

Equilibrium Price Dispersion and the Border Effect*

Ryan Chahrour
Boston College

Luminita Stevens
FRB Minneapolis and
University of Maryland

December 18, 2015

Abstract

We develop a model of equilibrium price dispersion via retailer search and show that the degree of market segmentation within and across countries cannot be separately identified by good-level price data alone. We augment a set of well-known empirical facts about the failure of the law of one price with data on aggregate intranational and international trade quantities, and calibrate the model to match price and quantity facts simultaneously. The calibrated model matches the data very well and implies that within-country markets are strongly segmented, while international borders contribute virtually no additional market segmentation.

Keywords: Law of one price, Border effect, Real exchange rate

JEL Codes: F41, F30, E30

*We would like to thank Emi Nakamura, Jaromir Nosal, Jón Steinsson, Rosen Valchev, Michael Woodford, and seminar participants at the Federal Reserve Bank of Boston, Boston College, and Columbia University for very helpful comments. We are grateful to João Ayres for excellent research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Contact: ryan.chahrour@bc.edu and stevens@econ.umd.edu.

1 Introduction

Increasing availability of price data at the product level has led to a resurgence of research into the failure of the law of one price across countries. The key feature of the recent literature, compared to the foundational work of [Engel and Rogers \(1996\)](#) and its immediate successors, is that the recent work explicitly considers price dispersion of identical products sold in particular cities or stores, rather than the behavior of aggregate price indices. In addition to circumventing the aggregation biases illustrated by [Broda and Weinstein \(2008\)](#), the use of disaggregated data has provided researchers with a richer set of empirical regularities against which to test potential theories. Following [Gorodnichenko and Tesar \(2009\)](#)'s criticism of reduced-form treatment of cross-border price data, this new empirical evidence has also been accompanied by a structural literature reinforcing the conclusion that high cross-country market segmentation is crucial for generating realistic patterns for good-level violations of the law of one price.

This paper demonstrates the existence of an identification problem in distinguishing between market segmentation that is induced by the international border and background market segmentation that exists between regions or markets within a country. Specifically, we show in the context of a model of retailer search that, even after imposing the structural restrictions of the model, detailed good-level price data between two countries cannot be used to determine whether international price dispersion arises from a friction inhibiting a retailer's search outside of her home region, or a friction that inhibits search outside of her home country. While price data and even data on international trade are insufficient to identify the source of market segmentation, we show that data on *intranational* trade is precisely what is needed to separately identify the parameters governing within- and across-country segmentation.

Our central finding is that the evidence on good-level prices, when combined with information on trade quantities, indicates that virtually all market segmentation occurs at the regional level. To arrive at this result, we calibrate our model of search to match a set of facts regarding good-level violations of the law of one price across the US-Canadian border, as well as both international and intranational levels of trade

between the two countries. We show that the search model is very successful in matching the empirical patterns of violations of the law of one price and can, simultaneously, match the observed levels of trade both within and across Canada and the United States. Matching both sets of facts simultaneously, however, requires a calibration of the economy in which regional markets are strongly segmented, but there is little difference in the segmentation between two regions in the same country and two regions each in different countries: the international border plays no special role in segmenting US and Canadian markets.

In our model, retailers engage in costly sequential search for the best price among producers in the economy. This search friction, combined with a distribution of good-specific and producer-specific productivity shocks, gives rise to endogenous equilibrium price dispersion, consistent with empirical evidence. Retailers search in a world of two countries, each with two regions. We define segmentation as the degree to which buyers' *access* to one market – their ability to sample prices from firms operating in that market – is lower than their access to another market. The model nests regionally and nationally segmented markets. Under regional segmentation, retailers located in a particular region are more likely to sample prices posted by producers located in their own region than in any of the other three regions. However, conditional on not sampling a price from their own region, retailers are equally likely to sample a price from any of the other three regions. Hence, their access to one of the markets located in the foreign country is no more limited than their access to the “away” market within their own country. On the other hand, under national segmentation, retailers in a particular region are more likely to sample prices posted by producers located in either region of their own country than in any of the two regions of the foreign country.

After setting out the model, we examine the degree to which it can match a set of stylized facts about cross-border violations of the law of one price. Specifically, recent empirical work¹ has shown that (1) good-level real exchange rates are more widely

¹For our pricing facts, we focus on the work of [Broda and Weinstein \(2008\)](#), [Gopinath et al. \(2011\)](#), [Burstein and Jaimovich \(2012\)](#), and [Crucini and Telmer \(2012\)](#) all of which compare prices across the US-Canadian border.

dispersed across than within countries; (2) the volatility of good-level real exchange rates is much greater than the volatility of the nominal exchange rate; (3) the aggregate real exchange rate closely follows the nominal exchange rate; and (4) changes in good-level prices are more correlated within countries than across countries. We show that the basic facts about price dispersion can be matched without taking a stand on the degree to which market segmentation is driven by the international border.

Although price data alone cannot identify within-country versus across-country market segmentation, we subsequently add information regarding intranational trade quantities to yield a decomposition of the relative importance of the two possible sources of market segmentation. Using data on state-to-state trade within the United States and interprovincial trade in Canada, we show that trade levels within countries are simply too low, relative to trade across the border, to be consistent with the view that the international border generates a major impediment to trade beyond that already imposed by the barriers to trade that appear across any regions in the economy. In particular, we find that – after controlling for country size – a retail firm in the United States is roughly 7 times more likely to search within its own region compared to the away region within the home country, but the probabilities of searching the neighboring US region and a region in Canada are roughly equal. Regional segmentation appears to be even stronger in Canada, with Canadian firms roughly 14 times more likely to search in their home region relative to their neighboring Canadian region while the international border adds only a trivial degree of segmentation beyond the level of regional segmentation.

Our paper relates to an extensive literature examining violations of the law of one price, which has been surveyed recently by [Burstein and Gopinath \(2014\)](#). [Alessandria \(2004, 2009\)](#) and [Alessandria and Kaboski \(2011\)](#) also use a search friction to motivate cross-border price differences, although these papers do not emphasize the distinction between frictions that occur across markets within countries versus those that occur at the border. Our modeling approach also resembles that of [Gopinath et al. \(2011\)](#) and [Burstein and Jaimovich \(2012\)](#) in that we consider a model with a real friction

in goods markets, coupled with country heterogeneity in the distribution of firm costs. In addition to the papers cited earlier, recent related papers using micro-level price data include [Baxter and Landry \(2012\)](#), [Goldberg and Hellerstein \(2013\)](#), [Fitzgerald and Haller \(2014\)](#), and [De Loecker, Goldberg, Khandelwal, and Pavcnik \(2015\)](#). A related literature, including [Atkeson and Burstein \(2008\)](#), [Drozd and Nosal \(2012\)](#) and [Candian \(2015\)](#), also studies failures of purchasing power parity in more aggregate – and dynamic – contexts.

[Burstein and Jaimovich \(2012\)](#) also introduce within-country regions in their economy and point out an identification problem using prices alone to distinguish between (i) differences in demand-shock correlations across within-country and across-country region pairs and (ii) differences in markup elasticities. However, the problem we highlight is distinct. The forces they discuss are the two forces in their model that can generate *within-country* price dispersion, and map to the sources of *within-region* price dispersion in our economy.² We implicitly break their identification problem by assuming perfect correlation of demand conditions within countries, but show there is nevertheless an identification problem regarding the *source* of market segmentation.

Methodologically, our paper is distinct from the earlier literature because it incorporates price and quantity data simultaneously. An exception is [Boivin, Clark, and Vincent \(2012\)](#), who consider price and quantity data using a good-level data set for a retailer in the online book market and find substantial deviations both within and across countries. Although their exercise is nonstructural, these authors also note a challenge in identifying market segmentation induced by the border from the segmentation that gives rise to price dispersion within countries.

The results in this paper are also related to an extensive literature studying the effects of national and subnational borders using gravity models of trade quantities. [McCallum \(1995\)](#) finds extremely that large effects of the international border are required to account for within- and across-country trades levels between the United States and Canada. [Anderson and van Wincoop \(2003\)](#) show that the estimated border effect

²Prices are dispersed within regions in our economy, but not in [Burstein and Jaimovich \(2012\)](#).

on trade is much smaller once theoretically motivated measures of multilateral resistance are added to the estimation. [Wolf \(2000\)](#) focuses particularly on the effect of intranational borders and finds evidence of substantial barriers to trade. His results are largely corroborated by [Millimet and Osang \(2007\)](#), who argue the effects can be accounted for by historical patterns of migration and past trade. Other related papers along this line include [Chen \(2004\)](#) and [Hillberry and Hummels \(2008\)](#).

Section 2 briefly reviews the existing empirical evidence on good-level violations of the law of one price. Section 3 outlines the search model and discusses the problem of identifying regional versus national market segmentation. Section 4 presents our model calibration and discusses the implications for the sources of segmentation. Section 5 concludes.

2 Pricing Facts

We evaluate the ability of our theory to match a set of facts that have emerged from the recent empirical literature, concerning both the level and growth rates of relative prices at the good level. For consistency, and reflecting the availability of data, we focus on papers that study relative prices across the US-Canada border.

Let $p_{i,t}(l)$ be the log price in local currency of a particular good i at time t in location l . Depending on the context, l may index a region, a city, or a specific store. Let $e_t(l, m)$ be the log of the nominal exchange rate between locations l and m , using the convention that an increase in $e_t(l, m)$ represents an appreciation of the currency in location m . For location pairs within countries, $e_t(l, m)$ is zero. For any two locations, the good-level real exchange rate is defined as

$$d_{i,t}(l, m) \equiv p_{i,t}(l) - p_{i,t}(m) - e_t(l, m). \quad (1)$$

The aggregate real exchange rate is the expenditure-weighted average of changes in good-level real exchange rates across the two locations, i and j :

$$\bar{d}_t(l, m) \equiv \sum_i \omega_i^{l,m} d_{i,t}(l, m). \quad (2)$$

In its strongest form, the law of one price posits that $d_{i,t} = 0$ for all t and i . This is referred to as the *absolute law of one price* and is the focus of [Gopinath et al. \(2011\)](#). A weaker hypothesis is the *relative law of one price*, $\Delta d_{i,t} = 0$ for all t and i , which is the main object of study in [Burstein and Jaimovich \(2012\)](#). One advantage of the relative law of one price is that it strips away price differences that may be due to transportation costs, tariffs, and other fixed physical costs of trading across space. [Broda and Weinstein \(2008\)](#) consider violations of price parity in both levels and changes.

