Accounting for Trade Patterns*

Stephen J. Redding
Princeton University, NBER and CEPR †

David E. Weinstein
Columbia University and NBER‡

March 7, 2018

Abstract

We develop a quantitative framework for exactly decomposing trade patterns into economically meaningful components. We derive price indexes that determine comparative advantage across countries and sectors and the aggregate cost of living. If firms and products are imperfect substitutes, we show that these price indexes depend on variety, average demand or quality, and the dispersion of demand or quality-adjusted prices, and are only weakly related to standard empirical measures of average prices, thereby providing insight for elasticity puzzles. Of the cross-section (time-series) variation in comparative advantage, 50 (90) percent is accounted for by variety and average demand or quality, with average prices contributing less than 10 percent.

JEL CLASSIFICATION: F11, F12, F14

KEYWORDS: comparative advantage, trade, prices, quality, variety

*We are especially grateful to Rob Feenstra, Keith Head, Pete Klenow, Kalina Manova, Thierry Mayer, Marc Melitz, Gianmarco Ottaviano and Daniel Xu for helpful comments. We would also like to thank colleagues and seminar participants at Carnegie-Mellon, Chicago, Columbia, Duke, NBER, Princeton, Stanford and UC Berkeley for helpful comments. Thanks to Mark Greenan, Ildiko Magyari, Charly Porcher and Dyanne Vaught for outstanding research assistance. Redding and Weinstein thank Princeton and Columbia respectively for research support. Weinstein would also like to thank the NSF (Award 1127493) for generous financial support. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the U.S. Census Bureau or any organization to which the authors are affiliated. Results have been screened to insure that no confidential data are revealed.

†Fisher Hall, Princeton, NJ 08544. Email: reddings@princeton.edu.

‡420 W. 118th Street, MC 3308, New York, NY 10027. Email: dew35@columbia.edu.

1 Introduction

International economists face a large number of modeling choices when constructing models aimed at matching international trade flows and welfare gains. In particular, it is not always obvious how to model the impact of international cost, quality, entry, and demand differences on trade volumes. Moreover, even if one did know the correct supply-side model, official price indexes are often constructed using very different formulas than those implied by the constant elasticity of substitution (CES) demand systems that dominate trade and macro. Thus, models may perform poorly because objects like import price indexes are defined by statistical agencies in very different ways than the theoretical objects in economic models.

This paper makes three contributions to our understanding of trade flows and aggregate welfare gains. First, we show how to use commonly available trade-transactions data to rigorously measure aggregate and import price indexes even when data (such as the prices of domestically produced tradables and non-tradables) are incomplete or missing. Second, we develop a method for exactly decomposing aggregate trade flows and revealed comparative advantage (RCA) into four factors that vary both in the cross section and the time series: average prices, the entry and exit of products and firms, demand or quality shifts, and the dispersion or heterogeneity of firm and product sales. Each of these factors has been emphasized in some form in many empirical implementations of trade models (c.f., (Crozet, Head, and Mayer 2012; Hallak and Schott (2011); Kehoe and Ruhl 2013; Manova and Zhang (2012); Mayer, Melitz, and Ottaviano 2014), but a key feature of our setup is that we can allow all of them to matter without needing to make any assumptions about the supply-side of the model. Finally, we develop novel moments for price and sales distributions that are necessary for any model to match trade patterns and import price changes in a CES demand system.

One can understand the challenge of linking import prices and import volumes by considering one of the earliest stylized facts in international economics: aggregate import prices and import volumes have correlations that are often the wrong sign or close to zero (*c.f.*, Adler 1945).[1] Figure 1 shows the relationship for the U.S. between the change in log of merchandise imports divided by GDP and the log in the change in the Import Price Index divided by the the GDP deflator over the last 30 years.[2] The slope of the line is 1.3, which, if we interpret through the lens of a CES demand system *in the absence demand shifts*, would imply that the import demand curve slopes upwards. Since Orcutt (1950), economists have long known the two main explanations for the relationship we see in Figure 1: "shifts of the demand schedule" and "faulty methods of [price] index construction." Demand shifts are typically ruled out *a priori* in almost all trade and macro models (which are based on time-invariant utility functions), and hence their ability to match actual flows requires us to believe in Orcutt's second explanation: that our price indexes are faulty and do not properly adjust for quality, variety, and other factors. However, this implies that even if we knew the relationship between imports and theoretically-correct price indexes, official price indexes do not capture important factors driving imports. Although Feenstra (1994) showed how to build price indexes that capture changes in variety, no import price

[1] For example, Adler writes, "While real income may account for most of the variation in imports, the remaining variation does not seem to be explained by changes in relative prices."

[2] We chose this particular specification because it can be motivated by a CES demand system. The lack of a clear negative relationship is not something that can be solved by simply changing the specification to long differences, levels, or just using real imports on the vertical axis.

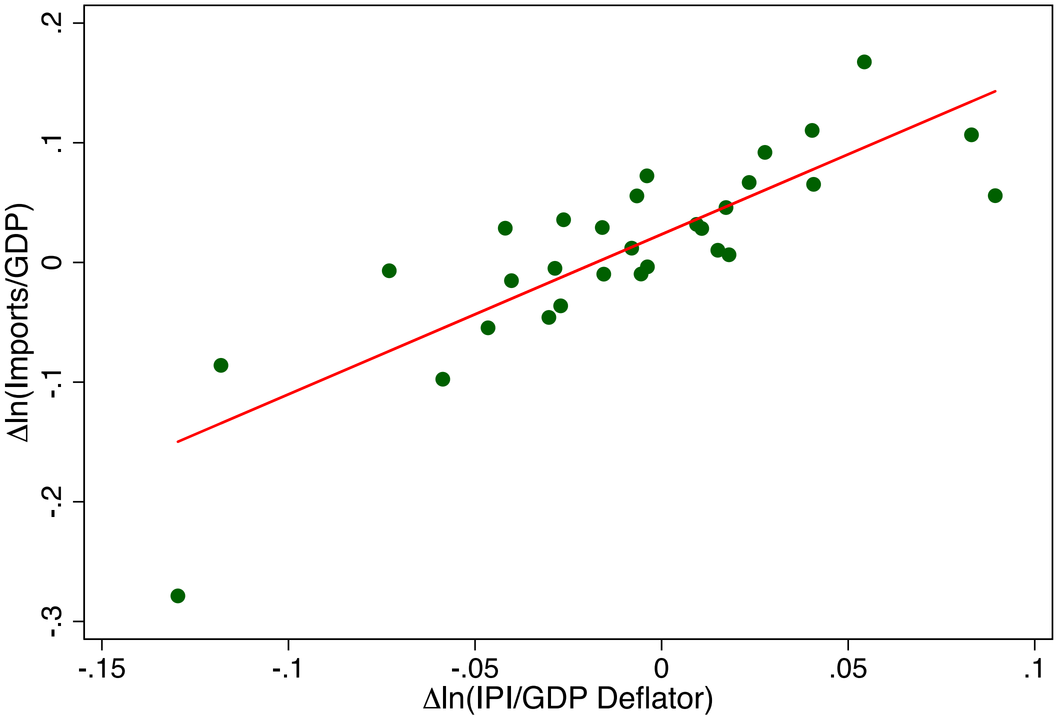

Note: Slope: 1.3367; SE: 0.2973; R-squared: 0.6817.

Figure 1: Changes in US merchandise imports as a share of GDP vs the BLS Import Price Index relative to GDP Deflator (1987-2017)

indexes rigorously allow for shifts in demand for or quality upgrading of existing products.[3]

One of our main contributions, therefore, is to demonstrate a method in which economists can account for these patterns in a way that allows for shifts in demand, prices, quality, variety, and firm heterogeneity or dispersion. We begin by noting that the demand system alone is sufficient to provide a framework for aggregating trade data from the micro to the macro level. The reason why the supply-side does not matter is that it only enters through the observed prices and expenditures and the estimated demand-side parameters. If the demand system is invertible, we can use these observables and the estimated parameters to solve for unique values for an unobserved demand or quality shifter for each good ("demand/quality"). Although this insight holds for any invertible demand system, we implement our analysis for CES preferences, as it is the most commonly-used specification in international trade. Our approach holds for *any* nesting structure within this demand system. We show that the value and price for each observed disaggregated trade transaction can be rationalized as the equilibrium of our model, without making any of the standard supply-side modeling assumptions. Furthermore, although the nested CES demand system is non-linear, we show that it admits an exact log-linear representation, which can be applied recursively across nests, and hence permits exact

[3]For example, Feenstra (1994) quantifies the entry and exit of goods over time, and Hummels and Klenow (2005) examine the contribution of variety and quality to a nation's exports, but neither paper examines all of these forces (and others) operating simultaneously. Hsieh et al. (2016) examine the contribution of this extensive margin to welfare using the Sato-Vartia price index and aggregate moments from U.S. and Canadian data. We show below that this Sato-Vartia price index cannot rationalize micro trade flows, because it assumes away idiosyncratic shifts in expenditure conditional on prices. Therefore, we use the unified price index (UPI) of Redding and Weinstein (2016), which enables us to both rationalize these micro trade flows and aggregate to the macro level.

additive decompositions of aggregate variables.

Our first main contribution is to use this trade nesting structure to derive exact *aggregate* price indexes at the national and exporter-sector level that are separable into average price, demand/quality, variety, and firm dispersion terms. The average price term has a functional form that corresponds closely to a conventional price index like the BLS import price index. The demand/quality term corrects the price index for shifts in an importer's demand for another country's exports, either due to a change of tastes for the exporter's output or quality upgrading of existing goods by exporters. The variety term captures how changes in the number of products in each firm and in the number of firms affect exports, and the dispersion term captures how changes in the size distribution of firms (e.g. as determined by demand and productivity differences across firms) affect exports when goods are imperfect substitutes. In addition, we also show how to aggregate to the national level even when detailed price and quantity data for certain sectors (e.g. non-traded sectors) is incomplete or missing.

Unlike prior import price indexes, our indexes allow not only for changes in prices and variety, but also demand and quality of existing goods. Thus, unlike frameworks built on time-invariant utility functions, it is possible in our setup that patterns like the one in Figure 1 can be driven by demand shifts. Importantly, we provide a quantification of each of the factors that are omitted when only prices are used to construct measures of import prices. Thus, while we show that the average price term has a similar functional form and turns out to be highly correlated with the BLS import price index ($\rho = 0.72$), the other terms in our price index have very different functional forms and turn out to be are negatively correlated with the BLS import price index. The importance of these terms helps explain the difficulty of trying to understand trade patterns using only conventional measures of import and domestic prices.

Our next theoretical contribution is to develop a rigorous measure of revealed comparative advantage (RCA) that is valid for the CES demand system and depends on the relative values of our price index across countries and sectors. This RCA measure is also additively separable into price, demand/quality, entry/exit, and dispersion terms in both the cross-section and time series. This property enables us to seamlessly move between the contribution of each factor to price indexes and how they affect import volumes. Thus, for example, when we observe that Chinese exports to the US rose dramatically between 1998 and 2011 in spite of a relative increase in Chinese export prices, we can do more than just note the factors that are left out of conventional indexes—our method enables us to exactly decompose this rise in Chinese imports into the amount attributable to each factor.

Our last theoretical contribution is to derive a general expression for understanding when distributional assumptions will and will not matter for understanding movements in import prices and trade patterns. There have been a number of proposed distributions for modeling the underlying firm productivity distributions that give rise to observed trade patterns: Pareto, Fréchet, and log normal being the most prominent ones. These distributional assumptions are made for reasons of tractability in theoretical models of international trade. However, it is less clear how much the assumed distributional assumptions matter for interpreting the results. Our paper makes both a destructive and a constructive contribution. On the negative side, we show that detailed trade-transactions data formally reject all three of these distributions, which means that although models based on these specifications may fit the aggregate data, they do not fully capture what is happening at the micro level. However, we also demonstrate an irrelevance result. In a CES setup, it does not matter what the underlying productivity distribution is for understanding trade patterns and import prices as long as the distribution is correctly centered such that it matches the geometric average price, geometric average quality, and geometric average sales of each firm. Thus, calibrated models will correctly capture trade flows and import price changes if their choice of distributional parameters enables them match these geometric averages in the data. For example, we show that parameters satisfying these criteria can easily be recovered from the QQ-estimation framework introduced into the trade literature by Head, Mayer, and Thoenig (2016).

We implement our approach using both U.S. data from 1997-2011 (reported in the main paper) and also Chilean data from 2007-14 (reported in the web appendix). Our decomposition reveals that firm entry/exit and average demand/quality each account for around 45 percent of the time-series variation in imports, with the dispersion of demand-adjusted prices making up most of the rest. We demonstrate that this pattern is robust across a range of alternative values for the elasticities of substitution. Indeed, for parameter values for which goods are imperfect substitutes, we show that the contributions from firm entry/exit and the dispersion of demand-adjusted prices to patterns of trade are invariant to these assumed elasticities. The fact that the prices of continuing goods are not negatively associated with the time-series variation in imports (as we saw in Figure 1) does not mean that prices do not matter for import volumes. Indeed, we estimate steeply negative demand curves. However, the data suggest that equilibrium price increases are small and associated with firm entry and demand/quality shifts that often more than offset whatever impact the price movement has. Similarly, in the cross-section, prices matter little in equilibrium with firm variety and dispersion (heterogeneity) accounting for close to 70 percent of the variation in trade patterns. In sum, these results suggest that models that emphasize small equilibrium movements in average prices and larger movements in other factors are more likely to capture the underlying forces at work.

Although there is an exact mapping between our price index and observed trade flows, the same need not be true for other approaches that impose stronger assumptions, such as the Feenstra (1994) price index, which corresponds to a special case in which there are no demand shifts for surviving goods. Thus, the difference between observed trade patterns and those predicted using alternative price indexes provides a metric for how successful models based on these assumptions are. By comparing actual RCA with counterfactual values of RCA based on different assumptions—e.g., with or without demand-shifts or variety corrections—we can directly assess the implications of these simplifying assumptions for understanding trade patterns. In particular, we find that models that assume no demand shifts and no changes in variety perform poorly on trade data. Models that incorporate variety changes while maintaining the assumption of no demand shifts do better, but still can only account for about ten percent of the changes in comparative advantage over time. These findings highlight the importance of changes in demand/quality within surviving varieties in understanding the changes in comparative advantage documented in Freund and Pierola (2015) and Hanson, Lind and Muendler (2015). They also point to the relevance of dynamic trade theories, in which comparative advantage evolves endogenously with process and product innovation, as in Grossman and Helpman (1991).

Our paper is related to several strands of existing research. First, we build on a long tradition in international trade that examines how to develop measures of prices in which quality and/or variety are changing (Feenstra 1994, Hallak and Schott 2011, and Feenstra and Romalis 2014). The papers developed important insights into how to adjust prices for variety and quality. However, none of them provide a method for exactly decomposing trade flows into the various factors explained by these forces. Our decomposition also builds on Hottman, Redding, and Weinstein (2016), which provides a methodology for doing a similar decomposition within consumer sectors using bar-code data. This paper extends that method to show how to aggregate across sectors even when some of the sectoral price data is missing. Thus, while Hottman *et al* found that within-firm demand/quality and product scope accounted for the vast majority of variation in sales across firms within sectors, they do not discuss variation across sectors or even aggregate variation in sales. Indeed, the two key factors that account for most of the variation in trade patterns—firm variety and firm dispersion—do not even appear in their earlier work. We also make use of the price index developed Redding and Weinstein (2016) but modify it by adding a nesting structure that not allows for firms, multiple sectors, and missing price data. These additions change both how the price indexes are constructed and how the parameters are estimated. Moreover, Redding and Weinstein (2016) do not consider how to use these indexes to construct aggregate real output measures.

Second, research has relaxed the constant-elasticity assumptions in neoclassical trade models by providing conditions under which they reduce to exchange models in which countries directly trade factor services (see Adao, Costinot and Donaldson 2017). In contrast, we assume a constant elasticity of demand, but relax the assumption of a constant elasticity of supply. By using additional structure on the demand-side, we are able to rationalize both micro and macro trade data, enabling us to aggregate up from the micro level and quantify the importance of different micro mechanisms for macro variables. As a result of imposing less structure on the supply-side, we can encompass non-neoclassical models with imperfect competition and increasing returns to scale (including Krugman 1980, Melitz 2003, and Atkeson and Burstein 2008).

Finally, our paper is related to the literature estimating elasticities of substitution between varieties and quantifying the contribution of new goods to welfare. As shown in Feenstra (1994), the contribution of entry and exit to the change in the CES price index can be captured using the expenditure share on common products (supplied in both periods) and the elasticity of substitution. Building on this approach, Broda and Weinstein (2006) quantify the contribution of international trade to welfare through an expansion on the number of varieties, and Broda and Weinstein (2010) examine product creation and destruction over the business cycle. Other related research using scanner data to quantify the effects of globalization includes Handbury (2013), Atkin and Donaldson (2015), and Atkin, Faber, and Gonzalez-Navarro (2015), and Fally and Faber (2016). Whereas this existing research assumes that demand/quality is constant for each surviving variety, we show that allowing for time-varying demand/quality is central to both rationalizing both aggregate and disaggregate patterns of trade.

The remainder of the paper is structured as follows. Section 2 introduces our theoretical framework. Section 3 outlines our structural estimation approach. Section 4 discusses our data. Section 5 reports our empirical results. Section 6 concludes. A web appendix contains technical derivations, additional empirical results for the U.S., and a replication of our U.S. results using Chilean data.

2 Theoretical Framework

We begin by showing that our framework exactly rationalizes observed micro trade data and permits exact aggregation, so that it can be used to quantify the importance of different micro mechanisms for macro variables. Although our approach could be implemented for any nested demand system that is invertible, we focus on CES preferences as the leading demand system in international trade, with a nesting structure guided by existing theories of trade, which distinguish countries, sectors, firms and products.

