

Everything all the time? Entry and Exit in U.S. Import Varieties*

Roc Armenter and Miklós Koren[†]

February 14, 2017

Abstract

We propose a new theory of the extensive margin of trade based on a standard random-utility, discrete choice model for import demand. Crucially, there are only a finite number of independent purchase decisions each period. Whereas traditional demand systems predict market shares, our model yields instead the probability that a purchase for a given good is supplied by any given country. The model has a rich set of predictions regarding the extensive margin across goods, countries, and time. The model naturally reconciles two commanding observations in the data: there is a large fraction of varieties that are not traded yet the entry and exit rates of commodities are very high. We pursue an exhaustive evaluation of the model's quantitative performance with data on U.S. imports at the HS10 product level over the period 1990-2001. The model reproduces faithfully the cross-section distribution of varieties traded per product along several dimensions. Regarding dynamic facts, the model is spot on with its predictions on the net change, gross entry and exit of commodities, both by count and weighted by value; as well as survival probabilities and hazard rates. We briefly explore the model's implications for price changes, using NAFTA as a case study, and welfare gains from new varieties.

*PRELIMINARY and INCOMPLETE. For useful comments we thank George Alessandria, Lukasz Drodz, Jonathan Eaton, and Etienne Gagnon as well as participants at the International Trade Workshop at Florida International University, Penn State, The Graduate Center at C.U.N.Y., the Federal Reserve Board, the Federal Reserve Bank of Saint Louis, and elsewhere. Joy Zhu, Donald Penn, and Patrick Doelp provided excellent research assistance. The views expressed here do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

[†]*Armenter*: Federal Reserve Bank of Philadelphia. E-mail: roc.armenter@phil.frb.org. *Koren*: Central European University, IE-CERS and CEPR. E-mail: korenm@ceu.hu

1 Introduction

International trade has expanded dramatically during the last half century, doubling every seven years. This wave of globalization has been driven by trade agreements, falling transportation costs, and increased specialization. A phenomenon of such sheer magnitude is surely worth dissecting. One of the questions to receive the most attention is how much newly-traded goods have contributed to overall trade growth—what is commonly referred to as the extensive margin of trade. The attention is not unwarranted. First, many of the benefits from increased trade integration proposed in the literature operate through newly-traded goods: by increasing the varieties available to consumers, through pro-competitive effects, and by improving aggregate productivity.¹ Second, the extensive margin may have its own idiosyncratic response to price or economic fluctuations, and possibly adjust at a distinct pace as well.² This, in turn, has important implications for the evaluation and design of trade policies.

In this paper we propose a new theory of the extensive margin of trade. We introduce a novel demand system for imports, building on a standard random-utility, discrete choice model of product differentiation.³ Each country is assumed to supply a differentiated variety for each good or product. Crucially, there are only a finite number of independent purchase decisions or transactions each period. As a result, the demand for any particular commodity is a random variable. Whereas traditional demand systems predict market shares, our model yields instead the probability that a purchase of a given good is supplied by any given country. As long as the number of purchase decisions is finite, there is a positive probability that a country variety for a given good is not traded—the exact probability being determined by the price vector and the structural parameters. In other words, even though all varieties for all goods are available in all periods, not all of them are purchased in every period. Our model collapses to the standard CES demand when the number of purchase decisions per period tends to infinity.⁴ By the law of large numbers, probabilities converge to market shares and all available commodities are traded in all periods.

The model yields a rich set of predictions regarding the extensive margin, both in the cross-section and across time. We obtain simple expressions for the expected

¹See among many others Feenstra (1994) and Broda and Weinstein (2006), Feenstra and Weinstein (forthcoming), and Melitz (2003), respectively.

²See Ruhl (2008), Alessandria and Choi (2007), among others.

³See Anderson et al. (1992) for an overview of these models.

⁴This result stems from the choice of a Gumbel distribution for the random-utility terms, as first pointed out by Anderson et al. (1987). The Gumbel's properties regarding maximal order statistics are crucial to keep the model tractable despite the underlying probabilistic structure.

number of traded commodities per period along the the dimension of choice, say, countries or product classifications. Across time, we derive formulas for the net change in newly-traded commodities, as well as the gross amounts of entering and disappearing varieties, their survival probabilities and the associated hazard rates. Our model also has implications for the intensive margin, which allow us to derive, among other things, the contribution of newly-traded goods weighted by value and how sales per variety change conditional on survival.

We evaluate the model’s quantitative performance with data on U.S. imports at the HS10 product level from 1990 through 2001. To this end we propose a parsimonious calibration of the model’s structural parameters. We assume a Cobb-Douglas preference structure across products, which we use to fit the average expenditures share across goods over the period. We assume the relative prices across varieties, for each good, are constant over time. This allows us to capture the relative prices through country fixed-effects or a standard gravity regression.⁵ Jointly with the number of purchase decisions, the country fixed-effects are used to match the average number of traded products per country. We also take advantage of the underlying relationship between the dispersion parameter of the random-utility terms in the discrete choice model and the elasticity of substitution to use the estimates at the HS10 product level provided by Broda and Weinstein (2006). The growth rate of purchase decisions for foreign varieties is set to reproduce import penetration from 1990 to 2001.

The model’s predictions line up with the data with astounding precision. We start by checking whether our model reproduces the cross-section distribution of varieties per product in a given year. It does, and it actually also traces very well the distribution conditional on the product expenditure shares or country size. These results are perhaps not too surprising since they are closely tied to our calibration. However, we also find in the data that varieties per product decrease with the elasticity of substitution—an important prediction of the model.

Our model performs excellently in matching time-series moments. The average growth rate of the number of trade country-good varieties, per year, is 2.18 percent in the model versus 2.23 percent in the data. The extensive margin contributes 0.3 percentage points to trade growth—both in the model and data. The model also approximates very well the gross entry and exit of varieties, both by count and weighted by value. We also find that entry and exit rates vary with good and country size as they do in the data. Finally we perform a survival analysis for the set of

⁵Our model demands the complete distribution for relative prices—including those for varieties that are not traded.

traded commodities in 1990. Once again the model predictions match the data very well. The model also reproduces how average sales per variety change with survival, i.e., years continuously being traded.

What is behind our model’s ability to reproduce the data? At the core the model posits that trade, for very detailed product classifications, is a rare event that should not be expected to be observed regularly at the annual frequency. This naturally reconciles, at once, two immediate observations in the data: there is a large fraction of varieties that are not traded yet the entry and exit rates of commodities are very high. The model’s quantitative performance, though, requires that in addition we get the distribution of the underlying probabilities right. For example, the amount of churning expected for a given variety is maximized when the probability of observing such commodity is exactly one half, and decreases symmetrically as the probability increases or decreases from there. The aggregate level of churning depends then crucially on where the mass of the distribution of probabilities lies.

We explore our model implications with two exercises. First we take a closer look at Mexico’s experience in the aftermath of the North American Free Trade Agreement (NAFTA). Because our baseline calibration features constant relative prices, the model misses the rapid expansion of U.S. imports from Mexico on both the extensive and intensive margin. Once we account for the tariff reductions, the model predicted response on the extensive margin lies side by side with the data. The reason is quite prosaic: as Mexican goods get cheaper, they get purchased more often. Thus a larger fraction of them will be observed in the course of a year. We see this result as suggestive that the model can also perform well in response to price changes.

Finally, we revisit the computation of variety-adjusted import price indexes in Feenstra (1994) and Broda and Weinstein (2006). We find that when purchase decisions grow over time there is the potential for a substantial downward bias in the import price index—and thus to overstate the welfare gains associated with new varieties. There are, obviously, no welfare gains in our model: all varieties are available at all times.

There is by now a very large literature on the extensive margin in trade that we do not hope or attempt to review here. Most trade models posit that economies of scale in exporting, at the level of the firm, are behind the extensive margin.⁶ Clearly

⁶Melitz (2003), Bernard et al. (2003), Chaney (2008) are some of the best known contributions. Ricardian models as Eaton and Kortum (2002) also have implications for the extensive margin but they have received somewhat less attention.

the extensive margin does not operate through economies of scale in our paper. We will consequently focus our review on the few alternatives in the literature.

Elsewhere we have developed an atheoretical benchmark designed to capture the sparsity that is commonplace in trade data.⁷ The balls-and-bins model matches several stylized facts regarding the cross-section of exports across goods, countries, and firms—suggesting there was little scope to identify the relevant theory of the extensive margin. While our model shares a similar underlying probabilistic structure, it is tied down to structural parameters allowing us, for example, to contrast the model’s implications for goods with different elasticities. We also emphasize that Armenter and Koren (2014) limited the analysis of sparsity to the cross-section while our model’s performance is especially accurate regarding dynamic facts—e.g., exit, entry, survival.

Eaton et al. (2013) amend a standard heterogeneous-firm model to allow for a finite number of firms to export and show how it can account for zeros in bilateral trade. As in our paper, removing the law of large numbers creates the desired sparsity in the data. There are two important differences. First, we argue for granularity on the demand side rather than on the supply side. This mainly reflects the difference in focus—firm selection in exporting or import composition—and we thus view both approaches as complementary. The second difference is more conceptual in nature. We posit a finite number of purchase decisions *per period*, which allows to explore the dynamic implications of our model. Eaton et al. (2013) instead view the population of firms as a single realization.

Some other work does not break the law of large numbers yet deviate from standard models in order to capture some of the sparsity in the trade data. Bekes and Murakozy (2012) tackle the phenomenon of temporary trade, that is, short trading spells at the firm level. This observation is closely related to our measures of churning. Kropf and Sauré (2014) introduce a fixed cost per trade shipment. Firms choose optimally to ship large amounts infrequently, balancing shipment and storage costs. It is thus possible then that an exporting firm is not observed to trade at an annual frequency. That said, such a model is not likely to be able to explain the extensive margin from the import side.