Fact 1: Good-level real exchange rates are more volatile across countries than within. [Burstein and Jaimovich \(2012\)](#) consider the quarterly growth, $\Delta d_{i,t}$, in *wholesale costs* of a single retailer with stores in multiple locations in Canada and in the US. Using a panel of matched goods, they find a standard deviation of roughly 7% in the United States, 6% in Canada, and 13% across countries from 2004 through 2006 depending on the precise set of goods included. Using weekly *retail price* data from the same retailer, [Gopinath et al. \(2011\)](#) find that over the period from 2004 through the middle of 2007, the median standard deviation of $|d_{n,t}|$, measured at the weekly frequency, is approximately 6% within the United States, 4% within Canada, but 24% between the two countries. In contrast, using buyer scanner data aggregated to the city level for the fourth quarter of 2003, [Broda and Weinstein \(2008\)](#) find smaller differences across and within countries: they document standard deviations of 22% within the United States, 19% within Canada, and 27% across the border.³ We view the two data sets as complementary: the former characterizes the distribution of wholesale prices sampled by a particular retailer across different markets and the distribution of prices posted by this retailer; the latter characterizes the distribution of prices sampled by consumers across different markets from all retailers.

Fact 2: Changes in cross-border good-level real exchange rates are significantly more volatile than changes in the aggregate real exchange rate and hence the nominal exchange rate. In [Burstein and Jaimovich \(2012\)](#), this corresponds to a fact about

³[Broda and Weinstein \(2008\)](#) have scanner data on prices paid by households in different cities in the United States and Canada. Their results refer to the price at which a particular good was purchased by a representative household in a given city, rather than the price at which that good was available in a particular store in that city.

relative unit labor costs: they find that the standard deviation of $\Delta d_{i,t}$ across borders is at least three times the standard deviation of changes in relative unit labor costs, which is virtually identical to the standard deviation of changes in the exchange rate over that period.

Fact 3: The aggregate real exchange rate closely follows the nominal exchange rate. This well-known fact has been reconstructed from micro-level data in various forms. [Burstein and Jaimovich \(2012\)](#) document a high correlation between changes in relative unit labor costs and the expenditure-weighted average of changes in good-level real exchange rates, $\Delta \bar{d}_t$, across the United States-Canada border over the period from 2004 through 2006, when variation in relative labor costs was almost entirely driven by the nominal exchange rate. Using the same data source, [Gopinath et al. \(2011\)](#) construct the median value of $d_{n,t}$ across products for each period t and show that this time series tracks the nominal exchange rate almost perfectly from 2004 through the middle of 2007. Similarly, [Broda and Weinstein \(2008\)](#) provide evidence that, controlling for the distance between markets, the United States-Canada expenditure weighted real exchange rate tracks the nominal exchange rate over the period 2001 through 2003.

Fact 4: Changes in good-level prices (in a common currency) are more correlated within countries than across countries. [Burstein and Jaimovich \(2012\)](#) find that, within countries, United States dollar denominated price changes have a correlation of 75% for the United States and 85% for Canada, while the correlation across countries is close to zero.

These facts have been taken as evidence that markets are segmented internationally, either completely or partially. In the next section, we lay out a theory that matches these facts without requiring any more segmentation across countries than within countries.

3 Model

We develop a model in which a search friction gives rise to endogenous price dispersion in the tradable good sector. The tradable sector consists of a continuum of good types,

each of which is produced by many producers with heterogeneous marginal costs. For each good type, retailers search among the set of producers, paying a fixed cost each time they wish to sample a price from the distribution of supplying firms.⁴ After settling on a supplier for the good type, retailers costlessly differentiate their supply of the purchased good into a retailer-specific variety and sell it to domestic consumers at a fixed markup.

The search friction, combined with heterogeneous producer costs and imperfectly elastic demand, yields a nondegenerate distribution of prices. We first present the basic setup, similar to [Reinganum \(1979\)](#), in which producers, retailers, and consumers operate in a single market, which we call a region. A distinguishing feature of the environment we adopt here is that it generates price dispersion within regions. We then extend the model to a two-region economy and then finally to a two-country world with two symmetric regions in each country.

3.1 The Single-Region Economy

Consumers: The representative consumer buys tradable and nontradable goods and supplies labor. For simplicity, we assume that consumers cannot borrow or save. The consumer’s problem is

$$\begin{aligned} \max_{C_T, C_N, L} \quad & \log \left(C_T^\phi C_N^{1-\phi} \right) - v(L) \\ \text{s.t.} \quad & P_T C_T + P_N C_N \leq wL + \Pi, \end{aligned} \tag{3}$$

where C_N denotes consumption of an aggregate nontradable consumption good, P_N its price, C_T denotes aggregate consumption of final goods from the tradable sector, P_T is the corresponding price index, L is the labor supply, w is the economy-wide nominal wage, Π is firm profits, and the disutility of labor $v(\cdot)$ satisfies standard assumptions. In the single region model, the components of the aggregate trade good will always be

⁴ We are motivated to focus on the problem of search at the retailer level by the evidence of [Gopinath et al. \(2011\)](#) that retail-level markups are stable, of [Goldberg and Hellerstein \(2013\)](#) that retail-level price changes always follow changes in wholesale costs, of [Berger et al. \(2012\)](#) that overall “distribution wedges” are stable over time, and the finding of [Hillberry and Hummels \(2003\)](#) that intranational trade bias is especially strong for wholesalers, who are the direct suppliers of retail outlets.

domestically produced but are subject to a search friction; in the multicountry model, these components may also be imported by retailers.

Consumers take the aggregate price indexes P_N and P_T as given, and expenditure is allocated across tradable and nontradable goods according to

$$\frac{P_T C_T}{P_N C_N} = \frac{\phi}{1 - \phi}. \quad (4)$$

The consumer thus spends a fixed fraction $\phi \in (0, 1)$ of total income on the purchase of tradable goods and a fraction $1 - \phi$ on the purchase of nontradable goods.

The tradable goods market is monopolistically competitive, with a continuum of varieties ν in each of a continuum of goods i . At each level, consumption is aggregated according to a constant elasticity aggregator,

$$C_T = \left(\int_0^1 c(i)^{\frac{\rho-1}{\rho}} di \right)^{\frac{\rho}{\rho-1}}, \quad (5)$$

$$c(i) = \left(\int_0^1 c(i, \nu)^{\frac{\eta-1}{\eta}} d\omega \right)^{\frac{\eta}{\eta-1}}, \quad (6)$$

where $\rho > 1$ is the constant elasticity of substitution between goods and $\eta > 1$ is the constant elasticity of substitution between the varieties of each good. The associated price indices are $P_T \equiv \left(\int_0^1 p(i)^{1-\rho} di \right)^{\frac{1}{1-\rho}}$ and $p(i) \equiv \left(\int_0^1 p(i, \nu)^{1-\eta} d\nu \right)^{\frac{1}{1-\eta}}$. We assume that the nontradable sector is perfectly competitive, with a homogeneous nontradable good and representative firm in that sector.

Following the standard derivation, demand for variety ν of the tradable good i is given by

$$c(i, \nu) = p(i, \nu)^{-\eta} p(i)^{\eta-\rho} P_T^\rho C_T. \quad (7)$$

Retailers: A unit mass of multiproduct retailers, indexed by ν , produce varieties by costlessly differentiating tradable intermediate goods. Each intermediate good is produced by a continuum of producers with heterogeneous costs, and retailers purchase these intermediate goods from producers subject to a search friction.⁵ Retailers know

⁵The multiproduct market structure for retailers is drawn from [Bernard, Redding, and Schott \(2011\)](#).

the distribution of prices posted by the producers of each good, but they do not know which producer sells at what price. Instead, they pay a fixed cost each time they wish to sample from the distribution of producer prices for each good. Each retailer then chooses either to purchase all of its demand for the good at the sampled price or to continue searching by paying the search cost and drawing a new price. Search continues until each retailer has settled on a single supplier for each good.

Having settled on a supplier for each good, the retailer differentiates its supply of good i into a retailer-specific good-variety, indexed by (i, ν) . The retailer is a price setter with elastic demand in the market for each good-variety and chooses a price schedule $p(i, \nu)$ to maximize total per-period profits

$$\pi^R(\nu) = \int_0^1 \pi^r(i, \nu) di \quad (8)$$

where good-level profits, $\pi^r(i, \nu)$, are given by

$$\pi^r(i, \nu) = [p(i, \nu) - \widehat{p}(i, \nu)] c(i, \nu) - \kappa n(i, \nu). \quad (9)$$

In equation (9), $\widehat{p}(i, \nu)$ is the producer price upon which the retailer settles after completing search for good i in the period, κ is the cost of search per producer searched and fixed in units of the nontraded good, and $n(i, \nu)$ is the total number of producers of good i that the retailer samples in the period. Since the search process is independent across goods and each retailer's variety competes with the varieties of the same good from other retailers, the retailer does not internalize any impact of changing the price of good-variety (i, ν) on demand for its variety of other goods, (i', ν) , and therefore solves its problem by maximizing expression (9) good-by-good.

The sequential nature of search implies that a retailer's choice to continue looking for a better price for good i is independent of the number of good i producers already sampled in the period. For each good, given the currently sampled producer, j_i^ν , and the current distribution of producer prices, f_i , the retailer's value function is

$$V(j_i^\nu; f_i) = \max\{V^s(f_i), V^{ns}(j_i^\nu)\}, \quad (10)$$

where $V^s(\cdot)$ denotes the value of continuing to search for good i and $V^{ns}(\cdot)$ denotes the value of stopping the search and purchasing all demand from the currently sampled

producer. The retailer continues to search as long as $V^s(f_i) \geq V^{ns}(j_i^\nu)$. The value of continuing to search for a supplier is

$$V^s(f) = E_f[V(j; f)] - \kappa, \quad (11)$$

where expectations are taken over the current distribution of producer prices. The value of halting the search is

$$V^{ns}(j_i^\nu) = \max_p \pi^r(j_i^\nu, \omega). \quad (12)$$

Since differentiation of each good at the retail level is costless, the retailer maximizes expression (12) by charging a constant markup over the producer price, $p = \mu \hat{p}$, where $\mu \equiv \eta / (\eta - 1)$.

The optimal search strategy is a stopping rule described by a unique reservation price $\hat{p}_r(i)$ for each good i that sets $V^s(f_i) = V^{ns}(j_i^\nu)$ for all retailers. All retailers sampling a price less than or equal to $\hat{p}_r(i)$ stop search and purchase all their demand for good i , $c(i, \nu)$, at the sampled price; all retailers sampling a price above $\hat{p}_r(i)$ continue to search for a better offer.

Producers: For each good i , there is a unit mass of producers j characterized by the production function

$$y(i, j) = A(i)\zeta(i, j)l(i, j), \quad (13)$$

where $A(i)$ is good-specific productivity, $\zeta(i, j)$ is good- and producer-specific productivity, and $l(i, j)$ is the labor input. The marginal cost of producer (i, j) is $mc(i, j) = \frac{w}{A_i \zeta(i, j)}$. The distribution of marginal costs for producers of good i therefore depends on the aggregate price of labor as well as on both good-level and producer-level shocks. The cumulative distribution of marginal costs across producers of good i is denoted as $G_i(mc)$.