Throughout the paper, we index importing countries ("importers") by j and exporting countries ("exporters") by i (where each country can buy its own output). Each exporter can supply goods to each importer in a number of sectors that we index by g (a mnemonic for "group"). We denote the set of sectors by Ω^G and we indicate the number of elements in this set by N^G. We denote the set of countries from which importer j sources goods in sector g at time t by Ω^I_{jgt} and we indicate the number of elements in this set by N^I_{jgt}. Each sector (g) in each exporter (i) is comprised of firms, indexed by f (a mnemonic for "firm"). We denote the set of firms in sector g that export from country i to country j at time t by Ω^F_{jigt}; and we indicate the number of elements in this set by N^F_{jigt}. Each active firm can supply one or more products that we index by u (a mnemonic for "unit," as our most disaggregated unit of analysis); we denote the set of products supplied by firm f at time t by Ω^U_{ft}; and we indicate the number of elements in this set by N^U_{ft}.[4]

2.1 Demand

The aggregate unit expenditure function for importer j at time t (P_{jt}) is defined over the sectoral price index (P^G_{jgt}) and "demand"[5] parameter (φ^G_{jgt}) for each sector $g \in \Omega^G$:

$$P_{jt} = \left[\sum_{g \in \Omega^G} \left(P^G_{jgt} / \varphi^G_{jgt} \right)^{1-\sigma^G} \right]^{\frac{1}{1-\sigma^G}}, \qquad \sigma^G > 1, \varphi^G_{jgt} > 0, \tag{1}$$

where σ^G is the elasticity of substitution across sectors and φ^G_{jgt} captures the relative demand for each sector. The unit expenditure function for each sector g depends on the price index (P^F_{ft}) and demand parameter (φ^F_{ft}) for each firm $f \in \Omega^F_{jigt}$ from each exporter $i \in \Omega^I_{jgt}$ within that sector:

$$P^G_{jgt} = \left[\sum_{i \in \Omega^I_{jgt}} \sum_{f \in \Omega^F_{jigt}} \left(P^F_{ft} / \varphi^F_{ft} \right)^{1-\sigma^F_g} \right]^{\frac{1}{1-\sigma^F_g}}, \qquad \sigma^F_g > 1, \varphi^F_{ft} > 0, \tag{2}$$

[4] We use the superscript G to denote a sector-level variable, the superscript F to represent a firm-level variable, and the superscript U to indicate a product-level variable. We use subscripts j and i to index individual countries, the subscript g to reference individual sectors, the subscript f to refer to individual firms, the subscript u to label individual products, and the subscript t to indicate time.

[5] There is no standard way to refer to this parameter. It is "taste" in Feenstra (1994), "quality" in Broda and Weinstein (2010), and "appeal" in Hottman, Redding and Weinstein (2016). We will refer to this parameter as either "demand" or "demand/quality" because it is isomorphic in the CES setup to have consumers demand more of a good conditional on price because it is higher quality in some objective sense or because they just like it more. We therefore refer to the parameter as "demand" or "demand/quality" because it corresponds to anything that shifts consumer demand conditional on price.

where σ_g^F is the elasticity of substitution across firms f for sector g and φ_{ft}^F controls the relative demand for each firm within that sector. We assume that the unit expenditure function within each sector takes the same form for both final consumption and intermediate use, so that we can aggregate both these sources of expenditure, as in Eaton and Kortum (2002) and Caliendo and Parro (2015).

We allow firm varieties to be horizontally differentiated and assume the same elasticity of substitution for domestic and foreign firms within sectors (σ_g^F).[6] The unit expenditure function for each firm f depends on the price (P_{ut}^U) and demand parameter (φ_{ut}^U) for each product $u \in \Omega_{ft}^U$ supplied by that firm:

$$P_{ft}^F = \left[\sum_{u \in \Omega_{ft}^U} \left(P_{ut}^U / \varphi_{ut}^U \right)^{1-\sigma_g^U} \right]^{\frac{1}{1-\sigma_g^U}}, \qquad \sigma_g^U > 1, \varphi_{ut}^U > 0, \tag{3}$$

where σ_g^U is the elasticity of substitution across products within firms for sector g and φ_{ut}^U captures the relative demand for each product within a given firm.

A few remarks about this specification are useful. First, we allow prices to vary across products, firms, sectors and countries, which implies that our setup nests models in which relative and absolute production costs differ within and across countries. Second, for notational convenience, we define the firm index $f \in \Omega_{jigt}^F$ by sector g, destination country j and source country i. Therefore, if a firm has operations in multiple sectors and/or exporting countries, we label these different divisions separately. As we observe the prices of the products for each firm, sector and exporting country in the data, we do not need to take a stand on market structure or the level at which product introduction and pricing decisions are made within the firm. Third, the fact that the elasticities of substitution across products within firms (σ_g^U), across firms within sectors (σ_g^F), and across sectors within countries (σ^G) need not be infinite implies that our framework nests models in which products are differentiated within firms, across firms within sectors, and across sectors. Moreover, our work is robust to collapsing one or more of these nests. For example, if all three elasticities are equal ($\sigma_g^U = \sigma_g^F = \sigma^G$), all three nests collapse, and the model becomes equivalent to one in which consumers only care about firm varieties. Alternatively, if $\sigma_g^U = \sigma_g^F = \infty$ and $\sigma^G < \infty$, only sectors are differentiated, and varieties are perfectly substitutable within sectors. Finally, if $\sigma_g^U = \sigma_g^F > \sigma^G$, firm brands are irrelevant, so that products are equally differentiated within and across firms for a given sector.

Fourth, the demand shifters (φ_{jgt}^G, φ_{ft}^F, φ_{ut}^U) capture anything that shifts the demand for sectors, firms and products *conditional* on price. Therefore, they incorporate both quality (vertical differences across varieties) and consumer tastes. We refer to these demand shifters as "demand/quality" to make clear that they can be interpreted either as shifts in consumer demand or product quality.[7] Finally, in order to simplify notation, we suppress the subscript for importer j, exporter i, and sector g for the firm and product demand shifters

[6]Therefore, we associate horizontal differentiation within sectors with firm brands, which implies that differentiation across countries emerges solely because there are different firms in different countries, as in Krugman (1980) and Melitz (2003). It is straightforward to also allow the elasticity of substitution to differ between home and foreign firms, which introduces separate differentiation by country, as in Armington (1969). Feenstra, Luck, Obstfeld and Russ (2014) find that they often cannot reject the same elasticity between home and foreign varieties as between foreign varieties.

[7]See, for example, the discussion in Di Comite, Thisse and Vandenbussche (2014). A large literature in international trade has interpreted these demand shifters as capturing product quality, including Schott (2004), Khandelwal (2010), Hallak and Schott (2011), Feenstra and Romalis (2008), and Sutton and Trefler (2016).

$(\varphi^F_{ft}, \varphi^U_{ut})$. However, we take it as understood that we allow these demand shifters for a given firm f and product u to vary across importers j, exporters i and sectors g, which captures the idea that a firm's varieties can be more appealing in some markets than others. For example, Sony products may be more appealing to Americans than Chileans, or may have more consumer appeal in the television sector than the camera sector, or even may be perceived to have higher quality if they are supplied from Japan rather than from another location.

2.2 Non-traded Sectors

We allow some sectors to be non-traded, in which case we do not observe products within these sectors in our disaggregated import transactions data, but we can measure total expenditure on these non-traded sectors using domestic expenditure data. We incorporate these non-traded sectors by re-writing the overall unit expenditure function in equation (1) in terms of the share of expenditure on tradable sectors (μ^T_{jt}) and a unit expenditure function for these tradable sectors ($\mathbb{P}^T_{jt}$):

$$P_{jt} = \left(\mu^T_{jt}\right)^{\frac{1}{\sigma^G - 1}} \mathbb{P}^T_{jt}. \tag{4}$$

The share of expenditure on the set of tradable sectors $\Omega^T \subseteq \Omega^G$ (μ^T_{jt}) can be measured using aggregate data on expenditure in each sector:

$$\mu^T_{jt} \equiv \frac{\sum_{g \in \Omega^T} X^G_{jgt}}{\sum_{g \in \Omega^G} X^G_{jgt}} = \frac{\sum_{g \in \Omega^T} \left(P^G_{jgt} / \varphi^G_{jgt}\right)^{1-\sigma^G}}{\sum_{g \in \Omega^G} \left(P^G_{jgt} / \varphi^G_{jgt}\right)^{1-\sigma^G}}, \tag{5}$$

where X^G_{jgt} is total expenditure by importer j on sector g at time t. The unit expenditure function for *tradable* sectors ($\mathbb{P}^T_{jt}$) depends on the price index for each tradable sector (P^G_{jgt}):

$$\mathbb{P}^T_{jt} \equiv \left[\sum_{g \in \Omega^T} \left(P^G_{jgt} / \varphi^G_{jgt}\right)^{1-\sigma^G}\right]^{\frac{1}{1-\sigma^G}}, \tag{6}$$

where we use the "blackboard" font $\mathbb{P}$ to denote price indexes that are defined over tradable goods.

Therefore, our assumption on demand allows us to construct an overall price index without observing entry, exit, sales, prices or quantities of individual products in non-tradable sectors. From equation (5), there is always a one-to-one mapping between the market share of tradable sectors and the relative price indexes in the two sets of sectors. In particular, if the price of non-tradables relative to tradables rises, the *share* of tradables (μ^T_{jt}) also rises. In other words, the share of tradables is a sufficient statistic for understanding the relative prices of tradables and non-tradables. As one can see from equation (4), if we hold fixed the price of tradables ($\mathbb{P}^T_{jt}$), a rise in the share of tradables (μ^T_{jt}) can only occur if the price of non-tradables sectors also rises, which means that the aggregate price index index (P_{jt}) must also be increasing in the share of tradables.

2.3 Domestic Versus Foreign Varieties Within Tradable Sectors

We also allow for domestic varieties within tradable sectors, in which case we again do not observe them in our import transactions data, but we can back out the implied expenditure on these domestic varieties using data on domestic shipments, exports and imports for each tradable sector. We incorporate domestic varieties within tradable sectors by re-writing the sectoral price index in equation (2) in terms of the share of expenditure on foreign varieties within each sector (the sectoral import share μ^G_{jgt}) and a unit expenditure function for these foreign varieties (a sectoral import price index $\mathbb{P}^G_{jgt}$):

$$P^G_{jgt} = \left(\mu^G_{jgt}\right)^{\frac{1}{\sigma^F_g - 1}} \mathbb{P}^G_{jgt}. \tag{7}$$

The sectoral import share (μ^G_{jgt}) equals total expenditure on imported varieties within a sector divided by total expenditure on that sector:

$$\mu^G_{jgt} \equiv \frac{\sum_{i\in\Omega^E_{jgt}} \sum_{f\in\Omega^F_{jigt}} X^F_{ft}}{X^G_{jgt}} = \frac{\sum_{i\in\Omega^E_{jgt}} \sum_{f\in\Omega^F_{jigt}} \left(P^F_{ft}/\varphi^F_{ft}\right)^{1-\sigma^F_g}}{\sum_{i\in\Omega^I_{jgt}} \sum_{f\in\Omega^F_{jigt}} \left(P^F_{ft}/\varphi^F_{ft}\right)^{1-\sigma^F_g}}, \tag{8}$$

where $\Omega^E_{jgt} \equiv \left\{\Omega^I_{jgt} : i \neq j\right\}$ is the subset of *foreign* countries $i \neq j$ that supply importer j within sector g at time t; X^F_{ft} is expenditure on firm f; and X^G_{jgt} is country j's total expenditure on all firms in sector g at time t. The sectoral import price index ($\mathbb{P}^G_{jgt}$) is defined over the foreign goods observed in our disaggregated import transactions data as:

$$\mathbb{P}^G_{jgt} \equiv \left[\sum_{i\in\Omega^E_{jgt}} \sum_{f\in\Omega^F_{jigt}} \left(P^F_{ft}/\varphi^F_{ft}\right)^{1-\sigma^F_g}\right]^{\frac{1}{1-\sigma^F_g}}. \tag{9}$$

In this case, the import share within each sector is the appropriate summary statistic for understanding the relative prices of home and foreign varieties within that sector. From equation (7), the sectoral price index (P^G_{jgt}) is *increasing* in the sectoral *foreign* expenditure share (μ^G_{jgt}). The reason is that our expression for the sectoral price index (P^G_{jgt}) conditions on the price of foreign varieties, as is captured by the import price index ($\mathbb{P}^G_{jgt}$). For a given value of this import price index, a higher foreign expenditure share (μ^G_{jgt}) implies that domestic varieties are less attractive, which implies a higher sectoral price index.[8]

2.4 Exporter Price Indexes

To examine the contribution of individual countries to trade patterns and aggregate prices, it proves convenient to rewrite the sectoral import price index ($\mathbb{P}^G_{jgt}$) in equation (9) in terms of price indexes for each foreign exporting country within that sector ($\mathbb{P}^E_{jigt}$):

$$\mathbb{P}^G_{jgt} = \left[\sum_{i\in\Omega^E_{jgt}} \left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}\right]^{\frac{1}{1-\sigma^F_g}}, \tag{10}$$

[8]In contrast, the expression for the price index in Arkolakis, Costinot and Rodriguez-Clare (2012) conditions on the price of *domestically-produced* varieties, and is *increasing* in the *domestic* expenditure share. The intuition is analogous. For a given price of domestically-produced varieties, a higher domestic trade share implies that foreign varieties are less attractive, which implies a higher price index.

where importer j's price index for exporter i in sector g at time t ($\mathbb{P}_{jigt}^{E}$) is defined over the firm price indexes (P_{ft}^{F}) and demand/qualities (φ_{ft}^{F}) for each of the firms f from that foreign exporter and sector:

$$\mathbb{P}_{jigt}^{E} \equiv \left[\sum_{f \in \Omega_{jigt}^{F}} \left(P_{ft}^{F} / \varphi_{ft}^{F} \right)^{1-\sigma_{g}^{F}} \right]^{\frac{1}{1-\sigma_{g}^{F}}}, \tag{11}$$

and we use the superscript E to denote a variable for a *foreign* exporting country.

This exporter price index (11) is a key object in our empirical analysis, because it summarizes importer j's cost of sourcing goods from exporter i within sector g at time t. We show below that the relative values of these exporter price indexes across countries and sectors determine comparative advantage. Note that substituting this definition of the exporter price index (11) into the sectoral import price index (10), we recover our earlier equivalent expression for the sectoral import price index in equation (9).

2.5 Expenditure Shares

Using the properties of CES demand, the share of each product in expenditure on each firm (S_{ut}^{U}) is given by:

$$S_{ut}^{U} = \frac{\left(P_{ut}^{U} / \varphi_{ut}^{U} \right)^{1-\sigma_{g}^{U}}}{\sum_{\ell \in \Omega_{ft}^{U}} \left(P_{\ell t}^{U} / \varphi_{\ell t}^{U} \right)^{1-\sigma_{g}^{U}}}, \tag{12}$$

where the firm and sector expenditure shares are defined analogously.

In the data, we observe product expenditures (X_{ut}^{U}) and quantities (Q_{ut}^{U}) for each product category. In our baseline specification in the paper, we assume that the level of disaggregation at which products are observed in the data corresponds to the level at which firms make product decisions. Therefore, we measure prices using unit values ($P_{ut}^{U} = X_{ut}^{U} / Q_{ut}^{U}$). From equation (12) above, demand-adjusted prices ($P_{ut}^{U} / \varphi_{ut}^{U}$) are uniquely determined by the expenditure shares (S_{ut}^{U}) and the elasticities (σ_{g}^{U}). Therefore, any multiplicative change in the units in which quantities (Q_{ut}^{U}) are measured, which affects prices ($P_{ut}^{U} = X_{ut}^{U} / Q_{ut}^{U}$), leads to an exactly proportionate change in demand/quality (φ_{ut}^{U}), in order to leave the demand-adjusted price unchanged ($P_{ut}^{U} / \varphi_{ut}^{U}$). It follows that the *relative* importance of prices and demand/quality in explaining expenditure share variation is unaffected by any multiplicative change to the units in which quantities are measured.

In Section A.7 of the web appendix, we show that our analysis generalizes to the case in which firms supply products at a more disaggregated level than the categories observed in the data. In this case, there can be unobserved differences in composition within observed product categories. However, we show that these unobserved compositional differences enter the model in exactly the same way as unobserved differences in demand/quality for each observed product category, and hence our analysis goes through in the sense that some of what we label product demand/quality may reflect compositional changes at a more disaggregate level than we can observe in the data.

2.6 Log-Linear CES Price Index

We now use the CES expenditure share to rewrite the CES price index in an exact log linear form that enables us to aggregate from micro to macro. We illustrate our approach using the product expenditure share within the firm tier of utility, but the analysis is analogous for each of the other tiers of utility. Rearranging the expenditure share of products within firms (12) using the firm price index (3), we obtain:

$$P_{ft}^{F} = \frac{P_{ut}^{U}}{\varphi_{ut}^{U}} \left(S_{ut}^{U} \right)^{\frac{1}{\sigma_g^U - 1}}, \tag{13}$$

which must hold for each product $u \in \Omega_{ft}^{U}$. Taking logarithms, averaging across products within firms, and adding and subtracting $\frac{1}{\sigma_g^U - 1} \ln N_{ft}^{U}$, we obtain the following exact log linear decomposition of the CES price index into four terms:

$$\ln P_{ft}^{F} = \underbrace{\mathbb{E}_{ft}^{U}\left[\ln P_{ut}^{U}\right]}_{\text{(i) Average log prices}} - \underbrace{\mathbb{E}_{ft}^{U}\left[\ln \varphi_{ut}^{U}\right]}_{\text{(ii) Average log demand}} + \underbrace{\frac{1}{\sigma_g^U - 1}\left(\mathbb{E}_{ft}^{U}\left[\ln S_{ut}^{U}\right] - \ln \frac{1}{N_{ft}^{U}}\right)}_{\text{(iii) Dispersion demand-adjusted prices}} - \underbrace{\frac{1}{\sigma_g^U - 1} \ln N_{ft}^{U}}_{\text{(iv) Variety}}, \tag{14}$$

where $\mathbb{E}\left[\cdot\right]$ denotes the mean operator such that $\mathbb{E}_{ft}^{U}\left[\ln P_{ut}^{U}\right] \equiv \frac{1}{N_{ft}^{U}} \sum_{u \in \Omega_{ft}^{U}} \ln P_{ut}^{U}$; the superscript U indicates that the mean is taken across products; and the subscripts f and t indicate that this mean varies across firms and over time.[9]

This expression for the firm price index in equation (14) has an intuitive interpretation. When products are perfect substitutes ($\sigma_g^U \to \infty$), the average of log demand-adjusted prices ($\mathbb{E}_{ft}^{U}\left[\ln\left(P_{ut}^{U}/\varphi_{ut}^{U}\right)\right]$) is a sufficient statistic for the log firm price index (as captured by terms (i) and (ii)). The reason is that perfect substitutability implies the equalization of demand-adjusted prices for all consumed varieties ($P_{ut}^{U}/\varphi_{ut}^{U} = P_{\ell t}^{U}/\varphi_{\ell t}^{U}$ for all $u, \ell \in \Omega_{ft}^{U}$ as $\sigma_g^U \to \infty$). Therefore, the mean of the log demand-adjusted prices is equal to the log demand-adjusted price for each product ($\mathbb{E}_{ft}^{U}\left[\ln\left(P_{ut}^{U}/\varphi_{ut}^{U}\right)\right] = \ln\left(P_{\ell t}^{U}/\varphi_{\ell t}^{U}\right)$ for all $u, \ell \in \Omega_{ft}^{U}$ as $\sigma_g^U \to \infty$).