There is also substantial work at measuring the extensive margin. Kehoe and Ruhl (2013), Hummels and Klenow (2005), and Bils and Klenow (2001). Feenstra (1994) and Broda and Weinstein (2006) go one step further and provide import price indexes that adjust for changes in the traded varieties. These, in turn, are used to compute welfare gains.

⁷Armenter and Koren (2014).

The paper is organized as follows. Section 2 describes our model and derives the main formulas for the extensive and intensive margin. We turn to the data description and model calibration in Section 3. Our results regarding the cross-section and time-series are presented in Section 4 and 5, respectively. NAFTA is the object of attention of Section 6; import price indexes that of Section 7. The last Section previsibly concludes.

2 Model

Our demand system is based on the discrete choice model detailed in Anderson et al. (1992). The commodity space is as follows. Each country $j = 1, 2, \dots, J$ supplies a differentiated variety of every product (or good) $g = 1, 2, \dots, G$. There is also a non-traded numeraire good, bringing the total number of differentiated commodities to $JG + 1$.

For each period $t = 0, 1, 2, \dots$ there is a finite number of purchase decisions or transactions for each good, denoted $n_t^g \in \mathbb{N}$. These decisions may originate at firms or consumers: the demand system is silent about the identity of the agents behind the purchases. The distribution of purchase decisions across goods is given by

$$n_t^g = \lceil \alpha^g n_t \rceil \quad (1)$$

where $\alpha^g > 0$, $\sum_G \alpha^g = 1$, and n_t is the overall trade intensity.⁸ We assume the latter evolves across time according to a linear time trend,

$$n_t = (1 + \gamma_n)^t n_0, \quad (2)$$

where $\gamma_n > 0$ is the net growth rate of purchase decisions. The assumption behind (1) is, not surprisingly, akin to a unit-elasticity demand across goods: as we would obtain under Cobb-Douglas preferences, the share of expenditures in good g will be constant across time and independent of prices.

The demand for the country j variety of a given good g is determined as follows. Each purchase decision represents an independent, discrete choice between each of the varieties $j = 1, 2, \dots, J$ governed by prices and a vector of random utility terms, $\theta \in \mathfrak{R}_+^J$. It is a *discrete* choice because each purchase decision must be satisfied by

⁸Note that n_t may not be equal to the total number of purchases due to the restriction that every n_t^g is an integer number. For the ease of the notation we will treat n_t as the actual number of purchase decisions: the difference is negligible once the number of purchases is high.

a single country—albeit in the quantity of choice. It is an *independent* choice each purchase for good g has its own idiosyncratic type θ , drawn from distribution F^g independently of the type of other purchases in the same or different good or period.

The non-negative vector θ determines the agent's preferences over the country-specific varieties $j = 1, 2, \dots, J$,

$$u(x, c; \theta) = \left(\sum_J \theta_j x_j \right)^\eta c^{1-\eta}$$

where x_j is the individual's purchase of each variety $j = 1, 2, \dots, J$ and c is the demand of the numeraire good. Each purchase is associated with a budget $y_t = Y_t/n_t$, where Y_t is the appropriate measure of aggregate income.⁹ We assume aggregate income grows at a constant rate $1 + \gamma_y$.

Let $p_t^g = \{p_{jt}^g : j = 1, 2, \dots, J\}$ be the vector of country prices for good g at date t . A single variety must be chosen for each purchase decision. First, we solve for the indirect utility function v_{jt}^g associated with choosing the country j variety. The optimal quantities x_j and c must then solve

$$v_{jt}^g = \max_{x_j, c} \{ \eta \ln x_j + (1 - \eta) \ln c : p_{jt}^g x_j + c \leq y_t \}. \quad (3)$$

The solution is

$$x_{jt}^g = \frac{\eta y_t}{p_{jt}^g}, \quad (4)$$

or, in terms of revenues, $y_{jt}^g = \eta y_t$. The Cobb-Douglas specification between the variety and the numeraire implies that the demand of the numeraire good, $c = (1 - \eta)y_t$, does not depend on the commodity chosen.

Given the vector $v_t^g = \{v_{jt}^g\}_J$, the optimal choice among country varieties associated with type θ purchase decision is given by

$$V_t^g(p_t^g; \theta) = \max \{ v_t(p_{jt}^g) + \ln \theta_j : j \in \{1, 2, \dots, J\} \}.$$

This is the usual structure of random utility choice models, featuring a common term v_t^g determined by prices (and perhaps other observable variables) and a random utility vector capturing the idiosyncratic nature of each choice.

⁹It is possible to introduce heterogeneity regarding y_t . In our specification, y_t is irrelevant for the choice of the variety, though it obviously determines the amount spent on the chosen variety as well as the numeraire good.

Note how the choice model lays out the foundations for the gains from new varieties. A new variety $J + 1$ comes with a new draw $\ln \theta_{J+1}$. Quite obviously, the maximum across draws, V_t^g , is increasing with the number of independent draws—that is, the number of varieties—as there is always the chance the random utility term $\ln \theta_{J+1}$ comes high enough that the new variety becomes the optimal choice. How likely this is will depend on the distribution F^g as well as the price differences across varieties.

Let $\ln \theta_j$ be independently and identically distributed according to a Gumbel (or type I extreme value) distribution F^g ,

$$F^g(m) = \exp \left(- \exp \left(- \frac{m}{\mu^g} \right) \right)$$

where $\mu^g > 0$ governs the dispersion of the random utility term. Note that the random utility term for *each* variety is independent of the other variety's terms as well as of the type of other purchases in the same or different time period.

The Gumbel distribution is a particular case of the generalized extreme value distribution. It thus arises naturally as the distribution of the maximum of a (properly normalized) sequence of i.i.d. random variables.¹⁰ Perhaps one can view these underlying random variables as the variety's value for different uses and the random utility term as the variety's value for its best use.

In any case, the Gumbel distribution keeps the demand system tractable due to its properties regarding the maximum order statistic. If $\{\ln \theta_1, \ln \theta_2, \dots, \ln \theta_J\}$ is a sequence of i.i.d. variables with a Gumbel distribution, the probability that country i supplies any given purchase decision for good g is then simply given by

$$s_{it}^g = \frac{(p_{it}^g)^{-1/\mu^g}}{\sum_J (p_{jt}^g)^{-1/\mu^g}} \quad (5)$$

It is immediate to see why the Gumbel distribution is commonly used in discrete choice models. It is perhaps less known that, as Anderson et al. (1987) first pointed out, the discrete choice model presented here serves as a foundation for the CES demand specification as the number of independent purchase decisions tends to infinity.

2.1 Retrieving the CES demand as $n_t = \infty$

For asymptotically large n_t we can invoke the law of large numbers and treat the probability that any country satisfies a given purchase decision for good g as the

¹⁰See David and Nagaraja (2003) for an introduction to extreme order statistics.

fraction of purchase decisions actually supplied by each country, s_{jt}^g . Each purchase decision is of the amount given by (4), and thus the demand of variety j of good g at date t is given by

$$X_{jt}^g = s_{jt}^g \frac{Y_t^g}{p_{jt}^g}.$$

where $Y_t^g = \alpha^g \eta Y_t$ are the total expenditures in good g .¹¹ Setting $\rho^g = \frac{1}{1+\mu^g}$ we recover the CES aggregate demand formulation,

$$X_{jt}^g = \left(\frac{p_{jt}^g}{P_t^g} \right)^{-\frac{1}{1-\rho^g}} \frac{Y_t^g}{P_t^g} \quad (6)$$

where

$$P_t^g = \left(\sum_J (p_{jt}^g)^{-\frac{\rho^g}{1-\rho^g}} \right)^{-\frac{1-\rho^g}{\rho^g}} \quad (7)$$

is the familiar price index.

The elasticity of substitution relates to the underlying distribution parameter by $\sigma^g = \frac{1+\mu^g}{\mu^g}$. When μ^g is low, the random utility terms show little dispersion and the optimal choice is mostly driven by the price, that is, varieties become close substitutes—as we will expect from a high elasticity σ^g . Conversely, when the random utility terms are highly dispersed (high μ^g), the choice of a variety is not very sensitive to prices. Note how the μ^g also links the elasticity of substitution with the magnitude of the gains from new varieties, even if prices are constant across old and new varieties. The higher μ^g , the more skewed to the right the Gumbel distribution is and new draws become more valuable.

A well-known property of the CES demand system is that we have that the reservation price for any variety is infinite, that is, there is positive trade $X_{jt}^g > 0$ for all g, j and periods t . In other words, all available varieties are traded all the time. This property is the key identification scheme used in Feenstra (1994) and Broda and Weinstein (2006): whenever a category is observed to have zero sales, it must be it is simply unavailable. We next show this not to be the case when the number of purchase decisions is finite.

2.2 Finite number of purchase decisions $n_t < \infty$

For a finite number of purchase decisions, the demand for each variety j is a non-degenerate random variable. Define z_{jt}^g as the number of purchase decisions satisfied

¹¹We thus confirm our model aggregates up to a Cobb-Douglas specification across goods.

by country j at date t in good g . The vector $z_t^g = \{z_{jt}^g : j = 1, \dots, J\}$ is a random variable distributed according to the multinomial distribution

$$\Pr(z) = \frac{n_t^g!}{z_1! z_2! \dots z_J!} (s_{1t}^g)^{z_1} (s_{2t}^g)^{z_2} \dots (s_{Jt}^g)^{z_J}, \quad (8)$$

with $\sum_J z_{jt}^g = n_t^g$.¹² Each purchase decision for variety j will demand x_{jt}^g units as given by (4), and thus the demand for country j variety of the good g adds up to

$$X_{jt}^g = \frac{z_{jt}^g \eta Y_t}{p_{jt}^g n_t}.$$

While the demand for each variety is a random variable, the total expenditures in good g is deterministic, as it is pinned down by n_t^g . Substituting n_t^g using (1) we confirm that a share $\eta \alpha^g$ of total expenditures is spent in good g . This is a convenient property, which ensures the market share of each good g does not depend on the realization of z_t^g . Letting $Y_t^g = \alpha^g \eta Y_t$, we can rewrite the demand for variety j as

$$X_{jt}^g = \frac{z_{jt}^g Y_t^g}{p_{jt}^g n_t^g}. \quad (9)$$

The market share for variety j within good g , $p_{jt}^g X_{jt}^g / Y_t^g$, is a random variable with mean equal to s_{jt}^g and variance $s_{jt}^g (1 - s_{jt}^g) / n_t^g$. Due to the underlying multinomial distribution of z_t^g , the demand of two varieties i, j of the same good g are negatively correlated, with covariance $-s_{jt}^g s_{it}^g / n_t^g$. The reason is one purchase decision supplied by variety i is one less purchase decision that may end up being supplied by variety j . Note how both variance and covariance converge to zero as n_t^g tends to infinity, that is, we recover the standard CES model.

