Accounting for Micro and Macro Patterns of Trade[*]

Stephen J. Redding
Princeton University and NBER [†]

David E. Weinstein
Columbia University and NBER[‡]

May 24, 2017

Abstract

We develop a framework for quantifying the forces underlying changes in trade patterns and aggregate prices that is analogous to growth accounting. Our approach both exactly rationalizes micro data on trade by firm, product, source and destination and also permits exact aggregation to the macro level. We use this approach to separate the contributions of different explanations for trade patterns and the aggregate cost of living proposed in prior research, including relative prices, demand/quality, firm variety and firm heterogeneity. Our approach encompasses most existing macroeconomic models of trade, because it uses only the demand system and its parameters and does not impose functional form restrictions on the supply-side (such as Fréchet or Pareto productivity distributions). We find that relative prices make a comparatively small contribution to levels and changes in trade patterns. Instead, we find that demand parameters (including product quality), firm entry and exit, and firm heterogeneity account for most of the observed variation in patterns of trade in general and China's exports in particular.

JEL CLASSIFICATION: F11, F12, F14
KEYWORDS: comparative advantage, firm heterogeneity, variety, trade, aggregate prices

1 Introduction

Research in international trade has traditionally focused heavily on two questions: what explains trade patterns and what are the implications of trade for domestic prices? Because trade data is often aggregated or only available in detail in non-representative samples, researchers have typically pursued these questions using one of two approaches. On the one hand, general equilibrium "macro" models of trade explain aggregate trade patterns based on *assumptions* about disaggregated economic activity. On the other hand, "micro" studies of firm export and import behavior use disaggregated datasets on firm-level trade decisions for particular industries and countries. While each line of research has informed the other in important ways, a substantial gap between the approaches remains. The reduced-form nature of the micro studies means that there is no clear mapping between what happens at the firm level and what happens at the aggregate level. Similarly, the supply-side assumptions of the macro models can be at odds with what researchers observe in the micro data. To bridge this gap, we develop a trade accounting framework that both exactly rationalizes observed disaggregated trade data and permits exact aggregation to the national level. This framework is analogous to growth accounting, in that can be used to decompose aggregate trade and prices into the contributions of the different mechanisms emphasized in existing trade theories.

Our first main insight is that the demand system and its parameters are sufficient to provide a framework for quantifying the relative importance of different mechanisms for observed trade patterns. All we require to implement this framework is data on prices and expenditures. The only way in which supply-side assumptions matter is through the estimation of the parameters of the demand system. While this insight applies for a range of demand systems, we implement our analysis for constant elasticity of substitution (CES) preferences, as the most common specification used in international trade. Our procedure is sufficiently general as to apply for *any* nesting structure within this demand system. This approach stands in contrast with much research in international trade theory, which has largely focused on different *supply-side* explanations for patterns of international trade. 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, such as iceberg trade costs, perfect or monopolistic competition, and Pareto or Fréchet productivity distributions. 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 log-linear aggregation from micro to macro.

We derive our exact aggregation approach using three key properties of CES demand. First, we use the invertibility of the demand system, which implies that the unobserved demand shifters for each good can be uniquely recovered (up to a choice of units) from the observed data on prices and expenditure shares and the model's substitution parameters. Second, we exploit the separability properties of CES, which enable us to partition the overall unit expenditure function into an expenditure share for a subset of goods and the unit expenditure function for that subset of goods. We use this property to incorporate entry and exit (as in Feenstra 1994), to allow for non-tradable sectors, and to incorporate domestic varieties within tradable sectors. Third, we build on the unified price index (UPI) of Redding and Weinstein (2016), which expresses the

CES unit expenditure function as a log-linear function of prices, expenditure shares and demand shifters. We extend this unified price index to a nested demand structure and develop an estimator of the parameters of this nested demand structure. We show that the combination of the UPI and the nesting structure permits exact log-linear aggregation from the micro to the macro level. While nested CES is the only demand system that jointly satisfies all three properties, the point that the demand system alone can be used to quantify the role of different theoretical mechanisms is more general. For other demand systems, the resulting decompositions would not be additive in logs, or only would be additive in logs up to a first-order approximation, whereas this property holds globally for nested CES.

Our demand-side approach naturally accommodates a number of features of disaggregated trade data. For example, we can undertake our analysis using standard datasets, in which researchers observe disaggregated foreign transactions, but cannot observe individual price and quantity data for domestic transactions. Our framework allows for the entry and exit of firms, products and countries, which are pervasive in micro data, and we use the structure of our model to correctly evaluate the implications of this entry and exit for aggregate trade and prices. We incorporate demand shocks for individual varieties, which we show are required to explain the observed trade data as an equilibrium of the model, because of idiosyncratic shifts in expenditure conditional on prices. We also provide a natural explanation for the sparsity of trade: the small fraction of possible observations on products, firms, sectors and countries with positive flows. Zero trade flows arise naturally in our framework, either as a result of prohibitive reservation prices, negligible consumer demand (since demand enters inversely to prices), or fixed costs.

Our second main insight is that this general approach can be applied to the specific nesting structures that have dominated thinking in international trade. Trade models are typically built on theoretical *constructs* that do not necessarily exist in the data, such as the price of a unit of consumption of a good. By contrast, actual trade *data* are objects like export values and quantities in a particular product category by a particular firm. The precise way the various outputs of a given firm enter into utility, and the way in which the outputs of different firms are aggregated, is based on assumptions about demand. In the CES setup, these assumptions are manifest in the nesting structure. For example, traditional neoclassical models of comparative advantage assume the existence of industry nests, through which the goods supplied by each industry enter utility (as in the classical variants of the Heckscher-Ohlin and Ricardian models). Other versions of these neoclassical models postulate products as a tier of utility (as in Eaton and Kortum 2002), which can be nested within industries (as in Costinot, Donaldson and Komunjer 2012). All of these formulations assume that firms supply undifferentiated output of a given good under conditions of perfect competition. In contrast, much of new trade theory postulates a firm tier of utility (as in Krugman 1980 or Melitz 2003), which again can be nested within industries (as in Helpman and Krugman 1985 and Bernard, Redding and Schott 2007). More recently, theories of multi-product firms have added products as an additional tier of utility within firms and industries (as in Bernard, Redding and Schott 2010, 2001 and Hottman, Redding and Weinstein 2016, among others). Since one of our objectives is to evaluate existing trade theories, we adopt a nesting structure that connects as closely as possible to these theories, in which products can be grouped by firm, firms can be grouped by industry, and industries can be grouped into aggregate welfare. Importantly, while our setup allows for this

rich nesting structure, the framework also allows for simpler nesting structures in which industry or firm nests are absent (as in Krugman 1980 or Eaton and Kortum 2002).

Our third main contribution is to use this nesting structure from existing trade theories to provide evidence on the relative importance of different mechanisms for trade patterns and the aggregate cost of living. We follow the influential development, growth and business cycle accounting literatures (as in Chari, Kehoe and McGratten 2007), and use the structure of our model to isolate different mechanisms. We show that measuring the relative cost of sourcing goods across countries and sectors (which determines comparative advantage) requires assumptions about demand if goods are differentiated (in the same way that measuring productivity with differentiated goods requires assumptions about demand). We derive a theoretically-consistent empirical measure of comparative advantage in an environment with many countries, heterogeneous goods, and differentiated firms using only our assumption of nested CES demand. We show how this measure can be exactly decomposed into the contributions of differences in relative prices (as in Armington 1969); demand/quality (as in Linder 1961); firm variety (as in Krugman 1980); firm heterogeneity (as in Melitz 2003); and multi-product firms (as in Bernard Redding and Schott 2010, 2011).

We implement our framework using Chilean import data from 2007-2014 and obtain a number of novel insights about the forces underlying trade patterns and the extent to which these forces differ across countries and sectors. One of our most striking findings is that relative prices, which underlie Armington models of specialization, are of little importance in understanding changes in import shares: they account for around ten percent of the variation in trade patterns. However, changes in firm variety and heterogeneity are much more important. Half of all export growth in our sample can be attributed to increases in the number of firms, and close to ten percent of export growth can be attributed to changes in firm heterogeneity. This pattern of results is consistent with the mechanisms emphasized in the literatures on new trade theory and heterogeneous firms following Krugman (1980) and Melitz (2003). Finally, shifts in firm demand/quality account for around forty percent of trade movements. Interestingly, these demand shifts, which may reflect quality upgrading, are not due to firms switching their product mix across narrowly-defined product categories. Instead, the shifts in demand occur for firms' existing product mix, which is more supportive of models that focus on the upgrading or marketing of existing products rather than the development of radically new products.

Although none of these factors are the causal determinants of trade patterns, any successful model of trade and aggregate prices (based on CES demand) must be consistent with the mechanisms we identify, in the same way that any empirically successful growth model must map into a standard growth accounting exercise. In this sense, our approach generates a set of empirical moments for disciplining theory and empirics. Our framework exactly rationalizes the observed data, which implies that there is an exact mapping between our price index and observed patterns of trade. However, the same need not be true for other approaches that impose stronger assumptions, such as the Feenstra price index, which corresponds to a special case of our framework that assumes no demand shocks 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 revealed comparative advantage with counterfactual values of RCA based on different assumptions—e.g., with or without demand-shifts or variety corrections—we

4

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. We show that standard distributional assumptions about productivity or firm size can be rejected statistically, but in our framework these assumptions do not matter for accounting for comparative advantage, as long as researchers choose distributions to match the means of the logs of firm prices, demand/quality, and market shares.

Finally, our approach provides insights for elasticity puzzles, which are based on the insensitivity of real economic variables to measured changes in prices. All trade models postulate a mapping between trade flows and (correctly measured) relative prices, but we show that conventional price indexes are a poor proxy for the theoretically-motivated price index of the CES demand system. For example, while the average relative price of Chinese export products rose, the negative impact that this had on Chinese exports was more than offset by substantial increases in the quality or demand for Chinese products as well as by increases in the number and heterogeneity of Chinese exporters. These "non-conventional" forces belong in a theoretically motivated price index but do not appear in conventional price indexes that are averages of price changes for a constant set of products. Therefore, researchers may obtain misleading estimates of the response of real economic variables to price changes if these non-conventional forces are correlated with conventional measures of average prices.

Our paper is related to several strands of existing research. First, we contribute to the literature on firm heterogeneity in international trade following Melitz (2003), as reviewed in Bernard, Jensen, Redding and Schott (2007) and Melitz and Redding (2014).[1] One strand of this literature has used micro data on plants and firms to examine performance differences between exporters and non-exporters following the early empirical work by Bernard and Jensen (1995, 1999). Another line of research has provided evidence on the extensive margin of firm entry into export markets, including Hummels and Klenow (2005) and Eaton, Kortum, and Kramarz (2004, 2011).[2] Other research has examined multi-product firms and the extensive margin of the number of products supplied by firms, including Bernard, Redding, and Schott (2010, 2011) and Hottman, Redding, and Weinstein (2016).[3]

Second, our research connects with the recent literature on quantitative trade models following Eaton and Kortum (2002). These models are rich enough to account for first-order features of the data (such as an aggregate gravity equation) but parsimonious enough to permit transparent parameterization and counterfactual analysis. Using constant elasticity assumptions for both demand and supply (including for example Fréchet or Pareto productivity distributions), these models exactly rationalize observed aggregate trade flows. Under

---

[1]See also the reviews of Helpman (2006), Bernard, Jensen, Redding and Schott (2017), Antràs (2015) and Melitz and Trefler (2015).

[2]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 the micro data, 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 the micro data and aggregate to the macro level.

[3]Other research on multi-product firms and trade includes Eckel and Neary (2010), Feenstra and Ma (2008), Dhingra (2013), Mayer, Melitz, and Ottaviano (2013), Arkolakis, Muendler, and Ganapati (2014), and Nocke and Yeaple (2014).

these assumptions, they can be used to evaluate the welfare gains from trade and undertake counterfactuals for changes in trade costs (as in Arkolakis, Costinot, and Rodriguez-Clare 2012). More recently, 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 decompose observed trade patterns into the contributions of different mechanisms. As a result of imposing less structure on the supply-side, we are able to encompass non-neoclassical models with imperfect competition and increasing returns to scale (including Krugman 1980, Melitz 2003, and Atkeson and Burstein 2008).

Third, 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 observed disaggregated trade data and explaining changes in aggregate trade patterns and prices.

More broadly, our contribution relative to all of these strands of research is to develop a quantitative trade model that exactly rationalizes observed disaggregated trade data by firm, product, and destination (using only the demand system and its parameters) and permits exact aggregation to enable us to explore the mechanisms underlying changes in aggregate trade and prices. By contrast, prior work often matches aggregate moments (e.g., trade flows), but does not match the disaggregated trade transactions data, or matches data at a particular level of aggregation (e.g., firms), but is silent on how to match either more or less aggregated data.

The remainder of the paper is structured as follows. Section 2 introduces our theoretical framework. Section 3 presents the structural estimation of the model. Section 4 discusses our data. Section 5 reports our empirical results. Section 6 concludes. A web appendix contains technical derivations and the proofs of the propositions.

2   Theoretical Framework

One of the key insights from our approach is that specifying a functional form for demand and its parameters is sufficient to exactly decompose trade and aggregate prices into a number of different mechanisms. This insight is quite general and holds for a range of functional forms and nesting structures. To illustrate this insight and quantify the importance of these different mechanisms, we focus on CES preferences, because of their prominence in the international trade literature. Our choice of nesting structure is guided by existing models of international trade, which distinguish countries, sectors, firms and products. In our empirical

analysis, we also examine the robustness of our results to alternative choices of nesting structures. Since we want to abstract from discussions of how intermediate input usage of a product can differ from final consumption, 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 sources of expenditure, as in Krugman and Venables (1995), Eaton and Kortum (2002), and Caliendo and Parro (2015).