In contrast, when products are imperfect substitutes ($1 < \sigma_g^U < \infty$), the firm price index also depends on both the number of varieties (term (iv)) and the dispersion of demand-adjusted prices across those varieties (term (iii)). The contribution from the number of varieties reflects consumer love of variety: if varieties are imperfect substitutes ($1 < \sigma_g^U < \infty$), an increase in the number of products sold by a firm (N_{ft}^{U}) reduces the firm price index. Keeping constant the price-to-quality ratio of each variety, consumers obtain more utility from firms that supply more varieties than others.

The contribution from the dispersion of demand-adjusted prices also reflects imperfect substitutability. If all varieties have the same demand-adjusted price, they all have the same expenditure share ($S_{ut}^{U} = 1/N_{ft}^{U}$). At this point, the mean of log-expenditure shares is maximized, and this third term is equal to zero. Moving away from this point and increasing the dispersion of demand-adjusted prices, by raising the demand-adjusted price for some varieties and reducing it for others, the dispersion of expenditure shares across varieties increases. As the log function is strictly concave, this increased dispersion of expenditure shares in turn implies a fall in the mean of log expenditure shares. Hence, this third term is negative when demand-adjusted prices differ across varieties ($\mathbb{E}_{ft}^{U}\left[\ln S_{ut}^{U}\right] < \ln\left(1/N_{ft}^{U}\right)$), which reduces the firm price index. Intuitively, holding constant average demand-adjusted prices, consumers prefer to source products from firms with more dis-

[9]This price index in equation (14) uses a different but equivalent expression for the CES price index from Hottman et al. (2016), in which the dispersion of sales across goods is captured using a different term from $\left(1/\left(\sigma_g^U - 1\right)\right)\mathbb{E}_{ft}^{U}\left[\ln S_{ut}^{U}\right]$.

persed demand-adjusted prices, because they can substitute away from products with high demand-adjusted prices and towards those with low demand-adjusted prices.

The decomposition in equation (14) can be undertaken in a sequence of steps. First, we can separate out the contribution of variety (term (iv)) and demand-adjusted prices (terms (i)-(iii)). Second, we can break down the demand-adjusted prices component (terms (i)-(iii)) into terms for average demand-adjusted prices (terms (i)-(ii)) and the dispersion of demand-adjusted prices (term (iii)). Third, we can disaggregate the demand-adjusted prices term into components for average prices (term (i)) and average demand (term (ii)). This sequential decomposition is useful, because it highlights the ways in which the model-based price indexes differ from standard empirical measures of average prices, since the change in average log prices (term (i) differenced) is the log of a conventional Jevons Price Index. Furthermore, as the decomposition in equation (14) is log additive, it provides the basis for exact log-linear decompositions of aggregate variables into different micro mechanisms in the model. Finally, the log-linear nature of this decomposition also implies that it is robust to measurement error in prices and/or expenditure shares that is mean zero in logs.

2.7 Entry, Exit and the Unified Price Index

One challenge in implementing this exact aggregation approach is the entry and exit of varieties over time in the micro data. To correctly take account of entry and exit between each pair of time periods, we follow Feenstra (1994) in using the share of expenditure on "common" varieties that are supplied in both of these time periods. In particular, we partition the set of firms from exporter i supplying importer j within sector g in periods $t-1$ and t (Ω^{F}_{jigt-1} and Ω^{F}_{jigt} respectively) into the subsets of "common firms" that continue to supply this market in both periods ($\Omega^{F}_{jigt,t-1}$), firms that enter in period t (I^{F+}_{jigt}) and firms that exit after period $t-1$ (I^{F-}_{jigt-1}). Similarly, we partition the set of products supplied by each of these firms in that sector into "common products" ($\Omega^{U}_{ft,t-1}$), entering products (I^{U+}_{ft}) and exiting products (I^{U-}_{ft-1}). A foreign exporting country enters an import market within a given sector when its first firm begins to supply that market and exits when its last firm ceases to supply that market. We can thus define analogous sets of foreign exporting countries $i \neq j$ for importer j and sector g: "common" ($\Omega^{E}_{jgt,t-1}$), entering (I^{E+}_{jgt}) and exiting (I^{E-}_{jgt-1}). We denote the number of elements in these common sets of firms, products and foreign exporters by $N^{F}_{jigt,t-1}$, $N^{U}_{ft,t-1}$ and $N^{E}_{jgt,t-1}$ respectively.

To incorporate entry and exit into the firm price index, we compute the shares of firm expenditure on common products in periods t and $t-1$ as follows:

$$\lambda^{U}_{ft} \equiv \frac{\sum_{u\in\Omega^{U}_{ft,t-1}} \left(P^{U}_{ut}/\varphi^{U}_{ut}\right)^{1-\sigma^{U}_{g}}}{\sum_{u\in\Omega^{U}_{ft}} \left(P^{U}_{ut}/\varphi^{U}_{ut}\right)^{1-\sigma^{U}_{g}}}, \qquad \lambda^{U}_{ft-1} \equiv \frac{\sum_{u\in\Omega^{U}_{ft,t-1}} \left(P^{U}_{ut-1}/\varphi^{U}_{ut-1}\right)^{1-\sigma^{U}_{g}}}{\sum_{u\in\Omega^{U}_{ft-1}} \left(P^{U}_{ut-1}/\varphi^{U}_{ut-1}\right)^{1-\sigma^{U}_{g}}}, \tag{15}$$

where recall that $\Omega^{U}_{ft,t-1}$ is the set of common products such that $\Omega^{U}_{ft,t-1} \subseteq \Omega^{U}_{ft}$ and $\Omega^{U}_{ft,t-1} \subseteq \Omega^{U}_{ft-1}$.

Using these common expenditure shares, the change in the log firm price index between periods $t-1$ and t ($\ln\left(P^{F}_{ft}/P^{F}_{ft-1}\right)$) can be exactly decomposed into four terms that are analogous to those for our levels decomposition in equation (14) above:

$$\ln\left(\frac{P_{ft}^{F}}{P_{ft-1}^{F}}\right)=\underbrace{\mathbb{E}_{ft}^{U*}\left[\ln\left(\frac{P_{ut}^{U}}{P_{ut-1}^{U}}\right)\right]}_{\text{(i) Average log prices}}-\underbrace{\mathbb{E}_{ft}^{U*}\left[\ln\left(\frac{\varphi_{ut}^{U}}{\varphi_{ut-1}^{U}}\right)\right]}_{\text{(ii) Average log demand}}+\underbrace{\frac{1}{\sigma_g^U-1}\mathbb{E}_{ft}^{U*}\left[\ln\left(\frac{S_{ut}^{U*}}{S_{ut-1}^{U*}}\right)\right]}_{\text{(iii) Dispersion demand-adjusted prices}}+\underbrace{\frac{1}{\sigma_g^U-1}\ln\left(\frac{\lambda_{ft}^{U}}{\lambda_{ft-1}^{U}}\right)}_{\text{(iv) Variety}},\tag{16}$$

as shown in Section A.2.7 of the web appendix; $\mathbb{E}_{ft}^{U*}\left[\ln\left(P_{ut}^{U}/P_{ut-1}^{U}\right)\right]\equiv\frac{1}{N_{ft,t-1}^{U}}\sum_{u\in\Omega_{ft,t-1}^{U}}\ln\left(P_{ut}^{U}/P_{ut-1}^{U}\right)$; the superscript U^* indicates that the mean is taken across common products; and the subscripts f and t indicate that this mean varies across firms and over time; S_{ut}^{U*} is the share of an individual common product in expenditure on all common products, which takes the same form as the expression in equation (12), except that the summation in the denominator is over the set of common products ($\Omega_{ft,t-1}^{U}$); if entering varieties are either more numerous or have lower demand-adjusted prices than exiting varieties, the common goods expenditure share at time t is lower than at time $t-1$, implying a fall in the price index ($\ln\left(\lambda_{ft}^{U}/\lambda_{ft-1}^{U}\right)<0$).

We refer to the exact CES price index in equation (16) as the "unified price index" (UPI), because the time-varying demand shifters for each product (φ_{ut}^{U}) ensure that it exactly rationalizes the micro data on prices and expenditure shares, while at the same time it permits exact aggregation to the macro level, thereby unifying micro and macro. This price index shares the same variety correction term $\left(\lambda_{ft}^{U}/\lambda_{ft-1}^{U}\right)^{1/\left(\sigma_g^U-1\right)}$ as Feenstra (1994). The key difference from Feenstra (1994) is the formulation of the price index for common goods, which we refer to as the "common-goods unified price index" (CG-UPI). Instead of using the Sato-Vartia price index for common goods, which assumes time-invariant demand/quality for each common good, we use the formulation of this price index for common goods from Redding and Weinstein (2016), which allows for changes in demand/quality for each common good over time.

2.8 Model Inversion

Given the observed data on prices and expenditures for each product $\{P_{ut}^{U}, X_{ut}^{U}\}$ and the substitution parameters $\{\sigma_g^U, \sigma_g^F, \sigma^G\}$, the model is invertible, such that unique values of demand/quality can be recovered from the observed data (up to a normalization or choice of units). We illustrate this inversion for the firm tier of utility, but the same approach holds for each of our tiers of utility. Dividing the share of a product in firm expenditure (12) by its geometric mean across common products within that firm, product demand can be expressed as the following function of data and parameters:

$$\frac{\varphi_{ut}^{U}}{\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^{U}\right]}=\frac{P_{ut}^{U}}{\mathbb{M}_{ft}^{U*}\left[P_{ut}^{U}\right]}\left(\frac{S_{ut}^{U}}{\mathbb{M}_{ft}^{U*}\left[S_{ut}^{U}\right]}\right)^{\frac{1}{\sigma_g^U-1}}.\tag{17}$$

where $\mathbb{M}\left[\cdot\right]$ is the geometric mean operator such that $\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^{U}\right]\equiv\left(\prod_{u\in\Omega_{ft,t-1}^{U}}\varphi_{ut}^{U}\right)^{1/N_{t,t-1}^{U}}$.[10]

As the CES expenditure shares are homogeneous of degree zero in the demand/quality parameters, we can only recover the product, firm and sector demand parameters $\{\varphi_{ut}^{U}, \varphi_{ft}^{F}, \varphi_{jgt}^{G}\}$ up to a choice of units

[10] Redding and Weinstein (2016) provide a micro-foundation for this normalization based on a theory of consumer behavior in which every product has a time-invariant, non-random component and a time-varying random component that varies by product and time. While one can imagine other normalizations, the fact that this one is micro-founded makes it particularly attractive.

in which to measure these parameters. We choose the convenient choice of units such that the geometric mean of product demand across common products within each foreign firm is equal to one ($\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^{U}\right]=1$ in equation (17)), the geometric mean of firm demand across common foreign firms within each sector is equal to one, and the geometric mean of sector demand across tradable sectors is equal to one. Under these normalizations, product demand (φ_{ut}^{U}) captures the relative demand/quality of products within foreign firms; firm demand (φ_{ft}^{F}) absorbs the relative demand/quality of foreign firms within sectors; and sector demand (φ_{jgt}^{G}) reflects the relative demand/quality of tradable sectors.[11]

Given this choice of units, we use the recursive structure of the model to solve for unique values of product, firm and sector demand $\{\varphi_{ut}^{U}, \varphi_{ft}^{F}, \varphi_{jgt}^{G}\}$, as shown in Section A.2.8 of the web appendix. First, we use the product expenditure share in equation (17) to solve for product demand/quality (φ_{ut}^{U}). Second, we use these solutions for product demand/quality to construct the firm price index for each foreign firm (P_{ft}^{F}). Third, we use the shares of individual foreign firms in expenditure on foreign imports within a sector to solve for demand/quality for each foreign firm (φ_{ft}^{F}). Fourth, we use these solutions for demand/quality for each foreign firm and the share of expenditure on foreign firms within the sector (μ_{jgt}^{G}) to compute the price index for each tradable sector (P_{jgt}^{G}). Fifth, we use the share of individual tradable sectors in expenditure on all tradable sectors (μ_{jt}^{T}) to solve for demand/quality for each tradable sector (φ_{gt}^{G}). Sixth, we use these solutions for sector demand/quality for each tradable sector and the share of aggregate expenditure on tradable sectors to compute the aggregate price index (P_{jt}).

Our decompositions of comparative advantage across countries and sectors are robust to alternative choices of units in which to measure product, firm and sector demand/quality. In particular, comparative advantage is based on relative comparisons across countries and sectors. Therefore, any common choice of units across firms within each sector differences out when we compare firms from different countries within that sector. Given the observed data on prices and expenditures $\{P_{ut}^{U}, X_{ut}^{U}\}$ and the substitution parameters $\{\sigma_{g}^{U}, \sigma_{g}^{F}, \sigma^{G}\}$, no supply-side assumptions are needed to undertake this analysis and recover the structural residuals $\{\varphi_{ut}^{U}, \varphi_{ft}^{F}, \varphi_{jgt}^{G}\}$. The reason is that we observe both prices (P_{ut}^{U}) and expenditures (X_{ut}^{U}). Therefore, we do not need to take a stand on the different supply-side forces that determine the observed prices (e.g. technology, factor prices, oligopoly, monopolistic competition or perfect competition). Hence, the only way in which supply-side assumptions enter our analysis is through the estimation of the elasticities of substitution, as discussed further in Section 3 below.

An important difference between our approach and standard exact price indexes for CES is that we allow the demand/quality parameters to change over time. Therefore, our framework captures demand/quality upgrading for individual foreign products (changes in φ_{ut}^{U}) for individual foreign firms (changes in φ_{ft}^{F}) and for individual tradable sectors (changes in φ_{jgt}^{G}). We also allow for proportional changes in the demand/quality for all foreign varieties relative to all domestic varieties within each sector, which are implicitly captured in

[11] For firms with no common products, we set the geometric mean of demand across all products equal to one ($\mathbb{M}_{ft}^{U}\left[\varphi_{ut}^{U}\right]=1$), which enables us to recover product demand (φ_{ut}^{U}) and construct the firm price index (P_{ft}^{F}) for these firms. This choice has no impact on the change in the exporter price indexes ($\mathbb{P}_{jigt}^{E}/\mathbb{P}_{jigt-1}^{E}$) and sectoral import price indexes ($\mathbb{P}_{jgt}^{G}/\mathbb{P}_{jgt-1}^{G}$), because firms with no common products enter these changes in price indexes through the variety correction terms ($\lambda_{jigt}^{F}/\lambda_{jigt-1}^{F}$ and $\lambda_{jgt}^{F}/\lambda_{jgt-1}^{F}$ respectively) that depend only on observed expenditures.

the shares of expenditure on foreign varieties within sectors (μ^{G}_{jgt}) in equation (7) for the sectoral price index (P^{G}_{jgt}). Similarly, we allow for proportional changes in the demand/quality for all tradable sectors relative to all non-tradable sectors, which are implicitly captured in the share of expenditure on tradable sectors (μ^{T}_{jt}) in equation (4) for the aggregate price index (P_{jt}). Finally, the only component of demand/quality that cannot be identified from the observed expenditure shares is proportional changes in demand/quality across all sectors (both traded and non-traded) over time. Nevertheless, our specification considerably generalizes the conventional assumption that demand/quality is time-invariant for all common varieties.

2.9 Exporter Price Movements

Having inverted the model to recover the unobserved demand/quality parameters that rationalize the observed data, we now show how to aggregate to the exporter price index that summarizes the cost of sourcing goods across countries and sectors. Recursively applying our log linear representation of the CES price index in equation (14) for the exporter and firm price indexes, we obtain the following exact log-linear decomposition of the exporter price index, as shown in Section A.2.9 of the web appendix:

$$\begin{aligned}\ln \mathbb{P}^{E}_{jigt} = \; & \underbrace{\mathbb{E}^{FU}_{jigt}\left[\ln P^{U}_{ut}\right]}_{\text{(i) Average log prices}} - \underbrace{\left\{\mathbb{E}^{F}_{jigt}\left[\ln \varphi^{F}_{ft}\right] + \mathbb{E}^{FU}_{jigt}\left[\ln \varphi^{U}_{ut}\right]\right\}}_{\text{(ii) Average log demand}} + \underbrace{\left\{\frac{1}{\sigma^{U}_{g}-1}\mathbb{E}^{FU}_{jigt}\left[\ln S^{U}_{ut} - \ln \frac{1}{N^{U}_{ft}}\right] + \frac{1}{\sigma^{F}_{g}-1}\mathbb{E}^{F}_{jigt}\left[\ln \mathbb{S}^{EF}_{ft} - \ln \frac{1}{N^{F}_{jigt}}\right]\right\}}_{\text{(iii) Dispersion demand-adjusted prices}} \\ & - \underbrace{\left\{\frac{1}{\sigma^{U}_{g}-1}\mathbb{E}^{F}_{jigt}\left[\ln N^{U}_{ft}\right] + \frac{1}{\sigma^{F}_{g}-1}\ln N^{F}_{jigt}\right\}}_{\text{(iv) Variety}},\end{aligned} \tag{18}$$

where $\mathbb{S}^{EF}_{ft}$ is the share of a firm in imports from an individual exporting country and sector, as defined in Section A.2.9 of the web appendix; $\mathbb{E}^{FU}_{jigt}\left[\ln P^{U}_{ut}\right] \equiv \frac{1}{N^{F}_{jigt}} \sum_{f\in\Omega^{F}_{jigt}} \frac{1}{N^{U}_{ft}} \sum_{u\in\Omega^{U}_{ft}} \ln P^{U}_{ut}$ is a mean across firms and products within that exporting country and sector; and $\mathbb{E}^{F}_{jigt}\left[\ln P^{F}_{ft}\right] \equiv \frac{1}{N^{F}_{jigt}} \sum_{f\in\Omega^{F}_{jigt}} \ln P^{F}_{ft}$ is a mean across firms for that country and sector.

Similarly, partitioning varieties into those that are common, entering and exiting, and taking differences over time, we obtain an analogous exact log linear decomposition for changes in the exporter price index:

$$\begin{aligned}\Delta \ln \mathbb{P}^{E}_{jigt} = \; & \underbrace{\mathbb{E}^{FU*}_{jigt}\left[\Delta \ln P^{U}_{ut}\right]}_{\text{(i) Average log prices}} - \underbrace{\left\{\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \varphi^{F}_{ft}\right] + \mathbb{E}^{FU*}_{jigt}\left[\Delta \ln \varphi^{U}_{ut}\right]\right\}}_{\text{(ii) Average log demand}} + \underbrace{\left\{\frac{1}{\sigma^{U}_{g}-1}\mathbb{E}^{FU*}_{jigt}\left[\Delta \ln S^{U*}_{ut}\right] + \frac{1}{\sigma^{F}_{g}-1}\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \mathbb{S}^{EF}_{ft}\right]\right\}}_{\text{(iii) Dispersion demand-adjusted prices}} \\ & + \underbrace{\left\{\frac{1}{\sigma^{U}_{g}-1}\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \lambda^{U}_{ft}\right] + \frac{1}{\sigma^{F}_{g}-1}\Delta \ln \lambda^{F}_{jigt}\right\}}_{\text{(iv) Variety}},\end{aligned} \tag{19}$$

as also shown in Section A.2.9 of the web appendix.

