose equation (4.30) into the contribution of each city j to the change in retailer r's price, denoting this contribution by $\hat{T}_{rj}$:

$$\hat{P}_{ri} = \sum_j \omega_{ri} \hat{T}_{rj} + \frac{s_{ri}}{(1 - s_{ri})\, \sigma_{ri}} \hat{s}_{ri} \tag{.30}$$

where the contribution of city j to the change in the price level of retailer r is given by

$$\hat{T}_{rj} = \frac{1}{1 - \sigma_g} \frac{1}{\alpha} \omega_{rj} \left[\left(1 - \frac{s_{rj}}{(1 - s_{rj})\, \sigma_{rj}} \right) \hat{s}_{rj} + \hat{X}_j \right] + \frac{1}{1 - \sigma_g} \hat{n}_c \tag{.31}$$

The contribution of demand growth in city j to the price level in city i can be expressed as the sum of city j's contributions $\hat{T}_{rj}$ to retailers in set R_{ij}, weighted by the share s_{ri} of each retailer in city i's demand. I denote the total contribution of city j to city i's price level by $\hat{T}_{ij}$. Combining equation (.29) with equations (.30) and (.31), the change in the price level in city i due to contributions to each retailer's price from city j is given by

$$\hat{T}_{ij} = \sum_{r \in R_{ij}} s_{ri} \omega_{rj} \hat{T}_{rj}. \tag{.32}$$

Finally, combining equation (.31) with equation (.32) expresses the connection between demand in city j and city i's price level:[11]

$$\hat{T}_{ij} = \frac{1}{1 - \sigma_g} \sum_{r \in R_{ij}} s_{ri} \omega_{rj} \left(\frac{1}{\alpha} \left[\left(1 - \frac{s_{rj}}{(1 - s_{rj})\, \sigma_{rj}} \right) \hat{s}_{rj} + \hat{X}_j \right] + \hat{n}_c \right) \tag{.33}$$

Equation (.33) expresses an intuitive relationship between demand in city pairs. City j has a larger impact on price level changes in city i whenever common retailers R_{ij} represent a large fraction of consumption in city i (s_{ri} is large), these retailers derive a significant fraction of their revenue from city j (ω_{rj} is large), or demand shocks in city j are particularly

[11] Because $\sum_j \omega_{rj} = 1$, the contribution of each city to trend growth $\hat{n}_c$can be divided proportionally across cities.

significant ($\hat{X}_j$ is large).

Calibrated Model

In this section, I recalibrate the parameters σ_g, α, $\tau_i$and $\hat{n}_c$ as well as the elasticity of substitution σ_r between retailers that governs the relationship between changes in the retailer's price and share as a result of demand shocks. Initial shares $\{s_{ri}, \omega_{ri}\}$ remain as characterized in section 5.1. I estimate an elasticity of substitution across retailers using the technique outlined in Appendix Section 7. I find a value of $\sigma_r = 1.2\,(0.04)$, implying relatively low substitutability across retailers.

Next, I describe how the values of the fixed cost parameter α from equation (.21), elasticity of substitution across goods σ_g, cost parameter τ_i and trend growth rate of product variety $\hat{n}_c$ are chosen. As in section 5.2, I measure the shock to city-level demand as total demand growth in each city-year. Because I relax the assumption that retailers are small and do not internalize their impact on the price level, the choice of variety is itself a function of the change in the retailer's share. I use the weighted average growth in a retailer's markets as an instrument for exogenous growth in the retailer's revenue:

$$\hat{X}_{rt-1} = \sum_i \omega_{rit} \hat{X}_{it-1} \tag{.34}$$

$$\hat{n}_{ri,t} = \beta_1 \left(\hat{X}_{r,t-1} + \sum_i \omega_{ri,t} \left[1 - \frac{s_{ri,t}}{(1 - s_{ri,t})\,\sigma_{ri}} \right] \hat{s}_{ri,t} \right) + \beta_2 (1 - LANYC_i) + \Gamma_c + \epsilon_{ri,t} \tag{.35}$$

$$\hat{s}_{ri} = (1 - \sigma_r)(\hat{P}_{ri} - \sum_{r' \in R_i} s_{ri} \hat{P}_{r'i}) \tag{.36}$$