We begin by defining some notation. 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] In our baseline specification in the paper, we assume that the level at which firms make product decisions is the same as the level observed in our data (u). In Section C of the web appendix, we allow firms to make product choice decisions at a more disaggregated level than the units observed in our data (e.g. firms may produce different varieties of goods classified into the same observed category).

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 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. To simplify notation, we suppress the subscripts for countries and sectors when we refer to firm and product-level variables: that is, we use X^F_{ft} (rather than X^F_{jigft}) for firm expenditure and X^U_{ut} (rather than X^U_{jigfut}) for product expenditure.

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 horizontal differentiation within sectors occurs across firms and that there is a single elasticity of substitution for both domestic and foreign firms (σ_g^F).[5] 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, as in the canonical Krugman (1980) and Melitz (2003) models. Alternatively, neoclassical trade models specify undifferentiated output by firms supplying the same good, which if sectors are interpreted as goods, corresponds to a special case in which varieties are perfectly substitutable within sectors ($\sigma_g^U = \sigma_g^F = \infty$ and $\sigma^G < \infty$). Finally, if firm brands are irrelevant, so that products are equally differentiated within and across firms for a given sector as in Eaton and Kortum (2002), this corresponds to the special case in which the firm and product elasticities are equal to one another and distinct from the sector elasticity ($\sigma_g^U = \sigma_g^F > \sigma^G$).

Fourth, the demand shifters ($\varphi_{jgt}^G, \varphi_{ft}^F, \varphi_{ut}^U$) 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.[6] Finally, in order to simplify notation,

---

[5]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).

[6]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).

we suppress the subscript for importer j, exporter i, and sector g for the firm and product demand shifters (φ_{ft}^{F}, φ_{ut}^{U}). 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 Canadians, 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_{jt}^{T}) and a unit expenditure function for these tradable sectors ($\mathbb{P}_{jt}^{T}$):

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

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

$$\mu_{jt}^{T} \equiv \frac{\sum_{g \in \Omega^{T}} X_{jgt}^{G}}{\sum_{g \in \Omega^{G}} X_{jgt}^{G}} = \frac{\sum_{g \in \Omega^{T}} \left(P_{jgt}^{G} / \varphi_{jgt}^{G} \right)^{1-\sigma^{G}}}{\sum_{g \in \Omega^{G}} \left(P_{jgt}^{G} / \varphi_{jgt}^{G} \right)^{1-\sigma^{G}}}, \tag{5}$$

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

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

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_{jt}^{T}) 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 ($\mathbb{P}_{jt}^{T}$), a rise in the share of tradables (μ_{jt}^{T}) 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_{jgt}^G) and a unit expenditure function for these foreign varieties (a sectoral import price index $\mathbb{P}_{jgt}^G$):

$$P_{jgt}^G = \left(\mu_{jgt}^G \right)^{\frac{1}{\sigma_g^F - 1}} \mathbb{P}_{jgt}^G. \tag{7}$$

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

$$\mu_{jgt}^G \equiv \frac{\sum_{i \in \Omega_{jgt}^E} \sum_{f \in \Omega_{jigt}^F} X_{ft}^F}{X_{jgt}^G} = \frac{\sum_{i \in \Omega_{jgt}^E} \sum_{f \in \Omega_{jigt}^F} \left(P_{ft}^F / \varphi_{ft}^F \right)^{1 - \sigma_g^F}}{\sum_{i \in \Omega_{jgt}^I} \sum_{f \in \Omega_{jigt}^F} \left(P_{ft}^F / \varphi_{ft}^F \right)^{1 - \sigma_g^F}}, \tag{8}$$

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

$$\mathbb{P}_{jgt}^G \equiv \left[\sum_{i \in \Omega_{jgt}^E} \sum_{f \in \Omega_{jigt}^F} \left(P_{ft}^F / \varphi_{ft}^F \right)^{1 - \sigma_g^F} \right]^{\frac{1}{1 - \sigma_g^F}}. \tag{9}$$

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_{jgt}^G) is *increasing* in the sectoral *foreign* expenditure share (μ_{jgt}^G). The reason is that our expression for the sectoral price index (P_{jgt}^G) conditions on the price of foreign varieties, as is captured by the import price index ($\mathbb{P}_{jgt}^G$). For a given value of this import price index, a higher foreign expenditure share (μ_{jgt}^G) implies that domestic varieties are less attractive, which implies a higher sectoral price index.[7]

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 ($\mathbb{P}_{jgt}^G$) in equation (9) in terms of price indexes for each exporting country within that sector ($\mathbb{P}_{jigt}^E$). These exporter price indexes capture the contribution of each

---

[7]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.

foreign trade partner to the sectoral import price index:

$$\mathbb{P}_{jgt}^{G} = \left[\sum_{i \in \Omega_{jgt}^{E}} \left(\mathbb{P}_{jigt}^{E} \right)^{1-\sigma_{g}^{F}} \right]^{\frac{1}{1-\sigma_{g}^{F}}}, \tag{10}$$

where we use the superscript E to denote a variable for a *foreign* exporting country; the exporter price index ($\mathbb{P}_{jigt}^{E}$) is defined over the firm price indexes (P_{ft}^{F}) and demand/qualities (φ_{ft}^{F}) for each of the firms f that supply importing country j from that foreign exporter and sector:

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

If we substitute 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, we can also obtain the following expressions for the share of each product in expenditure on each firm (S_{ut}^{U}), the share of each foreign firm in all expenditure on foreign firms ($\mathbb{S}_{ft}^{F}$), and the share of each traded sector in all expenditure on traded sectors ($\mathbb{S}_{jgt}^{T}$):

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

$$\mathbb{S}_{ft}^{F} = \frac{\left(P_{ft}^{F} / \varphi_{ft}^{F} \right)^{1-\sigma_{g}^{F}}}{\sum_{i \in \Omega_{jgt}^{E}} \sum_{m \in \Omega_{jigt}^{F}} \left(P_{mt}^{F} / \varphi_{mt}^{F} \right)^{1-\sigma_{g}^{F}}}, \tag{13}$$

$$\mathbb{S}_{jgt}^{T} = \frac{\left(P_{jgt}^{G} / \varphi_{jgt}^{G} \right)^{1-\sigma^{G}}}{\sum_{k \in \Omega^{T}} \left(P_{jkt}^{G} / \varphi_{jkt}^{G} \right)^{1-\sigma^{G}}}. \tag{14}$$

We use "blackboard" font $\mathbb{S}_{ft}^{F}$ for the firm expenditure share to emphasize that this variable is defined as a share of expenditure on *foreign* firms (since $\Omega_{jgt}^{E} \equiv \left\{ \Omega_{jgt}^{I} : i \neq j \right\}$ in the denominator of equation (13)). Similarly, we use the blackboard font $\mathbb{S}_{jgt}^{T}$ and superscript T for the sector expenditure share to signal that this variable is defined across *tradable* sectors (since $\Omega^{T} \subseteq \Omega^{G}$ in the denominator of equation (14)).

We observe product expenditures (X_{ut}^{U}) and quantities (Q_{ut}^{U}) for each Harmonized System (HS) 8-digit category in our data, where quantities are measured in consistent units for a given HS 8-digit category (e.g., counts or tons). In our baseline specification, we assume that the level of disaggregation at which products are observed in the data (brands within each HS 8-digit category) corresponds to the level at which firms make product decisions, and hence we measure prices by unit values ($P_{ut}^{U} = X_{ut}^{U} / Q_{ut}^{U}$). From equation (12), 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. We show in subsection 2.8 below that we can use the structure of the model to back out the demand parameters (φ_{ut}^U) that provide the theory-consistent weights for each variety from the observed expenditures (X_{ut}^U) and prices (P_{ut}^U).[8]

2.6 Log-Linear CES Price Index

We now use these expressions for expenditure shares to show that CES price index can be written in an exact log linear form. We illustrate our approach for the firm tier of utility, but the analysis takes the same form for each of the other tiers of utility. We proceed by rearranging the expenditure share of products within firms (12) using the firm price index (3) to obtain:

$$P_{ft}^F = \frac{P_{ut}^U}{\varphi_{ut}^U} \left(S_{ut}^U \right)^{\frac{1}{\sigma_g^U - 1}},$$ (15)

which must hold for each product $u \in \Omega_{ft}^U$. Taking logarithms, averaging across products within firms, and exponentiating, we obtain the following expression for the firm-level price index:

$$P_{ft}^F = \left(\frac{\mathbb{M}_{ft}^U \left[P_{ut}^U \right]}{\mathbb{M}_{ft}^U \left[\varphi_{ut}^U \right]} \right) \left(\mathbb{M}_{ft}^U \left[S_{ut}^U \right] \right)^{\frac{1}{\sigma_g^U - 1}},$$ (16)

where $\mathbb{M}_{ft}^U [\cdot]$ is the geometric mean operator such that:

$$\mathbb{M}_{ft}^U \left[P_{ut}^U \right] \equiv \left(\prod_{u \in \Omega_{ft}^U} P_{ut}^U \right)^{\frac{1}{N_{ft}^U}},$$

where the superscript U indicates that this geometric mean is taken across products; and the subscripts f and t indicate that it varies across firms and over time.[9]

This representation for the firm price index in equation (16) has an intuitive interpretation. When products are perfect substitutes ($\sigma_g^U \to \infty$), the geometric mean of demand-adjusted product prices ($\mathbb{M}_{ft}^U \left[P_{ut}^U / \varphi_{ut}^U \right] = \mathbb{M}_{ft}^U \left[P_{ut}^U \right] / \mathbb{M}_{ft}^U \left[\varphi_{ut}^U \right]$) is a sufficient statistic for the firm price index. The reason is that perfect substitutability implies the equalization of demand-adjusted prices ($P_{ut}^U / \varphi_{ut}^U = P_{\ell t}^U / \varphi_{\ell t}^U$ for all $u, \ell \in \Omega_{ft}^U$ as $\sigma_g^U \to \infty$). Therefore, the geometric mean of demand-adjusted prices is equal to the demand-adjusted price for each product ($\mathbb{M}_{ft}^U \left[P_{ut}^U / \varphi_{ut}^U \right] = P_{\ell t}^U / \varphi_{\ell t}^U$ 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 the geometric mean of product expenditure shares ($\mathbb{M}_{ft}^U \left[S_{ut}^U \right]$). This term summarizes the overall effect on the

---

[8]In Section C of the web appendix, we show that our analysis generalizes to the case in which firms make product decisions at a more disaggregated level (e.g. unobserved barcodes) within each observed category (e.g., HS 8 products). In this case, the product demand shifter for each observed category (φ_{ut}^U) controls for both demand/quality and unobserved differences in composition within each observed category.

[9]This price index in equation (16) 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 $\left(\mathbb{M}_{ft}^U \left[S_{ut}^U \right] \right)^{1/\left(\sigma_g^U - 1 \right)}$.

cost of living of the love of variety and heterogeneity forces emphasized by Krugman (1980) and Melitz (2003). If we take the logarithm of equation (16) and add and subtract $\frac{1}{\sigma_g^U - 1} \ln N_{ft}^U$, we can be further decompose the price index as follows:

$$\ln P_{ft}^F = \underbrace{\mathbb{E}_{ft}^U \left[\ln P_{ut}^U\right]}_{\text{Prices}} - \underbrace{\mathbb{E}_{ft}^U \left[\ln \varphi_{ut}^U\right]}_{\text{Demand}} - \underbrace{\frac{1}{\sigma_g^U - 1} \ln N_{ft}^U}_{\text{Variety}} + \underbrace{\frac{1}{\sigma_g^U - 1} \left(\mathbb{E}_{ft}^U \left[\ln S_{ut}^U\right] - \ln \frac{1}{N_{ft}^U}\right)}_{\text{Heterogeneity}}, \quad (17)$$

where $\mathbb{E}_{ft}^U [\cdot]$ denotes the mean operator such that $\mathbb{E}_{ft}^U \left[\ln P_{ut}^U\right] = \frac{1}{N_{ft}^U} \sum_{u \in \Omega_{ft}^U} \ln P_{ut}^U$; 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. This formulation of the CES price index has the advantage that it is robust to measurement error in prices and/or expenditure shares that is mean zero in logs.

Each of the terms in equation (17) corresponds to an underlying economic mechanism that influences the firm price index. The first two terms on the right-hand side capture average prices and demand/quality across products sold by a firm, as discussed above. The third term captures 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 fourth term captures heterogeneity across varieties. When all varieties within firms have the same expenditure share ($S_{ut}^U = 1/N_{ft}^U$), the mean of log-expenditure shares is maximized, and this fourth term equals zero. Starting from this point and increasing the dispersion of expenditure shares, by raising some expenditure shares and decreasing others, the mean of log expenditure shares will fall because the log function is strictly concave. Therefore, this fourth term is negative when market shares are heterogeneous ($\mathbb{E}_{ft}^U \left[\ln S_{ut}^U\right] < \ln \left(1/N_{ft}^U\right)$), reducing the firm price index. The intuition for this result is that varieties are substitutes ($1 < \sigma_g^U < \infty$). Hence, holding constant average prices ($\mathbb{E}_{ft}^U \left[\ln P_{ut}^U\right]$) and average demand/quality ($\mathbb{E}_{ft}^U \left[\ln \varphi_{ut}^U\right]$), consumers prefer to source products from firms with more dispersed demand-adjusted prices across products. The reason is that they can substitute from less to more attractive products within firms.