Equations (18) and (19) make explicit the three key features of our framework that allow exact aggregation from micro to macro. First, we can invert the model to recover the unobserved demand/quality parameters

$\{\varphi^U_{ut}, \varphi^F_{ft}, \varphi^G_{gt}\}$ that rationalize the observed data. Second, for each tier of utility, the CES price index can be written as a log linear form of these demand/quality parameters and the observed data. Third, demand is nested, such that the price index for utility tier K depends on the price index and demand/quality parameters for utility tier $K-1$. Combining these three properties, and noting that the mean for tier K of the means from tier $K-1$ remains linear, we obtain our exact log linear decomposition of aggregate variables into the contributions of different microeconomic mechanisms.

Each of the terms in these equations have an intuitive interpretation. The first term in equation (19) is the average log change in the price of common products sourced from exporting country i within sector g ($\mathbb{E}^{FU*}_{jigt}\left[\Delta \ln P^U_{ut}\right]$). This first component equals the log of a Jevons Index, which is a standard empirical measure of average prices, and is used to aggregate prices in the U.S. consumer price index.

The second term ($\mathbb{E}^{F*}_{jigt}\left[\ln \varphi^F_{ft}\right] + \mathbb{E}^{FU*}_{jigt}\left[\ln \varphi^U_{ut}\right]$) captures demand shifts or quality upgrading for common products and firms and its presence reflects the fact that consumers care about demand-adjusted prices rather than prices alone. Recall that our normalization in equation (A.2.7) implies that the average log change in common-product demand within foreign firms is equal to zero: $\mathbb{E}^{FU*}_{jigt}\left[\Delta \ln \varphi^U_{ut}\right] = 0$. Similarly, our normalization in equation (A.2.10) implies that the average log change in firm demand across *all* common foreign firms within a sector is equal to zero: $\mathbb{E}^{F*}_{jgt}\left[\Delta \ln \varphi^F_{ft}\right] = 0$. Nevertheless, the relative demand/quality of firms in different foreign countries within that sector can change, if demand/quality rises in some countries relative to others, in which case this second term is non-zero: $\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \varphi^F_{ft}\right] \neq \mathbb{E}^{F*}_{jgt}\left[\Delta \ln \varphi^F_{ft}\right] = 0$ for country $i \neq j$. Therefore, if one foreign exporter upgrades its demand/quality relative to another, this implies a fall in the cost of sourcing imports from that exporter relative to other foreign exporters.

The third term captures the dispersion of demand-adjusted prices across common products and firms for a given exporter and sector. Other things equal, if the dispersion of these demand-adjusted prices increases, this reduces the cost of sourcing goods from that exporter and sector ($\mathbb{E}^{FU*}_{jigt}\left[\Delta \ln S^{U*}_{ut}\right] < 0$ and $\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \mathbb{S}^{EF}_{ft}\right] < 0$). The reason is that this increased dispersion of demand-adjusted prices enhances the ability of consumers to substitute away from varieties with high demand-adjusted prices and towards varieties with low demand-adjusted prices.

The fourth term in equation (19) ($\frac{1}{\sigma^U_g - 1}\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \lambda^U_{ft}\right] + \frac{1}{\sigma^F_g - 1}\Delta \ln \lambda^F_{jigt}$) captures the effect of product turnover and firm entry and exit on the cost of sourcing imports from a given exporter and sector. If entering firms and products are more numerous or desirable than exiting firms and products, this again reduces the cost of sourcing goods from that exporter and sector ($\mathbb{E}^{F*}_{jigt}\left[\Delta \ln \lambda^U_{ft}\right] < 0$ and $\Delta \ln \lambda^F_{jigt} < 0$).

2.10 Patterns of Trade Across Sectors and Countries

Thus far, we have been focused on measuring the price indexes that determine the costs of sourcing goods from a given exporter and sector. The move from price indexes to trade patterns, however, is straightforward, because these patterns of trade are determined by relative price indexes. We can therefore translate our results for exporter price indexes into the determinants of patterns of trade across countries and sectors.

2.10.1 Revealed Comparative Advantage

We begin by deriving a theoretically-rigorous measure of revealed comparative advantage (RCA) that holds in all models based on a CES demand system. We start with importer j's expenditure on foreign exporter $i \neq j$ as a share of its expenditure on all foreign exporters within sector g at time t:

$$\mathbb{S}^{E}_{jigt} = \frac{\sum_{f \in \Omega^{F}_{jigt}} \left(P^{F}_{ft} / \varphi^{F}_{ft} \right)^{1-\sigma^{F}_{g}}}{\sum_{h \in \Omega^{E}_{jgt}} \sum_{f \in \Omega^{F}_{jhgt}} \left(P^{F}_{ft} / \varphi^{F}_{ft} \right)^{1-\sigma^{F}_{g}}} = \frac{\left(\mathbb{P}^{E}_{jigt} \right)^{1-\sigma^{F}_{g}}}{\left(\mathbb{P}^{G}_{jgt} \right)^{1-\sigma^{F}_{g}}}, \qquad i \neq j. \tag{20}$$

where the single superscript E is a mnemonic for exporter and indicates that this is the expenditure share for a foreign exporting country $i \neq j$; the numerator in equation (20) captures importer j's price index for exporting country i in sector g at time t ($\mathbb{P}^{E}_{jigt}$); and the denominator in equation (20) features importer j's overall import price index in sector g at time t ($\mathbb{P}^{G}_{jgt}$).

Using the definition of this exporter expenditure share (20), we measure RCA in sector g for import market j, by first taking the value of country i's exports relative to the geometric mean across countries for that sector ($\mathbb{X}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{X}^{E}_{jigt} \right]$), and then dividing by country i's geometric mean of this ratio across tradable sectors ($\mathbb{M}^{T}_{jit} \left[\mathbb{X}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{X}^{E}_{jigt} \right] \right]$):

$$RCA_{jigt} \equiv \frac{\mathbb{X}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{X}^{E}_{jigt} \right]}{\mathbb{M}^{T}_{jit} \left[\mathbb{X}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{X}^{E}_{jigt} \right] \right]} = \frac{\mathbb{S}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{S}^{E}_{jigt} \right]}{\mathbb{M}^{T}_{jit} \left[\mathbb{S}^{E}_{jigt} / \mathbb{M}^{E}_{jgt} \left[\mathbb{S}^{E}_{jigt} \right] \right]}, \tag{21}$$

where we use $\mathbb{X}^{E}_{jigt}$ to denote the value of bilateral exports from country i to importer $j \neq i$ within sector g at time t; $\mathbb{M}^{E}_{jgt} \left[\mathbb{X}^{E}_{jigt} \right] = \left(\prod_{h \in \Omega^{E}_{jgt}} \mathbb{X}^{E}_{jhgt} \right)^{1/N^{E}_{jgt}}$ is the geometric mean of these exports across all foreign exporters for that importer and sector; $\mathbb{M}^{T}_{jit} \left[\mathbb{X}^{E}_{jigt} \right] = \left(\prod_{k \in \Omega^{T}_{jit}} \mathbb{X}^{E}_{jikt} \right)^{1/N^{T}_{jit}}$ is the geometric mean of these exports across tradable sectors for that importer and foreign exporter; and $\mathbb{S}^{E}_{jigt} = \mathbb{X}^{E}_{jigt} / \sum_{h \in \Omega^{E}_{jgt}} \mathbb{X}^{E}_{jhgt} = \mathbb{X}^{E}_{jigt} / \mathbb{X}^{G}_{jgt}$ is the share of foreign exporter $i \neq j$ in country j's imports from all foreign countries within sector g at time t.

From equation (21), an exporter has a *revealed comparative advantage* in a sector within a given import market (a value of RCA_{jigt} greater than one) if its exports relative to the average exporter in that sector are larger than for its average sector. This RCA measure is similar to those in Costinot, Donaldson and Komunjer (2012) and Levchenko and Zhang (2016). However, instead of choosing an individual sector and country as the base for the double-differencing, we first difference relative to a hypothetical country within a sector (equal to the geometric mean country for that sector), and then second difference relative to a hypothetical sector (equal to the geometric mean across sectors).[12] We also derive our measure solely from our demand-side assumptions, without requiring a Ricardian supply-side to the model.

As we now show, these differences enable us to quantify the role of different economic mechanisms in understanding patterns of trade across countries and sectors. From equations (20) and (21), RCA captures the

[12] Our measure also relates closely to Balassa (1965)'s original measure of RCA, which divides a country's exports in a sector by the total exports of all countries in that sector, and then divides this ratio by the country's share of overall exports across all sectors. Instead, we divide a country's exports in a sector by the geometric mean exports in that sector across countries, and then divide this ratio by its geometric mean across sectors.

relative cost to an importer of sourcing goods across countries and sectors, as determined by relative price indexes and the elasticity of substitution ($\left(\mathbb{P}^{E}_{jigt}\right)^{1-\sigma^F_g}$):

$$RCA_{jigt} = \frac{\left(\mathbb{P}^{E}_{jigt}\right)^{1-\sigma^F_g} / \mathbb{M}^{E}_{jgt}\left[\left(\mathbb{P}^{E}_{jigt}\right)^{1-\sigma^F_g}\right]}{\mathbb{M}^{T}_{jit}\left[\left(\mathbb{P}^{E}_{jigt}\right)^{1-\sigma^F_g} / \mathbb{M}^{E}_{jgt}\left[\left(\mathbb{P}^{E}_{jigt}\right)^{1-\sigma^F_g}\right]\right]}. \tag{22}$$

Taking logarithms in equation (22), and using equation (18) to substitute for the log exporter price index ($\ln \mathbb{P}^{E}_{jigt}$), we can decompose differences in log RCA across countries and sectors into the contributions of average log prices ($\ln\left(RCA^{P}_{jigt}\right)$), average log demand ($\ln\left(RCA^{\varphi}_{jigt}\right)$), the dispersion of demand-adjusted prices ($\ln\left(RCA^{S}_{jigt}\right)$), and variety ($\ln\left(RCA^{N}_{jigt}\right)$):

$$\ln\left(RCA_{jigt}\right) = \underbrace{\ln\left(RCA^{P}_{jigt}\right)}_{\text{(i) Average log prices}} + \underbrace{\ln\left(RCA^{\varphi}_{jigt}\right)}_{\text{(ii) Average log demand}} + \underbrace{\ln\left(RCA^{S}_{jigt}\right)}_{\text{(iii) Dispersion demand-adjusted prices}} + \underbrace{\ln\left(RCA^{N}_{jigt}\right)}_{\text{(iv) Variety}}, \tag{23}$$

where each of these terms is defined in full in Section A.2.10.1 of the web appendix.

Each term is a double difference in logs, in which we first difference a variable for an exporter and sector relative to the mean across exporters for that sector (as in the numerator of RCA), before then second differencing the variable across sectors (as in the denominator of RCA). For example, to compute the average log price term ($\ln\left(RCA^{P}_{jigt}\right)$), we proceed as follows. In a first step, we compute average log product prices for an exporter and sector in an import market. In a second step, we subtract from these average log product prices their mean across all exporters for that sector and import market. In a third step, we difference these scaled average log product prices from their mean across all sectors for that exporter and import market. Other things equal, an exporter has a RCA in a sector if its log product prices relative to the average exporter in that sector are low compared to the exporter's average sector.

A key implication of equation (23) is that comparative advantage cannot be measured independently of demand when goods are differentiated ($\sigma^U_g < \infty$, $\sigma^F_g < \infty$, $\varphi^U_{ut} \neq \varphi^U_{\ell t}$ for $u \neq \ell$, and $\varphi^F_{ft} \neq \varphi^F_{mt}$ for $f \neq m$), in the same way that productivity cannot be measured independently of demand in this case.[13] The reason is that comparative advantage depends on relative price indexes, which cannot be inferred from relative prices alone if goods are differentiated. In such a setting, average demand/quality, the number of products and firms, and the dispersion of demand-adjusted prices across these products and firms (as captured by the dispersion of expenditure shares) are also important determinants of relative price indexes.

Similarly, partitioning varieties into those that are common, entering and exiting, and taking differences over time, we can decompose changes in RCA across countries and sectors into four analogous terms:

$$\Delta\ln\left(RCA^{*}_{jigt}\right) = \underbrace{\Delta\ln\left(RCA^{P*}_{jigt}\right)}_{\text{(i) Average log prices}} + \underbrace{\Delta\ln\left(RCA^{\varphi*}_{jigt}\right)}_{\text{(ii) Average log demand}} + \underbrace{\Delta\ln\left(RCA^{S*}_{jigt}\right)}_{\text{(iii) Dispersion demand-adjusted prices}} + \underbrace{\Delta\ln\left(RCA^{\lambda}_{jigt}\right)}_{\text{(iv) Variety}}, \tag{24}$$

[13]For a discussion of the centrality of demand to productivity measurement when goods are imperfect substitutes, see for example Foster, Haltiwanger and Syverson (2008) and De Loecker and Goldberg (2014).

where all four terms are again defined in full in subsection A.2.10.1 of the web appendix.

The interpretation of these four terms is similar to that for our decomposition of exporter price indexes above. Other things equal, an exporter's RCA in a sector rises if its prices fall faster than its competitors in that sector relative to other sectors. The second term incorporates the effects of average log demand/quality. All else constant, RCA increases in a sector if an exporter's demand/quality rises more rapidly than its competitors in that sector relative to other sectors. The third term summarizes the impact of the dispersion of demand-adjusted prices across varieties. Other things equal, RCA rises for an exporter in a sector if the dispersion of demand-adjusted prices increases relative to its competitors in that sector more than in other sectors. As its demand-adjusted prices become more dispersed, this enables consumers to more easily substitute from the exporter's less attractive varieties to its more attractive varieties, which increases the demand for its goods. Finally, the fourth term summarizes the contribution of entry/exit. All else constant, if entering varieties are more numerous or have lower demand-adjusted prices than exiting varieties, this increases the value of trade. Therefore, an exporter's RCA in a sector increases if the exporter's rate of net product and firm entry relative to its competitors in that sector exceeds its relative rate in other sectors.

2.10.2 Aggregate Trade

We now aggregate further to obtain an exact log linear decomposition of exporting countries' shares of total imports that can be used to examine the reasons for the large-scale changes in countries' import shares over our sample period (such as the dramatic rise in Chinese import penetration). At first sight, our ability to obtain log linear decompositions of both sectoral and aggregate trade is somewhat surprising, because aggregate trade is the sum of sectoral trade (rather than the sum of log sectoral trade). What makes this possible is that the structure of CES demand yields a closed-form solution for an exact Jensen's Inequality correction term that controls for the difference between the log of the sum and the sum of the logs.

Partitioning varieties into common, entering and exiting varieties, we show in Section A.2.10.2 of the web appendix that the log change in the share of foreign exporter i in importer j's total expenditure on all foreign importers can be exactly decomposed as follows:

$$\begin{aligned}
\Delta \ln \mathbb{S}^{E}_{jit} = & -\underbrace{\left\{\mathbb{E}^{TFU*}_{jit}\left[\left(\sigma^F_g-1\right)\Delta\ln P^U_{ut}\right]-\mathbb{E}^{TEFU*}_{jt}\left[\left(\sigma^F_g-1\right)\Delta\ln P^U_{ut}\right]\right\}}_{\text{(i) Average log prices}} \\
& +\underbrace{\left\{\mathbb{E}^{TFU*}_{jit}\left[\left(\sigma^F_g-1\right)\Delta\ln\varphi^U_{ut}\right]-\mathbb{E}^{TEFU*}_{jt}\left[\left(\sigma^F_g-1\right)\Delta\ln\varphi^U_{ut}\right]\right\}}_{\text{(ii) Average log product demand}}+\underbrace{\left\{\mathbb{E}^{TF*}_{jit}\left[\left(\sigma^F_g-1\right)\Delta\ln\varphi^F_{ft}\right]-\mathbb{E}^{TEF*}_{jt}\left[\left(\sigma^F_g-1\right)\Delta\ln\varphi^F_{ft}\right]\right\}}_{\text{(iii) Average log firm demand}} \\
& -\underbrace{\left\{\mathbb{E}^{TFU*}_{jit}\left[\frac{\sigma^F_g-1}{\sigma^U_g-1}\Delta\ln S^{U*}_{ut}\right]-\mathbb{E}^{TEFU*}_{jt}\left[\frac{\sigma^F_g-1}{\sigma^U_g-1}\Delta\ln S^{U*}_{ut}\right]\right\}}_{\text{(iv) Dispersion demand-adjusted product prices}}-\underbrace{\left\{\mathbb{E}^{TF*}_{jit}\left[\Delta\ln\mathbb{S}^{EF*}_{ft}\right]-\mathbb{E}^{TEF*}_{jt}\left[\Delta\ln\mathbb{S}^{EF*}_{ft}\right]\right\}}_{\text{(v) Dispersion demand-adjusted firm prices}} \\
& -\underbrace{\left\{\mathbb{E}^{TF*}_{jit}\left[\frac{\sigma^F_g-1}{\sigma^U_g-1}\Delta\ln\lambda^U_{ft}\right]-\mathbb{E}^{TEF*}_{jt}\left[\frac{\sigma^F_g-1}{\sigma^U_g-1}\Delta\ln\lambda^U_{ft}\right]\right\}}_{\text{(vi) Product Variety}}-\underbrace{\left\{\mathbb{E}^{T}_{jit}\left[\Delta\ln\lambda^F_{jigt}\right]-\mathbb{E}^{TE*}_{jt}\left[\Delta\ln\lambda^F_{jigt}\right]\right\}}_{\text{(vii) Firm Variety}}-\underbrace{\Delta\ln\left(\lambda^E_{jit}/\lambda^T_{jt}\right)}_{\text{(viii) Country-Sector Variety}} \qquad (25) \\
& +\underbrace{\Delta\ln\mathbb{K}^T_{jit}}_{\text{(ix) Country-sector Scale}}+\underbrace{\Delta\ln\mathbb{J}^T_{jit}}_{\text{(x) Country-sector Concentration}},
\end{aligned}$$

where the country-sector scale ($\Delta \ln \mathbb{K}^T_{jit}$) and country-sector concentration ($\Delta \ln \mathbb{J}^T_{jit}$) terms are defined in

Section A.2.10.2 of the web appendix; $\mathbb{E}_{jt}^{TEFU*}[\cdot]$, $\mathbb{E}_{jit}^{TFU*}[\cdot]$, $\mathbb{E}_{jit}^{TF*}[\cdot]$, $\mathbb{E}_{jt}^{TEF*}[\cdot]$, $\mathbb{E}_{jt}^{TE*}[\cdot]$ and $\mathbb{E}_{jit}^{T}[\cdot]$ are means across sectors, exporters, firms and products, as also defined in that section of the web appendix.