Let $\chi(i, j)$ denote the mass of retailers who settle on producer j in their search for a supplier of good i . Then

$$D(i, j) \equiv \chi(i, j) \mu^{-\eta} \hat{p}(i, j)^{-\eta} p(i)^\eta c(i) \quad (14)$$

gives the total demand from all the retailers searching for good i who settle on producer j and a producer setting price $\widehat{p}(i, j)$ faces demand

$$x(i, j) = \begin{cases} D(i, j) & \text{if } \widehat{p}(i, j) \leq \widehat{p}_r(i) \\ 0 & \text{if } \widehat{p}(i, j) > \widehat{p}_r(i) \end{cases}. \quad (15)$$

Equilibrium: An equilibrium in the producer-retailer market for good i is a retail reservation price $\widehat{p}_r(i)$ and a distribution of producer prices f_i such that (a) given f_i , retailers choose the optimal stopping rule governed by $\widehat{p}_r(i)$ and (b) given $\widehat{p}_r(i)$, producers maximizing profits generate f_i . The optimal price⁶ set by good i producers and the resulting cumulative distribution of producer prices are given by

$$\widehat{p}(i, j) = \min \{ \mu mc(i, j), \widehat{p}_r(i) \} \quad (16)$$

and

$$F_i(\widehat{p}) = \begin{cases} G_i(\frac{\widehat{p}}{\mu}) & \text{if } \widehat{p} \leq \widehat{p}_r(i) \\ 1 & \text{if } \widehat{p} > \widehat{p}_r(i) \end{cases}. \quad (17)$$

In the single-market model, the retailers' stopping rule implies that there is no search in equilibrium, since all producers post prices that are weakly below the retailers' reservation price. However, because demand below the reservation price is elastic, there is no [Diamond \(1971\)](#) paradox: producers with heterogeneous costs do not find it optimal to generate a single price equilibrium at $\widehat{p}_r(i)$ and, consistent with empirical evidence, the model generates equilibrium price dispersion. The degree of dispersion in prices is determined by the cross-sectional dispersion of producer costs, but the out-of-equilibrium threat of search leads to incomplete pass-through of marginal cost for high-cost producers: markups are constant for all good i producers with marginal costs less than $\widehat{p}_r(i)/\mu$, and are decreasing in cost for producers with costs larger than this threshold.

⁶For simplicity, our exposition ignores the shutdown choice of firms. Our numerical implementation of the model economy, however, ensures the shutdown of any firm that would otherwise earn negative profits and accounts for the effects of firm shutdown on the mass of searchers arriving at active firms.

3.2 The Two-Country Model

We now extend the model to a two-country setup, in which there are two regions in each country: a and b in the home country, H , and c and d in foreign country, F . We introduce asymmetry across countries along three dimensions: (i) the aggregate real wage, and therefore average marginal costs, may differ across countries; (ii) the realizations of good-specific productivity shocks may differ across countries, creating differences in average good-level costs even after controlling for wages; and (iii) the cross-sectional dispersion of idiosyncratic cost shocks are country-specific. In order to bring the model to the data, we also allow for the possibility that countries differ in size, which affects the search process and the probability of matches within and across markets. Importantly, we assume that the wage rate, good-specific productivities, and the distribution of idiosyncratic shocks are the same across regions within countries.⁷

With the introduction of regional and national borders, we now assume that retailers in each region sample producers in their region and other regions at potentially different rates. Specifically, let $f_i^r(\hat{p})$ denote the distribution of prices for good i in region r , and let $\omega(r, k)$ denote the relative region function, mapping region r to its k -th closest neighboring region.⁸ Moreover, let α_k^r represent the probability that a retailer in region r samples a firm in the r -relative region k . During search, a retailer in region a has probability α_1^a of drawing a price from the distribution of good i producer prices posted in its own region, $f_i^{\omega(a,1)} = f_i^a$, a probability α_2^a of drawing a price from the distribution in the neighboring region, $f_i^{\omega(a,2)} = f_i^b$, and so on. A retailer in region r therefore samples prices for good i from the distribution

$$f_{i,r}^{ret}(\hat{p}) \equiv \alpha_1^r f_i^{\omega(r,1)}(\hat{p}) + \alpha_2^r f_i^{\omega(r,2)}(\hat{p}) + \alpha_3^r f_i^{\omega(r,3)}(\hat{p}) + \alpha_4^r f_i^{\omega(r,4)}(\hat{p}), \quad (18)$$

with $\sum_k \alpha_k^r = 1$, $k \in \{1, 2, 3, 4\}$, and where $f_i^{\omega(r,k)}$ is the density of prices posted by the good i producers of the r -relative region k .

⁷The assumption that markets within the same country are symmetric is supported by the evidence in [Gopinath et al. \(2011\)](#) that price differentials are centered around zero within countries, and in [Burststein and Jaimovich \(2012\)](#) that average changes in relative prices within countries are zero.

⁸In our notation, we will let letters represent a region's *absolute* index, while integers represent each region's *relative* index. Thus, for example, $\omega(a, 1) = a$, $\omega(a, 2) = b$, and $\omega(b, 2) = a$. The relative indexes $k \in \{3, 4\}$ always index regions in the away country.




Figure 1: The two-country setup.

To incorporate a notion of country size, let s_r denote the mass of consumers, retailers, and producers in region r . We assume that the probability α_k^r that a retailer from region r samples a producer from its k -th neighbor depends on the size of the sampled region as well as a *pure bias* parameter β_k^r of the retailer toward that particular market, such that $\sum_k \beta_k^r = 1$, $k \in \{1, 2, 3, 4\}$. In the case of symmetrically sized markets, $\alpha_k^r = \beta_k^r$, while completely unsegmented markets imply $\beta_k^r = \frac{1}{4}, \forall r, k$. We assume the sampling probabilities are proportional to the country size, so that the probability of sampling from a region is given by

$$\alpha_k^r = \frac{\beta_k^r s_{\omega(r,k)}}{\sum_l \beta_l^r s_{\omega(r,l)}}, \quad l \in \{1, 2, 3, 4\}. \quad (19)$$

Figure 1 summarizes the general multicountry setup and displays relative search probabilities from the perspective of region a . Retailers in a particular region are now more likely to sample prices from their own region if their own region is relatively large or if they have a sampling bias $\beta_1^r > 1/4$ toward their own region. An important baseline case for us will be one in which $\beta_1^r > \beta_2^r$ but $\beta_2^r = \beta_3^r = \beta_4^r$. In this case, regional sampling bias implies a relatively high weight on domestic prices, even though it is unrelated to segmentation at the international border.

It is important to note that the notion of segmentation used earlier applies to one retail market with respect to producers in another market and does not imply symmetric frictions in crossing the border; it may be the case that retailers in region a import

easily from region b , while at the same time region b retailers encounter high frictions in importing from region a ($\alpha_2^a \gg \alpha_2^b$). For simplicity, we take the regional sampling bias parameters as exogenous and common across goods. They capture all the frictions and barriers to trade, either bilateral or unilateral, that may make transacting across regions less likely and may be motivated by informational advantages that ease access to the chain of production in one's own market, or by external barriers that make transacting with firms located outside one's own network more difficult. The exogeneity assumption can be relaxed, as long as the sampling bias remains independent of relative prices in the two regions. Otherwise, each retailer's problem is the same as in the single-region setting.

A crucial consequence of the supposed symmetry between regions, the distributions of prices at the producer, retailer, and consumer level for any good i are identical across the two regions within each country and, even if region-specific price distributions are available, they cannot be used to infer the degree of regional segmentation. This outcome foreshadows the challenge of separately identifying regional versus national segmentation in the two-country model. As shown later, pricing statistics generate an estimate of the overall segmentation between countries. Since the absence of regional shocks implies that regional segmentation cannot be identified using price data alone, we will show that an estimate of overall segmentation based only on price data will confound regional and national barriers.

In principle, both sizes and regional biases may differ across all four regions. To simplify the exposition and to reduce the number of free parameters in the quantitative analysis, we make the following normalizations. First, we assume that the two regions of each country are of equal size, $s_a = s_b \doteq s_H$ and $s_c = s_d \doteq s_F$. Second, we assume that bias parameters to relative regions within each country are identical, so that $\beta_k^a = \beta_k^b$ and $\beta_k^c = \beta_k^d$ for all k . Finally, we assume that retailers located in a particular country are symmetrically biased against both foreign regions: $\beta_3^r = \beta_4^r \doteq \beta_H^*$ for $r \in \{a, b\}$ and $\beta_3^r = \beta_4^r \doteq \beta_F^*$ for $r \in \{c, d\}$. The restrictions on the bias parameters reduce the set of free parameters to four, two for each country. In section 3.5, we show how the

remaining degrees of freedom correspond to within and across segmentation parameters for each country.

Because of cross-country heterogeneity, good i producers in each country generate different producer price distributions for good i , $f_i^a \neq f_i^c$, and retailers in each country may have different reservation prices, $\widehat{p}_{r,a}(i) \neq \widehat{p}_{r,c}(i)$. Differences in reservation prices arise across countries because of the three sources of cross-country asymmetry described earlier. Without loss of generality, let the home country be relatively less expensive, with $\widehat{p}_{r,a}(i) < \widehat{p}_{r,c}(i)$. In equilibrium, all producers in both countries post prices that are weakly lower than the high reservation price, $\widehat{p}_{r,c}(i)$. However, now the distribution of prices may contain mass points at one or both reservation prices, $\widehat{p}_{r,a}(i)$ and $\widehat{p}_{r,c}(i)$. High-cost producers in either country who post prices between the two reservation prices only sell to retailers in the foreign regions c and d . Conversely, producers charging prices weakly below $\widehat{p}_{r,a}(i)$ sell to retailers in both countries at a single monopoly price that takes into account the different price levels and relative demand in the two countries.

Unlike in the single-region version of the model, retailers may now repeat their search in equilibrium, and the demand faced by producers in the low-price country therefore depends on the probability that retailers do repeat their search. We leave a detailed derivation of the demand functions faced by producers to the appendix.

3.3 Exchange Rate Determination

We make assumptions that permit a simple model of the link between exchange rates and real labor costs. In particular, we assume that money demand follows a standard velocity equation, with fixed velocity normalized to one. Under these assumptions, $P_T^H C_T^H + P_N^H C_N^H = M^{H,s}$, $P_T^F C_T^F + P_N^F C_N^F = eM^{F,s}$, and

$$e = \frac{M^{H,s} P_T^F C_T^F + P_N^F C_N^F}{M^{F,s} P_T^H C_T^H + P_N^H C_N^H},$$

where $P_T^k C_T^k + P_N^k C_N^k$ gives the common-currency value of total output in country $k = H, F$ and an increase in the nominal exchange rate e represents a depreciation of the home currency.