Our exact aggregation approach combines this log-linear expression for the CES price index in equation (17) with the CES nesting structure. The log price index in each tier of utility equals the mean of the log prices in the lower tier of utility. Therefore, applying this log-linear representation recursively across each of our CES nests, we are able to write the log of aggregate prices in terms of means of the log prices of the disaggregated products observed in our data.

2.7 Entry, Exit and the Exact CES 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_{jigt}^F and Ω_{jigt-1}^F respectively) into the subsets of "common firms" that continue to supply this market in both periods ($\Omega_{jigt,t-1}^F$), firms that enter in period t (I_{jigt}^{F+}) and firms that exit after

13

period $t-1$ (I_{jigt-1}^{F-}). Similarly, we partition the set of products supplied by each of these firms in that sector into "common products" ($\Omega_{ft,t-1}^{U}$), entering products (I_{ft}^{U+}) and exiting products (I_{ft-1}^{U-}). 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 $i \neq j$ countries for importer j and sector g: "common" ($\Omega_{jgt,t-1}^{E}$), entering (I_{jgt}^{E+}) and exiting (I_{jgt-1}^{E-}). We denote the number of elements in these common sets of firms, products and foreign exporters by $N_{jigt,t-1}^{F}$, $N_{ft,t-1}^{U}$ and $N_{jgt,t-1}^{E}$ 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:

$$\lambda_{ft}^{U} \equiv \frac{\sum_{u \in \Omega_{ft,t-1}^{U}} \left(P_{ut}^{U}/\varphi_{ut}^{U}\right)^{1-\sigma_{g}^{U}}}{\sum_{u \in \Omega_{ft}^{U}} \left(P_{ut}^{U}/\varphi_{ut}^{U}\right)^{1-\sigma_{g}^{U}}}, \qquad \lambda_{ft-1}^{U} \equiv \frac{\sum_{u \in \Omega_{ft,t-1}^{U}} \left(P_{ut-1}^{U}/\varphi_{ut-1}^{U}\right)^{1-\sigma_{g}^{U}}}{\sum_{u \in \Omega_{ft-1}^{U}} \left(P_{ut-1}^{U}/\varphi_{ut-1}^{U}\right)^{1-\sigma_{g}^{U}}}, \tag{18}$$

where recall that $\Omega_{ft,t-1}^{U}$ is the set of common products such that $\Omega_{ft,t-1}^{U} \subseteq \Omega_{ft}^{U}$ and $\Omega_{ft,t-1}^{U} \subseteq \Omega_{ft-1}^{U}$. Note that λ_{ft}^{U} is equal to the total sales of continuing products in period t divided by the sales of all products available in time t evaluated at current prices. Its maximum value is one if no products enter in period t, and it falls as the share of new products rises. Similarly, λ_{ft-1}^{U} is equal to total sales of continuing products as share of total sales of all goods in the past period evaluated at $t-1$ prices. It equals one if no products cease being sold and falls as the share of exiting products rises.

Using these common expenditure shares, the change in the firm price index between periods $t-1$ and t (P_{ft}^{F}/P_{ft-1}^{F}) can be re-written as:

$$\frac{P_{ft}^{F}}{P_{ft-1}^{F}} = \left(\frac{\lambda_{ft}^{U}}{\lambda_{ft-1}^{U}}\right)^{\frac{1}{\sigma_{g}^{U}-1}} \left[\frac{\sum_{u \in \Omega_{ft,t-1}^{U}} \left(P_{ut}^{U}/\varphi_{ut}^{U}\right)^{1-\sigma_{g}^{U}}}{\sum_{u \in \Omega_{ft,t-1}^{U}} \left(P_{ut-1}^{U}/\varphi_{ut-1}^{U}\right)^{1-\sigma_{g}^{U}}}\right]^{\frac{1}{1-\sigma_{g}^{U}}} = \left(\frac{\lambda_{ft}^{U}}{\lambda_{ft-1}^{U}}\right)^{\frac{1}{\sigma_{g}^{U}-1}} \frac{P_{ft}^{F*}}{P_{ft-1}^{F*}}, \tag{19}$$

where the superscript asterisk indicates that a variable is defined for the common set of varieties.

The first term ($\left(\lambda_{ft}^{U}/\lambda_{ft-1}^{U}\right)^{\frac{1}{\sigma_{g}^{U}-1}}$) is the "variety-adjustment" term, which controls for the impact of the entry and exit of products on the firm price index. If new products have lower average prices relative to demand (lower $\left(P_{ut}^{U}/\varphi_{ut}^{U}\right)$) or have the same price-to-demand ratio but are more numerous than exiting products, then $\lambda_{ft}^{U}/\lambda_{ft-1}^{U} < 1$, and the firm price index (P_{ft}^{F}/P_{ft-1}^{F}) will fall due to the entering products being more desirable than the exiting products or an increase in variety. The second term ($P_{ft}^{F*}/P_{ft-1}^{F*}$) is the change in the firm price index for common products. Using the same notation of an asterisk for a variable that is defined over the set of common varieties, we can also define the share of expenditure on an individual common product in expenditure on all common products within the firm as:

$$S_{ut}^{U*} = \frac{\left(P_{ut}^{U}/\varphi_{ut}^{U}\right)^{1-\sigma_{g}^{U}}}{\sum_{\ell \in \Omega_{ft,t-1}^{U}} \left(P_{\ell t}^{U}/\varphi_{\ell t}^{U}\right)^{1-\sigma_{g}^{U}}} = \frac{\left(P_{ut}^{U}/\varphi_{ut}^{U}\right)^{1-\sigma_{g}^{U}}}{\left(P_{ft}^{F*}\right)^{1-\sigma_{g}^{U}}}. \tag{20}$$

Rearranging this common product expenditure share (20), and taking geometric means of both sides of the equation, the common goods firm price index (P_{ft}^{F*}) can be expressed in the log-linear form introduced in the

14

previous subsection. Using this log-linear representation in equation (19), we obtain the following expression for the overall price index:

$$\frac{P_{ft}^F}{P_{ft-1}^F} = \underbrace{\left(\frac{\lambda_{ft}^U}{\lambda_{ft-1}^U}\right)^{\frac{1}{\sigma_g^U-1}}}_{\text{Variety Correction}} \underbrace{\frac{\mathbb{M}_{ft}^{U*}\left[P_{ut}^U\right]}{\mathbb{M}_{ft}^{U*}\left[P_{ut-1}^U\right]} \left(\frac{\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^U\right]}{\mathbb{M}_{ft}^{U*}\left[\varphi_{ut-1}^U\right]}\right)^{-1} \left(\frac{\mathbb{M}_{ft}^{U*}\left[S_{ut}^{U*}\right]}{\mathbb{M}_{ft}^{U*}\left[S_{ut-1}^{U*}\right]}\right)^{\frac{1}{\sigma_g^U-1}}}_{\text{Common Goods Unified Price Index (CG-UPI)}}, \tag{21}$$

where $\mathbb{M}_{ft}^{U*}\left[P_{ut}^U\right] = \left(\prod_{u\in\Omega_{ft,t-1}^U} P_{ut}^U\right)^{1/N_{ft,t-1}^U}$ is the geometric mean across common products (superscript $U*$) within firm f between periods $t-1$ and t.

We refer to the exact CES price index in equation (21) 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.

Again this approach can be implemented for each tier of utility and applied recursively across these tiers of utility, which permits exact log-linear aggregation from micro to macro.

2.8 Model Inversion

We now show that there is a one-to-one mapping from the observed data on prices and expenditure shares $\{P_{ut}^U, S_{ut}^U, \mathbb{S}_{ft}^F, \mathbb{S}_{jgt}^T\}$ and the model's parameters $\{\sigma_g^U, \sigma_g^F, \sigma^G\}$ to the unobserved structural residuals $\{\varphi_{ut}^U, \varphi_{ft}^F, \varphi_{jgt}^G\}$. Therefore, the model can be inverted to recover these unobserved structural residuals from the observed data. 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:

$$\frac{\varphi_{ut}^U}{\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^U\right]} = \frac{P_{ut}^U}{\mathbb{M}_{ft}^{U*}\left[P_{ut}^U\right]} \left(\frac{S_{ut}^U}{\mathbb{M}_{ft}^{U*}\left[S_{ut}^U\right]}\right)^{\frac{1}{\sigma_g^U-1}}. \tag{22}$$

where $\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^U\right] \equiv \left(\prod_{u\in\Omega_{ft,t-1}^U} \varphi_{ut}^U\right)^{1/N_{t,t-1}^U}$. Similarly, dividing the share of a foreign firm in sectoral imports (13) by its geometric mean across common foreign firms within that sector, we obtain an analogous expression for firm demand:

$$\frac{\varphi_{ft}^F}{\mathbb{M}_{jgt}^{F*}\left[\varphi_{ft}^F\right]} = \frac{P_{ft}^F}{\mathbb{M}_{jgt}^{F*}\left[P_{ft}^F\right]} \left(\frac{\mathbb{S}_{ft}^F}{\mathbb{M}_{jgt}^{F*}\left[\mathbb{S}_{ft}^F\right]}\right)^{\frac{1}{\sigma_g^F-1}}. \tag{23}$$

where $\mathbb{M}_{jgt}^{F*}\left[\varphi_{ft}^F\right] \equiv \left(\prod_{i\in\Omega_{jgt,t-1}^E} \prod_{f\in\Omega_{jigt,t-1}^F} \varphi_{ft}^F\right)^{1/N_{jgt,t-1}^F}$. Finally, dividing the share of an individual tradable sector in all expenditure on tradable sectors by its geometric mean across these tradable sectors, we

obtain a similar expression for sector demand:

$$\frac{\varphi_{jgt}^G}{\mathbb{M}_{jt}^T\left[\varphi_{jgt}^G\right]} = \frac{P_{jgt}^G}{\mathbb{M}_{jt}^T\left[P_{jgt}^G\right]}\left(\frac{\mathbb{S}_{jgt}^T}{\mathbb{M}_{jt}^T\left[\mathbb{S}_{jgt}^T\right]}\right)^{\frac{1}{\sigma^G-1}},$$ (24)

where $\mathbb{M}_{jt}^T\left[\varphi_{jgt}^G\right] \equiv \left(\prod_{g\in\Omega^T}\varphi_{jgt}^G\right)^{1/N^T}$ and there is no asterisk in the superscript of the geometric mean operator across tradable sectors, because the set of tradable sectors is constant over time.

As expenditure shares are homogeneous of degree zero in the demand parameters, product demand (φ_{ut}^U), firm demand (φ_{ft}^F) and sector demand (φ_{jgt}^G) are only defined up to a choice of units in which to measure these parameters. We adopt the following convenient choice of units, such that the geometric mean of product demand across common products within each foreign firm is equal to one, 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:

$$\mathbb{M}_{ft}^{U*}\left[\varphi_{ut}^U\right] \equiv \left(\prod_{u\in\Omega_{ft,t-1}^U}\varphi_{ut}^U\right)^{\frac{1}{N_{t,t-1}^U}} = 1,$$ (25)

$$\mathbb{M}_{jgt}^{F*}\left[\varphi_{ft}^F\right] \equiv \left(\prod_{i\in\Omega_{jgt,t-1}^F}\prod_{f\in\Omega_{jigt,t-1}^F}\varphi_{ft}^F\right)^{\frac{1}{N_{jgt,t-1}^F}} = 1,$$ (26)

$$\mathbb{M}_{jt}^T\left[\varphi_{jgt}^G\right] \equiv \left(\prod_{g\in\Omega^T}\varphi_{jgt}^G\right)^{\frac{1}{N^T}} = 1.$$ (27)

Although for convenience we set each of these geometric means to one, our decompositions of changes over time are robust to any constant choice of units in which to measure each demand shifter. Furthermore, our decompositions of patterns of trade below are based on relative comparisons across firms in different exporters within a sector, which implies that any common choice of units across firms within each sector differences out. Under these normalizations in equation (25)-(27), 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.[10] Given the elasticities of substitution $\{\sigma_g^U, \sigma_g^F, \sigma^G\}$, 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 prices and expenditure shares $\{P_{ut}^U, S_{ut}^U, S_{ft}^F, S_{jgt}^T\}$ and hence do not need to take a stand on the different supply-side forces that determine prices (e.g. technology, factor prices, oligopoly, monopolistic competition or perfect competition). Therefore, the only way in which the supply-side can potentially enter our analysis is through the estimation of the elasticities of substitution, as discussed further in Section 3 below.

---

[10]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.