From the first term (i), an exporter's import share increases if the average prices of its products fall more rapidly than those of other exporters. In the second term (ii), our choice of units for product demand in equation (A.2.7) implies that the average log change in demand across common products within firms is equal to zero ($\mathbb{E}_{ft}^{U*}\left[\Delta \ln \varphi_{ut}^{U}\right]$), which implies that this second term is equal to zero. From the third term (iii), an exporter's import share also increases if the average demand/quality of its firms rises more rapidly than that of firms from other exporters within each sector (recall that our choice of units for firm demand only implies that its average log change equals zero across *all* foreign firms within each sector).

The fourth and fifth terms ((iv) and (v)) capture the dispersion of demand-adjusted prices across common products and firms. An exporter's import share increases if demand-adjusted prices become more dispersed across its products and firms compared to other foreign exporters. The sixth through eighth terms ((vi)-(viii)) capture the contribution of entry/exit to changes in country import shares. An exporter's import share increases if its entering products, firms and sectors are more numerous and/or have lower demand-adjusted prices compared to its exiting varieties than those for other foreign exporters.

The last two terms capture compositional effects across sectors. From the penultimate term (ix), an exporter's import share increases if its exports become more concentrated in sectors that account for large shares of expenditure relative to exports from other foreign countries. The final term (x) captures the concentration of imports across sectors for an individual exporter relative to their concentration across sectors for all foreign exporters. This final term is the exact Jensen's Inequality correction term discussed above.

2.11 Aggregate Prices

In addition to understanding aggregate trade patterns, researchers are often interested in understanding movements in the aggregate cost of living, since this is important determinant of real income and welfare. In Section A.2.11 of the web appendix, we show that the change in the aggregate price index in equation (4) can be exactly decomposed into the following five terms:

$$\underbrace{\Delta \ln P_{jt}}_{\substack{\text{Aggregate}\\ \text{Price Index}}} = \underbrace{\frac{1}{\sigma^{G}-1}\Delta \ln \mu_{jt}^{T}}_{\substack{\text{Non-Tradable}\\ \text{Competitiveness}}} + \underbrace{\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma_{g}^{F}-1}\Delta \ln \mu_{jgt}^{G}\right]}_{\substack{\text{Domestic}\\ \text{Competitiveness}}} + \underbrace{\mathbb{E}_{jt}^{T}\left[\Delta \ln \varphi_{jgt}^{G}\right]}_{\substack{\text{Average}\\ \text{Demand}}} + \underbrace{\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma^{G}-1}\Delta \ln \mathbb{S}_{jgt}^{T}\right]}_{\substack{\text{Dispersion demand-}\\ \text{adjusted prices across sectors}}} + \underbrace{\mathbb{E}_{jt}^{T}\left[\Delta \ln \mathbb{P}_{jgt}^{G}\right]}_{\substack{\text{Aggregate Import}\\ \text{Price Indexes}}}, \quad (26)$$

where $\mathbb{S}_{jgt}^{T}$ is the share of an individual tradable sector in expenditure on all tradable sectors. Recall that the set of tradable sectors is constant over time and hence there are no terms for the entry and exit of sectors in equation (26).

The first three terms capture shifts in aggregate prices that can be inferred from changes in market shares or demand. The first term ($\frac{1}{\sigma^{G}-1}\Delta \ln \mu_{jt}^{T}$) captures the relative attractiveness of varieties in the tradable and non-tradable sectors. Other things equal, a fall in the share of expenditure on tradable sectors ($\Delta \ln \mu_{jt}^{I} < 0$) implies that varieties in non-tradable sectors have become relatively more attractive, which reduces the cost of living. The second term ($\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma_{g}^{F}-1}\Delta \ln \mu_{jgt}^{G}\right]$) captures the relative attractiveness of domestic varieties within sectors. Other things equal, a fall in the average share of expenditure on foreign varieties within sectors ($\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma_g^F-1}\Delta\ln\mu_{jgt}^{G}\right] < 0$) implies that domestic varieties have become relatively more attractive within sectors, which again reduces the cost of living. The third term ($\mathbb{E}_{jt}^{T}\left[\Delta\ln\varphi_{jgt}^{G}\right]$) captures changes in the average demand/quality for tradable sectors, where the superscript T on the expectation indicates that this mean is taken across the subset of tradable sectors ($\Omega^{T} \subseteq \Omega^{G}$). Given our choice of units in which to measure sector demand/quality, this third term is equal to zero ($\mathbb{E}_{jt}^{T}\left[\Delta\ln\varphi_{jgt}^{G}\right] = 0$). The fourth term ($\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma^G-1}\Delta\ln\mathbb{S}_{jgt}^{G}\right]$) captures changes in the dispersion of demand-adjusted prices across tradable sectors. Intuitively, when sectors are substitutes ($\sigma^{G} > 1$), an increase in the dispersion of demand-adjusted prices across sectors reduces the cost of living, as consumers can substitute from less to more desirable sectors. The fifth and final term ($\mathbb{E}_{jt}^{T}\left[\Delta\ln\mathbb{P}_{jgt}^{G}\right]$) captures changes in aggregate import price indexes across all tradable sectors. Other things equal, a fall in these aggregate import price indexes ($\mathbb{E}_{jt}^{T}\left[\Delta\ln\mathbb{P}_{jgt}^{G}\right] < 0$) reduces the cost of living. We now show that this fifth term can be further decomposed.

Partitioning goods into common, entering and exiting varieties, Section A.2.11 of the web appendix shows that the change in aggregate import price indexes can be exactly decomposed as follows:

$$
\begin{aligned}
\underbrace{\mathbb{E}_{jt}^{T}\left[\Delta\ln\mathbb{P}_{jgt}^{G}\right]}_{\substack{\text{Aggregate Import}\\ \text{Price Indexes}}} = {} & \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\Delta\ln P_{ut}^{U}\right]}_{\text{Average log prices}} - \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\Delta\ln\varphi_{ft}^{F}\right]}_{\substack{\text{Average log}\\ \text{firm demand}}} - \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\ln\varphi_{ut}^{U}\right]}_{\substack{\text{Average log}\\ \text{product demand}}} \\
& \underbrace{\mathbb{E}_{jt}^{TE*}\left[\frac{1}{\sigma_g^F-1}\Delta\ln\mathbb{S}_{jigt}^{E}\right]}_{\substack{\text{Dispersion country-sector}\\ \text{demand-adjusted prices}}} + \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\frac{1}{\sigma_g^F-1}\Delta\ln\mathbb{S}_{ft}^{EF}\right]}_{\substack{\text{Dispersion firm}\\ \text{demand-adjusted prices}}} + \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\frac{1}{\sigma_g^U-1}\Delta\ln S_{ut}^{U}\right]}_{\substack{\text{Dispersion product}\\ \text{demand-adjusted prices}}} \\
& + \underbrace{\mathbb{E}_{jt}^{T*}\left[\frac{1}{\sigma_g^F-1}\Delta\ln\lambda_{jgt}^{E}\right]}_{\substack{\text{Country - Sector}\\ \text{Variety}}} + \underbrace{\mathbb{E}_{jt}^{TE*}\left[\frac{1}{\sigma_g^F-1}\Delta\ln\lambda_{jigt}^{F}\right]}_{\substack{\text{Firm}\\ \text{Variety}}} + \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\frac{1}{\sigma_g^U-1}\Delta\ln\lambda_{ft}^{U}\right]}_{\substack{\text{Product}\\ \text{Variety}}}.
\end{aligned} \tag{27}
$$

The interpretation of each of these components in equation (27) is analogous to the interpretation of the corresponding components of countries aggregate import shares in equation (25). Aggregate import price indexes fall with declines in average product prices, rises in average firm and product demand, increases in the dispersion of demand-adjusted prices across surviving countries, firms and products, and if entering countries, firms and products are more numerous or more desirable than those that exit.

3 Structural Estimation

In order to take our model to data, we need estimates of the elasticities of substitution $\{\sigma_g^U, \sigma_g^F, \sigma^G\}$. We now turn to our estimation of these elasticities, which is the only place in our approach where we are required to make assumptions about the supply-side. In particular, in the data, we observe changes in expenditure shares and changes in prices, which provides a standard demand and supply identification problem. In a CES demand system with N goods, this identification problem can be equivalently formulated as follows: we have N parameters, which include $N-1$ independent demand shifters (under a normalization) and one elasticity of substitution, but we have only $N-1$ independent equations for expenditure shares, resulting in underidentification.

In our baseline specification, we estimate these elasticities of substitution using an extension of the reverse-weighting (RW) estimator of Redding and Weinstein (2016). This reverse weighting estimator solves the above underidentification problem by augmenting the $N-1$ independent equations of the demand system with two additional equations derived from three equivalent ways of writing the change in the unit expenditure function. In the web appendix, we also report robustness checks, in which we compare our RW estimates of the elasticities of substitution to alternative estimates, and in which we examine the sensitivity of our results to alternative values of these elasticities of substitution using a grid search.

We extend the RW estimator to a nested demand system and show that the estimation problem is recursive. In a first step, we estimate the elasticity of substitution across products (σ_g^U) for each sector g. In a second step, we estimate the elasticity of substitution across firms (σ_g^F) for each sector g. In a third step, we estimate the elasticity of substitution across sectors (σ^G).

In this section, we illustrate the RW estimator for the product tier of utility, and report the full details of the nested estimation and the moment equation in Section A.3 of the web appendix. The RW estimator is based on three equivalent expressions for the change in the CES unit expenditure function: one from the demand system, a second from taking the forward difference of the unit expenditure function, and a third from taking the backward difference of the unit expenditure function. Together these three expressions imply the following two equalities

$$\Theta_{ft,t-1}^{U+}\left[\sum_{u\in\Omega_{ft,t-1}^U} S_{ut-1}^{U*}\left(\frac{P_{ut}^U}{P_{ut-1}^U}\right)^{1-\sigma_g^U}\right]^{\frac{1}{1-\sigma_g^U}} = \mathbb{M}_{ft}^{U*}\left[\frac{P_{ut}^U}{P_{ut-1}^U}\right]\left(\mathbb{M}_{ft}^{U*}\left[\frac{S_{ut}^{U*}}{S_{ut-1}^{U*}}\right]\right)^{\frac{1}{\sigma_g^U-1}}, \tag{28}$$

$$\Theta_{ft,t-1}^{U-}\left[\sum_{u\in\Omega_{ft,t-1}^U} S_{ut}^{U*}\left(\frac{P_{ut}^U}{P_{ut-1}^U}\right)^{-(1-\sigma_g^U)}\right]^{-\frac{1}{1-\sigma_g^U}} = \mathbb{M}_{ft}^{U*}\left[\frac{P_{ut}^U}{P_{ut-1}^U}\right]\left(\mathbb{M}_{ft}^{U*}\left[\frac{S_{ut}^{U*}}{S_{ut-1}^{U*}}\right]\right)^{\frac{1}{\sigma_g^U-1}}, \tag{29}$$

where the variety correction terms ($\left(\lambda_{ft}^U/\lambda_{ft-1}^U\right)^{\frac{1}{\sigma_g^U-1}}$) have cancelled because they are common to all three expressions; $\Theta_{ft,t-1}^{U+}$ and $\Theta_{ft,t-1}^{U-}$ are forward and backward aggregate demand shifters respectively, which summarize the effect of changes in the relative demand for individual products on the unit expenditure function (as defined in Section A.3 of the web appendix); finally the equalities in equations (28) and (29) are robust to introducing a Hicks-neutral shifter of demand/quality across all products within each firm, which would cancel from both sides of the equation (like the variety correction term).

The RW estimator uses equations (28) and (29) to estimate the elasticity of substitution across products (σ_g^U) under the identifying assumption that the shocks to relative demand/quality cancel out across products:

$$\Theta_{ft,t-1}^{U+} = \left(\Theta_{ft,t-1}^{U-}\right)^{-1} = 1. \tag{30}$$

The asymptotic properties of this estimator are characterized in Redding and Weinstein (2016). The RW estimator is consistent as demand shocks become small ($\varphi_{ut}^U/\varphi_{ut-1}^U \to 1$) or as the number of common goods becomes large and demand shocks are independently and identically distributed ($N^{U}_{t,t-1} \to \infty$). More generally, the identifying assumption in equation (30) is satisfied up to a first-order approximation. Therefore, the RW estimator can be interpreted as providing a first-order approximation to the data. In practice, we find that the RW estimated elasticities are similar to those estimated using other methods, such as the generalization of the Feenstra (1994) estimator used in Hottman et al (2016). More generally, a key advantage of our CES specification is that it is straightforward to undertake robustness checks to alternative values of these elasticities of substitution using a grid search.

4 Data Description

To undertake our empirical analysis of the determinants of trade patterns and aggregate prices, we use international trade transactions data that are readily available from customs authorities. In this section, we briefly discuss the U.S. trade transactions data that we use in the paper, and report further details in Section A.4.1 of the web appendix. In Section A.4.2 of the web appendix, we discuss the Chilean trade transactions data that we use in robustness tests in the web appendix.

For each U.S. import customs shipment, we observe the cost inclusive of freight value of the shipment in U.S. dollars (market exchange rates), the quantity shipped, the date of the transaction, the product classification (according to 10-digit Harmonized System (HS) codes), the country of origin, and a partner identifier containing information about the foreign exporting firm.[14] We concord the HS-digit 10-digit products to 4-digit sectors in the North American Industry Classification System (NAICS). We are thus able to construct a dataset for a single importer j (the U.S.) with many exporters i (countries of origin), sectors g (4-digit NAICs codes), firms f (foreign firm identifiers within exporters within sectors), and products u (10-digit HS codes within foreign firm identifiers, within exporters and within sectors) and time t (year). We standardize the units in which quantities are reported (e.g. we convert dozens to counts and grams to kilograms). We also drop any observations for which countries of origin or foreign firm identifiers are missing. Finally, we collapse the import shipments data to the annual level by exporting firm and product, weighting by trade value, which yields a dataset on U.S. imports by source country (exporter), foreign firm, product and year from 1997-2011. In our final year of 2011, we have over 3.7 million observations by exporter-firm-product.

Our measure of prices is the export unit value of an exporting firm within a 10-digit HS category. While these data necessarily involve some aggregation across different varieties of products supplied by the same exporting firm within an observed product category, Section A.7 of the web appendix shows that our framework generalizes to the case in which firms make product decisions at a more disaggregated level than observed in the data. In this case, the product demand shifter (φ^{U}_{ut}) captures unobserved compositional differences within each observed product category. Moreover, 10-digit HS categories are relatively narrowly defined, and the coverage of sectors is much wider than in datasets that directly survey prices. As a result, many authors—including those working for statistical agencies—advocate for greater use of unit value data in the

[14] See Kamal, Krizan and Monarch (2015) for further discussion of the U.S. trade transactions data and comparisons of these partner identifiers using import data for the U.S. and export data from foreign countries. In robustness checks, we show that we continue to find that the variety and average demand/quality terms dominate if we omit this partner identifier (and focus only on countries, sectors and products) or use Chilean trade transactions data that report foreign brands.

construction of import price indexes.[15] Furthermore, existing research comparing aggregate import price indexes constructed using unit values and directly surveyed prices finds only small differences between them, as reported using U.S. data in Amiti and Davis (2009). Similarly, in our data we find that the correlation between a Cobb-Douglas price index (using lagged import shares as weights) and the BLS import price index is 0.93, which suggests that unit value indexes capture much of the variation of import indexes based on price quotes.

In Section A.5.5 of the web appendix, we show that our U.S. trade transactions data exhibit the same properties as found by a number of existing studies in the empirical trade literature.[16] Two key features are the high concentration of trade across countries and the dramatic increase in Chinese import penetration. As shown in Figure A.5.5 of that section of the web appendix, the top 20 import source countries account for around 80 percent of U.S. imports in each year; China's import share more than doubles from 7 to 18 percent from 1997-2011; in contrast, Japan's import share more than halves from 14 to 6 percent over this period.

5 Empirical Results

We present our results in several stages. We begin in Section 5.1 by reporting our estimates of the elasticities of substitution (σ_g^U, σ_g^F, σ^G), which we use to invert the model and recover the values of product, firm and sector demand/quality (φ_{ut}^U, φ_{ft}^F, φ_{jgt}^G). In Section 5.2, we use these estimates to compute the exporter price indexes that determine the cost of sourcing goods across countries and sectors. In Section 5.3, we report our main results for comparative advantage, aggregate trade and aggregate prices. In Section 5.4, we compare the results of our framework with special cases that impose additional theoretical restrictions. In Section A.6 of the web appendix, we replicate all of these specification using our Chilean trade transactions data, and show that we find the same qualitative and quantitative pattern of results.

5.1 Elasticities of Substitution

In Table 1, we summarize our baseline estimates of the elasticities of substitution (σ_g^U, σ_g^F, σ^G). Since we estimate a product and firm elasticity for each sector, it would needlessly clutter the paper to report all of these elasticities individually. Therefore we report quantiles of the distributions of product and firm elasticities (σ_g^U, σ_g^F) across sectors and the single estimated elasticity of substitution across sectors (σ^G). The estimated product and firm elasticities are significantly larger than one statistically and always below eleven. We find a median estimated elasticity across products (σ_g^U) of 6.29, a median elasticity across firms (σ_g^F) of 2.66, and an elasticity across sectors (σ^G) of 1.36. These results imply that products within firms, firms within sectors and sectors are imperfect substitutes for one another, which has important implications for the measurement of comparative advantage below, because observed product prices are no longer sufficient statistics for the cost of sourcing goods across countries and sectors.

[15]For instance, Nakamura et al (2015) argue for the superiority of indexes based on disaggregated unit value data on theoretical grounds and "recommend alternatives to conventional price indexes that make use of unit values."

[16]For example, see Bernard, Jensen and Schott (2009) and Bernard, Jensen, Redding and Schott (2009) for the U.S.; Mayer, Melitz and Ottaviano (2014) for France; and Manova and Zhang (2012) for China.

Although we do not impose this restriction on the estimation, we find a natural ordering, in which varieties are more substitutable within firms than across firms, and firms are more substitutable within industries than across industries: $\hat{\sigma}_g^U > \hat{\sigma}_g^F > \hat{\sigma}^G$. We find that the product elasticity is significantly larger than the firm elasticity at the 5 percent level of significance for all sectors, and the firm elasticity is significantly larger than the sector elasticity at this significance level for all sectors as well.[17] Therefore, the data reject the special cases in which consumers only care about firm varieties ($\sigma_g^U = \sigma_g^F = \sigma^G$), in which varieties are perfectly substitutable within sectors ($\sigma_g^U = \sigma_g^F = \infty$), and in which products are equally differentiated within and across firms for a given sector ($\sigma_g^U = \sigma_g^F$). Instead, we find evidence of both firm differentiation within sectors and product differentiation within firms.