In order to capture the phenomenon of sticky nominal wages in our static setting, we assume that wages in each country incompletely adjust each period to their flexible wage value. That is, we assume that the prevailing wage in country C is given by

$$w^C = \mu_w w_{fix}^C + (1 - \mu_w) w_{flex}^C, \quad (20)$$

where w_{fix}^C represents the nominally fixed wage that is state invariant and w_{flex}^C represent the wage given by the household first order condition for labor:

$$w_{flex}^C = v'(L) C_T^C P_T^C / \phi. \quad (21)$$

By assuming values for ω_w near to one, shocks to relative money supply $\frac{M^{H,s}}{M^{F,s}}$, which we assume follows a persistent AR(1) process, generate both substantial fluctuations in the exchange rate and differences in real unit labor costs between countries, consistent with the evidence of [Burstein and Jaimovich \(2012\)](#).

3.4 Model Intuition

To demonstrate intuition for the model, we illustrate the properties of the model in three different settings with different parameterizations of search and relative marginal costs. Figure 2 plots the pricing policies and distribution of prices posted in each country for a version of the model in which markets are segmented ($\beta_1^r > 1/4$) but there are no cross-country differences in aggregate or product-specific marginal costs. In this case, the two-country model behaves identically to the single-region model, with a single reservation price and identical price-posting distributions in the two countries. In this example, the mass point at the single reservation price is represented by the vertical line in panel (b) of the figure, indicating that roughly 70% of firms post at the reservation price.

Next, figure 3 illustrates the case in which the home country experiences low relative unit labor costs ($w_H < w_F$), and the dispersion of marginal costs are asymmetric, but retailer search places equal weight on all regions ($\beta_k^r = \frac{1}{4}$). In this case, the distributions of prices posted by producers are different across the two countries, but the symmetry of search once again implies a single reservation price in equilibrium. In this case, the




Figure 2: Pricing functions and price distributions for symmetric, segmented economies.




Figure 3: Pricing functions and price distributions for unsegmented economies with different average and dispersion of marginal costs.

common-currency tradable goods price faced by consumers will be identical across the two countries.

Finally, figure 4 shows the case in which the two countries experience different unit labor costs ($w_H < w_F$) and face regional bias in search. In this case, the reservation price in the foreign country is substantially higher than in the home country because producers with high costs attempt to capitalize on “trapped” foreign retailers, rather than set a lower price that appeals to retailers in both regions. Under some param-




Figure 4: Pricing functions and price distributions for segmented economies with different average and dispersion of marginal costs.

eterizations, this may even be true in the low-cost home country to the extent that some retailers from the *foreign* country sample first in the home country and find it worthwhile to pay the higher reservation price rather than search again. In this case, the prices sampled by retailers - and, by extension, the tradable price index faced by consumers - differ across the two countries.

Figures 2 through 4 demonstrate an important point: in order to observe pricing-to-market (namely, firms with identical marginal costs charging different prices), markets must be at least partially segmented and also experience some asymmetry, either in country average costs or in the distributions of idiosyncratic productivities. Without both segmentation and differences between markets, firms with the same cost will charge the same price regardless of their location.

Finally, figure 5 breaks down the profit-maximizing policy for the home producers in the case of partial segmentation, showing profits as a function of marginal costs, given different pricing policies. As shown in panel (a), when marginal cost is less than $\hat{p}_{r,H}/\mu$ - the threshold set by the low reservation price - producers charge the monopoly price (indicated by the blue line) to retailers in both countries. At higher marginal costs, once the desired monopoly price exceeds the home reservation price, producers charge the home reservation price (green line, panel b), thereby maintaining market share in




Figure 5: Profit as a function of marginal cost under different pricing policies. When marginal cost is below p_r^H/μ , foreign firms charge the two-country monopoly price (panel a). When marginal cost is between p_r^H/μ and c^* , foreign firms charge the home reservation price (panel b). When marginal cost is between c^* and p_r^F/μ , foreign firms charge the foreign monopoly price (panel c). When marginal cost is greater than p_r^F/μ , foreign firms charge the foreign reservation price (panel d).

the home country. Under this policy, the producer more than makes up in volume from home retailers what he loses in pricing from foreign retailers. Marginal costs eventually reach a high enough critical point, c^* , where the producer no longer finds it worthwhile to keep selling to home retailers and forgo the profits of charging foreign retailers a higher price; instead, he starts charging the foreign monopoly price (purple line in panel c). Finally, with high enough marginal costs, the producer simply charges the foreign reservation price (red line, panel d) or drops out of the market. In this setting, retailer markups remain constant, while producer markups are heterogeneous, with a mass point at the maximum producer markup, μ .

3.5 Identification

In this section, we demonstrate that market segmentation within- and across-countries cannot be separately identified in our model using data on prices alone. The set of statistics regarding within and across country price dispersion, which have often been taken in the literature to be an indicator of strong market segmentation at the border, cannot tell us whether markets are indeed segmented at the international border or whether, instead, all regional markets are simply isolated from one another. After demonstrating the identification problem in the model, we then illustrate how introducing quantity data, in particular data on within-country trade levels, can resolve the identification problem.

3.5.1 Price Distributions and Segmentation

Under the assumptions made earlier, the distributions of both posted prices and reservation prices are the same within regions of the same country. Consider the distribution of prices sampled by a retailer from region a . Within-country symmetry then implies that the distribution of prices sampled depends only on the probability that a given price draw of a retailer from region a returns a price from the home distribution. Call this probability P_H^a . Imposing the symmetry of search probabilities across foreign regions

assumed in section 3.2, we have that⁹

$$\begin{aligned}
P_H^a &= \frac{(\beta_1^a + \beta_2^a)s_H}{(\beta_1^a + \beta_2^a)s_H + \left(\frac{1-\beta_1^a-\beta_2^a}{2}\right)s_F} \\
&= \frac{s_H}{s_H + \left(\frac{1-\beta_1^a-\beta_2^a}{\beta_1^a+\beta_2^a}\right)s_F/2} \\
&= P_H^b.
\end{aligned} \tag{22}$$

From equation (22), it follows that the probability that any retailer from home draws a home price depends only on the combination

$$\chi_H \equiv \left(\frac{1 - \beta_1^a - \beta_2^a}{\beta_1^a + \beta_2^a}\right). \tag{23}$$

Crucially, χ_H does not load independently on the parameters β_1^a and β_2^a but depends only on their sum. A similar argument applies to χ_F . Moreover, it is only through their effect on χ_H and χ_F that the bias parameters β_k^i influence the set of prices sampled by retailers, and therefore it is only through these values that the bias parameters can affect the pricing decisions of firms or influence quantities at the country level. Data on the distribution of prices and aggregate trade levels can therefore be used to identify χ_H and χ_F , but cannot distinguish the source of segmentation between bias toward the home region and bias toward the home country.

3.5.2 Quantities and Segmentation

While price distributions and aggregate trade quantities cannot pin down the degree of regional versus national segmentation, patterns of trade within countries can. To see this, compute the aggregate share of total tradable good demand in region a that is satisfied by region- a producers,

$$\begin{aligned}
\phi_a &= \int_0^1 \frac{\beta_1^a D^a(j)}{(\beta_1^a + \beta_2^a) D^a(j)} dj \\
&= \frac{\beta_1^a}{\beta_1^a + \beta_2^a},
\end{aligned} \tag{24}$$

⁹In our numerical implementation, firms earning negative profits are assumed to shut down, leading to good-specific measures of active firms in each country. The derivations that follow are extended in the appendix to incorporate this detail.

Within-country symmetry again implies that $\phi_a = \phi_b \equiv \phi_H$. An identical derivation applies to the fraction of domestically serviced demand in the foreign country, $\phi_c = \phi_d \equiv \phi_F$. Importantly, the fraction of demand satisfied in the home region is constant and does not depend on $D^a(j)$, the total regional demand for tradable good j , leaving a simple mapping from internal trade shares ϕ_H and ϕ_F and the bias parameters. Given values for χ_H and χ_F , internal trade shares ϕ_H and ϕ_F can therefore be used to separately identify the sources of market segmentation affecting the economy.

4 Results

In this section, we calibrate the two-country model to match the basic facts on within- and across-country real exchange rates described in section 2, as well as international trade shares. Since most of this earlier data was collected in the years prior to the financial crisis, and since the most recently available interstate United States trade data are from 2007, we use 2007 as the base year for our analysis. The calibration pins down the parameters χ_{us} and χ_{ca} but, following the earlier discussion in section 3.5, not the contribution of within- and across-country market segmentation to the observed price dispersion. We then use data on the internal trade flows from Canada and the United States to decompose this segmentation and show that it is driven almost exclusively by market segmentation within countries rather than across.

4.1 Model Calibration

Before calibrating to the data, we fix several model parameters to standard values in the literature. We fix the share of tradable consumption in the final good consumption aggregator $\phi = 0.66$, which matches the average trade share found by [Lombardo and Ravenna \(2012\)](#) across a panel of 25 countries. We set the elasticity of substitution at the good level to $\rho = 2$. Because it does not directly enter any individual optimization problems, this choice has a negligible effect on our results. We set the elasticity of substitution between varieties to $\eta = 5$. The latter value yields retail-level markups of 25% and also places a bound on producer level markups to be at or below the same

level. We select $v(L) = \psi \frac{L^{1+\frac{1}{\zeta}}}{1+\frac{1}{\zeta}}$, set ψ to target a nonstochastic value of $\bar{L} = 0.4$, and set the Frisch elasticity $\zeta = 1$. We then normalize the fixed component of the nominal wage to be consistent with the flexible price equilibrium of the economy when aggregate shocks are at their steady-state values and set the adjustment parameter $\mu_w = 0.15$ to capture substantial short-run stickiness in nominal wages. Finally, we fix the relative size of the two countries $s^{us}/s^{ca} = 9$, to match the relative levels of output in the United States and Canada.

On the producers' side, we assume that average relative productivity is distributed according to $\epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and that producer-specific productivity shocks are distributed within each period according to $\zeta_i \sim \mathcal{N}(0, \sigma_{\zeta,c}^2)$, with $c \in \{\text{us, can}\}$. The good-level cost shocks play an important role, along with monetary shocks, in delivering cross-border differences in marginal costs that, when combined with segmented markets, lead to differences in the price distributions sampled by retailers in each country. We calibrate the idiosyncratic cost shock parameters to match the main moments of the cost data for the United States and Canada considered by [Burstein and Jaimovich \(2012\)](#). Significant asymmetry in the distribution of idiosyncratic cost shocks plays a crucial part in delivering the asymmetric price dispersion observed across the two countries. Finally, the parameters σ_{ms} and ρ_{ms} determine both the persistence of the nominal exchange rate and, given our assumption regarding sticky wages, the degree to which monetary shocks generate aggregate differences in firms' production costs.