Importantly, the demand X_{jt}^g can be zero with positive probability. The variety is available, $s_{jt}^g > 0$, yet none of the purchase decisions may be supplied by country j at period t . A missing variety in our model constitutes a *sample* zero in a precise sense: we may observe all purchase decisions for some dates, but as long as we do not observe them *forever*, we find ourselves with a finite sampling period. That is, zero demand events have vanishing probability only as n_t^g tends to infinity—which requires an infinitely long sampling period.

¹²Formally, $\Pr(z) = 0$ if $\sum_J z_{jt}^g \neq n_t^g$

2.3 The extensive margin

Next we provide the model's key formulas regarding the extensive margin. Let d_{jt}^g be an indicator variable equal to one if $X_{jt}^g > 0$, zero otherwise. The probability that $d_{jt}^g = 0$, that is, variety j in good g at date t is not traded, is given by

$$\Pr(d_{jt}^g = 0) = (1 - s_{jt}^g)^{n_t^g}. \quad (10)$$

The probability (10) appears often in our computations so it is worth a brief discussion. The formula is quite intuitive: the probability country j does not supply one purchase decision is $1 - s_{jt}^g$. Since each purchase is associated with an independent draw of θ , the probability that none of the n_t^g purchase decisions were satisfied by country j is just the product of the probabilities of each independent event. The probability of a zero demand for variety j in good g increases with its price as s_{jt}^g decreases. The number of purchase decisions in such good g , n_t^g , also has the expected comparative statics: the more purchase decisions, the less likely that variety j ends up with zero demand.

We will be interested in the expected number of varieties imported for a given good; or the expected number of products that are imported from a given country. Both calculations can be obtained as the appropriate sum of the random variables, d_{jt}^g . While these are not independent of each other, the expectation operator is linear. For the expected number of varieties for good g at date t , we have

$$E \left\{ \sum_J d_{jt}^g \right\} = J - \sum_J (1 - s_{jt}^g)^{n_t^g} \quad (11)$$

where we have made use of the fact that $E d_{jt}^g = \Pr(d_{jt}^g = 1)$. The formula for the expected number of products that are imported from country j is similar, $E \left\{ \sum_G d_{jt}^g \right\} = G - \sum_G (1 - s_{jt}^g)^{n_t^g}$. Indeed the same derivation applies to any set of commodities we want to aggregate. We can also readily compute the expected number of periods with a positive import for a country-good pair.

For our analysis of the churning of varieties we need to compute the expected number of varieties that enter and exit between two periods, t and $t + 1$. Once again the indicator variable d_{jt}^g comes handy. Entry of variety j in good g at date $t + 1$ is equivalent to the event $(1 - d_{jt}^g)d_{jt+1}^g = 1$. Since purchase decisions are independent across time, the probability is simply given by

$$\Pr(d_{jt}^g = 0) \Pr(d_{jt+1}^g = 1) = (1 - s_{jt}^g)^{n_t^g} \left(1 - (1 - s_{jt+1}^g)^{n_{t+1}^g} \right). \quad (12)$$

We can then compute the expected number of varieties j that enter in good g as

$$E \left\{ \sum_J (1 - d_{jt}^g) d_{jt+1}^g \right\} = \sum_J (1 - s_{jt}^g)^{n_t^g} \left(1 - (1 - s_{jt+1}^g)^{n_{t+1}^g} \right). \quad (13)$$

The same steps lead to the expected number of varieties that exit in good g ,

$$E \left\{ \sum_J d_{jt}^g (1 - d_{jt+1}^g) \right\} = \sum_J \left(1 - (1 - s_{jt}^g)^{n_t^g} \right) (1 - s_{jt+1}^g)^{n_{t+1}^g}. \quad (14)$$

Both (13) and (14) can be easily written for the expected number of goods for a given country, or for two arbitrary dates. The expected net entry is simply the difference between the two.¹³

Finally, we show how to compute the survival function. The probability that a variety j in good g has positive sales for periods $t, t+1, \dots, t+k$ is equal

$$\Pr(d_{jl}^g = 1 : l = t, t+1, \dots, t+k) = \prod_{l=t}^{t+k} \left(1 - (1 - s_{jl}^g)^{n_l^g} \right)$$

since purchase decisions are independent across time. This property also implies the exit hazard rate at date $t+k$ is simply equal to $\Pr(d_{jt+k}^g = 0)$ or $(1 - s_{jt+k}^g)^{n_{t+k}^g}$. If prices and total purchases are constant over time, ours is a constant hazard model.

In practice, we will not be able to compute the survival function at the variety-good level from the data. Instead we will be aggregating across varieties, goods, or both. Unfortunately, there are no simple expressions for the survival function if there is price or good heterogeneity: the distribution *conditional* on survival for k periods within, say, good g is no longer given by (8). We can make some progress towards the model counterpart for the empirical survival function by computing the number of varieties (or goods) expected to last at least k periods. We rewrite the survival event as $\prod_{l=t}^{t+k} d_{jl}^g = 1$ to obtain

$$\begin{aligned} E \left\{ \sum_J \prod_{l=t}^{t+k} d_{jl}^g \right\} &= \sum_J E \left\{ \prod_{l=t}^{t+k} d_{jl}^g \right\} \\ &= \sum_J \prod_{l=t}^{t+k} \left(1 - (1 - s_{jl}^g)^{n_l^g} \right) \end{aligned}$$

¹³We have to be careful if we want to compute the expected entry and exit as a rate over previously traded varieties. The latter is a random variable, which can take value zero with positive probability. To avoid this, we will normalize entry and exit flows by the respective cardinality of the commodity space, J or G , or resort to simulation.

for adding up within a good g .¹⁴ We can also easily compute the expected number of varieties to exit for first time at date k ,

$$\sum_J \prod_{l=t}^{t+k-1} \left(1 - (1 - s_{jl}^g)^{n_l^g}\right) (1 - s_{jt+k}^g)^{n_{t+k}^g}.$$

We should make clear these expressions are not the survival function and exit hazard rate, respectively. They do, however, provide a good approximation by treating the expectations as the actual realization. We will return to this issue in Section 5 we will resort to simulation to compare with the data and confirm that the approximation is quite good.

2.4 The intensive margin

The model has also implications regarding the intensive margin, that is, the amount traded in a given variety-good pair conditional on such flow being positive. Conditional on $d_{jt}^g = 1$, the expected spending in variety j in good g can be easily calculated by noting that $E z_{tj}^g = \Pr(d_{jt}^g = 1) E \{z_{tj}^g | d_{jt}^g = 1\}$. We thus obtain

$$E \{Y_{jt}^g | d_{jt}^g = 1\} = \frac{s_{jt}^g}{1 - (1 - s_{jt}^g)^{n_t^g}} Y_t^g. \quad (15)$$

We can actually easily characterize the expectation on the additional condition that a set J^* of varieties has no sales:

$$E \{Y_{jt}^g | d_{jt}^g = 1, d_{it}^g = 0 : i \in J^*\} = \frac{s_{jt}^g Y_t^g}{(1 - s_{J^*t}^g) \left(1 - \left(1 - \frac{s_{jt}^g}{1 - s_{J^*t}^g}\right)^{n_t^g}\right)},$$

where $s_{J^*t}^g = \sum_{J^*} s_{jt}^g$. This formula captures the expected demand per traded variety—the intensive margin—conditional on observing no trade in a set of varieties—the extensive margin.¹⁵

We would also like to provide measures of exit and entry of new products in terms of spending. This can be done easily with (15) at hand. Say, for example, that we want to compute the expected spending in new varieties j in good g at date $t + 1$.

¹⁴The formulas for adding up within a country j are similar. We omit them for the sake of brevity.

¹⁵It is also possible to compute the expectation on the additional condition that all varieties $j \notin J^*$ display strictly positive sales.

As in (13), we rewrite the expectation as $E \left\{ \sum_J (1 - d_{jt}^g) d_{jt+1}^g Y_{jt+1}^g \right\}$. Using the law of iterated expectations, and noting that Y_{jt+1}^g is independent from d_{jt}^g , we obtain

$$E \left\{ \sum_J (1 - d_{jt}^g) d_{jt+1}^g Y_{jt+1}^g \right\} = \sum_J (1 - s_{jt}^g)^{n_t^g} s_{jt+1}^g Y_{t+1}^g.$$

Following the same steps we can compute the spending at date t for varieties disappearing at date $t + 1$, as well as equivalent formulas if we wish to aggregate across countries instead of goods. We can also express the spending associated with exit and entry of varieties in good g as a *rate* on total spending in good g in the previous and current period, respectively.

Finally we would like to compute the expected spending per variety j conditional on positive trade for k periods. Recalling that the survival event can be written as $\prod_{l=t}^{t+k} d_{jl}^g = 1$, and following similar steps as before, we find that

$$E \left\{ \sum_J \left(\prod_{l=t}^{t+k} d_{jl}^g \right) Y_{jt+k}^g \right\} = \sum_J \left(\prod_{l=t}^{t+k-1} \left(1 - (1 - s_{jl}^g)^{n_l^g} \right) \right) s_{jt+1}^g Y_{t+1}^g.$$

Unfortunately, we cannot derive an analytic expression for the expected rate over the initial spending, and we will have to resort to simulation to obtain the model's counterparts in the data.