Percentile	Elasticity Across Products (σ_g^U)	Elasticity Across Firms (σ_g^F)	Elasticity Across Sectors (σ^G)	Product-Firm Difference ($\sigma_g^U - \sigma_g^F$)	Firm-Sector Difference ($\sigma_g^F - \sigma^G$)
Min	5.14	1.97	1.36	1.51	0.60
5th	5.43	2.06	1.36	2.42	0.69
25th	5.85	2.36	1.36	3.13	1.00
50th	6.29	2.66	1.36	3.48	1.30
75th	6.99	3.41	1.36	3.94	2.04
95th	8.36	4.83	1.36	4.77	3.47
Max	10.59	7.66	1.36	5.51	6.30

Note: Estimated elasticities of substitution from the reverse-weighting estimator discussed in section 3 and in section A.3 of the web appendix. Sectors are 4-digit North American Industrial Classification (NAICS) codes; firms are foreign exporting firms within each foreign country within each sector; and products are 10-digit Harmonized System (HS) codes within foreign exporting firms within sectors.

Table 1: Estimated Elasticities of Substitution, Within Firms (σ_g^U), Across Firms (σ_g^F) and Across Sectors (σ^G) (U.S. Data)

Our estimated elasticities of substitution are broadly consistent with those of other studies that have used similar data but different methodologies and/or nesting structures. In line with Broda and Weinstein (2006), we find lower elasticities of substitution as one moves to higher levels of aggregation. Our estimates of the product and firm elasticities (σ_g^F and σ_g^U) are only slightly smaller than those estimated by Hottman et al. (2016) using different data (U.S. barcodes versus internationally-traded HS products) and a different estimation methodology based on Feenstra (1994).[18] Therefore, our estimated elasticities do not differ substantially from those obtained using other standard methodologies. As a check on the sensitivity of our estimated elasticities to the definition of categories, we re-estimated the product, firm, and sector elasticities using 6-digit instead of 4-digit NAICS codes as our definition of sectors. We find a similar pattern of results, with a median product elasticity of 6.20, a median firm elasticity of 2.70, and a sector elasticity of 1.47. As a check on the sensitivity of our results for comparative advantage to these estimated elasticities, report the results of a grid search over a range of alternative values for these elasticities in Section 5.3 below.

[17] In Figure A.5.1 in Section A.5.1 of the web appendix, we show the bootstrap confidence intervals for each sector.

[18] Our median estimates for the elasticities of substitution within and across firms of 6.3 and 2.7 respectively compare with those of 6.9 and 3.9 respectively in Hottman et al. (2016).

5.2 Exporter Price Indexes Across Sectors and Countries

We use these estimated elasticities ($\sigma_g^U, \sigma_g^F, \sigma^G$) to recover the structural residuals (φ_{ut}^U, φ_{ft}^F, φ_{jgt}^G) and solve for the exporter price indexes ($\mathbb{P}_{jigt}^E$) that summarize the cost of sourcing goods from each exporter and sector. A key implication of our framework is that these exporter price indexes depend not only on conventional average price terms, but also on the non-conventional forces of average demand/quality, variety and the dispersion of demand-adjusted prices. We now quantify the relative importance of each of these components in our data.

In the four panels of Figure 2, we display a bin scatter of the log of the exporter price index ($\ln \mathbb{P}_{jigt}^E$) and each of its components against average log product prices ($\mathbb{E}_{jigt}^{FU}\left[\ln P_{ut}^U\right]$), where the bins are twenty quantiles of each variable.[19] In each panel, we also show the regression relationship between the two variables based on the disaggregated (i.e., not binned) data. For brevity, we show results for the final year of our sample in 2011, but find the same pattern for other years in our sample. In the top-left panel, we compare the log exporter price index ($\ln \mathbb{P}_{jigt}^E$) to average log product prices ($\mathbb{E}_{jigt}^{FU}\left[\ln P_{ut}^U\right]$). In the special case in which firms and products are perfect substitutes within sectors ($\sigma_g^U = \sigma_g^F = \infty$) and there are no differences in demand/quality ($\varphi_{ft}^F = \varphi_{mt}^F$ for all f, m and $\varphi_{ut}^U = \varphi_{\ell t}^U$ for all u, ℓ), these two variables would be perfectly correlated. In contrast to these predictions, we find a positive but imperfect relationship, with an estimated regression slope of 0.59 and R^2 of 0.23. Therefore, the true cost of sourcing goods across countries and sectors can differ substantially from standard empirical measures of average prices.

In the remaining panels of Figure 2, we explore the three sources of differences between the exporter price index and average log product prices. As shown in the top-right panel, exporter sectors with high average prices (horizontal axis) also have high average demand/quality (vertical axis), so that the impact of higher average prices in raising sourcing costs is partially offset by higher average demand/quality. This positive relationship between average prices and demand/quality is strong and statistically significant, with an estimated regression slope of 0.41 and R^2 of 0.28. This finding of higher average demand for products with higher average prices is consistent with the quality interpretation of demand stressed in Schott (2004), in which producing higher quality incurs higher production costs.[20]

In measuring demand/quality as a residual that shifts expenditure shares conditional on price, we follow a long line of research in trade and industrial organization. This approach is similar to that taken to measure productivity in the growth literature, in which total factor productivity is a residual that shifts output conditional on inputs. The substantial variation in firm exports conditional on price is the underlying feature of the data that drives our finding of an important role for demand/quality in Figure 2. For plausible values of the elasticity of substitution, the model cannot explain this sales variation by price variation, and hence it is attributed to demand/quality. This result implies that the large class of trade models based on CES demand requires heterogenous demand/quality *and* cost shifts in order to rationalize the data.

In the bottom-left panel of Figure 2, we show that the contribution from the number of varieties to the

[19]We use a bin scatter, because U.S. Census disclosure requirements preclude showing results for each exporter-sector using the U.S. data. In Section A.6.2 of the web appendix, we show results by exporter-sector using publicly-available Chilean data.

[20]This close relationship between demand/quality and prices is consistent the findings of a number of studies, including the analysis of U.S. barcode data in Hottman et al. (2016) and the results for Chinese footwear producers in Roberts et al. (2011).

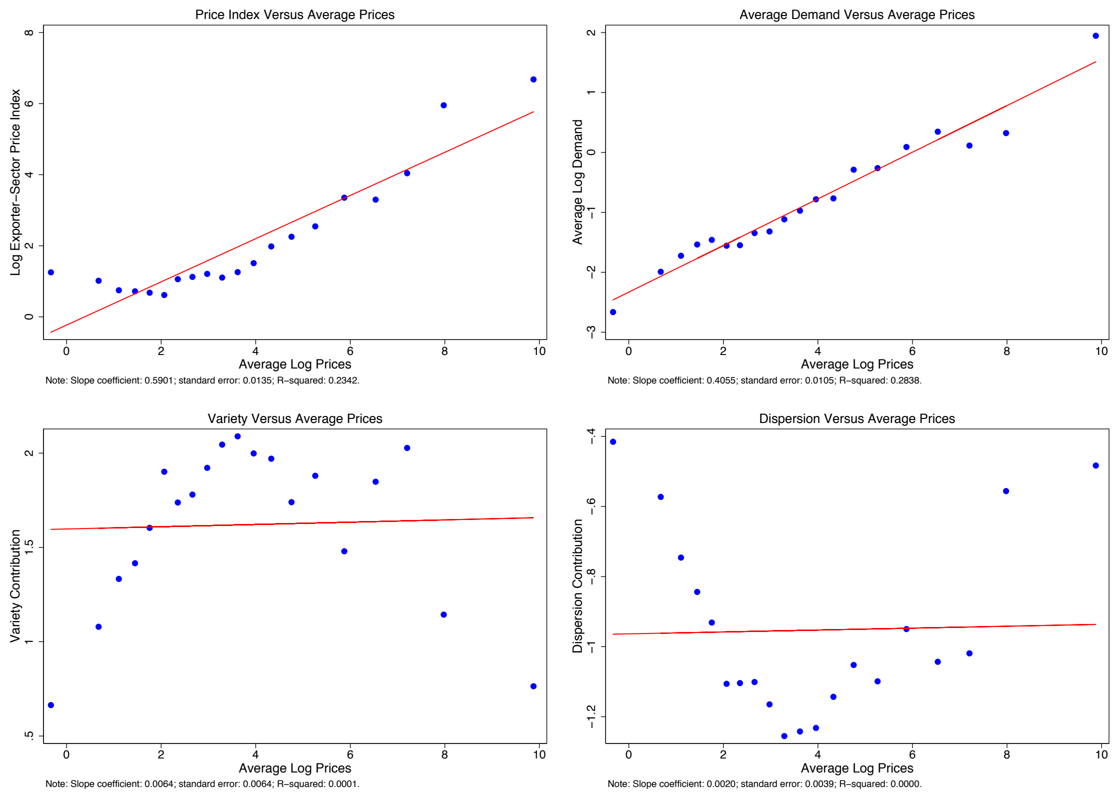

Figure 2: U.S. Exporter-Sector Price Indexes and their Components Versus Average Log Product Prices, 2011

exporter-sector price index exhibits an inverse U-shape, at first increasing with average prices, before later decreasing. This contribution ranges by more than two log points, confirming the empirical relevance of consumer love of variety. In contrast, in the bottom-right panel of Figure 2, we show that the contribution from the dispersion of demand-adjusted prices displays the opposite pattern of a U-shape, at first decreasing with average prices before later increasing. While the extent of variation is smaller than for the variety contribution, this term still fluctuates by around half a log point between its minimum and maximum value. Therefore, the imperfect substitutability of firms and products implies important contributions from the number of varieties and the dispersion in the characteristics of those varieties towards the true cost of sourcing goods across countries and sectors.

These non-conventional determinants are not only important in the cross-section but are also important for changes in the cost of sourcing goods over time. A common empirical question in macroeconomics and international trade is the effect of price shocks in a given sector and country on prices and real economic variables in other countries. However, it is not uncommon to find that measured changes in prices often appear to have relatively small effects on real economic variables, which has stimulated research on "elasticity puzzles" and the "exchange-rate disconnect." Although duality provides a precise mapping between prices and quantities, the actual price indexes used by researchers often differ in important ways from the formulas for price indexes from theories of consumer behavior. For example, as we noted earlier, our average price term is the log of the "Jevons Index," which is used by the U.S. Bureau of Labor Statistics (BLS) as part of its calculation of the consumer price index. Except in special cases, however, this average price term will not

US Aggregate Price Growth 1998-2011			
	Aggregate	Aggregate Terms	With Import Terms
Aggregate Price	0.22	-	-
Import Prices	-	0.12	-
Domestic Competitiveness	-	0.30	0.30
Non-Tradable Competitiveness	-	-0.19	-0.19
Sector Dispersion	-	0.00	0.00
Average Prices	-	-	0.39
Product Variety	-	-	-0.01
Product Dispersion	-	-	0.00
Firm Demand	-	-	-0.02
Firm Variety	-	-	-0.16
Firm Dispersion	-	-	-0.05
Country-Sector Variety	-	-	0.00
Country-Sector Dispersion	-	-	-0.03

Table 2: U.S. Aggregate Price Growth 1998-2011

equal the theoretically-correct measure of the change in the unit expenditure function.

We first demonstrate the importance of this point for aggregate prices. In Table 2, we decompose the log change in the U.S. aggregate cost of living from 1998-11 using equations (26) and (27). In the first column, we find that the aggregate U.S. price index increased by 0.22 log points over this time period. In the second column, we decompose this price change into four elements. First, the import price index rose by 0.12 log units which accounted for a little over half of the aggregate movement. Second, the value of imports rose as a share of tradables despite the rise in import prices, which implies that the exact price index of domestic tradables must have risen even more. This change in domestic competitiveness resulted in an increase in the price index by of an additional 0.3 log units. Offsetting this increase was a decline in the share of tradables in the US economy, which implies a relative decline in non-tradable prices that equaled a 0.19 log-unit decline in the U.S. aggregate price index. Finally, there was a negligible contribution from the dispersion of demand-adjusted prices across sectors. Thus, our decomposition enables us to capture not only the impact of import prices on aggregate prices, but also the impact of relative movements in the price indexes of domestic tradables and non-tradables.

Interestingly, the 0.12 log-point increase in aggregate import prices is much less than the 0.41 log point change in import prices between 1998 and 2011 reported in the BLS's U.S. Import Price Index for All Commodities. We can see the reason for the difference in the third column which expands our theoretical measure of the import price index into its components. The average log-price change, which equals the log of the Jevons index (the first term in equation (27)) rose by 0.39 log points over this time period: remarkably close to the 0.41 log point change reported in official series. Moreover, log changes in these two indexes are highly correlated in annual data as well ($\rho = 0.72$), which indicates that even at higher frequencies our Jevons index captures much of the variation in average import price changes as measured by the BLS. In other words, one obtains a very similar measure of import price increases regardless of whether one uses averages of log unit values or the price quotes used by the BLS in its Import Price Index. As we have been emphasizing, however, this index does not capture many of the other forces that matter for cost-of-living changes in more sophisticated models of consumer behavior. In particular, we find that the positive contribution from higher average prices of imported goods was offset by a substantial negative contribution from firm variety (see equation (27) for the definition of each term). This expansion in firm import variety reduced the cost of imported goods by around 0.16 log points. Changes in average firm demand and the dispersion of demand-adjusted prices across firms also acted to reduce aggregate import prices over this period. As a result, the true increase in the cost of imported goods from 1998-2011 was only 0.12 log points, less than one third the value implied by the conventional Jevons Index. In other words, a theory-based measure of aggregate import prices behaves very differently from one based only on average prices.

We next show that this point applies not only to aggregate import prices but also to the changes in exporter price indexes ($\Delta \ln \mathbb{P}^E_{jigt}$) that summarize the cost of sourcing goods across countries and sectors. Figure 3 displays the same information as in Figure 2, but for log changes from 1998-2011 rather than for log levels in 2011. In changes, the correlation between average prices and the true model-based measure of the cost of sourcing goods is much weaker (top-left panel) and the role for demand/quality is even greater (top-right panel). Indeed, the slope for the regression of average log changes in demand/quality on average log changes in prices is 0.92, indicating that most price changes are almost completely offset by demand/quality changes. This result suggests that the standard assumption of no shifts in demand/quality, which underlies standard price indexes such as the Sato-Vartia, is problematic. Price and demand/quality shifts are highly correlated.

5.3 Trade Patterns

We now use our results connecting RCA to relative exporter price indexes to examine the importance of the different components of these price indexes for comparative advantage across countries and sectors. We start with the decompositions of the level and change of RCA in equations (23) and (24) in Section 2.10.1 above. We use a variance decomposition introduced into the international trade literature by Eaton, Kortum and Kramarz (2004). We assess the contribution of each mechanism by regressing each component of RCA on the overall value of RCA. Therefore, for the level of RCA in equation (23), we have:

$$\begin{aligned}
\ln\left(RCA^P_{jigt}\right) &= \alpha_P + \beta_P \ln\left(RCA_{jigt}\right) + u^P_{jigt}, \\
\ln\left(RCA^\varphi_{jigt}\right) &= \alpha_\varphi + \beta_\varphi \ln\left(RCA_{jigt}\right) + u^\varphi_{jigt}, \\
\ln\left(RCA^S_{jigt}\right) &= \alpha_S + \beta_S \ln\left(RCA_{jigt}\right) + u^S_{jigt}, \\
\ln\left(RCA^N_{jigt}\right) &= \alpha_N + \beta_N \ln\left(RCA_{jigt}\right) + u^N_{jigt}.
\end{aligned} \tag{31}$$

where observations are exporters i and sectors g for a given importer j and year t. Since the sum of the dependent variables equals the independent variable, by the properties of OLS, $\beta_P + \beta_\varphi + \beta_S + \beta_N = 1$, and the relative value of each coefficient tells us the relative importance of each component of exporter price indexes. Similarly, we regress the log change in each component in equation (24) on the overall log change in RCA.

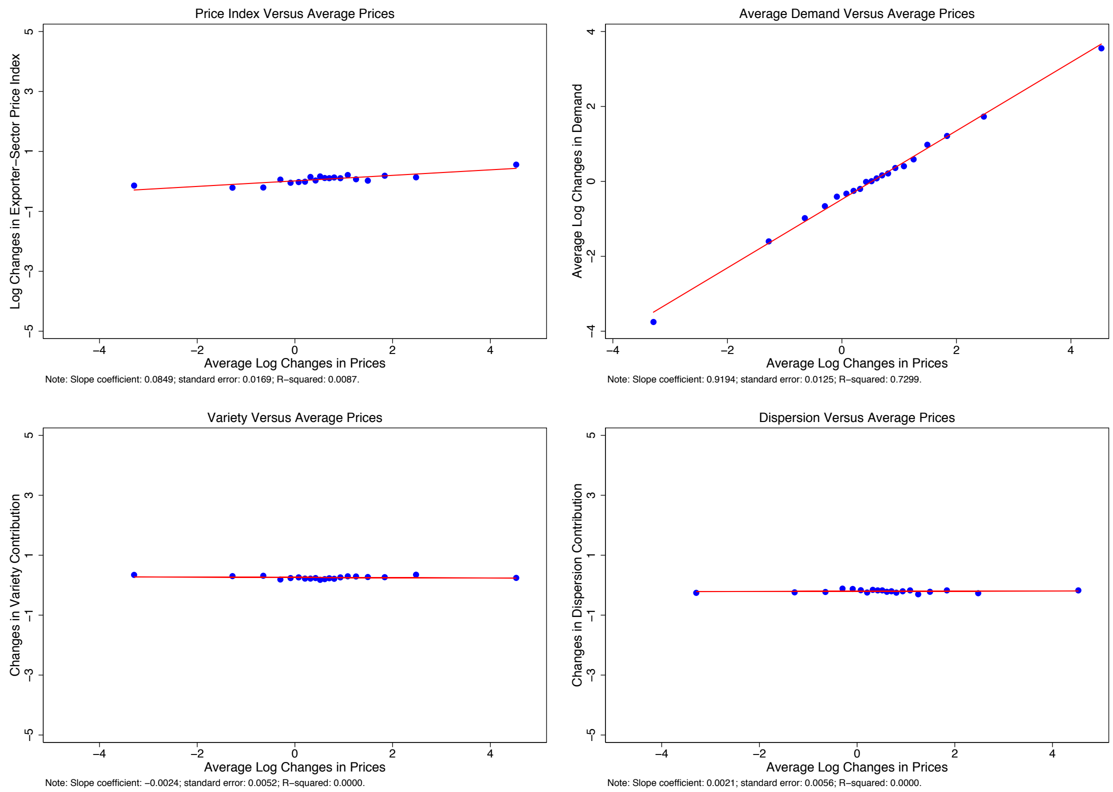

Figure 3: Log Changes in U.S. Sector-Exporter Price Indexes and their Components Versus Average Log Changes in Product Prices, 1998-2011

In Table 3, we report the results of these decompositions for both levels of RCA (Columns (1)-(2)) and changes of RCA (Columns (3)-(4)). In Columns (1) and (3), we undertake these decompositions down to the firm level. In Columns (2) and (4), we undertake them all the way down to the product level. For brevity, we concentrate on the results of the full decomposition in Columns (2) and (4). We find that average prices are comparatively unimportant in explaining patterns of trade. In the cross-section, average product prices account for 6.5 percent of the variation in RCA. In the time-series, we find that higher average prices are more than offset by higher demand/quality, resulting in a negative contribution of 4.8 percent from prices to changes in RCA. These results reflect the low correlations between average prices and exporter price indexes seen in the last section. If average prices are weakly correlated with exporter price indexes, they are unlikely to matter much for RCA, because RCA is determined by relative exporter price indexes.