We select parameter values to target eight moments for the model. In particular, we seek to match the three main statistics on the dispersion of relative price changes within and across countries presented in section 2. To be as model-consistent as possible, we focus on the price statistics provided by [Burstein and Jaimovich \(2012\)](#) for matched products produced in both countries. We select the volatility and persistence of monetary shocks to match an unconditional autocorrelation of 0.95 and a standard deviation of exchange rate changes of 2.9%. We also seek to match the share of imports from Canada in total domestic United States demand and the equivalent object from the Canadian perspective. We compute these values using 2007 National Income and

Table 1: Targeted moments and their values implied by the model.

	$\bar{\mu}_{us}$	$\sigma_{\Delta d_{t,us}}$	$\sigma_{\Delta d_{t,ca}}$	$\sigma_{\Delta d_{t,bord}}$	$\phi_{imp,us}$	$\phi_{imp,ca}$	$\sigma_{\Delta \log(e)}$	ρ_e
Target	0.150	0.080	0.060	0.140	0.021	0.175	0.029	0.950
Baseline	0.152	0.079	0.060	0.140	0.019	0.176	0.029	0.955
No PTM	0.223	0.191	0.237	0.229	0.023	0.218	0.032	0.955
Perfect Seg.	0.167	0.081	0.046	0.195	0.000	0.003	0.027	0.955

Note: Price moments in target data are from Table 5 panel B of [Burstein and Jaimovich \(2012\)](#). In the baseline calibration, the model is calibrated to match the data as described in section 2. In the no pricing pricing-to-market (No PTM) calibration, search costs are increased to a very high level so that producers charge constant monopoly markups. In the perfect segmentation (Perfect Seg.) calibration, $\chi_{us} = \chi_{ca} \approx 0$, eliminating international trade.

Product Accounts data from the United States and Canada to compute domestic demand, along with United States census data on gross trade flows in goods between the United States and Canada in 2007. As reported in table 1, imports of United States goods to Canada account for roughly $\phi_{imp,ca} = 17\%$ of domestic demand in Canada, while United States imports from Canada account for roughly $\phi_{imp,us} = 2\%$ of final United States demand. Finally, lacking strong evidence on producer-level markups, we target a modest $\bar{\mu}_{us} = 15\%$ average markup in the calibration. The target moments are summarized in the first line of table 1.

We calibrate eight model parameters, listed in table 2, to match the eight target moments. The search cost parameter κ largely determines the degree of pricing power held by producers, and the calibrated value of $\kappa = 0.004$ is sufficiently small to limit firm markups to a level well below their full-monopoly value of $\log(5/4) = .22$, but large enough to allow for the substantial price dispersion seen in the data. The parameters $\sigma_{\zeta,us} = 0.092$ and $\sigma_{\zeta,can} = 0.029$ are the crucial parameters for matching the levels of within-country price dispersion, while the good-level productivity shock $\sigma_e = 0.090$ provides an additional source of volatility in cross-country costs, beyond aggregate nominal shocks and idiosyncratic producer-level shocks, that is needed to drive the observed cross-country dispersion of price changes.

The crucial parameters for the question of the paper are the calibrated values of the search parameters, χ_{us} and χ_{ca} . The calibrated values of $\chi_{us} = 0.23$ and $\chi_{ca} = 0.13$ imply that, even after controlling for size, United States retailers are roughly five times

Table 2: Parameter values for the baseline model calibration.

Parameter	κ	χ_{us}	χ_{ca}	$\sigma_{\zeta,us}$	$\sigma_{\zeta,ca}$	σ_{ϵ}	σ_{ms}	ρ_{ms}
Value	0.004	0.230	0.130	0.092	0.029	0.090	0.055	0.950

Table 3: Other model moments not targeted.

	$\bar{\mu}_{ca}$	$\rho(\bar{d}_{t,bord}, \log(e_t))$	$\rho_{\Delta p_t,us}$	$\rho_{\Delta p_t,ca}$	$\rho_{\Delta p_t,bord}$
Baseline	0.199	0.936	0.830	0.876	0.401
No PTM	0.223	0.988	0.461	0.233	0.255
Perfect Seg.	0.217	0.928	0.845	0.941	0.022

more likely to search domestic rather than Canadian producers, while Canadian retailers are almost ten times more likely to search producers in their own country. Thus, the calibration is very much consistent with a strong home bias, even if we cannot yet say whether that bias is toward the home country or the home region. It perhaps somewhat surprising, given its relatively high degree of trade with the United States, that Canada appears to be more closed relative to the United States. This disconnect is only partially reconciled by the different in size of the two countries. Another part of the explanation lies in the relatively high dispersion of prices in the United States, which increases the probability that a Canadian retailer searching in the United States will land at producer with a price well below its reservation price and, therefore, purchase a relatively large quantity from that supplier.

The first and second rows of table 1 show that the baseline calibration of the model can almost perfectly match the targeted moments. In particular, changes in real relative prices are far more volatile across countries than within (fact 1), and relative prices across countries are approximately four times more volatile than the nominal exchange rate (fact 2).

Table 3 shows that the model has reasonable implications for other moments in the data as well. In particular, the average cross-border price is highly correlated with the exchange rate and, by extension, relative unit labor costs between the two countries

(fact 3). The model also qualitatively matches fact 4, that price changes are more correlated within countries than across. However, it cannot match the extremely low (near zero) correlation of cross-border price changes found in the data.

4.2 Regional versus International Segmentation

As discussed in section 3.5, decomposing the sources of home bias found earlier requires data on internal trade quantities. Data on interprovincial trade for Canada are provided by Statistics Canada, the governmental statistical agency for Canada, while analogous data on trade between United States states is provided by the Commodity Flow Survey produced by the United States Department of Transportation. Since state-level trade data available from that survey are from 2007, we again use 2007 as our base year for computing all trade quantities.

In order to map these data sources to our model economy, we must divide each country into two, ideally identically sized, regions. Geographic distinctions (i.e., East vs. West) provide a natural way to do this, but in order to ensure that our choice of region definitions is not influencing our results in a hidden way, we consider a much broader set of regional definitions. To do this for Canada, we first combine the data for the three territories of Canada (Northwest Territories, Nunavut, and Yukon) and then consider all possible two-bin (region c/region d) classifications of the 10 provinces + 1 combined territory. Among these, we keep only those region definitions for which the total economic output of each region is within 10 % of the other, thus aligning with our model of assumption of equal-size regions and avoiding cases in which a single small region drives the statistics.¹⁰ Thus, only 108 out of the possible 2^{10} region definitions survive. For the United States, we follow an identical procedure, after first combining the US states into 12 different regions according to a Department of Transportation sub-region classification. In this case, 464 of the 2^{11} different regional definitions survive.

For each country and admissible region definition, we then use equation (23) and (24) to back out the implied values of β_1^r , β_2^r , and β_3^r . Table 4 reports the average as

¹⁰Nevertheless, our computation of the underlying bias parameters does take into account the small size differences in the alternative region definitions.

Table 4: Estimated bias parameters.

	United States			Canada		
	Avg.	Min.	Max.	Avg.	Min.	Max.
β_1^r	0.73	0.71	0.76	0.83	0.82	0.84
β_2^r	0.10	0.08	0.13	0.06	0.04	0.07
β_3^r	0.08	0.08	0.08	0.06	0.06	0.06

well as the minimum and maximum values for each of these parameters. The results demonstrate that, regardless of particular region definitions, the market segmentation by regional trade levels is *almost entirely* within countries rather than across the international border, for both the United States and Canada. In the United States, a retailer is roughly 7 times more likely to sample a supplier within its own region rather than in the neighboring region of the same country, and only about 20% more likely to sample from a neighboring region than a region across the international border. For Canada, the results indicate an even stronger degree of within-country segmentation, with a retailer being roughly 14 times more likely to sample from her own region rather than the neighboring Canadian region, while the probabilities of searching the neighboring region compared to a region in the United States are nearly identical. Moreover, the particular region definitions do not appear to matter, as neither the minimum nor maximum values of the same parameters are substantially different from the average across all region definitions.

The previous results strongly indicate that market segmentation should be largely attributed to segmentation at the regional level; yet earlier literature has often assumed (either implicitly or explicitly) that market segmentation within countries is negligible. To understand the implications of this assumption, we consider a counterfactual experiment in which, rather than using internal trade data to pin down the fundamental search parameters, we instead achieve identification by assuming that retailers search across all within-country markets with the same probability. Table 5 indicates that this identification assumption leads to very misleading results regarding the importance of segmentation at the border, implying the conclusion that domestic firms are roughly 5

Table 5: Counter-factual bias parameters.

	United States			Canada		
	Avg.	Min.	Max.	Avg.	Min.	Max.
β_1^r	0.42	0.41	0.42	0.44	0.44	0.44
β_2^r	0.42	0.41	0.42	0.44	0.44	0.44
β_3^r	0.08	0.08	0.09	0.06	0.06	0.06

to 8 times more likely to search a neighboring domestic region, compared to a similarly-sized region across the border.

One possible objection to the approach taken here is that, in assuming away region-level shocks, we may have missed a potential source of identification that would arise if we could compare price differences *within and across regions*, as well as prices across the international border. Given the evidence cited in footnote 4, it seems unlikely to us that such shocks play an important role in generating the data. Nevertheless, the results regarding the relative degree of within- and across-country segmentation could not be overturned by relaxing this assumption. To see this, note that to the extent that the distribution of prices differs more across regions than within, firms drawing cross-region prices will optimally stop and purchase from across the regional border *more often* than in the same model without such differences. This will lead, other things equal, to a larger degree of within-country trade. In order to match the low observed level of internal trade, then, the model would require an *even greater* degree of within-country market segmentation. Thus, even if a within/across-region comparison of prices could offer some identifying information regarding the degree of internal segmentation, it would not lead to a qualitative change in our conclusion regarding the relative importance of within-country market segmentation.

A second possible objection to our general modeling strategy is our strong assumption regarding the exogeneity of search intensity with respect to the price distributions in each country. While this assumption plays an important role in keeping the environment tractable, it is important to notice that the model maintains two channels that allow for substantial expenditure switching as relative costs across the countries

fluctuate. First, conditional on settling on a supplier, elastic final demand implies that retailers who land at producers with relatively low costs purchase relatively large quantities. Second, the presence of good-level shocks, combined with the endogenous decision of firms to shut down, leads to a smaller mass of price-posting firms in the country with relatively higher costs, shrinking the probability that firms will sample suppliers from that country. To demonstrate that the channels lead to a qualitatively reasonable degree of expenditure switching, we regressed the log of the quantity of imports on a constant, the log of total domestic traded good absorption, and the log of the nominal exchange rate in each country. We find that a 1% depreciation of the United States dollar leads to a roughly 0.8% fall in United States imports of Canadian goods and a roughly 1.4% increase in Canadian imports of United States goods.