An important difference between our approach and standard exact price indexes for CES is that we allow the demand/quality parameters in equations (22)-(24) 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 captured in the shares of expenditure on foreign varieties within sectors (μ_{jgt}^G) in equation (7) for the sectoral price index (P_{jgt}^G). Similarly, we allow for proportional changes in the demand/quality for all tradable sectors relative to all non-tradable sectors, which are captured in the share of expenditure on tradable sectors (μ_{jt}^T) 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 Accounting for Exporter Price Movements

Having determined the structural residuals that rationalize the observed data on prices and expenditure shares $\{\varphi_{ut}^U, \varphi_{ft}^F, \varphi_{jgt}^G\}$, we have all the components needed for our decompositions. We start with the exporter price index, which summarizes the cost of sourcing imports from a foreign exporter and sector. We begin by using CES demand to express the share of an individual firm f in country j's imports from an exporting country $i \neq j$ within a sector g in terms of the exporter price index ($\mathbb{P}_{jigt}^E$):

$$\mathsf{S}_{ft}^{EF} = \frac{\left(P_{ft}^F / \varphi_{ft}^F\right)^{1-\sigma_g^F}}{\sum_{k \in \Omega_{jigt}^F} \left(P_{kt}^F / \varphi_{kt}^F\right)^{1-\sigma_g^F}} = \frac{\left(P_{ft}^F / \varphi_{ft}^F\right)^{1-\sigma_g^F}}{\left(\mathbb{P}_{jigt}^E\right)^{1-\sigma_g^F}}, \qquad i \neq j, \tag{28}$$

where the superscript EF is a mnemonic for exporter and firm, and indicates that this firm expenditure share is computed as a share of imports from a single foreign exporter.

Using the fact that the denominator in equation (28) is equal to $\left(\mathbb{P}_{jigt}^E\right)^{1-\sigma_g^F}$, and following the approach introduced in Section 2.6, we obtain the following log-linear expression for the exporter price index within a given sector in terms of the geometric means of prices, demand/quality and expenditure shares across firms from that country and sector:

$$\mathbb{P}_{jigt}^E = \frac{\mathbb{M}_{jigt}^F \left[P_{ft}^F\right]}{\mathbb{M}_{jigt}^F \left[\varphi_{ft}^F\right]} \left(\mathbb{M}_{jigt}^F \left[\mathsf{S}_{ft}^{EF}\right]\right)^{\frac{1}{\sigma_g^F - 1}}, \tag{29}$$

where $\mathbb{M}_{jigt}^F [\cdot]$ is the geometric mean across firms supplying importer j from exporter i within sector g at time t, as defined in Section B of the web appendix. Combining this expression for the exporter price index (29) with our earlier expression for the firm price index (16), we obtain:

$$\mathbb{P}_{jigt}^E = \left(\frac{\mathbb{M}_{jigt}^{FU} \left[P_{ut}^U\right]}{\mathbb{M}_{jigt}^F \left[\varphi_{ft}^F\right] \mathbb{M}_{jigt}^{FU} \left[\varphi_{ut}^U\right]}\right) \left(\mathbb{M}_{jigt}^{FU} \left[\mathsf{S}_{ut}^U\right]\right)^{\frac{1}{\sigma_g^U - 1}} \left(\mathbb{M}_{jigt}^F \left[\mathsf{S}_{ft}^{EF}\right]\right)^{\frac{1}{\sigma_g^F - 1}}, \tag{30}$$

where $\mathbb{M}_{jigt}^{FU}[\cdot]$ is the geometric mean across products within each firm (superscript U) and across firms (superscript F) for a given importer (subscript j), exporter (subscript i), sector (subscript g) and time (subscript t), as defined in Section B of the web appendix.

Taking logarithms and re-arranging terms, we obtain the following log-linear decomposition for the cost of importer j sourcing goods in sector g from an exporter i at time t:

$$\ln \mathbb{P}_{jigt}^E = \underbrace{\mathbb{E}_{jigt}^{FU}\left[\ln P_{ut}^U\right]}_{\text{Prices}} - \underbrace{\left\{\mathbb{E}_{jigt}^F\left[\ln \varphi_{ft}^F\right] + \mathbb{E}_{jigt}^{FU}\left[\ln \varphi_{ut}^U\right]\right\}}_{\text{Demand}} - \underbrace{\left\{\frac{1}{\sigma_g^U - 1}\mathbb{E}_{jigt}^F\left[\ln N_{ft}^U\right] + \frac{1}{\sigma_g^F - 1}\ln N_{jigt}^F\right\}}_{\text{Variety}} \quad (31)$$

$$+ \underbrace{\left\{\frac{1}{\sigma_g^U - 1}\mathbb{E}_{jigt}^{FU}\left[\ln S_{ut}^U - \ln \frac{1}{N_{ft}^U}\right] + \frac{1}{\sigma_g^F - 1}\mathbb{E}_{jigt}^F\left[\ln S_{ft}^{EF} - \ln \frac{1}{N_{jigt}^F}\right]\right\}}_{\text{Heterogeneity}},$$

where $\mathbb{E}_{jigt}^{FU}[\cdot]$ is a mean across firms and products and $\mathbb{E}_{jigt}^F[\cdot]$ is a mean across firms, as also defined in Section B of the web appendix; product prices (P_{ut}^U), numbers of products and firms (N_{ft}^U, N_{jigt}^F) and expenditure shares (S_{ut}^U, S_{ft}^{EF}) are directly observed; and product and firm demands (φ_{ft}^F, φ_{ut}^U) can be recovered from the observed data using the substitution parameters (σ_g^U, σ_g^F), as shown in Section 2.8 above.

Similarly, partitioning varieties into those that are common, entering and exiting, and taking differences over time, the log change in the exact CES price index for an importer j sourcing goods in sector g from an exporter i between periods $t-1$ and t can be decomposed exactly into the following four terms:

$$\Delta \ln \mathbb{P}_{jigt}^E = \underbrace{\mathbb{E}_{jigt}^{FU*}\left[\Delta \ln P_{ut}^U\right]}_{\text{Prices}} - \underbrace{\left\{\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \varphi_{ft}^F\right] + \mathbb{E}_{jigt}^{FU*}\left[\Delta \ln \varphi_{ut}^U\right]\right\}}_{\text{Demand}} + \underbrace{\left\{\frac{1}{\sigma_g^U - 1}\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \lambda_{ft}^U\right] + \frac{1}{\sigma_g^F - 1}\Delta \ln \lambda_{jigt}^F\right\}}_{\text{Variety}} \quad (32)$$

$$+ \underbrace{\left\{\frac{1}{\sigma_g^U - 1}\mathbb{E}_{jigt}^{FU*}\left[\Delta \ln S_{ut}^{U*}\right] + \frac{1}{\sigma_g^F - 1}\mathbb{E}_{jigt}^{F*}\left[\Delta \ln S_{ft}^{EF}\right]\right\}}_{\text{Heterogeneity}},$$

as shown in Section E of the web appendix.

Each of the terms in equation (32) has an intuitive interpretation. The first term is the average log change in the price of common products sourced from exporting country i within sector g, which equals the log of a Jevons Index index of import prices. The Jevons index is used to aggregate prices in the U.S. consumer price index and captures the price effects emphasized in Armington models. The second term ($\mathbb{E}_{jigt}^{F*}\left[\ln \varphi_{ft}^F\right] + \mathbb{E}_{jigt}^{FU*}\left[\ln \varphi_{ut}^U\right]$) is average log change in demand/quality of common products and firms sourced from country i within sector g. This term captures demand shifts or quality upgrading for common products and its presence reflects the fact that consumers care about demand-adjusted prices rather than prices alone. Recall that our normalization in equation (25) implies that the average log change in common-product demand within foreign firms is equal to zero: $\mathbb{E}_{jigt}^{FU*}\left[\Delta \ln \varphi_{ut}^U\right] = 0$. Similarly, our normalization in equation (26) implies that the average log change in firm demand across *all* common foreign firms within a sector is equal to zero: $\mathbb{E}_{jgt}^{F*}\left[\Delta \ln \varphi_{ft}^F\right] = 0$. 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: $\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \varphi_{ft}^F\right] \neq \mathbb{E}_{jgt}^{F*}\left[\Delta \ln \varphi_{ft}^F\right] = 0$ for country $i \neq j$. There-

fore, 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 ($\frac{1}{\sigma_g^U-1}\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \lambda_{ft}^U\right] + \frac{1}{\sigma_g^F-1}\Delta \ln \lambda_{jigt}^F$) captures the effect of product turnover and firm entry and exit on the cost of sourcing imports. If more products and firms enter from exporting country i within sector g than exit, this increases consumer utility and reduces the cost of sourcing goods from that country and sector, as reflected in a fall in the share of expenditure on common products and firms ($\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \lambda_{ft}^U\right] < 0$ and $\Delta \ln \lambda_{jigt}^F < 0$). Similarly, if the entering firms and products from country i within sector g are more attractive (have lower demand-adjusted prices) than those that exit, this again increases consumer utility and reduces the cost of sourcing goods from that country and sector, as reflected in a fall in the share of expenditure on common products and firms ($\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \lambda_{ft}^U\right] < 0$ and $\Delta \ln \lambda_{jigt}^F < 0$).

The fourth term captures the effect of heterogeneity across common products and firms on the cost of sourcing imports. Other things equal, if the dispersion of expenditure shares across common products and firms increases, this again raises consumer utility and reduces the cost of sourcing goods from that country and sector, as captured by a fall in the geometric mean of common expenditure shares across products and firms ($\mathbb{E}_{jigt}^{FU*}\left[\Delta \ln S_{ut}^{U*}\right] < 0$ and $\mathbb{E}_{jigt}^{F*}\left[\Delta \ln \mathbb{S}_{ft}^{EF}\right] < 0$). The reason is that this increased dispersion of expenditure shares reflects greater heterogeneity in demand-adjusted prices across varieties, which enhances the ability of consumers to substitute from less to more desirable varieties.

2.10   Patterns of Trade Across Sectors and Countries

Thus far, we have been focused on developing a consistent method to measure and decompose price indexes that determine the costs of sourcing imports 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. Moreover, we do not need to rely on a supply-side model, because such a model might explain why we observe the price and demand/quality parameters we do, but it is not necessary if we can observe or solve for these parameters. The key insight in this section is that in the CES-setup is possible to decompose a theoretically rigorous measure of revealed comparative advantage that is log-linear in exporter price indexes, which permits a simple translation of our exporter price index decomposition into a trade decomposition based on the same price, demand, variety and heterogeneity factors emphasized in existing trade theories.

2.10.1   Accounting for Revealed Comparative Advantage

Our main contribution in this section is to show that there exists an empirical measure of comparative advantage that holds in all models based on a CES demand system. Interestingly, this measure of revealed comparative advantage (RCA) takes a similar form to the original Balassa (1965) concept, except that it uses geometric averages of shares instead of shares of totals. We begin with the share of an individual foreign exporting country i in all foreign imports within a given sector g for importing country j at time t:

$$\mathbb{S}_{jigt}^E = \frac{\sum_{f \in \Omega_{jigt}^F}\left(P_{ft}^F/\varphi_{ft}^F\right)^{1-\sigma_g^F}}{\sum_{h \in \Omega_{jgt}^E}\sum_{f \in \Omega_{jhgt}^F}\left(P_{ft}^F/\varphi_{ft}^F\right)^{1-\sigma_g^F}} = \frac{\left(\mathbb{P}_{jigt}^E\right)^{1-\sigma_g^F}}{\left(\mathbb{P}_{jgt}^G\right)^{1-\sigma_g^F}}, \qquad i \neq j. \tag{33}$$

where the single superscript E is a mnemonic for exporter and indicates that this is the expenditure share for a foreign exporter; the numerator in equation (33) 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 (33) features importer j's overall import price index in sector g at time t ($\mathbb{P}^G_{jgt}$).

Using this exporter expenditure share (33), we measure Revealed Comparative Advantage (RCA) in sector g for import market j, by first taking 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 sectors ($\mathbb{M}^G_{jit}\left[\mathbb{X}^E_{jigt}/\mathbb{M}^E_{jgt}\left[\mathbb{X}^E_{jigt}\right]\right]$):

$$RCA_{jigt} \equiv \frac{\mathbb{X}^E_{jigt}/\mathbb{M}^E_{jgt}\left[\mathbb{X}^E_{jigt}\right]}{\mathbb{M}^G_{jit}\left[\mathbb{X}^E_{jigt}/\mathbb{M}^E_{jgt}\left[\mathbb{X}^E_{jigt}\right]\right]} = \frac{\mathbb{S}^E_{jigt}/\mathbb{M}^E_{jgt}\left[\mathbb{S}^E_{jigt}\right]}{\mathbb{M}^G_{jit}\left[\mathbb{S}^E_{jigt}/\mathbb{M}^E_{jgt}\left[\mathbb{S}^E_{jigt}\right]\right]}, \tag{34}$$

where we use the "blackboard" font $\mathbb{X}$ to denote bilateral trade with a foreign country; $\mathbb{X}^E_{jigt}$ is bilateral trade from exporter 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; $\mathbb{M}^G_{jit}\left[\mathbb{X}^E_{jigt}\right] = \left(\prod_{k\in\Omega^G_{jit}}\mathbb{X}^E_{jikt}\right)^{1/N^G_{jit}}$ is the geometric mean of these exports across 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 (34), 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 (2015). 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).[11] 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 decompose trade patterns into the key factors emphasized in much of trade theory. From equations (33) and (34), RCA captures the relative cost to an importer of sourcing goods across countries and sectors, as determined by relative price indexes and the elasticity of substitution ($\left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}$):

$$RCA_{jigt} = \frac{\left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}/\mathbb{M}^E_{jgt}\left[\left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}\right]}{\mathbb{M}^G_{jit}\left[\left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}/\mathbb{M}^E_{jgt}\left[\left(\mathbb{P}^E_{jigt}\right)^{1-\sigma^F_g}\right]\right]}. \tag{35}$$

A first insight from this relationship is that relative price indexes across countries and sectors determine

[11]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.

comparative advantage and trade. A second insight is that the demand-side of the model is central to computing the correct measures of relative price indexes. In the special case of Armington models, in which firms supply homogeneous products within sectors ($\sigma_g^U = \sigma_g^F = \infty$, $\varphi_{kt}^U = \varphi_{\ell t}^U$ and $\varphi_{ft}^F = \varphi_{mt}^F$ for all k, ℓ and f, m), relative price indexes can be directly measured using the price charged by any firm in each sector and exporter (since without differentiation the prices of all goods consumed within a given sector and exporter must be the same). Outside of this extreme special case, relative price indexes cannot be measured without taking a stand on the demand side. Furthermore, when the products supplied by firms are imperfect substitutes ($\sigma_g^U < \infty$, $\sigma_g^F < \infty$, $\varphi_{kt}^U \neq \varphi_{\ell t}^U$ and $\varphi_{ft}^F \neq \varphi_{mt}^F$ for some k, ℓ and f, m), product prices (P_{ut}^U) are only one of several determinants of relative price indexes. Demand/quality (φ_{ut}^U and φ_{ft}^F), the number of products and firms (N_{ft}^U and N_{jigt}^F) and the heterogeneity of demand-adjusted prices across these products and firms (as captured by the dispersion of expenditure shares S_{ut}^U and S_{ft}^{EF}) also influence relative price indexes and hence patterns of trade. Moreover, the relative contributions from each of these mechanisms are influenced by the elasticities of substitution across products (σ_g^U) and firms (σ_g^F). Therefore, just as productivity cannot be measured separately from demand when goods are imperfect substitutes, comparative advantage also cannot be measured independently of demand in such a differentiated goods environment.[12] Furthermore, we now show that our demand-side assumptions are sufficient to separate out a number of different mechanisms through which patterns of trade are determined in existing trade theories.