2.5 Missing domestic varieties

In principle, the set of J of country-specific varieties should include the *domestic* variety as well. Unfortunately, there are no data on U.S. spending on domestic products at the desired level of detail. This data shortcoming is ever-present in the literature. To get around it, domestic varieties are commonly assumed to enter the agent's preferences as a composite at the highest level of aggregation, instead of at the product level.¹⁶ With this assumption it is then possible to compute the welfare gains associated with new import varieties within the data restrictions.

While our model allows this approach, we will not pursue it. We will instead simply treat the demand for domestic varieties as an unobserved parameter. Doing so, the model implications for the foreign varieties are isomorphic to those for the full set of varieties. Let $j = 1$ be the domestic variety and n_{1t}^g the number of purchase decisions being captured by the domestic variety. The distribution over the $J - 1$

¹⁶For example, Broda and Weinstein (2006).

foreign varieties, conditional on n_{1t}^g , is still a multinomial with parameters $n_t^g - n_{1t}^g$ and probabilities equal to

$$\hat{s}_{jt}^g = \frac{s_{jt}^g}{1 - s_{1t}^g}$$

for $j > 1$. Thus the probability of a given vector $\hat{z}$ over $J-1$ is still given by (8) simply by substituting the conditional set of probabilities. The formulas for the extensive and intensive margin carry over after the appropriate substitutions.

3 Data and calibration

3.1 Data

Our main data source are the U.S. imports merchandise trade files provided by the U.S. Census Bureau for the years 1990 to 2001.¹⁷ The data contain all goods valued at more than \$2,000 per commodity shipped by individuals and organizations (including importers and customs brokers) into the U.S. from other countries. The records are compiled in terms of commodity classification, values and country of origin as well as several other dimensions.

Regarding the product classification, we will work with the Harmonized Tariff Schedules of the United States at the ten-digit level (HS10 henceforth).¹⁸ Since their introduction in January 1989, the HS10 classification spans more than 18,000 codes. The classification, though, is dynamic as either the World Customs Organization or the United States International Trade Commission introduced new product codes or declared some obsolete. To address this, we keep only those HS10 product codes that have positive trade in 1990-2001.¹⁹ We make two additional adjustments. We drop chapters 98 and 99 containing special classification provisions and estimates of low-value imports, among others. We also drop products under the four-digit headline

¹⁷Data and documentation are available at <http://www.census.gov/foreign-trade/about/index.html>.

¹⁸The World Customs Organization assigns 6-digit codes for general categories. Individual countries can then define their own codes at more detailed levels. The United States International Trade Commission administers imports codes for the U.S. See <http://www.usitc.gov/> for detailed information.

¹⁹Some of the remaining product codes may have seen their definition change as they merged or split into other HS10 codes. Using the algorithm provided by Pierce and Schott (2012), we found these occurrences to be rare and involve a very small amount of trade in our selection of product codes.

numbers 2710 to 2716: these codes include petroleum, fuels, and electric energy. These goods constitute an important part of total U.S. imports but are clearly out of place in our analysis. A total of 10537 HS10 product codes are in the final data set we use.

We also restrict our list of countries as follows. As done with product codes, we drop all countries that drop or appear during the period 1990-2001. We do make an exemption with Germany, assigning it the imports from West Germany in 1990.²⁰ We further restrict our data to the largest 120 U.S. import partners over all the sample, accounting for more than 99 percent of total imports. We also append some gravity and GDP data, obtained from the CEPII website and completed with some additional GDP data from World Bank.

3.2 Calibration

We aim for a calibration as parsimonious as possible. We start with the parameters governing spending across goods, $\{\alpha^g : g = 1, 2, \dots, G\}$. The counterpart in the data is easily available, so for each HS10 product code we pin down α^g to its average share in total U.S. imports.²¹ The defining feature of the product share distribution is its skewness: there is a very large fraction of very small products. It is not clear these are economically meaningful classifications, yet we lack a systematic method to establish which are so.

Recall there is a one-to-one correspondence between our dispersion parameters μ^g and the elasticity of substitution, σ^g . We exploit this and obtain the good-specific elasticities of substitution from Broda and Weinstein (2006). There are several product codes for which there are no estimates available: we assign them the average elasticity for their respective four-digit headline classification. The mean and median elasticities are 19.4 and 4, respectively. It should be mentioned that close to 40 percent of all products have an elasticity of substitution below 3. There is also a substantial amount of products with very high elasticities too, in excess of 100, for example.

Our model requires taking a stand on the full distribution of prices. The data, at best, contains only average unit prices for the observed trade flows in a given period. To bridge both, we posit a parsimonious yet quite flexible specification for prices,

²⁰East Germany is present in 1990 but has no recorded trade.

²¹Most product shares are stable in the period 1990-2001: only very small products display considerable variation.

allowing for country, good, and time fixed-effects:

$$\ln p_{jt}^g = \ln \bar{p} + \ln \tau_j + \ln \tau_g + \ln \tau_t. \quad (16)$$

While (16) can be used to fit the available data, it turns out that the only source of price heterogeneity that is relevant for the model are the country fixed-effects. Inspecting (5) is clear that both good and time fixed-effects cancel in the calculation of the probabilities s_{jt}^g . These, in turn, determine every other moment in the model. Inverting the process, it is possible to recoup the county fixed effects from the share of products imported from each country.²² Fortunately, we have a good handle on the sources of import price differences across countries, both due to trade costs and real wage variation. We check our calibrated fixed effects against a gravity-style regression featuring to distance, size, a dummy for the presence of a common border and per-capita income differences. We find that the fit is very good, and all coefficients have the expected sign.²³

Finally we turn to the key parameter of the model: the number of purchase decisions, n_t . The defining feature of a purchase decision is that it is an independent event. This makes it difficult to map n_t directly to the data. For example, a large firm may decide to buy a specific input from a foreign country, resulting in a large order that may be spread over many shipments over the course of several months. At the same time, it would be naive to equate n_t to the number of firms or households: clearly a retail chain may decide to stock its shelves with a several computer models, possibly produced in as many countries.

We decided instead to set n_0 to match the total count of good-variety pairs with positive trade in 1990. There were approximately 128,000 good-variety pairs with positive trade—a little more than 10 percent of all possible good-variety pairs. We find that we need to set n_0 in the neighborhood of 1.6 million. Together with our calibration of the country fixed-effects, we have thus matched the data regarding the count of varieties per country in 1990, that is, the extensive margin across countries in the first year in our sample.

We are left with three parameters: the expenditure share of the numeraire η and the growth rate of purchase decisions and aggregate income, γ_n and γ_y respectively.²⁴ We set $\eta = .3$ so 70 percent of expenditures are in domestic, non-tradeable goods;

²²We target the average share over the period 1990-2001. Country shares are very stable over time, with the notable exemption of Mexico—which shall receive its due attention in Section 6.

²³Indeed, we also computed the model using the fitted values from the gravity regression as the country fixed-effects. We obtain very similar results.

²⁴The initial level of income Y_0 is normalized to 1.

Description	Parameter	Target/Source
Preferences across goods	$\{\alpha^g\}_G$	HS10 expenditure shares
Elasticities across goods	$\{\sigma^g\}_G$	Broda and Weinstein (2006)
Country fixed effects	$\{\tau_j\}_J$	Products traded per country
Purchases decisions in 1990	$n_0 = 16 \times 10^6$	Good-varieties traded in 1990
Growth rate income	$\gamma_y = .035$	Real GDP growth 1990-2001
Growth rate purchases	$\gamma_n = .065$	Import penetration 1990-2001
Preference non-tradeables	$\eta = .3$	Expenditure share non-tradeables

Table 1: **Summary calibration**

and set γ_y to the average growth rate of real GDP over the period, 3.5 percent. It turns out neither of these parameters have virtually any impact on our results. It is instead the growth rate of purchase decisions, γ_n , that is key. Recall that n_t should measure the number of purchase decisions supplied abroad, as we do not have access to detailed domestic varieties. The natural target for the calibration is thus the sharp increase in import penetration in the U.S. over the period 1990-2001. Real imports grew at a torrid rate over the period, in excess of 10 percent on average. We match the observed real import growth then by setting $\gamma_n = .1 - \gamma_y = .065$.

Table 1 summarizes our parameter choices and the targeted fact or source which we link the parameter to. We should note that all country fixed-effects and the number of purchase decisions have to be determined jointly by solving for the model. The link, though, with the stated targets is tight enough for the purpose of the exposition.

4 Cross-sectional results

We have calibrated the parameters to target the extensive margin across countries and the intensive margin across products. The first step is to check that the model replicates successfully the cross-section distribution of the extensive margin across products we observe in the data.

Let us start by looking at the distribution of the number of varieties per product. By construction, the model exactly matches the average number of varieties per product in 1990. Figure 1 shows the model's fit extends to the overall distribution. The dashed line plots the data by 5 percentile intervals between the 10th and the 90th.²⁵ The solid line corresponds to the distribution of expected number of varieties

²⁵Data are for 1990: there is little variation across years.

across products. Model and data are rarely off by more than 1 variety: since the data will necessary come at discrete values, one could argue that the fit is essentially spot on. The model does come a tad below the data around the median (7.8 versus 9) and makes it up by over-predicting the right tail.




Figure 1: **Distribution of number of varieties per product**

While the model reproduces very well the overall distribution, we want to make sure it is the right products that have the right amount of expected varieties. We will look then at the number of varieties by product market share and the elasticity of substitution. Regarding the former, data display a clear positive relationship: products with larger market share feature more countries supplying the good. We sorted the products in 20 categories with an equal number of products, according to their market shares.²⁶ Thus the first bin contains products with a market share equal or smaller than the 5th percentile, the second bin contains products with a market share bigger than the 5th percentile but smaller than the 10th percentile, and so on. Products in the median category have an average of ten varieties, about double of what products in the first quartile do, and 50 percent less than what products in the last quartile do. It is not surprising that the model reproduces this very well: in

²⁶There are thus more than 500 products per category. For all calculations we have used data for the year 1990.

our calibration, larger products simply have more purchase decisions, and thus the expected number of varieties is increasing in the product market share.