4.3 The Importance of Pricing-to-Market and Cross-Border Competition

We next investigate to what extent our results are driven by pricing-to-market. Since producers in our model set a single price for each good-variety, we define pricing-to-market as the tendency of producers with *equal marginal costs* to set different prices depending on the market in which they are located. To study the role of pricing-to-market, we consider two counterfactual parameterizations of the model that capture, respectively, the minimal and maximal degrees of pricing-to-market that the model can deliver, given the baseline distributions of relative marginal cost shocks.

In the first exercise, we parameterize the model so that search costs are high enough that retailers always purchase from the first producer they search. This parameterization shuts down the pricing-to-market created by the presence of different reservation prices across countries and ensures that producers always charge a constant monopoly markup relative to their marginal cost. The third row of table 1 shows that, under this calibration, the degree of additional price dispersion created by the border essentially disappears, with cross-border price dispersion falling somewhere between the levels for the United States and Canada. This calibration indicates that pricing-to-market plays a very important role in explaining cross-border price differentials, a finding that is

consistent with [Gopinath et al. \(2011\)](#) and [Burstein and Jaimovich \(2012\)](#).

Next, we consider a calibration of the economy in which we set the parameters $\chi_{us} = \chi_{ca} \approx 0$, essentially eliminating international trade.¹¹ The final row of [table 1](#) describes the model-implied moments from this counterfactual exercise, in which market conditions in the foreign country are irrelevant for the pricing decisions of home firms. In this case, cross-border pricing differentials are much larger than in the baseline calibration, and they are also much larger relative to within-country price differentials. The contrast with the baseline economy demonstrates that cross-border market conditions do indeed have a substantial effect on firms' pricing decisions and serve to compress what, according to our model, would otherwise be substantially larger cross-border pricing differentials.

Taken together, these two counterfactual exercises suggest a modest but nontrivial effect of cross-border competition: pricing-to-market in the baseline economy is substantial but smaller than it would be if national markets were completely isolated. Our finding of partial market segmentation contrasts to some degree with the finding of [Gopinath et al. \(2011\)](#) that international markets are nearly fully segmented.

5 Conclusion

We have demonstrated that a model of price dispersion via retailer search can replicate the most prominent facts about good-level real exchange rates without relying on market segmentation induced by the international border. Evidence on intranational trade from the United States and Canada strongly indicates that, indeed, the border plays no special role in segmenting markets. Instead, regional segmentation seems the far larger source of market segmentation.

Beyond the application to the particular case of US-Canada bilateral trade, identifying the degree of market segmentation in other contexts is also likely to require data on variables other than prices, most notably trade quantities. Although the model

¹¹A small amount of trade is needed to pin down the nominal exchange rate in the economy, since it is determined by the balanced trade condition.

we have presented here incorporates a reduced-form wage friction, we leave for future work the introduction of a micro-founded nominal price stickiness. Although it cannot account for the data on its own, it is possible that in a dynamic setting, price stickiness interacts in an important way with our search friction, supporting the persistence of price dispersion.

References

- Alessandria, G. (2004). International Deviations from the Law of One Price: The Role of Search Frictions and Market Share. *International Economic Review* 45(4), 1263–1291.
- Alessandria, G. (2009). Consumer Search, Price Dispersion, and International Relative Price Fluctuations. *International Economic Review* 50(3), 803–829.
- Alessandria, G. and J. P. Kaboski (2011). Pricing-to-Market and the Failure of Absolute PPP. *American Economic Journal: Macroeconomics* 3(1), 91–127.
- Anderson, J. E. and E. van Wincoop (2003). Gravity With Gravitas: A Solution to the Border Puzzle. *American Economic Review* 93(1), 170–192.
- Atkeson, A. and A. Burstein (2008). Pricing-to-Market, Trade Costs, and International Relative Prices. *American Economic Review* 98(5), 1998–2031.
- Baxter, M. and A. Landry (2012). IKEA: Product, Pricing, and Pass-Through. Working Paper, Federal Reserve Bank of Dallas.
- Berger, D., J. Faust, J. H. Rogers, and K. Steverson (2012). Border Prices and Retail Prices. *Journal of International Economics* 88(1), 62–73.
- Bernard, A. B., S. J. Redding, and P. K. Schott (2011). Multiproduct Firms and Trade Liberalization. *Quarterly Journal of Economics* 126(3), 1271–1318.
- Boivin, J., R. Clark, and N. Vincent (2012). Virtual Borders. *Journal of International Economics* 86(2), 327–335.

- Broda, C. and D. E. Weinstein (2008, May). Understanding International Price Differences Using Barcode Data. Working Paper 14017, National Bureau of Economic Research.
- Burstein, A. and G. Gopinath (2014). International Prices and Exchange Rates. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Amsterdam, pp. 391–451. Elsevier.
- Burstein, A. and N. Jaimovich (2012). Understanding Movements in Aggregate Product-Level Real-Exchange Rates. Working paper, UCLA.
- Candian, G. (2015). Information Frictions and Real Exchange Rate Dynamics. Working paper, Boston College.
- Chen, N. (2004). Intra-national Versus International Trade in the European Union: Why Do National Borders Matter? *Journal of International Economics* 63(1), 93–118.
- Crucini, M. J. and C. I. Telmer (2012, April). Microeconomic Sources of Real Exchange Rate Variability. Working Paper 17978, National Bureau of Economic Research.
- De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2015). Prices, Markups and Trade Reform. *Econometrica* forthcoming.
- Diamond, P. A. (1971). A Model of Price Adjustment. *Journal of Economic Theory* 3(2), 156–168.
- Drozd, L. A. and J. B. Nosal (2012). Understanding International Prices: Customers As Capital. *American Economic Review* 102(1), 364–395.
- Engel, C. and J. H. Rogers (1996). How Wide is the Border? *American Economic Review* 86(5), 1112–1125.
- Fitzgerald, D. and S. Haller (2014). Pricing-to-Market: Evidence from Plant-Level Prices. *Review of Economic Studies* 81(2), 761–786.

- Goldberg, P. K. and R. Hellerstein (2013). A Structural Approach to Identifying the Sources of Local Currency Price Stability. *The Review of Economic Studies* 80(1), 175–210.
- Gopinath, G., P.-O. Gourinchas, C.-T. Hsieh, and N. Li (2011). International Prices, Costs, and Markup Differences. *American Economic Review* 101(6), 2450–2486.
- Gorodnichenko, Y. and L. L. Tesar (2009). Border Effect or Country Effect? Seattle May Not Be So Far from Vancouver After All. *American Economic Journal: Macroeconomics* 1(1), 219–241.
- Hillberry, R. and D. Hummels (2003). Intranational Home Bias: Some Explanations. *Review of Economics and Statistics* 85(4), 1089–1092.
- Hillberry, R. and D. Hummels (2008). Trade Responses to Geographic Frictions: A Decomposition Using Micro-data. *European Economic Review* 52(3), 527–550.
- Lombardo, G. and F. Ravenna (2012). The Size of the Tradable and Non-tradable Sectors: Evidence from Input-Output Tables For 25 Countries. *Economics Letters* 116(3), 558–561.
- McCallum, J. (1995). National Borders Matter: Canada-U.S. Regional Trade Patterns. *American Economic Review* 85(3), 615–623.
- Millimet, D. L. and T. Osang (2007). Do State Borders Matter For U.S. Intranational Trade? The Role of History and Internal Migration. *Canadian Journal of Economics* 40(1), 93–126.
- Reinganum, J. F. (1979). A Simple Model of Equilibrium Price Dispersion. *Journal of Political Economy* 87(4), 851–858.
- Wolf, H. C. (2000). Intranational Home Bias in Trade. *Review of Economics and Statistics* 82(4), 555–563.

A Detailed Derivation of the Two-Country, Two-Region Economy [For Online Publication]

Incorporating producer shutdown, the measures of active producers in each country-region for good i are given by $\left(m_a^{prod}(i), m_b^{prod}(i), m_c^{prod}(i), m_d^{prod}(i)\right)$ where $m_r^{prod}(i) = s_r(1 - G_i(mc_r^*(i)))$ and $mc_r^*(i)$ denotes the level of marginal costs at which firms producing good i in region r no longer find it profitable to operate. All regions produce the same goods, which is a continuum represented by $i \in [0, 1]$.

We are assuming that the regions within a country are identical (except with respect to the search bias), so we the following:

$$\begin{aligned} m_a^{prod}(i) &= m_b^{prod}(i) \equiv m_H^{prod}(i) \\ m_c^{prod}(i) &= m_d^{prod}(i) \equiv m_F^{prod}(i) \\ m_a^{ret} &= m_b^{ret} \equiv s_H \\ m_c^{ret} &= m_d^{ret} \equiv s_F. \end{aligned}$$

A.1 Households

Households in region $k \in \{a, b, c, d\}$ solve the following problem:¹²

$$\begin{aligned} \max_{C_{k,T}, C_{k,N}, L_k} & \log \left(C_{k,T}^\phi C_{k,N}^{1-\phi} \right) - v(L_k) \\ s.t. & P_{k,T} C_{k,T} + P_{k,N} C_{k,N} \leq w_k L_k + \Pi_k, \end{aligned}$$

where $C_{k,T}$ is the consumption basket

$$\begin{aligned} C_{k,T} &= \left(\int_0^1 c_k(i)^{\frac{\rho-1}{\rho}} di \right)^{\frac{\rho}{\rho-1}}, \\ c_k(i) &= \left(\int_{\Omega_k} c_k(i, \nu)^{\frac{\eta-1}{\eta}} d\nu \right)^{\frac{\eta}{\eta-1}}, \end{aligned}$$

and Ω_k denotes the set of retailers in region k (measure m_k^{ret}). Prices are expressed in terms of currency of country 1. Each retailer represents one variety (ν), and it can offer a continuum of goods (i).

¹²Again, we are considering that a and b are in country 1, and c and d are in country 2.

The optimal allocation of the household must satisfy the following conditions:

$$\begin{aligned}
C_{k,T} & : \frac{\phi}{C_{k,T}} = \lambda_k P_{k,T}, \\
C_{k,N} & : \frac{1 - \phi}{C_{k,N}} = \lambda_k P_{k,N}, \\
L_k & : v'(L_k) = \lambda_k w_k, \\
BC & : P_{k,T} C_{k,T} + P_{k,N} C_{k,N} = w_k L_k + \Pi_k,
\end{aligned}$$

where λ_k is the Lagrange multiplier with respect to the budget constraint. These four equations allow us to solve for $C_{k,T}$, $C_{k,N}$, L_k , and λ_k as functions of prices $P_{k,T}$, $P_{k,N}$, w_k , and profits Π_k (profits of retailers and producers of region k). The problem is similar for each country-region k .