In particular, using equation (30) to substitute for the exporter price index ($\mathbb{P}_{jigt}^E$) in equation (35), we can decompose differences in RCA across countries and sectors into the contributions of average prices ($\ln\left(RCA_{jigt}^P\right)$), average demand ($\ln\left(RCA_{jigt}^\varphi\right)$), variety ($\ln\left(RCA_{jigt}^N\right)$) and heterogeneity ($\ln\left(RCA_{jigt}^S\right)$):

$$\ln\left(RCA_{jigt}\right) = \underbrace{\ln\left(RCA_{jigt}^P\right)}_{\text{Prices}} + \underbrace{\ln\left(RCA_{jigt}^\varphi\right)}_{\text{Demand}} + \underbrace{\ln\left(RCA_{jigt}^N\right)}_{\text{Variety}} + \underbrace{\ln\left(RCA_{jigt}^S\right)}_{\text{Heterogeneity}}, \qquad (36)$$

where each of these terms is defined in full in Section F 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 price term ($\ln\left(RCA_{jigt}^P\right)$), 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. This price term captures the conventional role of relative prices in determining trade patterns from Armington trade models. 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.

The second term ($\ln\left(RCA_{jigt}^\varphi\right)$) captures the role of demand/quality in shaping patterns of trade, as emphasized by the literature following Linder (1961).[13] Other things equal, an exporter has a RCA in a sector

---

[12]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).

[13]See, in particular, Schott (2004), Fajgelbaum, Grossman and Helpman (2011), Hallak and Schott (2011), Feenstra and Romalis

if the demand/quality for its goods relative to the average exporter in that sector is high compared to the exporter's average sector. The third term ($\ln \left(RCA_{jigt}^{N} \right)$) captures firm variety, as emphasized in the new trade literature following Krugman (1980) and Helpman and Krugman (1985). Other things equal, an exporter has a RCA in a sector if its number of varieties relative to the average exporter in that sector is large compared to the exporter's average sector. The reason is that the products supplied by firms are imperfect substitutes and hence the value of bilateral trade is increasing in the number of varieties. Finally, the fourth term ($\ln \left(RCA_{jigt}^{S} \right)$) summarizes the role of heterogeneity across varieties, as emphasized in the heterogeneous firm literature following Melitz (2003). Other things equal, an exporter has a RCA in a sector if the heterogeneity across its varieties relative to the average exporter in that sector is large compared to the exporter's average sector. The explanation is again that the products supplied by firms are imperfect substitutes, and the greater the heterogeneity across varieties, the greater the ability of the consumer to substitute from less to more appealing varieties.

Our framework also permits an exact decomposition of changes over time in patterns of RCA across countries and sectors. We are therefore able to shed light on the different theoretical mechanisms underlying the turbulence in trade patterns over time reported in Proudman and Redding (2000), Freund and Pierola (2015), and Hanson, Lind and Muendler (2016). Taking differences over time in equation (35), we obtain the following exact decomposition of changes in RCA over time:

$$\Delta \ln \left(RCA_{jigt}^{*} \right) = \underbrace{\Delta \ln \left(RCA_{jigt}^{P*} \right)}_{\text{Prices}} + \underbrace{\Delta \ln \left(RCA_{jigt}^{\varphi*} \right)}_{\text{Demand}} + \underbrace{\Delta \ln \left(RCA_{jigt}^{\lambda} \right)}_{\text{Variety}} + \underbrace{\Delta \ln \left(RCA_{jigt}^{S*} \right)}_{\text{Heterogeneity}}, \quad (37)$$

where all four terms are again defined in full in subsection F of the web appendix. We compute these log changes for all common exporter-sector pairs with non-zero values of RCA in both periods, as indicated by the asterisks in the superscripts.

The first term again captures the role of prices as emphasized in Armington models of trade. Other things equal, an exporter gains a RCA in a sector if its prices fall faster relative to its competitors in that sector compared to other sectors. The second term incorporates the effects of demand/quality. All else constant, an exporter gains a RCA in a sector if its demand/quality rises more rapidly relative to its competitors in that sector compared to other sectors. The third term summarizes the contribution of entry/exit. Other things equal, if entering varieties are more numerous or have lower demand-adjusted prices than exiting varieties, this increases the value of trade. An exporter gains a RCA in a sector if this contribution from entry/exit is large relative to its competitors in that sector compared to other sectors. Finally, the fourth term summarizes the impact of heterogeneity across varieties. All else constant, an exporter gains a RCA in a sector if its varieties become more heterogeneous relative to its competitors in that sector compared to other sectors.

2.10.2 Aggregate Trade Accounting

In addition to decomposing revealed comparative advantage, we can also use the structure of our model to decompose aggregate trade flows. Although aggregate imports are the sum across sectors of imports (rather

---

(2014), Fieler, Eslava and Xu (2014), and Sutton and Trefler (2016).

than the sum of log imports), we derive an exact log-linear decomposition of the share of each foreign exporter in aggregate imports. We use this decomposition to shed light on the different mechanisms underlying the large-scale changes in countries' shares of aggregate imports observed over our sample period.

Partitioning varieties into common and entering/exiting varieties, and using the share of each country in sectoral imports from equation (33), the log change in a country's share of aggregate trade can be exactly decomposed into the following terms:

$$
\begin{aligned}
\Delta \ln S_{jit}^{E} = &-\underbrace{\left\{\mathbb{E}_{jit}^{GFU*}\left[\left(\sigma_g^F - 1\right)\Delta \ln P_{ut}^U\right] - \mathbb{E}_{jt}^{GEFU*}\left[\left(\sigma_g^F - 1\right)\Delta \ln P_{ut}^U\right]\right\}}_{\text{Prices}} \\
&+\underbrace{\left\{\mathbb{E}_{jit}^{GFU*}\left[\left(\sigma_g^F - 1\right)\Delta \ln \varphi_{ut}^U\right] - \mathbb{E}_{jt}^{GEFU*}\left[\left(\sigma_g^F - 1\right)\Delta \ln \varphi_{ut}^U\right]\right\}}_{\text{Product Demand}} \\
&+\underbrace{\left\{\mathbb{E}_{jit}^{GF*}\left[\left(\sigma_g^F - 1\right)\Delta \ln \varphi_{ft}^F\right] - \mathbb{E}_{jt}^{GEF*}\left[\left(\sigma_g^F - 1\right)\Delta \ln \varphi_{ft}^F\right]\right\}}_{\text{Firm Demand}} \\
&-\underbrace{\left\{\mathbb{E}_{jit}^{GF*}\left[\frac{\sigma_g^F - 1}{\sigma_g^U - 1}\Delta \ln \lambda_{ft}^U\right] - \mathbb{E}_{jt}^{GEF}\left[\frac{\sigma_g^F - 1}{\sigma_g^U - 1}\Delta \ln \lambda_{ft}^U\right]\right\}}_{\text{Product Entry/Exit}} - \underbrace{\left\{\mathbb{E}_{jit}^{G}\left[\Delta \ln \lambda_{jigt}^F\right] - \mathbb{E}_{jt}^{GE*}\left[\Delta \ln \lambda_{jigt}^F\right]\right\}}_{\text{Firm Entry/Exit}} - \underbrace{\Delta \ln \left(\lambda_{jit}^E / \lambda_{jt}^T\right)}_{\text{Country-Sector Entry/Exit}} \\
&-\underbrace{\left\{\mathbb{E}_{jit}^{GFU*}\left[\frac{\sigma_g^F - 1}{\sigma_g^U - 1}\Delta \ln S_{ut}^{U*}\right] - \mathbb{E}_{jt}^{GEFU*}\left[\frac{\sigma_g^F - 1}{\sigma_g^U - 1}\Delta \ln S_{ut}^{U*}\right]\right\}}_{\text{Product Heterogeneity}} - \underbrace{\left\{\mathbb{E}_{jit}^{GF*}\left[\Delta \ln S_{ft}^{EF*}\right] - \mathbb{E}_{jt}^{GEF*}\left[\Delta \ln S_{ft}^{EF*}\right]\right\}}_{\text{Firm Heterogeneity}} \\
&+\underbrace{\Delta \ln \mathbb{K}_{jit}^T}_{\substack{\text{Country-sector} \\ \text{Scale}}} + \underbrace{\Delta \ln \mathbb{J}_{jit}^T}_{\substack{\text{Country-sector} \\ \text{Concentration}}} ,
\end{aligned}
\tag{38}
$$

as shown in Section H of the web appendix; where $\mathbb{E}_{jt}^{GEFU*}[\cdot]$ denotes the mean across common sectors (superscript G), common exporters within sectors (superscript E), common firms within an exporter and sector (superscript F) and common products within a firm (superscript U) for a given importer (subscript j) and time period (subscript t); and the other means are defined analogously, as reported in Section B of the web appendix.

From the first term, 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, our choice of units for product demand in equation (25) 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, 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 through sixth terms capture the contribution of entry/exit to changes in country import shares. An exporter's import share increases if on average its entering products, firms and sectors are more numerous and/or have lower demand-adjusted prices compared to its exiting varieties than for other foreign exporters. The seventh through eighth terms terms capture the impact of changes in the heterogeneity in demand-adjusted prices across products and firms. An exporter's import share increases if on average demand-adjusted prices become more dispersed across its products and firms compared to other foreign exporters.

The last two terms capture compositional effects across sectors. From the penultimate term, an exporter's import share increases if its exports become more concentrated in sectors that account for large shares of

23

expenditure relative to exports from other foreign countries. The final term captures the concentration of imports across sectors for an individual exporter relative to their concentration across sectors for all foreign exporters. This final term can be interpreted as an exact Jensen's Inequality correction term that controls in this log-linear decomposition for the fact that aggregate imports are the sum across sectors rather than the sum of the logs across sectors.

2.11 Accounting for Aggregate Price Movements

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. We now show that our exact aggregation approach can be used to separate out the contributions of different theoretical mechanisms to changes in the aggregate cost of living. Combining the aggregate price index in equation (4), with the tradable sector expenditure share (14), and the sectoral price index in equation (7), the change in the aggregate cost of living can be decomposed into the following five terms:

$$
\Delta \ln P_{jt} = \underbrace{\frac{1}{\sigma^G - 1} \Delta \ln \mu_{jt}^T}_{\substack{\text{Relative} \\ \text{Tradable Prices}}} + \underbrace{\mathbb{E}_{jt}^T \left[\frac{1}{\sigma_g^F - 1} \Delta \ln \mu_{jgt}^G \right]}_{\substack{\text{Domestic} \\ \text{Competitiveness}}} - \underbrace{\mathbb{E}_{jt}^T \left[\Delta \ln \varphi_{jgt}^G \right]}_{\substack{\text{Sector} \\ \text{Demand}}} + \underbrace{\mathbb{E}_{jt}^T \left[\frac{1}{\sigma^G - 1} \Delta \ln \mathbb{S}_{jgt}^G \right]}_{\substack{\text{Sector} \\ \text{Heterogeneity}}} + \underbrace{\mathbb{E}_{jt}^T \left[\Delta \ln \mathbb{P}_{jgt}^G \right]}_{\substack{\text{Import} \\ \text{Price Indexes}}}, \tag{39}
$$

as shown in Section G of the web appendix. Recall that the set of tradable sectors is constant over time and hence there are no terms for the entry and exit of sectors.

The first term, "Relative Tradable Prices," 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, "Domestic Competitiveness," 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, "Sector Demand," 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 in equation (27), this third term is equal to zero ($\mathbb{E}_{jt}^T \left[\Delta \ln \varphi_{jgt}^G \right] = 0$). Recall that we implicitly capture changes in demand/quality in tradable sectors relative to non-tradable sectors in the share of expenditure on tradable sectors ($\Delta \ln \mu_{jt}^T$) in the first term for "Relative Tradable Prices."