Perhaps a more striking data feature is that there remains a large amount of dispersion in the number of varieties in each size category. For example, the size category corresponding to the median number of varieties (about 10) has an inter-quartile range of 5 and 15 varieties per product. The dispersion is also clearly increasing in the product share.




Figure 2: **Distribution of number of varieties per product market share**

Figure 2 plots the first to last quartiles for each product size category for both data (dotted line) and model (solid line). With the exception of the right tail for very small products, the model tracks the data with uncanny precision. As in the data, the inter-quartile range expands with the product market share. Note that because the calibration matches the product market shares in the data, each size category contains exactly the same set of products in the model and in the data.

In the model, the only source of variation in the expected number of varieties per product, given its market share, is the elasticity of substitution. There is no correlation between the elasticities and market shares. However, we find that products with higher elasticities tend to have fewer varieties with positive trade. The dispersion in the number of varieties per product also has a (weak) negative relationship with the elasticity of substitution.




Figure 3: **Distribution of number of varieties per product market share**

Figure 3 reproduces the plot in Figure 2, but this time products have been classified according to their elasticity of substitution.²⁷ The dotted lines display the first to last quartiles. Both the median and the top quartile are decreasing with the elasticity of substitution—though the bottom quartile displays only a slight downward slope.

The model performance is striking. The median and bottom quartile are virtually spot on with the data thorough the range of elasticities. The model overpredicts the amount of skewness for low elasticities and has a substantially sharper fall in the dispersion for higher elasticities. Still, the top quartile does not wander too far from the data. Recall the elasticities were estimated by Broda and Weinstein (2006) using unit price data.

What is behind the decreasing relationship between elasticities and the extensive margin? As it is well-known to anyone familiar with the standard CES demand system, a higher elasticity given prices will lead to a more skewed distribution of market shares. Intuitively, suppliers with lower prices capture a larger market share the closer substitutes the commodities are. In our framework, a high elasticity corresponds to a low dispersion parameter μ^g and the same comparative statics lead to a more skewed

²⁷We only display the range from 2 to 16, covering 70 percent of the total number of goods. The elasticity estimates are very skewed themselves, with close to 500 products having elasticities above 100. For higher elasticities there are not enough products to acquire a clear picture of the conditional distribution.

probability distribution s_t^g , as it can be easily confirmed from (5). For the next step, we reproduce the formula for the expected number of varieties for a given good g , (11):

$$E \left\{ \sum_j^J d_{jt}^g \right\} = J - \sum_j^J (1 - s_{jt}^g)^{n_t^g}.$$

A more skewed probability distribution means consecutive purchase decisions are more likely to repeat suppliers. For example, imagine that one particular variety enjoys an overwhelming advantage, say $s_{1t}^g = .99$. As most of the purchase decisions will choose variety 1, we will not expect many of the remaining $J - 1$ to have any positive sales. At the other extreme, we would expect two given purchase decisions to be satisfied by a different suppliers in an hypothetical symmetric case, $s_{jt}^g = 1/J$. Thus, given a constant number of purchases n_t^g , the more skewed the probability distribution is, the less number of varieties are expected. We believe this to be an excellent example of the kind of insights that are to be gained by considering a finite number of purchase decisions.

Finally, we also confirm that the model's prediction for the intensive margin—that is, the country's market share in revenues—line up with the data very well. The model's prediction for both quartiles and the median are within half a percentage point from the data. The model also reproduces the observed skewness in the distribution of country's market shares: the ranking of the largest 5 trading partners coincides in both data and mode, and together they account for the same fraction of total trade—about half.

5 Entry, exit and survival

We now turn to the model's performance over time and compare it with the data. The first question has to be whether the model correctly predicts the long-run growth of the extensive margin in trade. In the decade from 1990 to 2001, we saw the number of traded country-product pairs increase from 127,500 to 162,500.

Figure 4 plots the expected number of good-variety pairs traded in the model and the data, in thousands, from 1990 to 2001. We did not target this data series directly in any year other than 1990, yet the model is virtually on the top of the data. The model predicts the average growth rate of the extensive margin within half a tenth of a percentage point: 2.18 percent versus 2.23 percent in the data.

This result is more nuanced than it may appear at first sight. The predicted growth rate of traded varieties is only about one third of the growth rate of purchases, 6.5




Figure 4: **Good-variety pairs with positive trade**

percent. The difference in growth rates is due to some of the new purchase decisions in each good being assigned to varieties that had already positive trade, and thus do not add to the total count of traded varieties. Given that only about 10 percent of the possible varieties were traded, a tempting back-of-the-envelope calculation would lead us to conclude that only 10 percent of the new purchase orders would end up in previously-traded varieties—and the growth rate of traded varieties would have to be near 5 percent. Indeed, this would be an accurate calculation if, for each good, *all varieties had the same price*. But there is instead a large amount of price variation across countries. Some varieties are thus more attractive than others, that is, have a higher probability of capturing any given purchase. We are to expect that previously-traded varieties are those that are more attractive and thus bound to absorb a large fraction of the new purchases. The more skewed the underlying probability distribution, the slower the growth rate of the extensive margin. Reproducing the overall growth rate of the extensive margin is a strong reassurance for our calibration.

We performed some robustness analysis by varying some of the parameters. Alternative values for the growth rate of purchase decisions between 4 and 8 percent implies that the extensive growth rate is between 1.71 and 2.6 percent. We should also emphasize that model's predictions regarding the growth rate depend on the level of n_0 as well. For example, we set n_0 such that only 6 percent of the varieties are traded in 1990: the predicted growth rate of the extensive margin increases to 2.7

percent. Going the opposite way, setting n_0 such that 15 percent of the varieties are traded in 1990 cuts down the predicted growth rate to 1.7 percent. The reason is quite straightforward. The more varieties are previously traded, the harder it is that the new purchases count towards the extensive margin.

5.1 Entry and exit

A striking feature of the data is the large number of varieties that start or stop being traded in any given year. As mentioned earlier, the total number of traded varieties increases a little bit more than 2 percent per year. On average, though, a little bit less of 23 percent of the previously-traded varieties stop being so, and the amount of new varieties almost adds up to a quarter of the number of traded varieties. The large count numbers do not have continuity in terms of expenditures: it is mainly varieties with small value that enter and exit. The net change in varieties accounts only for a small 0.3 percent of the previous value, with gross entry and exit being just around 1 percent.

	Data		Model	
	By count	By value	By count	By value
Entry	24.6 %	1.1 %	27.0 %	1.2 %
Exit	22.4 %	0.8 %	24.8 %	1.0 %
Net	2.2 %	0.3 %	2.2 %	0.3 %

Table 2: **Entry and exit**

Table 2 presents the rates for data and model, both by count and by value.²⁸ The model predictions are very much line with the data. Gross flows are very large by count but relatively small by value. We have already seen the model matched the data regarding net entry by count: it is also spot on regarding the trade growth, in value, contributed by net entry. Since we calibrated the model to match total trade value growth, the model fully matches the decomposition between intensive and extensive margin.

The model slightly overstates the amount of churning by count. This is perhaps not surprising: we have no intrinsic persistence mechanism, as purchase decisions are

²⁸Rates by count use the previous period total number of traded varieties; rates by value use the previous period total trade value. All rates are the average across 1990-2001. For the model we have to resort to simulation to compute the exact expected rate: we find, though, that using that taking the ratio of expectations returns virtually identical values.

	Data		Model	
	Goods	Countries	Goods	Countries
Entry	.69	.91	.72	.93
Exit	.54	.91	.68	.93
Net	-.13	-.11	-.08	-.01

Table 3: **Correlations of entry and exit with good and country size**

completely independent across periods. The “persistence” is instead driven by the probability a variety has positive trade: varieties that are either very likely or very unlikely to have positive trade count little toward churning. To see this, consider the probability a variety j in good g enters assuming that there is no growth in purchase decisions so varieties are identically distributed across periods. From (12), the probability is simply $\Pr(d_j^g = 0) (1 - \Pr(d_j^g = 0))$.²⁹ Clearly the expression is concave in $\Pr(d_j^g = 0)$ and thus the probability of entry is maximized at $\Pr(d_j^g = 0) = .5$. It is also symmetric, that is, two varieties with $\Pr(d_j^g = 0) = 1 - \Pr(d_i^g = 0)$ have identical probability of entry. Once we average across varieties—see (13)—the amount of entry will be determined by the fraction of varieties whose likelihood $\Pr(d_j^g = 0)$ is near .5, that is, how many varieties are very uncertain to be observed.³⁰

What is the effect of growth in purchase decisions? By inspecting (13) and (14) we confirm what we would guess, namely, that faster growth leads to more entry and less exit. A perhaps more intriguing comparative static is regard the level of purchase decisions, n_0 . More purchase decisions reduce *both* exit and entry, so it unambiguously reduces churning. We explore some higher values of n_0 and find the model can match the observed entry and exit rates spot on while only slightly overstating the number of traded varieties in 1990.³¹

The model’s predictions for total entry and exit, both measured by count and by value, are very much in line with the data. We now go one step further and ask whether entry and exit across goods and countries are similar in model and data. To

²⁹Due to the no growth assumption, this is also the probability of exit and thus net entry will be expected to be zero.

³⁰The amount of churning is actually very closely related to the entropy of the underlying distribution for z_t^g . Indeed, $\Pr(d_j^g = 0) = .5$ is the point of maximum entropy.

³¹We conjecture the model performance is also slightly hindered by the arbitrary cut-off of December 31st when the data are aggregated by calendar years. If purchase decisions span several months, shipments late and early in their respective years will not be independent as we assumed.

this end we compute entry, exit, and net entry at the good and country level.³² We then see how the gross and net changes correlate with the market share of the good or country.