A.2 Nontradable Sector

We assume perfect competition in the nontradable sector. The technology is

$$y = \bar{A}_k l_k.$$

This implies that the following condition must hold in equilibrium:

$$P_{k,N} = \frac{w_k}{\bar{A}_k}.$$

A.3 Demand for Differentiated Final Goods

All regions produce all the goods, indexed by i . Households buy differentiated final goods from *local* retailers. Each retailer, indexed by ν , solves the profit maximization problem for each good independently.¹³ The demand for good i sold by retailer ν in

¹³This follows from the assumption that the search process is independent across goods.

region k is given by¹⁴

$$\begin{aligned}
y_k(i) &= p_k(i)^{-\rho} P_{k,T}^\rho C_{k,T}, \\
y_k(i, \nu) &= p(i, \nu)^{-\eta} p_k(i)^\eta y_k(i), \\
P_{k,T} &= \left(\int_0^1 p_k(i)^{1-\rho} di \right)^{\frac{1}{1-\rho}}, \\
p_k(i) &= \left(\int_{\Omega_k} p_k(i, \nu)^{1-\eta} d\nu \right)^{\frac{1}{1-\eta}}.
\end{aligned}$$

Note that up to this point, we consider only local demand $C_{k,T}$, since retailers sell only to local consumers.

A.4 Retailers

A retailer in region r samples prices for good i from the distribution

$$f_{i,r}^{ret}(\hat{p}) \equiv \alpha_1^r f_i^{\omega(r,1)}(\hat{p}) + \alpha_2^r f_i^{\omega(r,2)}(\hat{p}) + \alpha_3^r f_i^{\omega(r,3)}(\hat{p}) + \alpha_4^r f_i^{\omega(r,4)}(\hat{p}),$$

with $\sum_k \alpha_k^r = 1$, $k \in \{1, 2, 3, 4\}$, and where $f_i^{\omega(r,k)}$ is the density of prices posted by the good i producers of the r -relative region k . It is important to remember that retailers from all regions face the same distribution of producer prices in any given region; what differs is the probability of sampling from these distributions, the α 's.

Retailer ν in region r must choose whether to purchase the good upon drawing producer j_i^ν and its respective price or whether to search again at a cost κ . Assume the price of producer j_i^ν is \hat{p} . If the retailer decides to **purchase the good**, her profit is

$$V^{ns}(j_i^\nu; \hat{p}) = \max_{p_r(i, \nu)} p_r(i, \nu)^{1-\eta} p_r(i)^{\eta-\rho} P_{r,T}^\rho C_{r,T} - \hat{p} p_r(i, \nu)^{-\eta} p_r(i)^{\eta-\rho} P_{r,T}^\rho C_{r,T}.$$

The marginal cost is \hat{p} , so the optimal price is

$$p_r(i, \nu) = \frac{\eta}{\eta - 1} \hat{p} = \mu \hat{p}.$$

The quantity demanded by the retailer at the optimal price is

$$x_r(i, \nu; \hat{p}) = (\mu \hat{p})^{-\eta} p_r(i)^{\eta-\rho} P_{r,T}^\rho C_{r,T},$$

¹⁴This is the solution to the household problem.

and profits are:

$$V^{ns}(j_i^\nu; \widehat{p}) = \widehat{p}^{1-\eta} (\mu)^{-\eta} (\mu - 1) p_r(i)^{\eta-\rho} P_{r,T}^\rho C_{r,T}.$$

If the retailer chooses to **search again**, her expected profit is

$$V^s(f_i) = \int_{\Omega_p} p^{1-\eta} (\mu)^{-\eta} (\mu - 1) p_r(i)^{\eta-\rho} P_{r,T}^\rho C_{r,T} f_{i,r}^{ret}(p) dp - \kappa,$$

where Ω_p is the set of posted prices, that is, the support of $f_{i,r}^{ret}(p)$.

So the decision is,

$$V(j_i^\nu; f_i, \widehat{p}) = \max\{V^s(f_i), V^{ns}(j_i^\nu; \widehat{p})\}.$$

When making her decision, $V^s(f_i)$ is a constant for the retailer, while $V^{ns}(j_i^\nu; \widehat{p})$ is decreasing in \widehat{p} . This implies that the solution is characterized by a reservation price $\overline{p}_r(i)$ such that if the sampled price is above $\overline{p}_r(i)$, the retailer will choose to search again. Otherwise, she will purchase the good from producer j_i^ν . Note that the reservation price is product-specific.

A.5 Producers

Let $(i, j)_k$ denote the producer j in region k that produces good i . She has the following technology to produce good i :

$$y_k(i, j) = A_k(i) \zeta_k(i, j) l_k(i, j),$$

where $A_k(i)$ is the product-specific productivity shock, and $\zeta_k(i, j)$ is the producer-specific productivity shock. It is assumed that $A_k(i)$ is country-specific, that is, for a given product i , both regions within a country have the same productivity $A_k(i)$. Both distributions are assumed to be country-specific.

The producer posts her price after observing the shocks. She can be matched to retailers from every country-region, so she takes that into account. We assume that $\overline{p}_k(i)$ differs only across countries.¹⁵ So we have $\overline{p}_a(i) = \overline{p}_b(i)$ and $\overline{p}_c(i) = \overline{p}_a(i)$ (country 1 and country 2, respectively). Without loss of generality, we assume $\overline{p}_a(i) < \overline{p}_c(i)$.

¹⁵This is a guess and verify approach, since we know that given the symmetry, it will be the case in equilibrium.

Assumption: Country 1 has the lower reservation price, $\bar{p}_a(i) < \bar{p}_c(i)$.

The producer can choose to post a price $p \leq \bar{p}_c(i)$ or to shut down.¹⁶ If $\bar{p}_a(i) < p \leq \bar{p}_c(i)$, the producer sells **only to retailers from country 2** (both regions). If $p < \bar{p}_a(i)$, **it sells to retailers from both countries** (and both regions within the countries). So we have to solve for the optimal price:

$$\max_{\hat{p}} (\mu \hat{p})^{-\eta} \left(\hat{p} - \frac{w_k}{A_k(i) \zeta_k(i, j)} \right) \left[\begin{array}{l} I\{p \leq \bar{p}_a(i)\} \omega_a(i, j)_k p_a(i)^{\eta-\rho} P_{a,T}^\rho C_{a,T} \\ + I\{p \leq \bar{p}_a(i)\} \omega_b(i, j)_k p_b(i)^{\eta-\rho} P_{b,T}^\rho C_{b,T} \\ + \omega_c(i, j)_k p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T} \\ + \omega_d(i, j)_k p_d(i)^{\eta-\rho} P_{d,T}^\rho C_{d,T} \end{array} \right],$$

where $I\{p \leq \bar{p}_a(i)\}$ is the indicator function, equal to zero if the price is above $\bar{p}_a(i)$ and equal to one otherwise. The term $\omega_{k'}(i, j)_k$ is the mass of retailers from region k' sampling producer $(i, j)_k$. Note that without the indicator function, and without the limit on p (that it must be lower than $\bar{p}_c(i)$), the producer would just choose to charge a constant markup over marginal cost. However, when the constant markup price is just above $\bar{p}_a(i)$, the producer will compare this price to the reservation price $\bar{p}_a(i)$, since the demand is discontinuous. In this region, if she chooses to charge the reservation price $\bar{p}_a(i)$, she would have a lower markup but would face a larger demand. The same applies when the constant markup price is above $\bar{p}_c(i)$. But in this case, charging a price above $\bar{p}_c(i)$ means zero profits. So the producer will choose to charge $\bar{p}_c(i)$ until profits are zero.

¹⁶Demand will be zero if the price is above the highest reservation price.

Let $\zeta_k^1(i, j)$ be the solution to

$$\begin{aligned} (\mu \bar{p}_a(i))^{-\eta} \left(\bar{p}_a(i) - \frac{w_k}{A_k(i)\zeta_k(i, j)} \right) & \begin{bmatrix} \omega_a(i, j)_k p_a(i)^{\eta-\rho} P_{a,T}^\rho C_{a,T} \\ + \omega_b(i, j)_k p_b(i)^{\eta-\rho} P_{b,T}^\rho C_{b,T} \\ + \omega_c(i, j)_k p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T} \\ + \omega_d(i, j)_k p_d(i)^{\eta-\rho} P_{d,T}^\rho C_{d,T} \end{bmatrix} = \\ \left(\mu^2 \frac{w_k}{A_k(i)\zeta_k(i, j)} \right)^{-\eta} (\mu - 1) \left(\frac{w_k}{A_k(i)\zeta_k(i, j)} \right) & \begin{bmatrix} \omega_c(i, j)_k p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T} \\ + \omega_d(i, j)_k p_d(i)^{\eta-\rho} P_{d,T}^\rho C_{d,T} \end{bmatrix}. \end{aligned}$$

It is the productivity level where the producer is indifferent between charging the constant markup price and the lower reservation price $\bar{p}_a(i)$.

The solution to the producer problem is the following:

$$\hat{p}_k(i, j) = \left\{ \begin{array}{l} \frac{\mu w_k}{A_k(i)\zeta_k(i, j)}, \quad \text{if } \zeta_k(i, j) > \frac{\mu w_k}{A_k(i)\bar{p}_a(i)} \\ \bar{p}_a(i), \quad \text{if } \zeta_k^1(i, j) < \zeta_k(i, j) \leq \frac{\mu w_k}{A_k(i)\bar{p}_a(i)} \\ \frac{\mu w_k}{A_k(i)\zeta_k(i, j)}, \quad \text{if } \frac{\mu w_k}{A_k(i)\bar{p}_c(i)} < \zeta_k(i, j) \leq \zeta_k^1(i, j) \\ \bar{p}_c(i), \quad \text{if } \frac{w_k}{A_k(i)\bar{p}_c(i)} < \zeta_k(i, j) \leq \frac{\mu w_k}{A_k(i)\bar{p}_c(i)} \end{array} \right\}$$

Note that for productivity levels below $\frac{w_k}{A_k(i)\bar{p}_c(i)}$, the producer would make negative profits, so it chooses not to operate. Also, note that the distribution of prices within each country will be the same, since wages and product-specific productivity are the same within countries.