Table B.1: Parameter Calibration Regression

Parameters	Value
α	6.51
	(0.61)
$\hat{\tau}$	-0.007
	(0.001)
Number of Obs	2,007
Retailers	70
R^2	0.39

Note: The table reports the results of regression equation (.35), which is used to calibrate parameters α and $\hat{\tau}$.

$$\hat{P}_{ri} = \frac{1}{1-\sigma_g}\hat{n}_{ri} + \frac{s_{ri}}{(1-s_{ri})\,\sigma_{ri}}\hat{s}_{ri} \tag{.37}$$

Because $\hat{n}_{ri,t}$ depends on $\hat{s}_{ri}$, which in turn depends on the shocks faced by retailer r relative to other retailers in the set R_i, I estimate the parameters in this system of equations using a two-step numerical procedure. In an inner loop, I solve the system of equations for a given vector of parameters, while the outer loop minimizes mean squared error over the parameter space. Table B.1 reports the parameters τ_i and α calculated using this method. The values are almost identical to those found in Table B.1.

Next, I characterize the contribution of changes in market share to changes in product variety. The growth in product variety at the retailer level depends on a function, denoted $\hat{\gamma}_{r,t}$, of the change in the retailer's share $s_{ri,t}$ in each market:

$$\hat{\gamma}_{r,t} = \sum_i \omega_{ri,t}\left[1 - \frac{s_{ri,t}}{(1-s_{ri,t})\,\sigma_{ri}}\right]\hat{s}_{ri,t} \tag{.38}$$

The changes in share $\hat{s}_{ri,t}$ predicted by the model are small: the minimum change in retailer-city share is -0.11 percent and the maximum is 0.23 percent. For most retailers, equation (.38) suggests that the weighted net effect of changes in shares $\hat{\gamma}_{r,t}$ will be smaller because the coefficient on $\hat{s}_{ri,t}$ will be between zero and one. For retailers with large shares $s_{ri,t}$ such

Table B.2: Parameter Values in Baseline Calibration

		Value
Expenditure and Revenue Shares	$\{s_{ri}, \omega_{ri}\}$	Table I.6
City Trend Growth Adjustment	$\hat{\tau}$	-0.007
Retailer Elasticity of Substitution	σ_r	1.20
Goods Elasticity of Substitution	σ_g	4.09
Stocking Cost Shape Parameter	α	6.51

Notes: The table summarizes the parameter values used in the calibration. The distribution of shares, which implicitly defines γ_{ri}, was discussed previously.

that $s_{ri,t} > 2(1 - s_{ri})\sigma_{ri}$, the net effect of a change in share may not only be negative but have an elasticity greater than one with respect to the change in share. For this subset of firms, the negative effect on variety generated by an increase in market share outweighs the positive effect of an increase in revenue. In practice, the potential for amplification of changes in share is relatively unimportant: $\hat{\gamma}_{r,t}$ is smaller than the weighted change in shares for 97 percent of retailers in the sample.

I also consider the implication of equation (.27) for changes in the price on the intensive margin. The correlation between the model-implied change in the markup and measured changes in city-retailer price is -0.026, suggesting that the small changes in inflation suggested by the model are not important in explaining changes in price on the intensive margin.

The value of all parameters in the model with variable markups are summarized in Table B.2. To evaluate the model's predictions, I use the calibration of the parameters in equations (.24) and (.27) to predict city-level welfare gains based on demand growth in each city. I compare the model's predictions to the welfare gains calculated in section 3.1 using microdata at the UPC level. The predictions of the model are plotted against the gains based on micro data in Figure B.1. Within the period 2008-2010, the correlation between the baseline model and the microdata estimate is 0.75, while within the period 2011-2014 it is 0.86. Overall, the correlation between the baseline model and the microdata estimates is 0.92.