Table 3 documents the set of correlations for both model and data.³³ Let us first discuss the data. Both entry and exit gross flows are positively correlated with the market share of the underlying good or country. In particular, the correlation is very high with the country market share. Thus, we observe more churning for larger good categories and with large trade partners. Net entry, though, is negatively, albeit weakly, correlated. Hence, there is a substantial amount of net entry being driven by smaller products and trade partners.

The same patterns arise in the model: the predicted correlations are very close to the data, with perhaps the exceptions of exit flows across goods and the net entry across countries. As in the data, correlations are stronger for country market shares. The correlation with net entry is negative, albeit it is virtually zero for country market shares.

What is behind the positive relationship between churning and country/good market share? As discussed earlier, the amount of churning is highest when the probability of observing a given variety is close to .5. The vast majority of varieties in our model have a much lower probability than .5: after all, the model is calibrated such that just 10 percent of them are expected to have positive trade. Now, the distribution of probabilities shifts to the right for goods or countries with larger market share and thus most varieties get closer to the maximum entropy level of .5. As a result, churning increases for these categories. Note we could have ended up with a negative correlation if all varieties were very likely to be observed: then a rightward shift would have set most probabilities away from .5 and thus decreased churning.

The explanation for the negative correlation of net entry is different. Our model features a strong non-linearity regarding the number of purchase decisions. If there are very few purchases initially, every new one is virtually guaranteed to imply a new variety being traded, as the likelihood the new purchase order ends up being supplied by a previously-traded is quite low. At the other extreme, when most varieties are

³²We compute the flows as rate over the total good/country possible pairs (J or G) instead of over existing varieties. The latter are simply too volatile at the disaggregated level, with rates often hitting 100 percent.

³³To compute the correlations in the model we have to resort to simulation. This is unfortunately a quite costly endeavor since the whole set of goods must be simulated, simultaneously, from a multinomial with 120 categories (countries). To cut computing time we just use two consecutive periods, instead of averaging over 10 periods. Data are presented correspondingly for 1990-1991 but there is little variation across year pairs. We use 1000 draws.

already traded, new purchases are very likely to be satisfied by a previously traded category. Thus, despite all categories having the same growth in purchase decisions, the expected growth in varieties is faster in smaller categories.³⁴

5.2 Survival analysis

We complete our analysis by tracking the 1990 cohort of traded varieties over the entire sample. We will then compare the “survival” probability and exit hazard rates in the data with their model counterparts. Special attention will be paid to the composition of the varieties over time that are continuously traded over time.

Let us start with the survival analysis with the data. Figure 5 displays the empirical survival probability and hazard function, in a dashed line, in the top and bottom panels, respectively. The survival probability drops steeply in the first and second years, and then slopes down only very gently. After ten years, about half of the varieties with positive trade in 1990 have been traded continuously for all the years in between. The exit hazard rate also falls steeply. A quarter of the traded varieties in 1990 are dropped in 1991 but, among those surviving for five years, less than 6 percent do. The hazard rate seems to settle down to a very low 2 percent after 8 years.

Figure 5 also plots the model’s predictions with a solid line. Once again model and data lie very close together. The model replicates the steep drop in both survival probabilities and the hazard function. The survival probability also becomes practically flat after five years. Looking at the hazard function we see the model slightly overstates the exit rate in the first year but then tracks the data quite well. In the long run, the model’s hazard function dips below 1 percent, undershooting the data.

In Section 2 we described how to approximate the survival probability. The forces behind the downward sloping survival probability and hazard function are purely those of selection: the underlying likelihood of a variety of being trade is essentially constant in our model.³⁵ We should expect the less likely varieties to drop earlier, and thus the remaining varieties have a decreasing probability of exit. These dynamics are well known and we would not dwell further on them.

The selection mechanism also appears when we look at the intensive margin of the surviving varieties. Figure 6 plots the average value associated with varieties that have been continuously traded, for data and the model. It should not be surprising that it is increasing both in the data and the model. Note the magnitudes are substantial,

³⁴The correlation is weak simply because net entry has a lot of variance in the model.

³⁵It does increase, slightly, because of the growth in purchase decisions.




Figure 5: **Survival probability and hazard function**

and the model does an excellent job at matching the data quantitatively. After five years the surviving variety is about 32 percent larger than the original cohort in 1990—in the model, it is just below 35 percent. After ten years, the data shows an increase of 44 percent versus 46 percent in the model.³⁶

Finally we look at the cumulative entry and exit since 1990. That is, for every year in the sample after 1990 we compute the number of varieties present in that year which were not traded in 1990 for entry; and similarly for exit. Note the main difference with the previous survival analysis is that we do not require the disappearing varieties to have been traded until the current year, or new varieties to appear for first time then. For example, a variety that was traded, say, in 1990 and 1991 as well as in 1994 but not in 1992, would count toward exit in 1995 if it is not traded in 1995.

Figure 7 plots the data—in a dashed line—and the model—in a solid line. Model and data still line up quite well, but we can no longer claim the model delivers a good quantitative fit. This is, in part, because 1991 turns out to be quite an outlier in the sample, with very low entry—actually the net change in varieties was zero. As a result the model predicts too much entry early on and too little later on.

It is also apparent from Figure 7 that while the model predicts a sloping-down

³⁶Growth in purchases also contributes to the upward slope in Figure 6, yet its impact is quite small since it also contributes to reduce the likelihood that less likely varieties exit fast.




Figure 6: **Survivor average trade value**




Figure 7: **Cumulative entry and exit since 1990**

cumulative exit, in the data the cumulative exit rate seems to stabilize.³⁷ In contrast,

³⁷Using later years as the benchmark shows the cumulative exit rate does slope down, but substantially less than in the model.

the model seems to predict the cumulative entry should grow at a slightly slower rate than it does in the data. Recall that the model is pretty much spot on regarding survival, as shown in Figure 5. What the model is having some trouble reproducing are the probabilities of re-entry and re-exit. In particular, the data features a substantial amount of varieties that permanently exit after 1990, as well as a fraction of varieties that do not exit after seeing positive trade for first time after 1990.

We see this as indicative of underlying price movements. Recall our calibration assumed constant relative prices across countries. This is clearly a simplification that served our aim to keep the calibration as parsimonious as possible, yet it is blatantly at odds with the data. Temporary price movements will naturally bring the model more in line with the data by introducing mean-reversion effects. Varieties that had a (temporarily) low price in 1990 would be likely to appear in 1990 but, once their price returned to the mean, may not be expected to re-appear. Similarly, varieties with a temporary price increase would likely return permanently to positive trade later in the sample once their prices return to the mean. While the model allows for price dynamics, we would need substantial additional structure in order to infer the counter-factual prices. In the next Section we will instead look at a particular episode for which we do have a good grasp of the price dynamics.

6 Entry and exit during NAFTA

The North American Free Trade Agreement (NAFTA) is one of the most studied trade liberalizations episodes in the literature.³⁸ NAFTA called for the phasing out of virtually all restrictions on trade among the United States, Canada, and Mexico. The biggest impact was regarding Mexico, since U.S. and Canada were already well into removing trade barriers between themselves in accordance to their Free Trade Agreement signed five years prior. U.S. imports from Mexico grew a staggering 70 percent in just three years, virtually doubling Mexico's share of U.S. imports and eventually consolidating its position as the second largest U.S. trade partner, only behind Canada.³⁹

A natural question is what fraction of the trade expansion following NAFTA was due to the introduction of new goods from Mexico and Canada. The answer for Mexico is substantial. In three years 30 percent more goods were imported from Mexico

³⁸See, for example, Romalis (2007), Kehoe and Ruhl (2013), Hilberry and McDaniel (2002), and Treffer (2004).

³⁹To the dismay of every trade economist, the implementation of NAFTA coincided with a large depreciation of the peso, entangling the immediate impact of the tariff reduction.

into the U.S., adding to more than 10 percent in trade value. Recall that the average growth rate for varieties (including all countries) over the period is substantially more modest, of the tune of 2 percent product growth per year, adding to little value overall. The extensive-margin response for U.S. imports from Canada was quite muted in comparison.

We now approach the response to the extensive margin to NAFTA through the lens of our model. We start by comparing the model predictions under the baseline calibration, featuring constant relative prices. In other words, there is no NAFTA-related reduction in tariffs so we can think of the model results as a counterfactual exercise in which only trend growth in trade expands the set of goods with positive trade. Figures 8 and 9 compare the data with the model output for the net changes by count and value, as well as the cumulative entry and exit rates, relative to 1993—the last year before NAFTA.




Figure 8: **Canada after NAFTA: Model and Data**

The model is spot on with Canada, as Figure 8 shows. The net change, by count, is right as predicted; the contribution, by value, is a little bit more noisy but the model remains in the ballpark of the data. Canada is also one of the countries for which the model does an excellent job with the cumulative entry rate. The model is also very close regarding cumulative exit, though the data does not display a downward slope, an issue previously encountered. We conclude from Figure 8 that NAFTA did not have a substantial impact regarding Canada imports to the U.S. This is not really




Figure 9: **Mexico after NAFTA: Model and Data**

surprising, as U.S. tariffs for Canadian goods were already low or nonexistent since 1989.

The contrast with Mexico is remarkable, as shown in Figure 9. The model does not predict anything close to what we observed: net change, both by count and by value, are much larger in the data than in the model. Similarly the cumulative entry rates virtually triple those predicted in the model over the period 1994-2001. Only the cumulative exit rate seems to be in line.

In short, the model is completely missing the impact NAFTA had on U.S. imports from Mexico. We set ourselves to fix this omission by tweaking Mexico's trade costs starting in 1994. In order to match the increase in Mexico's share of total U.S. imports over the sample period we find we have to decrease the country-specific trade costs by 10 percentage points.⁴⁰ We also have to decide on the transition path: a substantial fraction of the import tariffs were phased out in five years or more, as the NAFTA agreement allowed to, yet many tariffs on Mexico exports to the U.S. were immediately eliminated in January 1994. By the end of our sample period, more than 85 percent of Mexico goods entered the U.S. duty free.⁴¹ We settle on a three year transition.