The fourth term, "Sector Heterogeneity," captures changes in the distribution of expenditure shares across tradable sectors. As the log function is concave, an increase in the dispersion of expenditure shares across tradable sectors necessarily reduces the average log sectoral expenditure share ($\mathbb{E}_{jt}^T \left[\frac{1}{\sigma^G - 1} \Delta \ln \mathbb{S}_{jgt}^G \right]$), which reduces the cost of living. Intuitively, when sectors are substitutes ($\sigma^G > 1$), an increase in the dispersion of demand-adjusted prices across sectors (as reflected in an increase in the dispersion of sectoral expenditure shares) reduces the cost of living, as consumers can substitute from less to more desirable sectors.

The fifth and final term, "Import Price Indexes," captures changes in average import price indexes across all tradable sectors. Other things equal, a fall in these average 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, and using the share of a foreign exporter in imports within a sector from equation (33), the share of a firm in imports from a foreign exporter and sector from equation (28), and the share of a product in firm imports from equation (12), we obtain:

$$
\underbrace{\mathbb{E}_{jt}^{T}\left[\Delta \ln \mathbb{P}_{jgt}^{G}\right]}_{\substack{\text{Import} \\ \text{Price Indexes}}} = \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\Delta \ln P_{ut}^{U}\right]}_{\text{Average Prices}} - \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\Delta \ln \varphi_{ft}^{F}\right]}_{\substack{\text{Firm} \\ \text{Demand}}} - \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\ln \varphi_{ut}^{U}\right]}_{\substack{\text{Product} \\ \text{Demand}}} \tag{40}
$$

$$
+ \underbrace{\mathbb{E}_{jt}^{T*}\left[\frac{1}{\sigma_g^F - 1}\Delta \ln \lambda_{jgt}^{E}\right]}_{\substack{\text{Country - Sector} \\ \text{Variety}}} + \underbrace{\mathbb{E}_{jt}^{TE*}\left[\frac{1}{\sigma_g^F - 1}\Delta \ln \lambda_{jigt}^{F}\right]}_{\substack{\text{Firm} \\ \text{Variety}}} + \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\frac{1}{\sigma_g^U - 1}\Delta \ln \lambda_{ft}^{U}\right]}_{\substack{\text{Product} \\ \text{Variety}}}
$$

$$
+ \underbrace{\mathbb{E}_{jt}^{TE*}\left[\frac{1}{\sigma_g^F - 1}\Delta \ln S_{jigt}^{E}\right]}_{\substack{\text{Country-Sector} \\ \text{Heterogeneity}}} + \underbrace{\mathbb{E}_{jt}^{TEF*}\left[\frac{1}{\sigma_g^F - 1}\Delta \ln S_{ft}^{EF}\right]}_{\substack{\text{Firm} \\ \text{Heterogeneity}}} + \underbrace{\mathbb{E}_{jt}^{TEFU*}\left[\frac{1}{\sigma_g^U - 1}\Delta \ln S_{ut}^{U}\right]}_{\substack{\text{Product} \\ \text{Heterogeneity}}},
$$

as shown in Section G of the web appendix.

The first term, "Average Prices," captures changes in the average price of common imported products that are supplied in both periods t and $t-1$. Other things equal, a fall in these average prices ($\mathbb{E}_{jt}^{TEFU*}\left[\Delta \ln P_{ut}^{U}\right] < 0$) reduces average import price indexes and hence the cost of living. The second and third terms incorporate changes in average firm demand (φ_{ft}^{F}) across common firms and average product demand (φ_{ut}^{U}) across common products. Our choice of units for product demand in equation (25) implies that the second term for the average log change in demand across common products within each firm is zero: $\mathbb{E}_{jt}^{TEFU*}\left[\ln \varphi_{ut}^{U}\right] = 0$. Our choice of units for firm demand in equation (26) implies that the unweighted average log change in demand across common foreign firms within each sector is zero: $\mathbb{E}_{jt}^{TF*}\left[\Delta \ln \varphi_{ft}^{F}\right] = 0$. However, the average of firm demand in the third term ($\mathbb{E}_{jt}^{TEF*}\left[\Delta \ln \varphi_{ft}^{F}\right]$) involves first averaging across firms within a given foreign exporter, and then averaging across foreign exporters, which corresponds to a weighted average across firms. Although in principle the weighted and unweighted averages across firms could differ from one another, we find that in practice they take similar values, which implies that the third term is close to zero. While the average in this third term is taken across foreign firms, we capture changes in the demand/quality for domestic firms relative to foreign firms within sectors in the term for the share of expenditure on foreign varieties within sectors ($\mathbb{E}_{jt}^{T}\left[\frac{1}{\sigma_g^F - 1}\Delta \ln \mu_{jgt}^{G}\right]$) in equation (39) above.

The fourth to sixth terms summarize the effect of the entry/exit of exporter-sector pairs, firms and products respectively. "Firm Variety" accounts for the entry and exit of foreign firms when at least one foreign firm from an exporter and sector exports in both time periods. "Country-Sector Variety" is an extreme form of foreign firm entry and exit that arises when the number of firms from a foreign exporter rises from zero to a positive value or falls to zero. Finally, "Product Variety" accounts for changes in the set of products within continuing foreign firms. For all three terms, the lower the shares of expenditure on common varieties at time t relative to those at time $t-1$ (the smaller values of $\Delta \ln \lambda_{jgt}^{E}$, $\Delta \ln \lambda_{jigt}^{F}$ and $\Delta \ln \lambda_{ft}^{U}$), the more attractive are entering varieties relative to exiting varieties, and the greater the reduction in the cost of living between the two time periods.

The seventh to ninth terms summarize the impact of the heterogeneity in expenditure shares across common exporter-sector pairs, common firms and common products, respectively. "Country-Sector Heterogene-

ity" reflects the fact that consumers are made better off if exporters improve performance in their most successful sectors. For example, consumers are better off if Japanese car makers and Saudi oil drillers become more relatively more productive (raising heterogeneity) than if Saudi car makers and Japanese oil drillers are the relative winners (lowering heterogeneity). Similarly at the firm-level, consumers benefit more from relative cost reductions or quality improvements of large sellers, which serve to raise firm heterogeneity. Since varieties are substitutes ($\sigma_g^U > 1$ and $\sigma_g^F > 1$), increases in the dispersion of these expenditure shares reduce the cost of living, as consumers can substitute away from less to more desirable varieties.

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 where assumptions about the supply-side become relevant. 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. 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 decompositions to alternative values of the 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). We report bootstrap standard errors that take into account that the estimates for each subsequent step depend on those in the preceding step.

In this section, we illustrate the RW estimator for the product tier of utility, and report the full details of the nested estimation in Section D 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

$$\Theta_{ft,t-1}^{U+} \left[\sum_{u \in \Omega_{ft,t-1}^U} S_{ut-1}^{U*} \left(\frac{P_{ut}^U}{P_{ut-1}^U} \right)^{1-\sigma_g^U} \right]^{\frac{1}{1-\sigma_g^U}} = \mathbb{M}_{ft}^{U*} \left[\frac{P_{ut}^U}{P_{ut-1}^U} \right] \left(\mathbb{M}_{ft}^{U*} \left[\frac{S_{ut}^{U*}}{S_{ut-1}^{U*}} \right] \right)^{\frac{1}{\sigma_g^U-1}}, \tag{41}$$

$$\Theta_{ft,t-1}^{U-} \left[\sum_{u \in \Omega_{ft,t-1}^{U}} S_{ut}^{U*} \left(\frac{P_{ut}^{U}}{P_{ut-1}^{U}} \right)^{-\left(1-\sigma_g^U\right)} \right]^{-\frac{1}{1-\sigma_g^U}} = \mathbb{M}_{ft}^{U*} \left[\frac{P_{ut}^{U}}{P_{ut-1}^{U}} \right] \left(\mathbb{M}_{ft}^{U*} \left[\frac{S_{ut}^{U*}}{S_{ut-1}^{U*}} \right] \right)^{\frac{1}{\sigma_g^U-1}}, \qquad (42)$$

where the variety correction terms $\left(\left(\lambda_{ft}^U / \lambda_{ft-1}^U \right)^{\frac{1}{\sigma_g^U-1}} \right)$ 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 the web appendix); finally the equalities in equations (41) and (42) 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 (41) and (42) 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:

$$\Theta_{ft,t-1}^{U+} = \left(\Theta_{ft,t-1}^{U-} \right)^{-1} = 1. \qquad (43)$$

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 ($N_{t,t-1}^U \to \infty$). More generally, the identifying assumption in equation (43) 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, an advantage of our CES specification is that the supply-side only enters through these estimated elasticities of substitution, and it is straightforward to undertake robustness checks to these elasticities 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. We currently report results using Chilean imports data from 2007-14, although future versions of the paper will report results using United States imports data. For each import customs shipment, the data report 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 8-digit Harmonized System (HS) codes), the country of origin, and the brand of the exporter (e.g. Nestlé, Toyota).

Using this information on import shipments, we construct a dataset for a single importer j (Chile) with many exporters i (countries of origin), sectors g (2-digit HS codes), firms f (foreign brands within exporter within sector), and products u (8-digit HS codes within foreign brands 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 drop the small number of HS8 codes that do not use consistent units after this standardization (e.g. we drop any HS8 code that switches from counts to kilograms). We also drop any observations for which

countries of origin or brands are missing. We collapse the import shipments data to the annual level by exporting firm and product, weighting by trade value, which yields around 5 million observations on Chilean imports by exporter-firm-product-year. We choose our baseline definitions of sectors and firms to remain close to previous empirical research in international trade. We also report robustness checks in which we consider alternative definitions (such as interpreting sectors as 4-digit rather than 2-digit HS codes).

Our measure of prices is the export unit value of a particular firm in an 8-digit HS category. While these data necessarily involve some aggregation across different varieties of products produced by the same firm, Section C 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_{ut}^U) captures unobserved compositional differences within each observed category. Moreover, 8-digit 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 construction of import price indexes.[14] 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).

One of the challenges of using trade-transactions data is that we need to identify firms based on their name and country of origin, and customs officials sometimes make typos or use non-standard abbreviations. For example, although it is likely that "Toyota Motor", "Toyoda Motor", "Toyota 7TR Motor", "Toyota Motor Corp", "Toyota Motor Corp.", "ToyotaMotor", and "Toyota Motor Corporation" refer to the same firm, they are all spelled differently. We therefore developed a name-matching algorithm to ensure that we correctly grouped different spellings of the same company together. The exact procedure is reported in the web appendix, but we provide a sense of it here. We start with around 1.7 million unique firm names in the raw data. First, we undertook some basic cleaning that resolved various data entry problems by eliminating extraneous strings (e.g., "-F" or "S.A."), non-numeric and non-alphabetic characters (e.g., "."), words that started with numbers, and uninformative entity names (e.g., "LLC" or "LTD"). We also standardized common words, so that "corporation" became "corp." Similarly, "technology" and "technologies" became "tech". This initial cleaning reduced the number of unique firm names to 1.4 million. Following with our example, this would have reduced our initial set of seven firm spellings to "Toyota Motor", "Toyoda Motor", and "ToyotaMotor". Our next step was to us a string similarity algorithm to measure the "distance" between any two spellings (measured in terms of how many letters needed to be changed to move from one firm name to another) to merge or distinguish between the remaining firm names. Thus, firms whose names differed by only one character would be grouped together. When we completed this and a few other minor cleaning procedures, we generated our final sample of just over one million unique firm identifiers. We then checked how well our procedure worked by manually checking the results of this algorithm for the 1,249 raw firm names in the Japanese steel sector (which we had not looked at when developing the procedure). Our cleaning algorithm and manual checking grouped firms in the same way 99.9 percent of the time.[15]

---

[14]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."

[15]As a robustness check, we also replicate our estimation and decompositions using the firm names before applying our name

To provide a check on the quality of the trade transactions data for Chile, we confirm that they exhibit similar properties as those for other countries examined in the empirical trade literature.[16] For example, as shown in Figure I.1 in Section I of the web appendix, we find high rates of turnover of exporting firms and exporting products, as well as selection conditional on firm and product survival. Around 50 percent of the exporting-firm-product observations in 2014 have been present for one year or less, but the just over 10 percent of these observations that have survived for at least seven years account for over 40 percent of import value. Additionally, as shown in Figure I.2 in Section I of the web appendix, Chile's imports are dominated by multi-product exporters. Although less than 30 percent of exporting firms are multi-product, they account for more than 70 percent of import value. Finally, as shown in Figure I.3 in Section I of the web appendix, we find that the extensive margins of exporting firms and exported products account for most of the cross-section variation in Chile's imports across countries, leaving a relatively small role for the intensive margin of average exports per firm and product. Across these and other empirical moments, we find that the Chilean data are representative of findings from other countries.

Another important similarity between Chilean data and that of other countries is the rapid expansion in imports from China. Chilean imports are highly concentrated across countries and characterized by a growing role of China over time. As shown in Figure 1, Chile's six largest import sources in 2007 were (in order of size) China, the U.S., Brazil, Germany, Mexico, and Argentina, which together accounted for more than 60 percent of its imports. Between 2007 and 2014, China's import share grew by over 50 percent, with all other major suppliers except Germany experiencing substantial declines in their market shares.

---

matching procedure, and find a similar pattern of results, confirming that this procedure is not consequential for our results.