Notice that the fit between the model with variable markups and the gains calculated from microdata is almost exactly the same as the fit between the baseline model and the microdata-based welfare gains. In fact, the correlation between the two models is close to

Figure B.1: Baseline Model Predicted Welfare Gains vs. Microdata-Based Welfare Gains

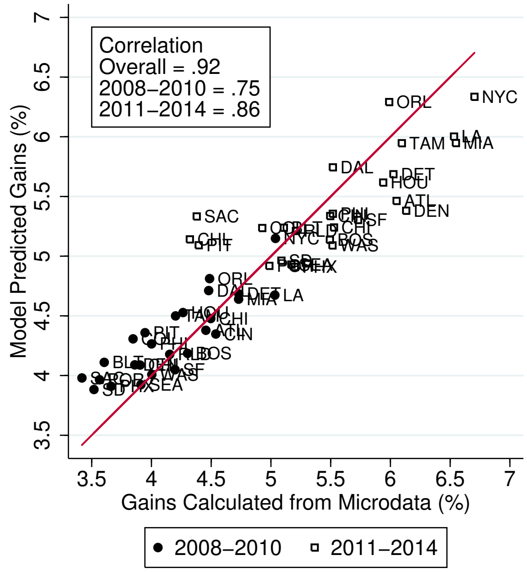

Notes: The figure compares the prediction of the calibrated model to the welfare gains estimated using UPC-level microdata for each retailer.

one, suggesting there are essentially no improvements in the prediction generated by allowing for variation in the amount of product entry generated by changes in market power.

The irrelevance of market power in explaining net product entry is not surprising. Within the subset of cities in the sample, retailers rarely have significant market power, especially considering that the shares as measured in the Nielsen data represent an upper bound on the true market share including non-participating retailers. Furthermore, the model only allows for changes in market share generated by demand shocks. In the data, there may be retailer or retailer-city specific shocks affecting consumer taste. Modeling these shocks could result in a somewhat larger role for changes in market share and markups in explaining changes in product variety, but the finding above that $\gamma_{r,t} < \sum_i \omega_{ri,t}\hat{s}_{ri,t}$ for at least 97 percent of retailers in the data suggests that no change in market share is likely to result in large differences in the predicted growth of variety in this model.

Appendix C: Supplementary Figures and Tables

Figure C.1: Welfare Gains from New Products and City-Level GDP Growth

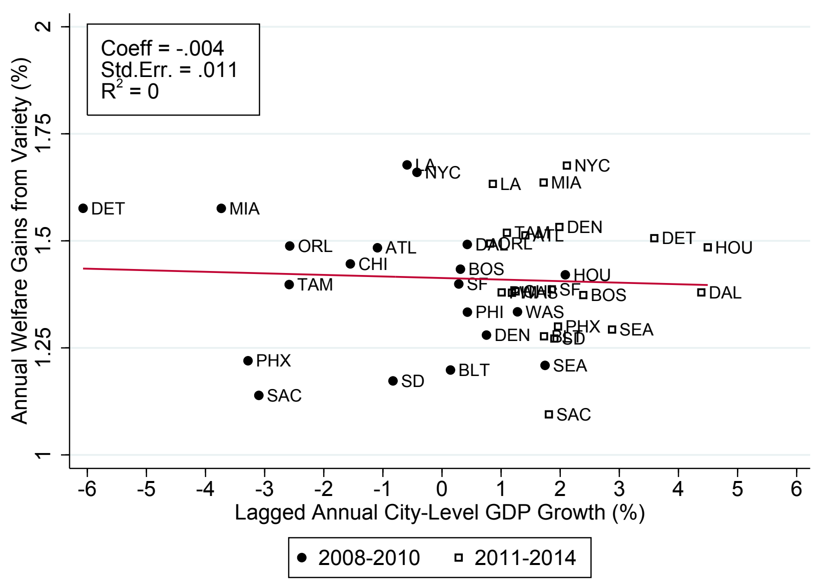

Notes: The figure plots welfare gains derived from a four-tier CES price index for the periods 2008-2010 and 2011-2014 against average annual lagged GDP growth from the Bureau of Economic Analysis.