One potential explanation for the relative unimportance of average prices in explaining trade patterns arises in the neoclassical Heckscher-Ohlin model. In an international trade equilibrium characterized by factor price equalization, relative goods prices are the same across all countries, and patterns of trade across countries and sectors are entirely explained by relative factor endowments. However, we find substantial differences in average prices across countries within sectors. Conditional on observing these price differences, the demand-side of our model implies that patterns of trade must be explained by some combination of average prices, average demand, the dispersion of demand-adjusted prices, and variety. Of these four components of exporter price indexes, we find that average prices are the least important.

	Log Level RCA 2011		Log Change RCA 1998-11	
	Firm-Level Decomposition	Product-Level Decomposition	Firm-Level Decomposition	Product-Level Decomposition
Firm Price Index	0.094	-	-0.005	-
Firm Demand	0.220	0.220	0.422	0.422
Firm Variety	0.324	0.324	0.501	0.501
Firm Dispersion	0.362	0.362	0.082	0.082
Product Prices	-	0.065	-	-0.048
Product Variety	-	0.014	-	0.037
Product Dispersion	-	0.014	-	0.007

Note: Variance decomposition for the log level of RCA in 2011 and the log change in RCA from 1998-2011 (from equation (31)).

Table 3: Variance Decomposition U.S. RCA

By contrast, we find that average demand/quality is over three times more important than average prices, with a contribution of 22 percent for levels of RCA and 42 percent for changes in RCA in Table 3. This empirical finding for the relative importance of these two terms for patterns of trade is the reverse of the relative amount of attention devoted to them in existing theoretical research. In principle, one could reinterpret the predictions of neoclassical trade models as predictions for relative prices after adjusting for demand/quality. However, we find marked differences in these average demand/quality-adjusted prices across countries within sectors, implying substantial departures from goods price equalization even after making these adjustments. What sustains these differences in adjusted prices in the model is the imperfect substitutability of firms and products within sectors, which implies that both the number of varieties and the dispersion of demand-adjusted prices also matter for patterns of trade. Finally, it is not obvious that the determinants of quality/demand are the same as those of prices, with, for example, a large literature in industrial organization emphasizing the importance of sunk costs for quality (e.g. Sutton 1991).

By far the most important of the different mechanisms for trade in Table 3 is firm variety, which accounts for 32 and 50 percent of the level and change of RCA respectively. We also find a substantial contribution from the dispersion of demand-adjusted prices across firms, particularly in the cross-section, where this term accounts for 36 percent of the variation in RCA. In the time-series, this firm dispersion term is relatively less important, though it still accounts for 8 percent of the variation in RCA. On the one hand, our findings for firm variety are consistent with research that emphasizes the role of the extensive margin in understanding patterns of trade (e.g. Hummels and Klenow 2005, Chaney 2008). On the other hand, our findings for the dispersion of demand-adjusted prices across common varieties imply that the intensive margin is also important (consistent with the analysis for a log normal distribution in Fernandes et al. 2015). In particular, we find quantitatively relevant differences in the second moment of the distribution of demand-adjusted prices across common products and firms within exporters and sectors.

More broadly, this pattern of empirical results is consistent with theoretical frameworks in which comparative advantage operates not only through prices but also through the mass of firms and the distributions of productivity and demand across firms, such as Bernard, Redding, and Schott (2007). While recent empirical studies have documented substantial churning in patterns of comparative advantage over time, as in Freund and Pierola (2015) and Hanson, Lind and Muendler (2016), our findings imply that this churning largely occurs through changes in average demand/quality and firm entry/exit. The dominance of these two components of changes in average demand/quality and firm entry/exit points towards the relevance of theoretical frameworks in which comparative advantage arises from endogenous investments in product and process innovation, as in Grossman and Helpman (1991).

We find that our results for comparative advantage are robust across a number of different specifications. As a check on the sensitivity of our findings to the definition of categories, we replicated our entire analysis using a definition of sectors based of 6-digit instead of 4-digit NAICS codes. Using this different definition, we find a similar pattern of results as in our baseline specification, with average demand/quality accounting for 23 and 46 percent of the level and change of RCA, and firm variety making up 34 and 47 percent. As a further robustness check, we undertook a grid search over the range of plausible values for the elasticities of substitution across firms and products. In particular, we consider values of σ_g^F from 2 to 8 (in 0.5 increments) and values of σ_g^U from $(\sigma_g^F + 0.5)$ to 20 in 0.5 increments, while holding σ^G constant at our estimated value, which respects our estimated ranking that $\sigma_g^U > \sigma_g^F > \sigma^G$. As shown in Section A.5.3 of the web appendix, the contributions from firm variety and the dispersion of demand-adjusted prices across firms are invariant across these parameter values, because the elasticities of substitution cancel from these two terms. In contrast, the contributions from average prices and average demand/quality are increasing and decreasing in σ_g^F respectively. Nevertheless, across the grid of parameter values, we find that average prices account for less than 25 percent of the level of the RCA and less than 10 percent of the changes in RCA. Therefore, our finding that the relative price indexes that determine comparative advantage depart substantially from average prices is robust across the range of plausible elasticities of substitution.

We now show that the non-conventional forces of variety, average demand/quality, and the dispersion of demand-adjusted prices are also important for understanding aggregate U.S. imports from its largest suppliers. In Figure 4, we show the time-series decompositions of aggregate import shares from equation (25) for the top-five trade partners of the United States. We find that most of the increase in China's market share over the sample period occurs through increases in the number of firm varieties (orange), average firm demand/quality (dark gray) and the dispersion of demand-adjusted prices across firm varieties (light blue).[21] In contrast, average product prices (green) increased more rapidly for China than for the other countries in our sample, which worked to reduce China's market share. Therefore, the reasons for the explosive growth of Chinese exports were not cheaper Chinese exports, but rather substantial firm entry (variety), demand/quality upgrading, and improvements in the performance of leading firms relative to lagging firms (the dispersion of demand-adjusted prices). For Canada, we find that firm exit (orange) makes the largest contribution to the decline in its import share. For Germany, Japan and Mexico, we find substantial contributions from average firm demand/quality (gray) and the dispersion of demand-adjusted prices across firms (light blue), which are large relative to the contributions from average prices. Therefore, consistent with our results for sectoral

[21]Our finding of an important role for firm entry for China is consistent with the results for export prices in Amiti, Dai, Feenstra, and Romalis (2016). However, their price index is based on the Sato-Vartia formula, which abstracts from changes in demand/quality for surviving varieties, and they focus on Chinese export prices rather on than trade patterns.

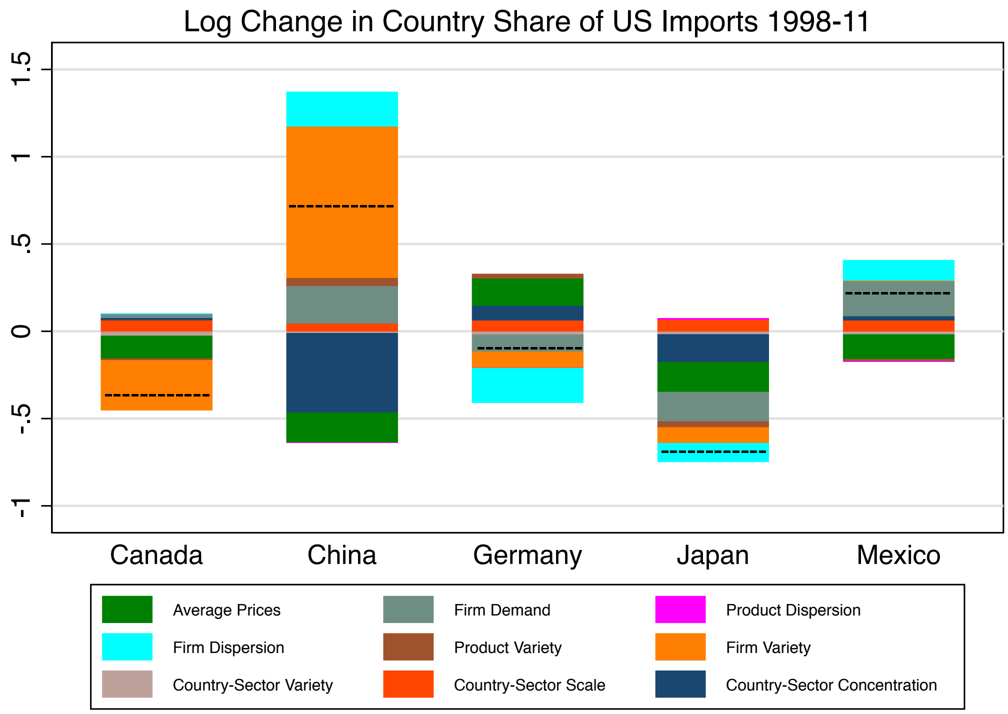

Figure 4: Country Aggregate Shares of U.S. Imports

patterns of trade above, we find that most of the change in aggregate import shares is explained by forces other than standard empirical measures of average prices.

Taken together, the results of this section highlight the role of imperfect substitutability across firms and products for comparative advantage and the aggregate volume of trade. Both are determined by relative price indexes that summarize the cost of sourcing goods from each country and sector. In a world in which goods are imperfect substitutes, these relative price indexes cannot be inferred solely from conventional measures of average prices. Instead, they also depend on the non-conventional forces of the number of varieties, demand/quality upgrading, and the performance of leading relative lagging varieties. Empirically, we find that these non-conventional forces are the dominant ones at work in the data.

5.4 Additional Theoretical Restrictions

We now compare our approach, which exactly rationalizes both micro and macro trade data, with special cases of this approach that impose additional theoretical restrictions. As a result of these additional theoretical restrictions, these special cases no longer exactly rationalize the micro trade data, and we quantify the implications of these departures from the micro data for macro trade patterns and prices.

Almost all existing theoretical research with CES demand in international trade is encompassed by the Sato-Vartia price index, which assumes no shifts in demand/quality for common varieties. Duality suggests that there are two ways to assess the importance of this assumption. First, we can work with a price index and examine how a CES price index that allows for demand shifts (i.e., the UPI in equation (16)) differs from a CES price index that does not allow for demand shifts (i.e., the Sato-Vartia index). Since the common goods component of the UPI (CG-UPI) and the Sato-Vartia indexes are identical in the absence of demand shifts, the difference between the two is a metric for how important demand shifts are empirically. Second, we can substitute each of these price indexes into equation (22) for revealed comparative advantage (RCA), and examine how important the assumption of no demand shifts is for understanding patterns of trade. Since the UPI perfectly rationalizes the data, any deviation from the data arising from using a different price index must reflect the effect of the restrictive assumption used in the index's derivation. In order to make the comparison fair, we need to also adjust the Sato-Vartia index for variety changes, which we do by using the Feenstra (1994) index, which is based on the same no-demand-shifts assumption for common goods, but adds the variety correction term given in equation (16) to incorporate entry and exit.

In Figure 5, we report the results of these comparisons. The top two panels consider exporter price indexes, while the bottom two panels examine RCA. In the top-left panel, we show a bin scatter of the Sato-Vartia exporter price index (on the vertical axis) against the common goods exporter price index (the CG-UPI on the horizontal axis), where the bins are twenty quantiles of each variable.[22] We also show the regression relationship between the two variables based on the disaggregated (i.e. not binned) data. If the assumption of time-invariant demand/quality were satisfied in the data, these two indexes would be perfectly correlated with one another and aligned on the 45-degree line. However, we find little relationship between them. The reason is immediately apparent if one recalls the top-right panel of Figure 3, which shows that price shifts are strongly positively correlated with demand shifts. The Sato-Vartia price index fails to take into account that higher prices are typically offset by higher demand/quality. In the top-right panel, we compare the Feenstra exporter price index (on the vertical axis) with our overall exporter price index (the UPI on the horizontal axis). These two price indexes have exactly the same variety correction term, but use different common goods price indexes (the CG-UPI and Sato-Vartia indexes respectively). The importance of the variety correction term as a share of the overall exporter price index accounts for the improvement in the fit of the relationship. However, the slope of the regression line is only around 0.5, and the regression R^2 is about 0.1. Therefore, the assumption of no shifts in demand/quality for existing goods results in substantial deviations between the true and measured costs of sourcing goods from an exporter and sector.

In the bottom left panel, we compare predicted changes in RCA based on relative exporter Sato-Vartia price indexes (on the vertical axis) against actual changes in RCA (on the horizontal axis). As the Sato-Vartia price index has only a weak correlation with the UPI, we find that it has little predictive power for changes in RCA, which are equal to relative changes in the UPI across exporters and sectors. Hence, observed changes in trade patterns are almost uncorrelated with the changes predicted under the assumption of no shifts in demand/quality and no entry/exit of firms and products. In the bottom right panel, we compare actual changes in RCA (on the horizontal axis) against predicted changes in RCA based on relative exporter Feenstra price indexes (on the vertical axis). The improvement in the fit of the relationship attests to the importance of adjusting for entry and exit. However, again the slope of the regression line is only around

[22] Again we use a bin scatter, because U.S. Census disclosure requirements preclude showing results for each exporter-sector using the U.S. data. In Section A.6.4 of the web appendix, we show results by exporter-sector using publicly-available Chilean data.

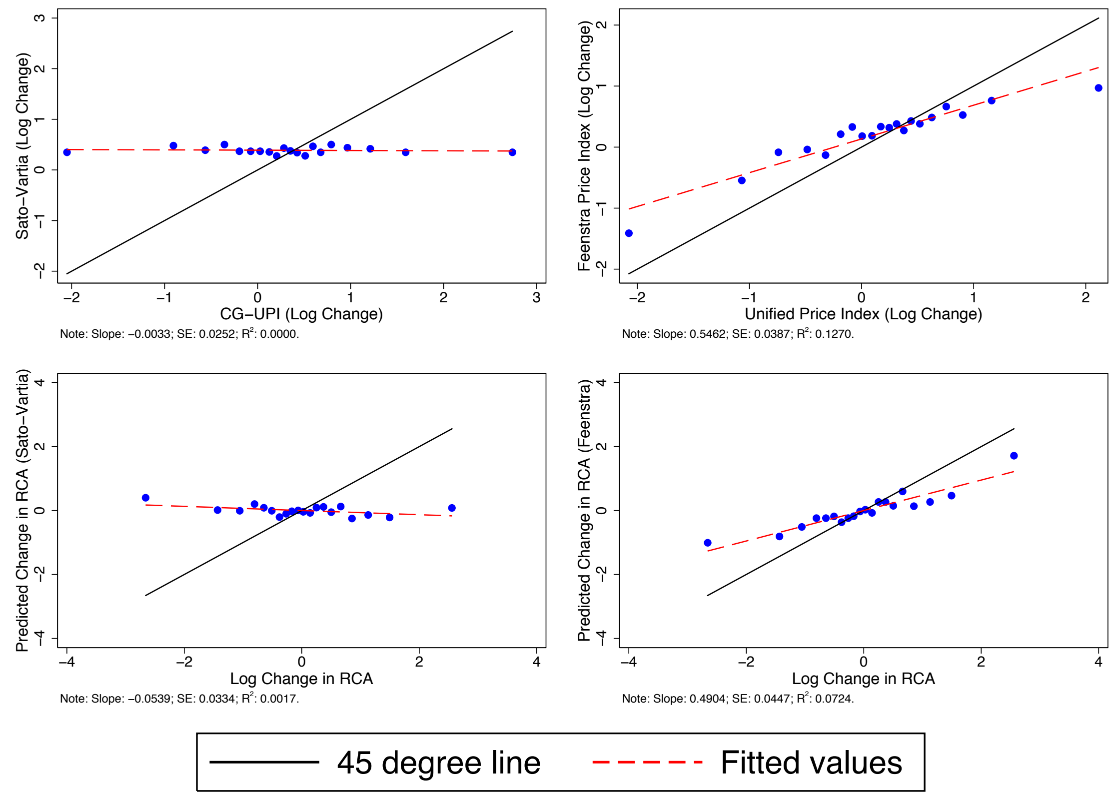

Figure 5: Sector-exporter Price Indexes with Time-Invariant Demand/Quality (Vertical Axis) Versus Time-Varying Demand/Quality (Horizontal Axis) for the U.S.

0.5 and the regression R^2 is less than 0.1. Therefore, even after adjusting for the shared entry and exit term, the assumption of no demand shifts for existing goods can generate predictions for changes in trade patterns that diverge substantially from those observed in the data.

Although the Sato-Vartia price index assumes no shifts in demand/quality for surviving varieties, it does not impose functional form restrictions on the cross-sectional distributions of prices, demand/quality and expenditure shares. We now examine the implications of imposing additional theoretical restrictions on these cross-sectional distributions. In particular, an important class of existing trade theories assumes not only a constant demand-side elasticity but also a constant supply-side elasticity, as reflected in the assumption of Fréchet or Pareto productivity distributions.[23] As our approach uses only demand-side assumptions, we can examine the extent to which these additional supply-side restrictions are satisfied in the data. In particular, we compare the observed data for firm sales and our model solutions for the firm price index and firm demand/quality ($\ln V^F_{ft} \in \left\{\ln \mathbb{X}^F_{ft}, \ln P^F_{ft}, \ln \varphi^F_{ft}\right\}$) with their theoretical predictions under alternative supply-side distributional assumptions.

[23] Ricardian trade models following Eaton and Kortum (2002) frequently assume a Fréchet productivity distribution, as in Costinot, Komunjer and Donaldson (2012). The firm heterogeneity literature following Melitz (2003) often assumes a Pareto productivity distribution, as in Chaney (2008) and Bernard, Redding and Schott (2011). Eaton, Kortum and Kramarz (2011) combines the assumption of a Pareto productivity distribution with stochastic shocks to demand and fixed costs that are log normally distributed. Fernandes et al. (2015) assumes a log normal distribution of productivity. Arkolakis, Costinot and Rodriguez-Clare (2012) provides macro restrictions on preferences, technology and market structure under which the import demand system exhibits a constant elasticity with respect to trade costs.