A.6 Price Distributions

Next, let F_k denote the cumulative distribution of prices in region k , that is, $H_k(p, i) = \Pr(\hat{p}_k(i, j) \leq p)$. We then have

$$F_k(p, i) = \left\{ \begin{array}{l} \Pr \left(\zeta_k(i, j) \geq \frac{\mu w_k}{A_k(i)p} \right), \text{ if } p < \bar{p}_a(i) \\ \Pr \left(\zeta_k(i, j) \geq \zeta_k^1(i, j) \right), \text{ if } p = \bar{p}_a(i) \\ \Pr \left(\zeta_k(i, j) \geq \frac{\mu w_k}{A_k(i)p} \right), \text{ if } \bar{p}_a(i) < p < \bar{p}_c(i) \\ 1, \text{ if } p = \bar{p}_c(i). \end{array} \right.$$

Note that the maximum observed price is $\bar{p}_c(i)$, so the cumulative distribution is equal to one at that point. Also, remember that we are taking $A_k(i)$ as given. We assume that $\ln \zeta_k(i, j)$ is independently normally distributed, $\ln \zeta_k(i, j) \sim N(0, \sigma_k^\zeta)$. So we have

$$F_k(p, i) = \left\{ \begin{array}{l} \frac{1 - \Phi(\ln \mu + \ln w_k - \ln A_k(i) - \ln p; 0, \sigma_k^\zeta)}{1 - \Phi(\ln w_k - \ln A_k(i) - \ln p; 0, \sigma_k^\zeta)}, \text{ if } p < \bar{p}_a(i) \\ \frac{1 - \Phi(\ln \zeta_k^1(i, j); 0, \sigma_k^\zeta)}{1 - \Phi(\ln w_k - \ln A_k(i) - \ln p; 0, \sigma_k^\zeta)}, \text{ if } p = \bar{p}_a(i) \\ \frac{1 - \Phi(\ln \mu + \ln w_k - \ln A_k(i) - \ln p; 0, \sigma_k^\zeta)}{1 - \Phi(\ln w_k - \ln A_k(i) - \ln p; 0, \sigma_k^\zeta)}, \text{ if } \bar{p}_a(i) < p < \bar{p}_c(i) \\ 1, \text{ if } p = \bar{p}_c(i), \end{array} \right.$$

where $\Phi(x; \mu, \sigma)$ denotes the cumulative distribution function of a variable that is normally distributed with mean μ and standard deviation σ . Let $F_{i,r}^{ret}(\hat{p})$ denote the cumulative distribution of $f_{i,r}^{ret}(\hat{p})$, that is, the cumulative distribution of prices from which a retailer in region r samples prices. It is given by

$$F_{i,r}^{ret}(\hat{p}) \equiv \alpha_1^r F_i^{\omega(r,1)}(\hat{p}) + \alpha_2^r F_i^{\omega(r,2)}(\hat{p}) + \alpha_3^r F_i^{\omega(r,3)}(\hat{p}) + \alpha_4^r F_i^{\omega(r,4)}(\hat{p}),$$

where the cumulative distributions in each region, F 's, are the ones characterized earlier.

A.7 Search Process (for Good i)

Next, we determine the mass of retailers sampling **each producer**. Retailers from country 2 search only once. We assume that producers within a region have the same probability of been sampled. So the mass of retailers from country 2 drawing producer

$(i, j)_k$ and purchasing the good from it is¹⁷

$$\frac{\alpha_k^c m_c^{ret} + \alpha_k^d m_d^{ret}}{m_k^{prod}}.$$

Retailers from country 1 search multiple times, since there is a positive probability that they will draw prices above their reservation price. The term $F_{i,a}^{ret}(\bar{p}_a(i))$ also represents the fraction of retailers in region a who find a seller (producer) with price less than $\bar{p}_a(i)$ (law of large numbers).

In this case, the mass of retailers from country 1 in region a visiting producer $(i, j)_k$ is

$$\frac{\alpha_k^a}{m_k^{prod}} (m_a^{ret}) + \frac{\alpha_k^a}{m_k^{prod}} [1 - F_{i,a}^{ret}(\bar{p}_a(i))] (m_a^{ret}) + \frac{1}{m_k^{prod}} \alpha_k^a [1 - F_{i,a}^{ret}(\bar{p}_a(i))]^2 (m_a^{ret}) + \dots = \frac{m_a^{ret}}{m_k^{prod}} \frac{\alpha_k^a}{F_{i,a}^{ret}(\bar{p}_a(i))},$$

and search costs are

$$m_a^{ret} \kappa + m_a^{ret} [1 - F_{i,a}^{ret}(\bar{p}_a(i))] \kappa + m_a^{ret} [1 - F_{i,a}^{ret}(\bar{p}_a(i))]^2 \kappa + \dots = \frac{m_a^{ret}}{F_{i,a}^{ret}(\bar{p}_a(i))} \kappa.$$

We assume that the measure of retailers and producers is the same within countries, so the previous expressions also hold for region b in country 1.

A.8 Demand Faced by Producers

For producers of region k , we have: $\omega_a(i, j)_k$ is the mass of retailers from region a settling on producer j in region k , and $F_{i,a}^{ret}(\bar{p}_a(i))$ is the mass of retailers in region a who find a seller with a price less than $\bar{p}_a(i)$. Following this notation, we have

$$\omega_a(i, j)_k = \left\{ \begin{array}{ll} \frac{m_a^{ret}}{m_k^{prod}} \frac{\alpha_k^a}{F_{i,a}^{ret}(\bar{p}_a(i))}, & \text{if } p \leq \bar{p}_a(i) \\ 0, & \text{otherwise} \end{array} \right\}$$

$$\omega_b(i, j)_k = \left\{ \begin{array}{ll} \frac{m_b^{ret}}{m_k^{prod}} \frac{\alpha_k^b}{F_{i,a}^{ret}(\bar{p}_a(i))}, & \text{if } p \leq \bar{p}_a(i) \\ 0, & \text{otherwise} \end{array} \right\}$$

$$\omega_c(i, j)_k = \left\{ \begin{array}{ll} \frac{\alpha_k^c m_c^{ret}}{m_k^{prod}}, & \text{if } p \leq \bar{p}_c(i) \end{array} \right\}$$

$$\omega_d(i, j)_k = \left\{ \begin{array}{ll} \frac{\alpha_k^d m_d^{ret}}{m_k^{prod}}, & \text{if } p \leq \bar{p}_c(i) \end{array} \right\}.$$

¹⁷We assume that within a region, producers are sampled with equal probability.

A.9 Final Demand for Producers

Producer $(i, j)_k$ faces the following demand

$$x(p; i, j)_k = \left\{ \begin{array}{ll} (\mu p)^{-\eta} \left[\omega_a(i, j)_k p_a(i)^{\eta-\rho} P_{a,T}^\rho C_{a,T} + \omega_b(i, j)_k p_b(i)^{\eta-\rho} P_{b,T}^\rho C_{b,T} + \right. \\ \left. \omega_c(i, j)_k p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T} + \omega_d(i, j)_k p_d(i)^{\eta-\rho} P_{d,T}^\rho C_{d,T} \right], & \text{if } p \leq \bar{p}_a(i) \\ (\mu p)^{-\eta} \left[\omega_c(i, j)_k p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T} + \omega_d(i, j)_k p_d(i)^{\eta-\rho} P_{d,T}^\rho C_{d,T} \right], & \text{if } \bar{p}_a(i) < p \leq \bar{p}_c(i) \\ 0, & \text{otherwise.} \end{array} \right\}$$

Finally, using the previous results, we have

$$\begin{aligned} \omega_a(i, j)_k &= \frac{m_a^{ret}}{m_k^{prod}} \frac{\alpha_k^a}{F_{i,a}^{ret}(\bar{p}_a(i))}, \\ \omega_b(i, j)_k &= \frac{m_b^{ret}}{m_k^{prod}} \frac{\alpha_k^b}{F_{i,b}^{ret}(\bar{p}_a(i))}, \\ \omega_c(i, j)_k &= \frac{\alpha_k^c m_c^{ret}}{m_k^{prod}}, \\ \omega_d(i, j)_k &= \frac{\alpha_k^d m_d^{ret}}{m_k^{prod}}. \end{aligned}$$

If we use the symmetry within countries, we reach

$$x(p; i, j)_k = \left\{ \begin{array}{ll} (\alpha_k^a + \alpha_k^b) \frac{m_a^{ret}}{m_k^{prod} F_{i,a}^{ret}(\bar{p}_a(i))} (\mu p)^{-\eta} p_a(i)^{\eta-\rho} P_{a,T}^\rho C_{a,T} \\ + (\alpha_k^c + \alpha_k^d) \frac{m_c^{ret}}{m_k^{prod}} (\mu p)^{-\eta} p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T}, & \text{if } p < \bar{p}_a(i) \\ (\alpha_k^c + \alpha_k^d) \frac{m_c^{ret}}{m_k^{prod}} (\mu p)^{-\eta} p_c(i)^{\eta-\rho} P_{c,T}^\rho C_{c,T}, & \text{if } \bar{p}_a(i) < p \leq \bar{p}_c(i) \\ 0, & \text{otherwise} \end{array} \right\}$$

The demand just derived is for a single producer.

A.10 Price of Composite Goods

We have

$$P_{k,T} = \left(\int_0^1 p_k(i)^{1-\rho} di \right)^{\frac{1}{1-\rho}}$$

and

$$p_k(i) = \left(\int_{\Omega_k} p_k(i, \nu)^{1-\eta} d\nu \right)^{\frac{1}{1-\eta}}.$$

The distribution of prices $p_k(i, \nu)$ follows from the distribution of producer prices. Each retailer in country k draws good i prices from $dF_{i,k}^{ret}(p; A_a(i), A_c(i))$. Retailers will charge a markup over the producer price:

$$p_k(i, \nu) = \mu \widehat{p}.$$

Retailers from country 2 search just once. We have a mass $m_c^{ret} = m_d^{ret}$ in each region of country 2. So the distribution of prices (law of large numbers) is

$$p_c(i) = p_d(i) = [m_c^{ret}]^{\frac{1}{1-\eta}} \left(\int_{\Omega_p} [\mu p]^{1-\eta} dF_{i,c}^{ret}(p; A_a(i), A_c(i)) dp \right)^{\frac{1}{1-\eta}}.$$

Retailers from country 1 draw prices from $dF_a(p; A_a(i), A_c(i))$, but conditional on $p \leq \overline{p}_a(i)$, since they will choose to search again if $p > \overline{p}_a(i)$. So we have

$$p_a(i) = [m_a^{ret}]^{\frac{1}{1-\eta}} \left(\int_{\Omega_p} [\mu p]^{1-\eta} dF_a^{ret}(p; A_a(i), A_c(i) | p \leq \overline{p}_a(i)) dp \right)^{\frac{1}{1-\eta}}.$$

Next, the composite good $P_{k,T} = \left(\int_0^1 p_k(i)^{1-\rho} di \right)^{\frac{1}{1-\rho}}$ reflects the distributions over $A_k(i)$. Again, using the Law of Large Numbers and given the fact that we are considering a measure 1 of goods, we have

$$P_{k,T} = \left(\int_{\Omega_{A_a} \times \Omega_{A_c}} p_k(i; A_a(i), A_c(i))^{1-\rho} f(A_a) f(A_c) dA_a dA_c \right)^{\frac{1}{1-\rho}}.$$