⁴⁰The reduction in Mexico's trade costs due to NAFTA is very much in line with estimates elsewhere, which should not be surprising given that Broda and Weinstein (2006) estimate the elasticities we used over the same period.

⁴¹See CBO (2003).

Note we are not adjusting the underlying growth in purchase decisions. Implicitly, we are assuming no domestic varieties were substituted by Mexican imports: while we doubt this was really the case, we prefer an agnostic calibration regarding this margin as we have no data to contrast any alternative.




Figure 10: **Mexico after NAFTA with tariff reduction**

Figure 10 shows the model's path for the net change in imported products from Mexico after incorporating the tariff cut (solid line), as well as the data (dashed line) and the baseline calibration (dotted line). The model's predicted response captures very well the magnitude in the expansion in the good mix imported from Mexico. It also broadly captures the speed of adjustment. We should note, though, that this is entirely due to the assumed pace of tariff phase outs: the model does not have intrinsic persistence. The model also still underpredicts the importance of the extensive margin by value, only closing the gap in Figure 9 by about half.

We find it surprising that the model matches the data with such ease once tariff cuts are accounted for in the calibration. After all, the extensive margin in our framework adjusts for very prosaic reasons. As Mexican goods get cheaper, they get purchased more often. Thus a larger fraction of them will be observed in the course of a year. We view our results as strongly suggestive that a complete analysis of NAFTA or similar trade liberalization episodes has to specify a micro-founded model of demand at the transaction level.

7 Re-evaluating welfare gains from new varieties

TBC

Feenstra (1994) and Broda and Weinstein (2006) have pioneered the estimation of the welfare gains due to new import varieties under a CES demand system. Available commodities are identified as those observed to have positive trade, as dictated by the CES demand. As we have already seen, all varieties are available in our model yet most are not traded. In this Section we re-evaluate the computation of the welfare gains through the lens of our model.

Feenstra (1994) shows how to derive an exact import price index under a CES demand when the set of varieties available changes across periods. The price index can then be used to compute the welfare gains associated with trade.⁴² For the ease of notation, we drop the superscript g and compare two consecutive periods, $t-$ and t . Let I_t be the set of varieties with positive trade at date t , $I_t = \{i : X_{it} > 0\}$, and $I = I_t \cap I_{t-1}$. Given that prices are constant across periods, the import price index when $I \neq \emptyset$ is given by

$$\pi(X_{t-1}, X_t) = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma-1}} \quad (17)$$

where

$$\lambda_t = \frac{\sum_{j \in I} p_j X_{jt}}{\sum_{j \in I_t} p_j X_{jt}}.$$

When $I = \emptyset$ it is not possible to define an import price index. However for our purposes here we can just set $\pi = 1$ whenever $I = \emptyset$.⁴³

When we have the exact CES demand system, that is, when $n \rightarrow \infty$ in both periods, we have that all commodities are traded in all periods, $I_t = I_{t-1}$. Since there are no price changes, $\lambda_t = \lambda_{t-1} = 1$, and as one would expected the import price index is unchanged, $\pi = 1$. We will obviously find no welfare gains associated with new varieties, since there are none.

If the number of purchase decisions is finite, then λ_t, X_t, I_t are random variables, and so it is π . Thus we should treat the resulting import price index as an estimate, subject to sampling error. We will next ask whether it is a unbiased estimator, and if it is not, what can be said about the sign and size of the bias.

Let us take a closer look at the import price index with finite purchase decisions. Since a constant fraction of income is always spent on the variety of choice, we have

⁴²See Broda and Weinstein (2006) for an application to U.S. imports.

⁴³Feenstra (1994) provides the more elaborated formulas for the case of non-constant prices.

that

$$\lambda_t = \frac{\sum_{j \in I} z_{jt}}{n_t}. \quad (18)$$

When $n_t = \infty$ all commodities would have positive sales, and thus $\lambda_t = 1$. For finite n_t the intersection subset I is a random variable. Conditional on a realization of I , we have that

$$E\{\lambda_t | I\} = \frac{\sum_{j \in I} E\{z_{jt} | I\}}{n_t} \leq s_I$$

where $s_I = \sum_{j \in I} s_j$ is the probability a purchase decision is supplied by a variety in I . Thus $E\{\lambda_t | I\} \leq 1$, with strict inequality whenever a commodity is not included in the intersection subset I . Hence λ_t is biased downwards,

$$E\{\lambda_t\} < 1$$

once we take the expectation across all possible intersection sets I . This simply reflects that there is a positive probability that a variety is not included in the intersection set.

Because of the intersection set I , λ_t and λ_{t-1} are not independent variables. Moreover, the import price index π is a non-linear function of both,

$$\pi = \left(\frac{n_{t-1} \sum_{j \in I} z_{jt-1}}{n_t \sum_{j \in I} z_{jt}} \right)^{\frac{1}{\sigma-1}}.$$

The exact probability distribution of the import price index is thus quite difficult to characterize. Fortunately we are able to characterize the estimator properties under certain conditions. For the general case we will have to resort to simulation.

7.1 No growth $n_{t-1} = n_t = n$

We first analyze the case when the number of purchase decisions is identical in both periods, $n_{t-1} = n_t = n$. In this case the particular realizations of z in each period are identically and independently distributed events. This allows us to characterize the bias even as the import price index π is a non-linear function of λ_t and λ_{t-1} , which are not independent.

Let z and z' be two realizations from the multinomial distribution (8). Because the realizations in each period are independent, it is equally likely that $z_{t-1} = z$ and

$z_t = z'$ than $z_{t-1} = z'$ and $z_t = z$. In a slight abuse of notation let $\pi(z, z')$ and $\pi(z', z)$ be the import price index in each case. Clearly,

$$\pi(z, z') = \frac{1}{\pi(z', z)}.$$

We also have that for any $\pi > 0$,

$$\pi + \frac{1}{\pi} \geq 2$$

with strict equality only if $\pi = 1$. Now we collect all the possible realizations of z_{t-1} and z_t in unordered pairs. Since there exist realizations such that $\pi(z_{t-1}, z_t) \neq 1$, we conclude that import price index will overstate its true value

$$E\{\pi\} > 1$$

whenever $n_{t-1} = n_t = n < \infty$. As the import price index is biased upwards, welfare gains from new varieties will be biased downwards.

The upward bias, though, tends to be very small. Since z_t is i.i.d., variable λ has the same expectation in each period, $E\{\lambda_t\} = E\{\lambda_{t-1}\}$. For the convexity to produce a sizeable bias, we would need λ_t to be quite dispersed. Instead the variance in λ_t actually converges to zero quite fast with n . Given intersection set I , $n * \lambda_t$ is distributed according to a Binomial distribution with parameters s_I and n . The variance of λ_t is then

$$E_I \lambda_t^2 - (E_I \lambda_t)^2 = \frac{s_I(1 - s_I)}{n}.$$

Moreover, the probability distribution over intersection sets I rapidly becomes skewed towards large sets as n grows, and thus the variance falls faster than rate n . We will confirm these results in our simulations below.

7.2 Growth $n_{t-1} < n_t$

Next we consider the case where the number of purchase decisions grows over time, $n_{t-1} < n_t$, which constitutes our baseline calibration. An exact characterization of $E\{\pi\}$ proves to be particularly difficult in this case, since z_{t-1} and z_t are no longer identically distributed. Fortunately, we can make considerable progress by ignoring the non-linearity,

$$\hat{\pi} = \left(\frac{E \lambda_t}{E \lambda_{t-1}} \right)^{\frac{1}{\sigma-1}}.$$

Obviously we have that $\hat{\pi} \neq E\{\pi\}$. Yet $\hat{\pi}$ proves to be an excellent approximation.

We start by solving for $E\lambda_t$. We must keep in mind that I is itself a random variable, that depends on realizations in both periods. To make the argument clear, consider the indicator function $\chi_I(i)$, which takes value 1 if $i \in I$, 0 otherwise. We can then write

$$E \left\{ \sum_{j \in I} z_{jt} \right\} = E \left\{ \sum_j^k z_{jt} \chi_I(j) \right\} = \sum_j^k E \{ z_{jt} \chi_I(j) \}.$$

Since $z_{jt-1} > 0$ and $z_{jt} > 0$ are independent events, we have that

$$E \{ z_{jt} \chi_I(j) \} = \Pr(z_{jt} > 0) \Pr(z_{jt-1} > 0) E \{ z_{jt} | z_{jt} > 0 \}.$$

Noting that $E \{ z_{jt} \} = \Pr(z_{jt} > 0) E \{ z_{jt} | z_{jt} > 0 \}$ we have then that

$$E \{ z_{jt} \chi_I(j) \} = \Pr(z_{jt-1} > 0) E \{ z_{jt} \}.$$

We can easily solve for each of these terms:

$$\begin{aligned} \Pr(z_{jt-1} > 0) &= (1 - (1 - s_j)^{n_{t-1}}) \\ E \{ z_{jt} \} &= s_j n_t. \end{aligned}$$

Thus

$$E\lambda_t = \sum_j^k (1 - (1 - s_j)^{n_{t-1}}) s_j.$$

Note how the expectation $E\lambda_t$ increases with the number of purchase decisions for the *other* period, n_{t-1} . As n_{t-1} grows, the set I is likely to be larger and thus include more purchase decisions in the current period.

Combining $E\lambda_t$ and $E\lambda_{t-1}$ we obtain

$$\hat{\pi} = \left(\frac{\sum_j^k (1 - (1 - s_j)^{n_{t-1}}) s_j}{\sum_j^k (1 - (1 - s_j)^{n_t}) s_j} \right)^{\frac{1}{\sigma-1}}.$$

Whenever $n_{t-1} < n_t$ we have that the approximated import price index is biased downwards, $\hat{\pi} < 1$.