To derive these theoretical predictions, we use the QQ estimator of Kratz and Resnick (1996), as introduced into the international trade literature by Head, Mayer and Thoenig (2016). This QQ estimator compares the empirical quantiles in the data with the theoretical quantiles implied by alternative distributional assumptions. As shown in Section A.5.4 of the web appendix, under the assumption that V_{ft}^{F} has an untruncated Pareto distribution, we obtain the following theoretical prediction for the quantile of the logarithm of that variable:

$$\ln\left(V_{ft}^{F}\right) = \ln \underline{V}_{jigt}^{F} - \frac{1}{a_g^V} \ln\left[1 - \mathcal{F}_{jigt}\left(V_{ft}^{F}\right)\right], \tag{32}$$

where $\mathcal{F}_{jigt}\left(\cdot\right)$ is the cumulative distribution function; $\ln \underline{V}_{jigt}^{F}$ is the lower limit of the support of the untruncated Pareto distribution, which is a constant across firms f for a given importer j, exporter i, sector g and year t; a_g^V is the shape parameter of this distribution, which we allow to vary across sectors g.

We estimate equation (32) by OLS using the empirical quantile for $\ln\left(V_{ft}^{F}\right)$ on the left-hand side and the empirical estimate of the cumulative distribution function for $\mathcal{F}_{jigt}\left(V_{ft}^{F}\right)$ on the right-hand side, as discussed further in the web appendix. We estimate this regression for each sector across foreign firms (allowing the slope coefficient a_g^V to vary across sectors) and including fixed effects for each exporter-year-sector combination (allowing the intercept $\ln \underline{V}_{jigt}^{F}$ to vary across exporters, sectors and time). The fitted values from this regression correspond to the predicted theoretical quantiles, which we compare to the empirical quantiles observed in the data. Under the null hypothesis of a Pareto distribution, there should be a linear relationship between the theoretical and empirical quantiles that coincides with the 45-degree line.

To assess the empirical validity of this theoretical prediction, we estimate equation (32) for two separate subsamples: firms with values below the median for each exporter-sector-year cell and firms with values above the median for each exporter-sector-year cell. Under the null hypothesis of a Pareto distribution, the estimated slope coefficient $1/a_g^V$ should be the same for firms below and above the median.[24] As shown in Section A.5.4 of the web appendix, we strongly reject this null hypothesis of a Pareto distribution for all three variables, with substantial differences in the estimated coefficients below and above the median, which are statistically significant at conventional levels.[25]

To provide a point of comparison, we also consider the log-normal distributional assumption. As shown in Section A.5.4 of the web appendix, we obtain the following theoretical prediction for the quantile of the logarithm of a variable V_{ft}^{F} under this distributional assumption:

$$\ln\left(V_{ft}^{F}\right) = \kappa_{jigt}^{V} + \chi_g^V \Phi^{-1}\left(\mathcal{F}_{jigt}\left(V_{ft}^{F}\right)\right). \tag{33}$$

where $\Phi^{-1}\left(\cdot\right)$ is the inverse of the normal cumulative distribution function; κ_{jigt}^{V} and χ_g^V are the mean and standard deviation of the log variable, such that $\ln\left(V_{ft}^{F}\right) \sim \mathcal{N}\left(\kappa_{jigt}^{V}, \left(\chi_g^V\right)^2\right)$; we make analogous assumptions about these two parameters as for the untruncated Pareto distribution above; we allow the parameter controlling the mean (κ_{jigt}^{V}) to vary across exporters i, sectors g and time t for a given importer j; we allow the parameter controlling dispersion (χ_g^V) to vary across sectors g.

[24] U.S. Census disclosure requirements preclude showing the quantiles for individual foreign firms using the U.S. data. In Figures A.6.7 and A.6.8 in Section A.6.4 of the web appendix, we show firm quantiles using publicly-available Chilean data.

[25] A similar analysis can be undertaken for a Fréchet distribution. We find a similar pattern of statistically significant departures from the predicted linear relationship between the theoretical and empirical quantiles under this distributional assumption.

Again we estimate equation (33) by OLS using the empirical quantile for $\ln\left(V_{ft}^{F}\right)$ on the left-hand side and the empirical estimate of the cumulative distribution function for $\mathcal{F}_{jigt}\left(V_{ft}^{F}\right)$ on the right-hand side. We estimate this regression for each sector across foreign firms (allowing the slope coefficient χ_{g}^{V} to vary across sectors) and including fixed effects for each exporter-year-sector combination (allowing the intercept κ_{jigt}^{V} to vary across exporters, sectors and time). As shown in Section A.5.4 of the web appendix, we find that the log-normal distributional assumption provides a closer approximation to the data than the Pareto distributional assumption. Consistent with Bas, Mayer and Thoenig (2017), we find smaller departures from the predicted linear relationship between the theoretical and empirical quantiles for a log-normal distribution than for a Pareto distribution. Nevertheless, we reject the null hypothesis of a log-normal distribution at conventional significance levels for all three variables for the majority of industries, with substantial differences in estimated coefficients above and below the median. Instead of imposing such supply-side distributional assumptions, our demand-side approach uses the observed empirical distributions of prices and expenditure shares, and the resulting implied distribution of demand/quality under our assumption of CES demand.

As a concluding point, we examine the implications of these departures from a Pareto and log-normal distributions for understanding trade patterns across countries and sectors. Here, we demonstrate a surprising result. If one rationalizes the data using the unified price index, distributional assumptions about the underlying parameters do not matter, as long as these distributions are centered on the correct mean of the logs of each variable. To see this, we take the mean of the predicted values for log firm import shares, log firm-price indexes and log firm-demand/quality, and use our estimated elasticities of substitution to construct the predicted log common-goods unified price index for each exporter and sector:

$$\widehat{\ln \mathbb{P}_{jigt}^{E}} = \mathbb{E}_{jigt}^{F}\left[\widehat{\ln P_{ft}^{F}}\right] - \mathbb{E}_{jigt}^{F}\left[\widehat{\ln \varphi_{ft}^{F}}\right] + \frac{1}{\sigma_{g}^{F}-1}\mathbb{E}_{jigt}^{F}\left[\widehat{\ln \mathbb{S}_{ft}^{EF}}\right], \tag{34}$$

where a hat above a variable denotes a predicted value (recall that $\mathbb{S}_{ft}^{EF}$ is the share of each firm $f \in \Omega_{jigt}^{F}$ in county j's imports from a given exporter $i \neq j$ in sector g at time t).

A notable feature of this equation is that if we remove the hats, we obtain the exporter price index, which rationalizes revealed comparative advantage exactly. In this case, each of the terms on the right-hand side correspond the means of the logs of each variable. It follows immediately from this that that *any* distribution of the log of prices, demand/quality parameters, and shares that has the same means as in the data will produce the correct exporter price index and match RCA. Given our inclusion of exporter-sector-year fixed effects in equations (32) and (33), both of the estimated distributions are centered on the correct means of the logs of each variable for each exporter-sector-year. Therefore, these distributional assumptions do not matter for our conclusions about the sources of variation of RCA, as long as they implemented in such a way as to preserve the correct means of the logs of each variable for each exporter-sector-year.

6 Conclusions

One of the earliest stylized facts in international economics is that aggregate import prices and import volumes have correlations that are often the wrong sign or close to zero. Economists have long known the two main reasons for this finding: "shifts in the demand schedule" and "faulty methods of price index construction." We develop a method for accounting for movements in aggregate import price indexes and trade flows that allows for shifts in demand, prices, quality, variety and firm heterogeneity or dispersion. Our approach uses the invertibility of the CES demand system to recover unobserved values of demand/quality that exactly rationalize disaggregated trade data. We show how to aggregate to the national level even when data are incomplete or missing using a log-linear representation that permits exact decompositions of aggregate variables.

Unlike prior import price indexes, our indexes allow not only for changes in prices and variety, but also innovations in the demand for and quality of existing goods. We show that our average price term has a similar functional form and turns out to be highly correlated with the BLS import price index. But the other terms in our price index have very different functional forms and turn out to be negatively correlated with the BLS import price index. The importance of these other terms helps to explain the difficulty of trying to understand trade patterns using only conventional measures of import and domestic prices.

We derive a theoretically-rigorous measure of revealed comparative advantage (RCA) that depends on the relative values of our price indexes across countries and sectors. This RCA measure is also additively separable into price, demand/quality, entry/exit and dispersion terms in both the cross-section and time-series. We show that firm entry/exit and average demand/quality each account for around 45 percent of the time-series variation in imports, with the dispersion of demand-adjusted prices making up most of the rest. We show that these same factors account for most of the growth of China's share of aggregate U.S. imports, in spite of a relative increases in Chinese export prices.

Although there is an exact mapping between our price indexes and observed trade flows, the same need not be true for other approaches that impose stronger assumptions. We show that models that assume no demand shifts and no changes in variety perform poorly on international trade data. Models that account for variety changes while maintaining the assumption of no demand shifts do better, but still account for only about ten percent of the changes in comparative advantage over time. Finally, we provide new evidence on the empirical relevance of standard distributional assumptions (such as Pareto, Fréchet or log normal) for understanding patterns of international trade and import prices. Remarkably, we show that under CES demand it does not matter what the underlying productivity distribution is as long as it is correctly centered such that it matches the geometric average price, quality and sales of each firm.

References

[1] Adler, J. Hans (1945) "United States Import Demand During the Interwar Period," *The American Economic Review,* 35(3), 418-430.

[2] Adao, Rodrigo, Arnaud Costinot and Dave Donaldson (2017) "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, 107(3), 633-689.

[3] Amiti, Mary, Mi Dai, Robert Feenstra, and John Romalis (2016) "How Did China's WTO Entry Benefit U.S. Consumers?", UC Davis, mimeograph.

[4] Amiti, Mary and Donald R. Davis (2009) "What's Behind Volatile Import Prices from China?," *Current Issues in Economics and Finance*, 15(1), 1-7.

[5] Arkolakis, Costas, Arnaud Costinot and Andres Rodriguez-Clare (2012) "New Trade Models, Same Old Gains," *American Economic Review*, 102(1), 94-130.

[6] Armington, Paul S. (1969) "A Theory of Demand for Products Distinguished by Place of Production," *International Monetary Fund Staff Papers*, 16(1), 159–78.

[7] Atkeson, Andrew, and Ariel Burstein (2008) "Pricing to Market, Trade Costs and International Relative Prices," *American Economic Review*, 98, 1998-2031.

[8] Atkin, David, and Dave Donaldson (2015) "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper*, 21439.

[9] Atkin, David, Benjamin Faber and Marco Gonzalez-Navarro (2016) "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy*, forthcoming.

[10] Balassa, Bela (1965) "Trade Liberalization and Revealed Comparative Advantage," *Manchester School of Economic and Social Studies*, 33, 99-123.

[11] Bas, Maria, Thierry Mayer and Mathias Thoenig (2017) "From Micro to Macro: Demand, Supply and Heterogeneity in the Trade Elasticity," *Journal of International Economics*, 108, 1-19.

[12] Bernard, Andrew. B., and J. Bradford Jensen (1995) "Exporters, Jobs, and Wages in US Manufacturing: 1976-87," *Brookings Papers on Economic Activity: Microeconomics*, 67-112.

[13] Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott (2009) "Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. that Trade Goods," in *Producer Dynamics: New Evidence from Micro Data*, ed. by T. Dunne, J. B. Jensen, and M. J. Roberts. Chicago University Press, Chicago.

[14] Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2007) "Firms in International Trade," *Journal of Economic Perspectives*, 21(3), 105-130.

[15] Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2009) "The Margins of U.S. Trade," *American Economic Review*, Papers and Proceedings, 99(2), 487-493.

[16] Bernard, Andrew B., Stephen J. Redding and Peter K. Schott (2007) "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 74(1), 31-66.

[17] Bernard, Andrew B., Stephen J. Redding and Peter K. Schott (2010) "Multi-product Firms and Product Switching," *American Economic Review*, 100(1), 70– 97.

[18] Bernard, Andrew B., Stephen J. Redding and Peter K. Schott (2011) "Multiproduct Firms and Trade Liberalization," *Quarterly Journal of Economics*, 126(3), 1271-1318.

[19] Broda, Christian and David E. Weinstein (2006) "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121, 541-86.

[20] Broda, Christian and David E. Weinstein (2010) "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 100, 691-723.

[21] Caliendo, Lorenzo and Fernando Parro (2015) "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studie*s, 82 (1), 1-44.

[22] Chaney, Thomas (2008) "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707-1721.

[23] Costinot, Arnaud, Dave Donaldson and Ivana Komunjer (2012) "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas," *Review of Economic Studies*, 79, 581-608.

[24] Crozet, Matthieu, Keith Head, and Thierry Mayer (2012) "Quality Sorting and Trade: Firm-level Evidence for French Wine," *Review of Economic Studies,* 79(2): 609-644.

[25] De Loecker, Jan and Penny Goldberg (2014) "Firm Performance in a Global Market," *Annual Review of Economics*, 6, 201-227.

[26] Dhingra, Swati (2013) "Trading Away Wide Brands for Cheap Brands," *American Economic Review*, 103(6), 2554-84.

[27] Di Comite, F., Thisse, J., Vandenbussche, H. (2014). "Verti-zontal Differentiation in Export Markets", *Journal of International Economics*, 93, 50-66.

[28] Eaton, Jonathan and Samuel S. Kortum (2002) "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741-1779.

[29] Eaton, Jonathan, Samuel S. Kortum and Francis Kramarz (2004) "Dissecting Trade: Firms, Industries, and Export Destinations," *American Economic Review*, Papers and Proceedings, 94(2), 150–154.

[30] Eaton, Jonathan, Samuel S. Kortum and Francis Kramarz (2011) "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79(5), 1453–1498.

[31] Eckel, Carsten and J. Peter Neary (2010) "Multi-product Firms and Flexible Manufacturing in the Global Economy," *Review of Economic Studies*, 77, 188-217.

[32] Fally, Thibault and Benjamin Faber (2016) "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," University of California, Berkeley, mimeograph.

[33] Feenstra, Robert C. (1994) "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84, 157-177.

[34] Feenstra, Robert C. and Hong Ma "Optimal Choice of Product Scope for Multiproduct Firms," in *The Organization of Firms in a Global Economy*, in (eds) Elhanan Helpman, Dalia Marin and Thierry Verdier, chapter 6, 173-199. Cambridge MA: Harvard University Press.

[35] Feenstra, Robert C., Philip Luck, Maurice Obstfeld, and Katheryn N. Russ (2014) "In Search of the Armington Elasticity," *Review of Economics and Statistics*, forthcoming.

[36] Feenstra, Robert C. and John Romalis (2014) "International Prices and Endogenous Quality," *Quarterly Journal of Economics*, 129(2), 477-527.

[37] Fernandes, Ana, Peter Klenow, Martha Denisse Pierola, Andres Rodriguez-Clare and Sergii Meleshchuk (2015) "The Intensive Margin in Trade," Stanford University, mimeograph.

[38] Foster, Lucia, John Haltiwanger and Chad Syverson (2008) "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1), 394-425.

[39] Freund, Caroline and Martha Denisse Pierola (2015) "Export Superstars," *Review of Economics and Statistics*, 97(5), 1023-32.

[40] Grossman, Gene and Elhanan Helpman (1991) *Innovation and Growth in the Global Economy*, Cambridge: MIT Press.

[41] Hallak, Juan Carlos and Peter K. Schott (2011) "Estimating Cross-Country Differences in Product Quality," *Quarterly Journal of Economics*, 126, 417-474.

[42] Handbury, Jessie (2013) "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," the Wharton School, University of Pennsylvania.

[43] Hanson, Gordon H., Nelson Lind and Marc-Andreas Muendler (2015) "The Dynamics of Comparative Advantage," *NBER Working Paper*, 21753.

[44] Head, Keith, Thierry Mayer and Mathias Thoenig (2014) "Welfare and Trade Without Pareto," *American Economic Review*, Papers and Proceedings, 104(5), 310-316.

[45] Hottman, Colin, Stephen J. Redding and David E. Weinstein (2016) "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131(3), 1291-1364.

[46] Hsieh, Chang-Tai, Nicholas Li, Ralph Ossa, and Mu-Jeung Yang (2016) "Accounting for the Gains from Trade Liberalization," University of Chicago, mimeograph.

[47] Hummels, David and Peter J. Klenow (2005) "The Variety and Quality of a Nation's Exports," *American Economic Review*, 95(3), 704-723.

[48] Kamal, Fariha, C. J. Krizan and Ryan Monarch (2015) "Identifying Foreign Suppliers in U.S. Merchandise Import Transactions", *Center for Economic Studies Working Paper*, 15-11, Bureau of the Census.

[49] Kehoe, Timothy J. and Ruhl, Kim J. (2013) "How Important is the New Goods Margin in International Trade," *Journal of Political Economy,* 121(2) 2013.

[50] Khandelwal, Amit K. (2010) "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, 77, 1450-1476.

[51] Kratz, Marie, and Sidney I Resnick (1996) "The QQ-estimator and Heavy Tails," *Stochastic Models*, 12(4), 699-724.

[52] Krugman, Paul (1980) "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950-959.

[53] Levchenko, Andrei A. and Jing Zhang (2016) "The Evolution of Comparative Advantage: Measurement and Welfare Implications," *Journal of Monetary Economics*, 78, 96-111.

[54] Manova, Kalina and Zhiwei Zhang (2012) "Export Prices across Firms and Destinations," *Quarterly Journal of Economics*, 127, 379-436.

[55] Mayer, Thierry, Marc Melitz and Gianmarco I. P. Ottaviano (2014) "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104(2), 495-536.

[56] Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695-1725.

[57] Melitz, Marc J. and Stephen J. Redding (2014) "Heterogeneous Firms and Trade," in *Handbook of International Economic*s, ed. by E. Helpman, G. Gopinath, and K. Rogoff, vol. 4, chapter 1, 1-54, Elsevier, North Holland: Amsterdam.

[58] Nakamura, Alice, W. Erwin Diewert, John S. Greenlees, Leonard I. Nakamura, and Marshall B. Reinsdorf (2015) "Sourcing Substitution and Related Price Index Biases," in Susan N. Houseman and Michael Mandel (eds.), *Measuring Globalization: Better Trade Statistics for Better Policy*, Volume 1, Biases to Price, Output, and Productivity Statistics from Trade W.E. Upjohn Institute for Employment Research, 21-88.

[59] Redding, Stephen J. and David E. Weinstein (2016) "A Unified Approach to Estimating CES Demand and Welfare," *NBER Working Paper*, 22479.

[60] Roberts, Mark J., Xiaoyan Fan, Daniel Xu and Shengxing Zhang (2011) "A Structural Model of Demand, Cost, and Export Market Selection for Chinese Footwear Producers," *NBER Working Paper*, 17725.

[61] Schott, Peter K. (2004) "Across-product Versus Within-product Specialization in International Trade," *Quarterly Journal of Economics*, 119, 647-678.

[62] Sutton, John (1991) *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*, Cambridge MA: MIT Press.

[63] Sutton, John and Daniel Trefler (2016) "Capabilities, Wealth, and Trade," *Journal of Political Economy*, 124(3), 826-878.