Figure C.2: Welfare Gains from New Products and House Price Growth

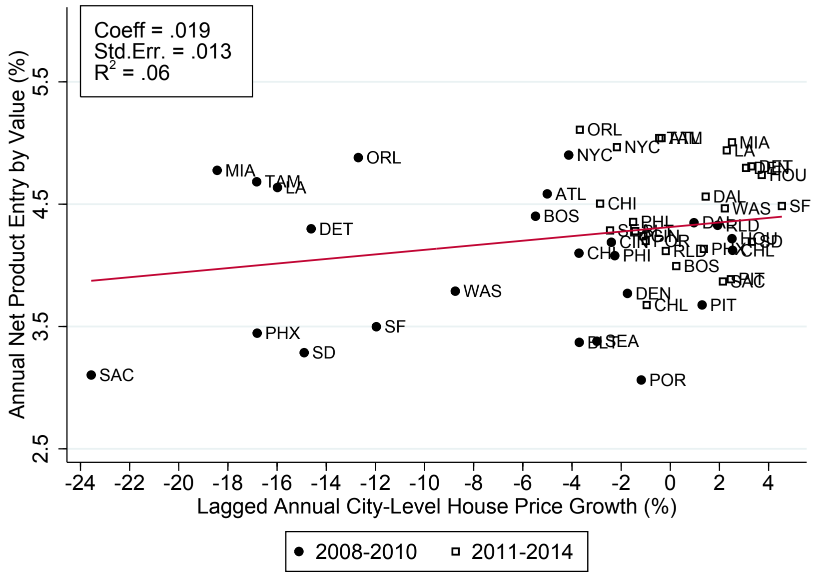

Notes: The figure plots welfare gains derived from a four-tier CES price index for the periods 2008-2010 and 2011-2014 against average annual lagged house price growth from the Federal Housing Finance Agency.

Table C.1: Extensive Margin Entry and Exit Patterns by Number of Initial Markets

Markets 2008	Product Enters All Retailer's Markets (%)	Product Exits All Retailer's Markets (%)	Share of Value (%)
1	100	100	16
2-13	86	72	63
14-25	80	58	3
26	86	64	17

Notes: In the entry (exit) grid, the 'Markets' column is the total number of markets into (out of) which the retailer has introduced (discontinued) the product relative to 2008. The horizontal bins reflect the number of markets into which the product enters within one year (four quarters) of entering its first market. Percentages represent the share of average quarterly revenue of all exiting products over the time period that falls into each bin.

Table C.2: Extensive Margin Entry and Exit Patterns by Number of Initial Markets

Markets 2010	Product Enters All Retailer's Markets (%)	Product Exits All Retailer's Markets (%)	Share of Value (%)
1	100	100	20
2-13	90	74	62
14-25	81	52	4
26	87	59	12

Notes: In the entry (exit) grid, the 'Markets' column is the total number of markets into (out of) which the retailer has introduced (discontinued) the product relative to 2010. The horizontal bins reflect the number of markets into which the product enters within one year (four quarters) of entering its first market. Percentages represent the share of average quarterly revenue of all exiting products over the time period that falls into each bin.

Table C.3: Log Sales by Exporter Experience, 2008-2012, HS2 Products

	(1)	(2)	(3)
New Exporters			
Year 1	-0.267	-0.223	-0.411
	(0.386)	(0.443)	(0.743)
Year 2	-0.0896	-0.0734	-0.247
	(0.556)	(0.588)	(0.746)
Year 3	-0.370	-0.410	-0.596
	(0.730)	(0.724)	(0.663)
Year 4	0.133	0.0974	-0.0872
	(0.711)	(0.738)	(0.853)
Year 5	-0.0761	-0.100	-0.299
	(0.988)	(1.000)	(1.008)
Experienced Exporters			
Year 1	-1.445***	-1.339***	-1.494***
	(0.472)	(0.498)	(0.424)
Year 2	-1.280	-1.202	-1.343**
	(0.945)	(0.951)	(0.613)
Year 3	-0.829	-0.807	-0.960*
	(0.887)	(0.897)	(0.511)
Year 4	-0.658	-0.631	-0.783
	(0.966)	(0.984)	(0.566)
Year 5	-0.657	-0.620	-0.786**
	(0.952)	(0.929)	(0.375)
Fixed Effects	Year	Year, Dest	Year, Dest, HS2
Observations	161,924	161,924	161,924
R-squared	0.001	0.054	0.198

Standard errors are clustered at the firm (NIT) level. ** signifies p<0.01, * signifies p<0.05