[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.
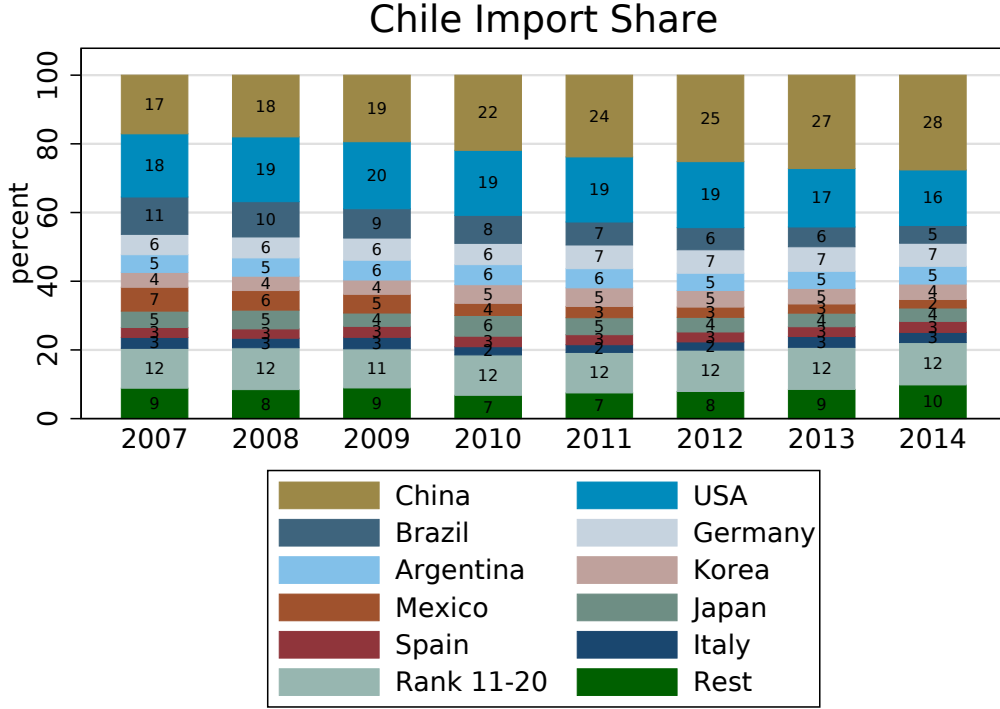
Figure 1: Chilean Shares of Total Imports over Time

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) that rationalize the observed data as an equilibrium. We use these structural estimates to solve for exporter price indexes, which in turn determine import shares in equation (33) and revealed comparative advantage (RCA) in equation (35). In Section 5.2, we use these solutions to decompose levels and changes in exporter price indices. In Section 5.3, we use the fact that RCA depends on relative exporter price indices across countries and sectors to decompose levels and changes in patterns of trade. We show that demand/quality is an important factor in understanding trade patterns in both the cross section and time series. We also document the importance of two other key forces determining trade: variety and heterogeneity. Importantly, we also show that demand-shifts, variety, and heterogeneity matter not only for trade patterns, but also for the measurement of prices as well.

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 ten. We find a median

estimated elasticity across products (σ_g^U) of 4.9, a median elasticity across firms (σ_g^F) of 2.7 and an elasticity across sectors (σ^G) of 1.69. These results imply that products within firms, firms within sectors and sectors are imperfect substitutes for one another, which has important implications for interpreting the data through the lens of existing trade theories. In particular, these results suggest that caution should exercised in taking models in which firms supply homogeneous outputs within sectors ($\sigma_g^F \to \infty$) and applying them directly to the standard statistical definitions of sectors used in the empirical trade literature.

A second important result concerns the nesting structure. 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 98 percent of sectors, and the firm elasticity is significantly larger than the sector elasticity at this significance level for all sectors. 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.

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).[17] Similarly, if we apply this alternative methodology to our data, we also obtain quite similar estimates, with median elasticities of 4.2 at the product level and 1.8 at the firm level, which are close to the 5.0 and 2.7 obtained here. Thus, our estimated elasticities do not differ substantially from those obtained using other standard methodologies. Finally, as a robustness check, we re-estimated the product, firm and sector elasticities using 4-digit HS categories as our definition of sectors instead of 2-digit HS categories. We find a similar pattern of results, with a somewhat larger median product elasticity of 5.2, a median firm elasticity of 2.6, and a sector elasticity of 1.7.

5.2 Exporter Price Indexes Across Sectors and Countries

We use these estimated elasticities (σ_g^U, σ_g^F, σ^G) to recover the structural residuals (φ_{ut}^U, φ_{ft}^F, φ_{jgt}^G) and solve for the exporter price indices ($\mathbb{P}_{jigt}^E$). A key implication of this section is measuring the cost of sourcing goods from an exporter and sector (as captured by these exporter price indices) involves making assumptions about demand in settings where goods are differentiated. In such environments, exporter price indices depend not only on conventional price terms, but also on the non-conventional forces of demand/quality, variety and heterogeneity. We now quantify the relative importance of each of these components in our data.

In the four panels of Figure 2, we display the log of the exporter price index ($\ln \mathbb{P}_{jigt}^E$) for against its

---

[17]Our median estimates for the elasticities of substitution within and across firms of 5.0 and 2.9 respectively compare with those of 6.9 and 3.9 respectively in Hottman et al. (2016).

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	4.34	1.80	1.69	1.36	0.11
5th	4.44	2.09	1.69	1.63	0.40
25th	4.63	2.40	1.69	2.06	0.71
50th	5.01	2.68	1.69	2.39	0.99
75th	5.54	3.02	1.69	2.82	1.34
95th	6.88	3.40	1.69	4.33	1.71
Max	8.47	4.14	1.69	4.43	2.45

Note: Estimated elasticities of substitution from the reverse-weighting estimator discussed in section 3 and in section D of the web appendix. Sectors are 2-digit Harmonized System (HS) codes; firms correspond to foreign exported brands within each foreign country within each sector; products; and products u reflect 8-digit HS codes within exported brands within sectors.

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

components, where each observation is an exporter-sector pair. For brevity, we show results for 2014, but find the same pattern for the 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 ($\mathbb{E}_{jigt}^{FU} [\ln P_{ut}^U]$). 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 only a weak positive correlation, with an estimated slope of around 0.21 and a R^2 of essentially zero. In other words, average prices are weakly correlated with the true CES price index, which underscores the problem of using average prices as a proxy for the CES price index.

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 elasticity of above 0.6 and regression R^2 of above 0.40. This finding of a tight connection between higher demand and higher prices is consistent with the quality interpretation of demand stressed in Schott (2004), in which producing higher quality incurs higher production costs.[18]

We follow a long line of research in trade and industrial organization in measuring demand/quality as a residual that shifts expenditure shares conditional on price, much like total factor productivity in the growth literature is a residual that shifts output conditional on inputs. The underlying feature of the data driving the importance of demand/quality in Figure 2 is the substantial variation in firm sales conditional on price. 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. Although we derive our results for a CES setup, we conjecture that this underlying feature of the data would generate a substantial role for demand/quality for a range of plausible demand systems.

---

[18]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).
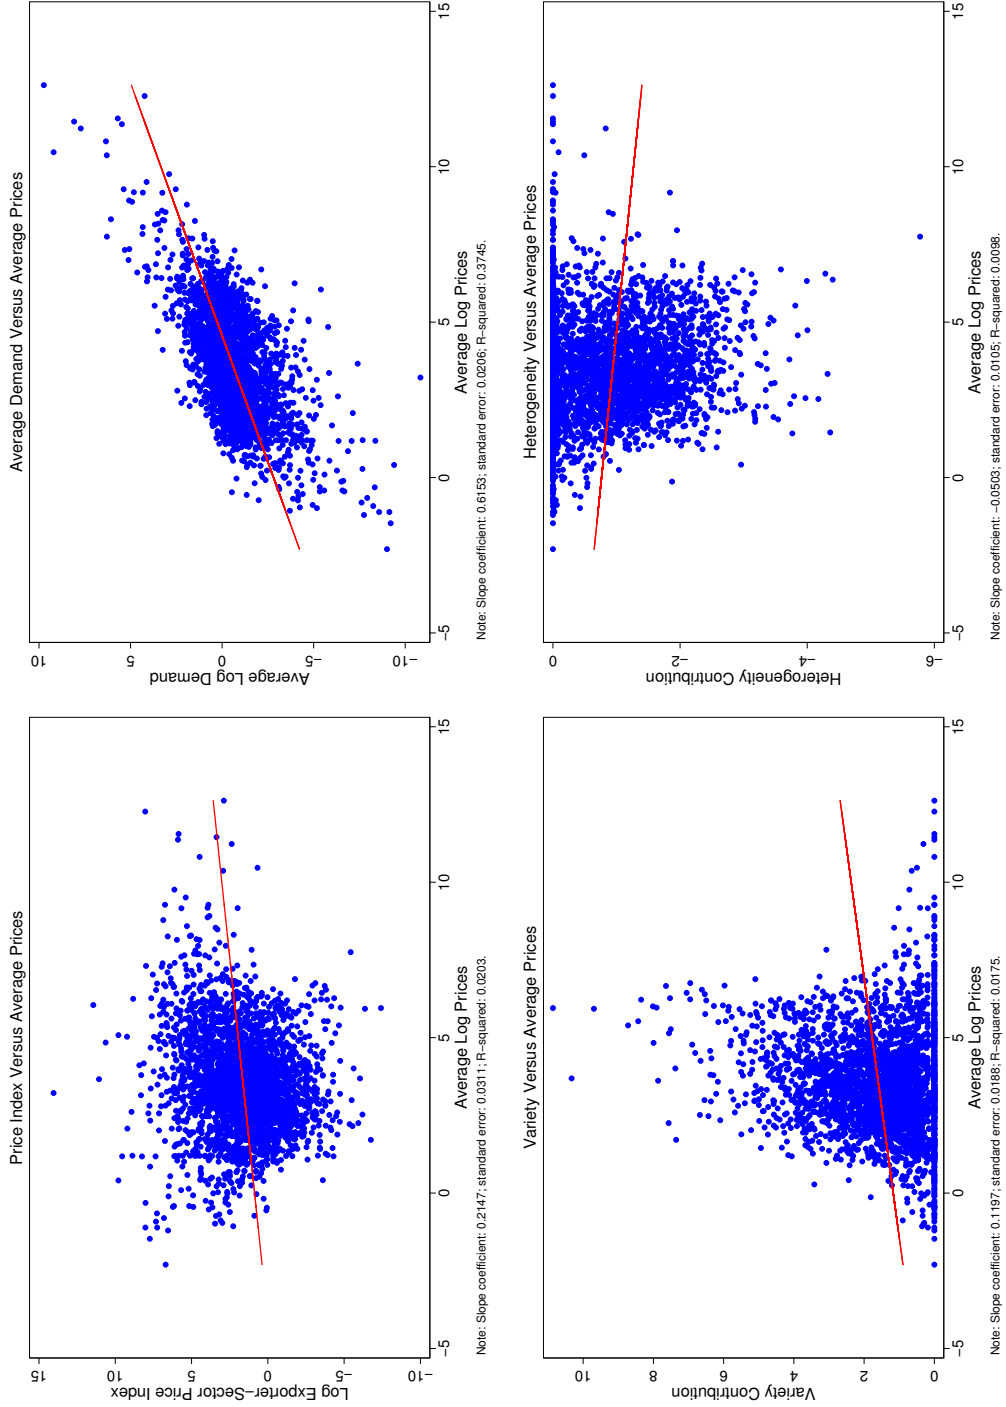
Figure 2: Exporter-Sector Price Indexes and their Components Versus Average Log Product Prices, 2014

In the bottom-left panel Figure 2, we show that exporter-sectors with high average prices (horizontal axis) also have many exporting firms and products (vertical axis), so that the increase in sourcing costs induced by higher average prices is also ameliorated by greater variety. This positive relationship is again strong and statistically significant (with an estimated elasticity of 0.12), although noisier (with a regression R^2 of less than 0.10). This finding highlights the empirical relevance of the love of variety forces emphasized by Krugman (1980). For our estimated elasticities of substitution across firms and products (and for other empirically plausible values of these parameters), we find these love of variety forces to be substantial relative to the observed differences in average prices.

As shown in the bottom-right panel, exporter-sectors with high average prices (horizontal axis) also exhibit greater heterogeneity in demand-adjusted prices across firms and products, as reflected in lower mean log expenditure shares (vertical axis). Therefore, the impact of higher average prices in raising sourcing costs is also mitigated by more scope to substitute from less to more attractive varieties. Although this relationship is precisely estimated, it is less strong than for demand/quality (with an estimated elasticity of 0.06 and a regression R^2 of less than 0.10). These results provide support for the mechanism of heterogeneity across goods emphasized in Melitz (2003), even after controlling for the overall number of varieties. For empirically plausible values of the elasticities of substitution across firms and products, we find that these heterogeneity forces are large relative to the observed differences in average prices.

These non-conventional determinants of the costs of sourcing goods across countries and sectors are also important in the time-series. 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 "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 equal the theoretically-correct measure of the change in the unit expenditure function.

We first demonstrate this point for aggregate import prices. In Figure 3, we use equation (40) to decompose the log change in aggregate import price indexes from 2008-14. The figure provides some important insights into why it is difficult to link import behavior to conventional price measures. If one simply computed the change in the cost of imported goods using a conventional Jevons index of the prices of those goods (the first term in equation (40)), one would infer a substantial increase in the cost of imported goods of around 9.2 percent over this time period (prices are measured in current price U.S. dollars). However, this positive contribution from higher prices of imported goods was offset by a substantial negative contribution from firm entry (variety). This expansion in firm import variety reduced the cost of imported goods by around 11.7 percent. By contrast, country-sector and firm heterogeneity fell over this period, which served to raise the CES price index and offset some of the variety effects. As a result, the true increase in aggregate import prices

Figure 3: Growth of Aggregate Import Prices 2008-14

from 2008-14 was only 4.4 percent, less than half of the value implied by a conventional geometric average of import prices. In other words, the true measure of aggregate import prices is strongly affected by factors other than movements in average prices.