Not surprisingly, the size of the downward bias depends on how fast the number of purchase decisions grows. It is easy to see that $\hat{\pi}$ decreases with the difference $n_t - n_{t-1}$ for a fixed n_t . Moreover even a small increase in purchase decisions can result in a large downward bias if the number of purchase decisions in the base period n_{t-1} is small. This is relevant since most studies on variety gains compare years often a decade or more apart for finely disaggregated categories that can be expected to have little trade to start with.

7.3 A sample of import products

We take a look at the quantitative importance of the biases discussed above. We choose to focus on the 8 import goods studied in Feenstra (1994) as a basis for the simulations, now comparing years 1967 and 1987. We take the point estimate for the elasticity from Feenstra (1994) and use the same country fixed-effects computed in Section 3.⁴⁴ We pin down the number of purchase decisions in each period simply by matching the observed number of varieties traded. Note we do not impose a common growth rate or constant expenditure shares across goods. By letting the data speak, good by good, we get a chance to explore our results under different but informed parameters.

Table 4 includes a description of each import good, the number of suppliers in 1967 and 1987, and the estimated elasticity for each category. There is substantial variation on both counts. All but one good (portable typewriters) saw the number of varieties increase yet some good categories almost tripled the variety count, while others had very modest growth. The elasticity estimates range from a low of 3 to highs of 27 and 43 for gold and silver bullion, respectively. The next column reports the good-specific growth rate in purchase decisions, at an annual rate, needed to match the data. The growth rate are quite dispersed, as one would expect by matching point to point differences: recall the number of varieties in a given year is a random variable from the point of view of the model. Note it is not necessarily the case that goods that saw the largest increase in varieties have the highest growth rate in purchase decisions. Both the elasticity of substitution as well as the level of varieties impact the calibrated growth rate.

The last three columns in Table 4 finally reports our results regarding the import price index. Our “null hypothesis” is that all varieties were available at both dates—but not necessarily observed to have positive trade. For each product we report $\hat{\pi}$ and contrast this linear approximation with $E\{\pi\}$, computed by simulating the model. From the latter we also obtain the probability that the import price index is biased downwards, $\Pr(\pi \leq 1)$.

Let us start with the first four goods. They all display a substantial downward bias, ranging from close to 4 percent for men’s leather athletic shoes to 11 percent for carbon steel sheets. The linear approximation turns out to be pretty good. For three out of these four products we are almost certain that the point estimate will come below 1. For stainless steel bars the estimates for the import price index are very dispersed but there is more than a 75 percent probability of them being below

⁴⁴The country mix is somewhat different, but we do not re-calibrate the model.

1. Note the bias is larger for goods with lower elasticities and higher growth: the good with the largest bias, carbon steel sheets, features one of the lowest elasticities in our sample as well as above average growth. We have already discussed how the bias depends on the growth rate of purchase decisions. The relationship between the bias and the elasticity is also quite straightforward, as lower elasticities makes the price index simply more sensitive to variation in the ratio of expenditures λ_t/λ_{t-1} . If varieties are close to perfect substitutes, then the price index naturally gives little importance to new goods. This is perfectly illustrated by our results regarding gold and silver bullion. For these two goods, the bias is essentially negligible.

Let us turn our discussion to color television receivers and portable typewriters. These two goods experienced a substantial transformation over the period, albeit with very different fates. Portable typewriters were progressively outdated by personal computers, while color televisions started as a luxurious novelty and ended up in every American household. Their results stand out in Table 4 for their own reasons.

First, portable typewriters saw negative growth in purchase decisions and, as a result, the import price index is biased *upwards*. Yet in close to one third of our simulations the import price index came below 1, as the estimate is very noisy. Second, color television receivers show a gap between the linear approximation and the simulated expectation. This reflects that more often than not there was simply no overlap between varieties between 1967 and 1987.⁴⁵ This points to a potential problem with goods that have experienced a quick product cycle driven by technological transformations: as the common base of varieties supplied at both ends of the time sample becomes narrow, or non-existent, the resulting import price indexes are less precisely estimated.

8 Conclusions

Conspicuously absent from our theory of the extensive margin are economies of scale—the most common approach in the literature to capture missing trade links has been to introduce fixed costs on the supply side. We find these trade models most appealing for the export decision margin at the firm level. Our results indicate, though, that one does not need economies of scale to explain the extensive margin at a more aggregated level. That said, we should note that economies of scale may be important for wholesale retailers which, we suspect, would be a key determinant of the number of independent purchase decisions we need to match the data.

⁴⁵Recall that whenever this happens we imputed a value of $\pi = 1$.

Our model also has implications regarding at what frequency we should approach the data. For example, given a ratio of purchase decisions per unit of time, we can derive the length of time we would need to have a certain level of confidence that a given commodity would have been observed if available. Unfortunately, at high levels of detail, our model indicates these sampling periods will far exceed anything practical. It is thus perhaps the right decision to work with likelihoods at the usual frequencies, annual or bi-annual. Ideally, a complete model would allow some probability that a given commodity is indeed not available. It is an open question what additional data can be used to inform this margin given that our model matches the data very well while assuming all varieties are available.

References

- Alessandria, George and Horag Choi**, “Do Sunk Costs of Exporting Matter for Net Export Dynamics?,” *Quarterly Journal of Economics*, 2007, *122* (1), 289–336.
- Anderson, Simon P., Andre De Palma, and Jacques-Francois Thisse**, “The CES is a discrete choice model?,” *Economic Letters*, 1987, *24*, 139–140.
- , —, and —, *Discrete Choice Theory of Product Differentiation*, The MIT Press, 1992.
- Armenter, Roc and Miklós Koren**, “A Balls-and-Bins Model of Trade,” *American Economic Review*, July 2014, *104* (7), 2127–51.
- Bekes, Gabor and Balasz Murakozy**, “Temporary Trade and heterogeneous firms,” *Journal of International Economics*, 2012, *87*, 232–246.
- Bernard, A.B., J. Eaton, J.B. Jensen, and S. Kortum**, “Plants and Productivity in International Trade,” *American Economic Review*, 2003, *93* (4), 1268–90.
- Bils, Mark and Peter Klenow**, “The Acceleration of Variety Growth,” *American Economic Review*, 2001, *91*, 274–280.
- Broda, Christian and David E. Weinstein**, “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 2006, *121* (2), 541–585.
- CBO**, “The Effects of NAFTA on U.S. Mexican trade and GDP,” Technical Report, Congressional Budget Office 2003.
- Chaney, Thomas**, “Distorted Gravity: Heterogeneous Firms, Market Structure, and the Geography of International Trade,” *American Economic Review*, 2008, *98*, 1707–1721.
- David, H.A. and H.N. Nagaraja**, *Order Statistics*, John Wiley and Sons, 2003.
- Eaton, J. and S. Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, *70* (5), 1741–1779.
- Eaton, Jonathan, Samuel S. Kortum, and Sebastian Sotelo**, “International Trade: Linking Micro and Macro,” in Daron Acemoglu, Manel Arellano, and Eddie Dekel, eds., *Advances in Economics and Econometrics*, Vol. 2 Cambridge University Press May 2013, pp. 329–372.

- Feenstra, R. C. and D. E. Weinstein**, “Globalization, Markups and U.S. Welfare,” *Journal of Political Economy*, forthcoming.
- Feenstra, Robert C.**, “New Product Varieties and the Measurement of International Prices,” *The American Economic Review*, 1994, *84* (1), 157–177.
- Hilberry, R. H. and C.A. McDaniel**, “A Decomposition of North American Trade Growth since NAFTA,” July 2002. U.S. International Trade Commission, Working paper 2002-12.
- Hummels, David and P. J. Klenow**, “The Variety and Quality of a Nation’s Exports,” *American Economic Review*, 2005, *95* (3), 704–23.
- Kehoe, Timothy J. and Kim J. Ruhl**, “How Important Is the New Goods Margin in International Trade?,” *Journal of Political Economy*, 2013, *121* (2), 358–392.
- Kropf, Andreas and Philip Sauré**, “Fixed costs per shipment,” *Journal of International Economics*, 2014, *92* (1), 166–184.
- Melitz, Marc J.**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, *71* (6), 1695–1725.
- Pierce, Justin R. and Peter K. Schott**, “Concording U.S. Harmonized System Codes Over Time,” *Journal of Official Statistics*, 2012, *28* (1), 53–68.
- Romalis, J.**, “NAFTA’s and CUSFTA’s Impact on International Trade,” *Review of Economics and Statistics*, 2007, *89* (3), 416–35.
- Ruhl, K. J.**, “The International Elasticity Puzzle,” March 2008. Working paper, NYU.
- Trefler, Daniel**, “The Long and Short of the Canada-U.S. Free Trade Agreement,” *American Economic Review*, 2004, *94* (4), 870–895.

Import	Varieties ^a		Elasticity ^a	Growth purchases		Import Price Index	
	1967	1987		σ	γ_n^g	$\hat{\pi}$	$E\{\pi\}$ $\Pr(\pi \leq 1)$
Men's Leather Athletic Shoes	16	27	6.2	6.2	4.4 %	.9620	.9623 .9103
Men's and boys' cotton knit shirts	24	51	5.8	5.8	8.0%	.9592	.9590 .9974
Stainless steel bars	10	15	3.6	3.6	3.1 %	.9241	.9293 .7601
Carbon steel sheets	11	26	4.2	4.2	7.2 %	.89	.8896 .9678
Color television receivers	6	15	8.4	8.4	6.8 %	.9317	.982 .9231
Portable typewriters	17	14	3	3	-1.7 %	1.0484	1.0549 .3277
Gold bullion ^b	19	32	27.2	27.2	9.5 %	.9931	.9931 .9389
Silver bullion	11	15	42.9	42.9	2.5 %	.9963	.9962 .7131

^aAs reported in Feenstra (1994).

^bFor gold bullion we compare 1977 and 1987.

Table 4: Bias in import price indexes