Table C.4: Log Sales by Exporter Experience, 2008-2012, HS6 Products

	(1)	(2)	(3)
New Exporters			
Year 1	0.494***	0.590***	0.430***
	(0.152)	(0.149)	(0.150)
Year 2	1.055***	1.119***	0.990***
	(0.161)	(0.164)	(0.158)
Year 3	1.340***	1.357***	1.205***
	(0.154)	(0.156)	(0.153)
Year 4	1.381***	1.402***	1.239***
	(0.155)	(0.155)	(0.153)
Year 5	1.282***	1.314***	1.136***
	(0.161)	(0.162)	(0.159)
Experienced Exporters			
Year 1	-0.235***	-0.145***	-0.199***
	(0.0471)	(0.0459)	(0.0444)
Year 2	0.216***	0.273***	0.250***
	(0.0463)	(0.0450)	(0.0433)
Year 3	0.503***	0.513***	0.467***
	(0.0468)	(0.0460)	(0.0431)
Year 4	0.679***	0.694***	0.637***
	(0.0471)	(0.0459)	(0.0434)
Year 5	0.638***	0.664***	0.591***
	(0.0488)	(0.0475)	(0.0446)
Fixed Effects	Year	Year, Dest	Year, Dest, HS4
Observations	269,605	269,605	269,605
R-squared	0.004	0.050	0.188

Standard errors are clustered at the firm (NIT) level. ** signifies p<0.01, * signifies p<0.05

Table C.5: Extended Gravity Regressions

	Effect on probability of exporting to a new country in group		
	Latin America	European Union	US & Canada
First Destination Group		Without industry controls	
Latin America	0.06***	-0.02	-0.1***
European Union	-0.14***	0.15***	0.13**
US & Canada	0.01	0.13***	-0.21***
First Destination Group		With industry controls	
Latin America	-0.02	0.01	-0.02
European Union	-0.04	-0.09**	-0.18***
US & Canada	0.11***	0.11***	-0.40***
Observations	1,296	1,296	1,296

The table reports results of regressing an indicator for entering a Latin American market, an EU market, or either the US or Canada, conditioning on a subsample of new exporters who do enter some foreign market. The first set of results do not condition on HS2 industry, while the second set conditions on the modal HS2 industry exported by the firm. Asterisks indicate coefficient significance: *** if $p<0.01$, ** if $p<0.05$, * if $p<0.1$.

Bibliography

[1] Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. J., and Li, H. Missing growth from creative destruction. Working paper, 2017.

[2] Albornoz, F., Calvo Pardo, H. F., Corcos, G., and Ornelas, E. Sequential exporting. *Journal of International Economics 88*, 1 (2012), 17–31.

[3] Alessandria, G., Choi, H., and Ruhl, K. Trade adjustment dynamics and the welfare gains from trade. Working Paper, 2015.

[4] Argente, D., and Lee, M. Cost of living inequality during the great recession. Working paper, 2016.

[5] Asdrubali, P., Sorensen, B., and Yosha, O. Channels of interstate risk sharing: United states 1963-1990. *The Quarterly Journal of Economics 111*, 4 (1996), 1081–1110.

[6] Bartelme, D. Trade costs and economic geography: Evidence from the us. Working paper, 2015.

[7] Bernard, A., Redding, S., and Schott, P. Multiple-product firms and product switching. *American Economic Review 100*, 1 (January 2010), 70–97.

[8] Bernard, A. B., Boler, E., Massari, R., Reyes, J.-D., and Taglioni, D. Exporter dynamics, firm size and growth, and partial year effects. *American Economic Review 107*, 10 (2017), 3211–28.

[9] Bernard, A. B., and Jensen, J. B. Why some firms export. *Review of Economics and Statistics 86*, 2 (2006), 561–569.

[10] Bernard, A. B., Moxnes, A., and Ulltveit-Moe, K. H. Two-sided heterogeneity and trade. Working Paper, 2014.

[11] Broda, C., and Weinstein, D. Globalization and the gains from variety. *Quarterly Journal of Economics 121*, 2 (2006), 541–585.

[12] Broda, C., and Weinstein, D. Product creation and destruction: Evidence and price implications. *American Economic Review 100*, 3 (June 2010), 691–723.

[13] Broda, C., and Weinstein, D. E. Globalization and the gains from variety. *Quarterly Journal of Economics 121*, 2 (2006), 541–585.

[14] Caliendo, L., Parro, F., Rossi-Hansberg, E., and Sarte, P.-D. The impact of regional and sectoral productivity changes on the u.s. economy. *Review of Economic Studies* ((forthcoming)).

[15] Coibion, O., Gorodnichenko, Y., and Hong, G. H. The cyclicality of sales, regular and effective prices: Business cycle and policy implications. 993–1029.

[16] Das, S., Roberts, M. J., and Tybout, J. R. Market entry costs, producer heterogeneity, and export dynamics. *Econometrica 75*, 3 (2007), 837–873.

[17] Eaton, J., Eslava, M., Kugler, M., and Tybout, J. Export dynamics in colombia: Firm-level evidence. Working Paper, 2007.

[18] Eaton, J., Kortum, S., and Kramarz, F. An anatomy of international trade: Evidence from french firms. *Econometrica 79*, 5 (2011), 1453–1498.

[19] Erickson, T., and Pakes, A. An experimental component index for the cpi: From annual computer data to monthly data on other goods. *American Economic Review 101*, 5 (2011), 1707–1738.

[20] Feenstra, R. New product varieties and the measurement of international prices. *American Economic Review 84*, 1 (March 1994), 157–177.

[21] Gagnon, E., and Lopez-Salido, D. Small price responses to large demand shocks. Working paper, 2014.

[22] Gil-Pareja, S., Llorca, R., and Serrano, J. A. M. Measuring the impact of regional export promotion: The spanish case. *Papers in Regional Science 87*, 1 (2008), 139–146.

[23] Griffith, R., Leibtag, E., Leicester, A., and Nevo, A. Consumer shopping behavior: How much do consumers save? *Journal of Economic Perspectives 23*, 2 (2009), 99–120.

[24] Handbury, J., and Weinstein, D. Goods prices and availability in cities. *Review of Economic Studies* (January 2015).

[25] Hanson, G. Market potential, increasing returns and geographic concentration. *Journal of International Economics 67*, 1 (2005), 1–24.

[26] Hausman, J. Valuation of new goods under perfect and imperfect competition. *The Economics of New Goods* (1996), 207–248.

[27] Head, K., and Ries, J. Do trade missions increase trade? *Canadian Journal of Economics 43*, 3 (2010), 754–775.

[28] Hottman, C. Retail markups, misallocation, and store variety across u.s. cities. Working paper, 2015.

[29] Jaimovich, N., Rebelo, S., and Wong, A. Trading down and the business cycle. Working paper, 2017.

[30] Leamer, E. Is it a demand curve, or is it a supply curve? partial identification through inequality constraints. *Review of Economics and Statistics 63*, 3 (1981), 319–327.

[31] Morales, E., Sheu, G., and Zahler, A. Extended gravity. Working Paper, 2017.

[32] Moretti, E. Local labor markets. *Handbook of Labor Economics* (2011).

[33] Nevo, A., and Wong, A. The elasticity of substitution between time and market goods: Evidence from the great recession. Working paper, 2015.

[34] Nitsch, V. State visits and international trade. *The World Economy 30*, 12 (2007), 1797–1816.

[35] Redding, S., and Weinstein, D. A unified approach to estimating demand and welfare. Working paper, 2016.

[36] Rose, A. K. The foreign service and foreign trade: embassies as export promotion. *The World Economy 30*, 1 (2011), 22–38.

[37] Ruhl, K. J., and Willis, J. L. New exporter dynamics. *International Economic Review 58*, 3 (2017), 703–726.

[38] Sato, K. The ideal log-change index number. *Review of Economics and Statistics 58*, 2 (1976), 223–228.

[39] Vartia, Y. Ideal log-change index number. *Scandinavian Journal of Statistics 3*, 3 (1976), 121–126.

[40] Volpe Martincus, C., and Carballo, J. Exporting and firm performance: Evidence from a randomized experiment. *Journal of International Economics 76*, 1 (2008), 89–106.

[41] Volpe Martincus, C., and Carballo, J. Entering new country and product markets: does export promotion help? *Review of World Economics 146* (2010), 437–467.