We next show that this point applies not only to aggregate import prices but also to changes in the cost of sourcing goods from individual exporters and sectors (as captured by changes in exporter price indices $\Delta \ln \mathbb{P}^{E}_{jigt}$). Figure 4 displays the same information as in Figure 2, but for log changes from 2008-2014 rather than for log levels in 2014. In changes, the correlation between average prices and the true model-based measure of the cost of sourcing goods is even weaker and the role for demand/quality is even greater. Indeed, the slope for the regression of average log changes in prices on average log changes in quality is almost one, indicating that most price changes are almost completely offset by quality changes. This result implies a problem for standard price indexes that assume no demand or quality shifts for commonly available goods, such as the Sato-Vartia price index.

5.3 Decomposing Trade Patterns

In the model, trade patterns (as captured by RCA) depend on the relative costs of sourcing goods across countries and sectors (as captured by relative exporter price indexes). Therefore, we now build on the results for exporter price indexes in the previous subsection to examine the contribution of each our mechanisms towards patterns of trade. We start with the decompositions of the level and change of RCA in equations (36) and (37) in Section 2.10.1 above. We use a variance decomposition that is employed in another context
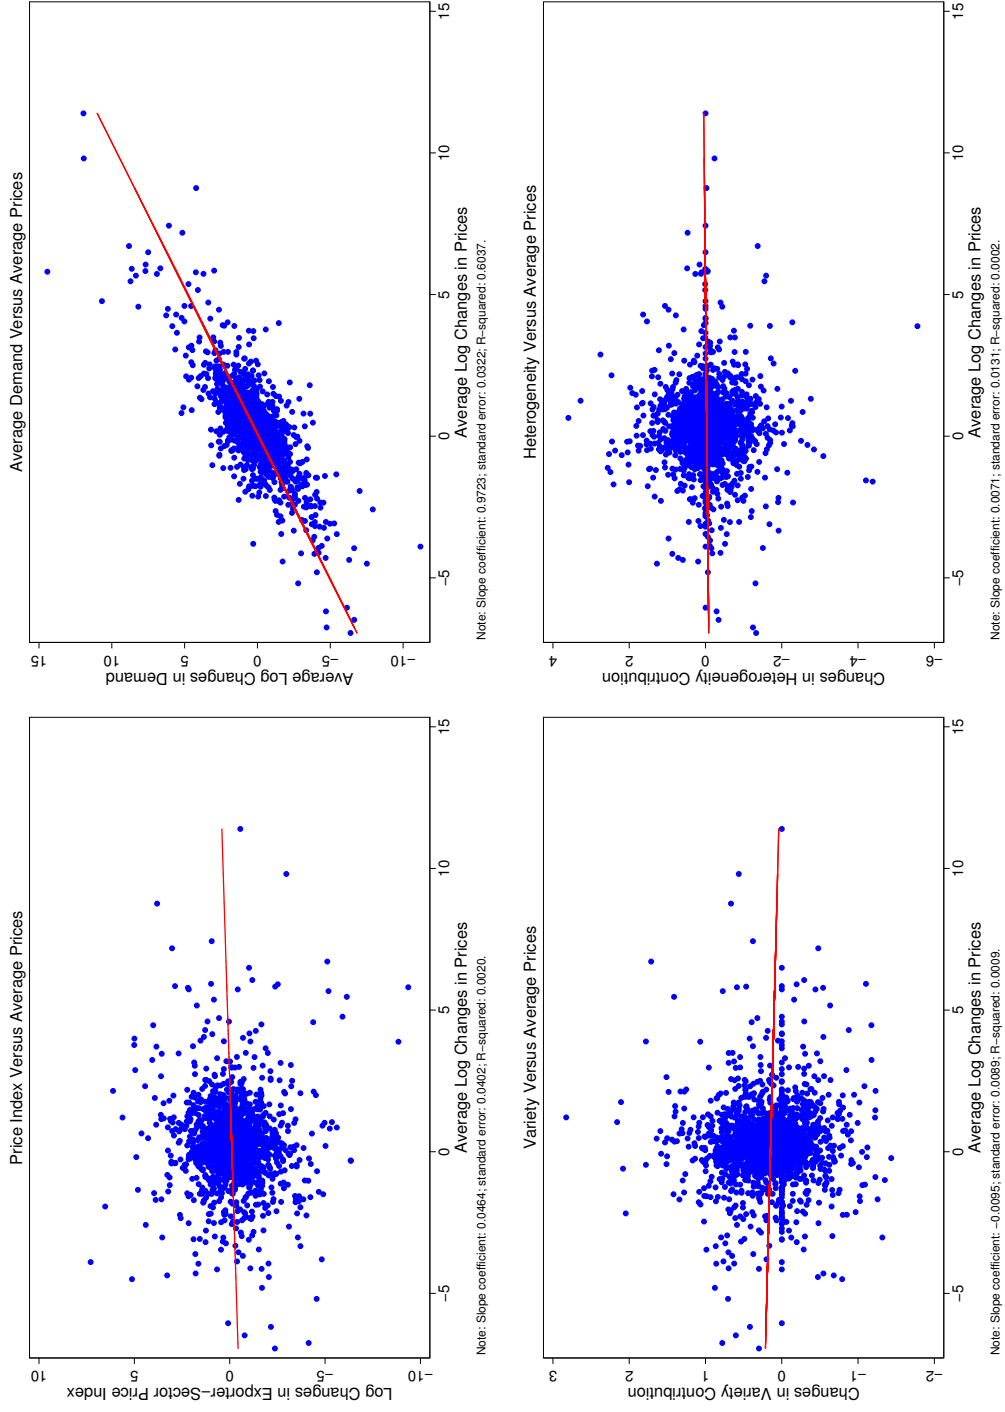
Figure 4: Log Changes in Sector-Exporter Price Indexes and their Components Versus Average Log Changes in Product Prices, 2008-2014

36

in 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 (36), we have:

$$\ln\left(RCA_{jigt}^{P}\right) = \alpha_P + \beta_P \ln\left(RCA_{jigt}\right) + u_P, \tag{44}$$

$$\ln\left(RCA_{jigt}^{\varphi}\right) = \alpha_\varphi + \beta_\varphi \ln\left(RCA_{jigt}\right) + u_\varphi,$$

$$\ln\left(RCA_{jigt}^{N}\right) = \alpha_N + \beta_N \ln\left(RCA_{jigt}\right) + u_N.$$

$$\ln\left(RCA_{jigt}^{S}\right) = \alpha_S + \beta_S \ln\left(RCA_{jigt}\right) + u_S,$$

where observations are exporters i and sectors g for a given importer j and year t. By the properties of OLS, $\beta_P + \beta_\varphi + \beta_N + \beta_S = 1$, and the relative value of each coefficient tells us the relative importance of each mechanism (prices, demand/quality, variety and heterogeneity). Similarly, we regress the log change in each component in equation (37) on the overall log change in RCA.

In Table 2, 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 the decomposition down to the firm level. In Columns (2) and (4), we undertake the decomposition all the way down to the product-level. In the interests of brevity, we concentrate on the results of the full decomposition in Columns (2) and (4). In contrast to Armington models of trade, we find that relative prices are comparatively unimportant in explaining patterns of trade. In the cross-section, average product prices account for 12.6 percent of the cross-section variation in RCA. In the time-series, we find that higher average prices account for only 9 percent of the variation. The results reflect the low correlations between average prices and exporter price indices seen in the last section. If average prices are weakly correlated with exporter price indices, they are unlikely to matter much for RCA, because RCA is determined by relative exporter price indices.

By contrast, we find that average demand/quality is two to three times more important than average prices, with a contribution of 23 percent for the levels of RCA and 36 percent for the changes in RCA. This empirical finding for the relative importance of these two determinants of patterns of international 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 trade models for relative prices as predictions for demand/quality-adjusted relative prices. However, it is not at all obvious that the determinants of quality/demand are exactly the same as those of prices, with, for example, a large literature in industrial organization emphasizing the importance of endogenous sunk costs for quality (e.g. Sutton 1991, 1998).

By far the most important of the different mechanisms for trade is firm variety, which accounts for 34 and 46 percent of the level and change of RCA respectively. Firm heterogeneity also makes a substantial contribution, particularly in the cross-section, where this term accounts for 30 percent of the variation in RCA. In the time-series, changes in the dispersion of expenditure shares across common firms are relatively less important, although they still account for 9 percent of the changes in RCA. Taken together, these findings support the empirical relevance of the mechanisms of love of variety and firm heterogeneity emphasized in Krugman (1980) and Melitz (2003) respectively. In principle, one could also interpret these results as consistent with neoclassical models such as Eaton and Kortum (2002). Although firm boundaries are indeterminate in

37

	Log Level RCA 2014		Log Change RCA 2008-14	
	Firm-Level Decomposition	Product-Level Decomposition	Firm-Level Decomposition	Product-Level Decomposition
Firm Price Index	0.126	-	0.091	-
Firm Demand	0.233	0.233	0.357	0.357
Firm Variety	0.344	0.344	0.464	0.464
Firm Heterogeneity	0.297	0.297	0.089	0.089
Product Prices	-	0.107	-	0.059
Product Variety	-	0.013	-	0.030
Product Heterogeneity	-	0.010	-	0.002

Note: Variance decomposition for the log level of RCA in 2014 and the log change in RCA from 2008-14 (from equation (44)).

Table 2: Variance Decomposition

that model, one could argue that each of the firms observed in the data specializes in a distinct disaggregated product within each sector. However, as we saw earlier, the fact that the elasticity of substitution within firms is larger than that across firms indicates that the data supports the theoretical assumption that demand is differentiated by firm.[19]

Differences in product variety within firms and heterogeneity across products within firms account for less than five percent of the variation in both the cross-section and time-series. In other words, within-firm differences in variety and heterogeneity do not account for much of the variation in import patterns. In line with the literature on multi-product firms such as Bernard, Redding and Schott (2010, 2011), these multi-product firms can account for a substantial share of expenditure within sectors. However, consistent with the empirical results in Hottman, Redding and Weinstein (2016), we find that much of the observed size differences across firms can be explained by firm demand/quality. Once we control for these overall size differences between multi-product firms and other firms, we find a relative small contribution from the number of products within firms and the sales distribution across products within firms.[20]

We now show that the non-conventional forces of demand/quality, variety and heterogeneity are also important for understanding changes in aggregate trade volumes. In Figure 5, we show the time-series decompositions of aggregate import shares from equation (38) for Chile's top six trade partners. As apparent from the figure, we find that we can account for the substantial increase in China's market share over the sample period by focusing mostly on increases in firm demand/quality, variety, and heterogeneity.[21] In con-

---

[19]In a robustness check defining sectors as 4-digit rather than 2-digit HS categories, we find a similar pattern of results. Demand/quality accounts for 28 and 42 percent of the level and change of RCA respectively, while firm variety accounts for 35 and 45 percent respectively.

[20]As a further robustness check, we undertake a grid search over σ_g^F from 2 to 8 (in 0.5 increments) and over σ_g^U from $(\sigma_g^F + 0.5)$ to 20 in 0.5 increments, 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 F of the web appendix, the firm variety and firm heterogeneity contributions are invariant to these elasticities (because they cancel from these terms). A higher value for σ_g^F raises the price contribution and reduces the demand/quality contribution. However, across the grid of values for these parameters, we find that average prices account for less than 35 percent of the level and less than 25 percent of the changes in RCA.

[21]This 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, so that they cannot implement the other decompositions developed in this paper, and they focus on Chinese export prices rather than international trade patterns.

Figure 5: Country Aggregate Import Shares

trast, average product prices increased more rapidly for China than for the other countries in our sample, reducing the extent of the increase in China's market share. In other words, our decomposition indicates that the reason for the explosive growth of Chinese exports was not due to cheaper Chinese exports, but rather substantial firm entry (variety), product upgrading (demand/quality), and improvements in the performance of leading firms relative to lagging firms (heterogeneity). By contrast the dramatic falls in import shares from Argentina and Brazil were driven by a confluence of factors that all pushed in the same direction: higher average product prices, firm exit (variety), a deterioration in the performance of leading firms relative to lagging firms (firm heterogeneity), and falls in average demand/quality relative to other countries.

Taken together, the results of this section suggest that empirically-successful trade theories based on CES demand should predict that most of the variation in patterns of trade occurs not through relative prices but rather through relative demand/quality, entry and exit, and heterogeneity in supplier characteristics across markets. Shifts in demand and variety account for 80 percent of movements in revealed comparative advantage with heterogeneity accounting for ten percent.

5.4 Additional Theoretical Restrictions

We have shown that our approach exactly rationalizes the micro data and permits exact log linear aggregation to the macro level. Therefore, our approach provides a coherent and internally-consistent framework for quantifying the relative importance of different mechanisms proposed in existing trade theories for aggregate trade patterns and prices. In this section, we compare this approach with existing special cases that impose

additional theoretical restrictions. As a result of these additional theoretical restrictions, these special cases no longer exactly rationalize the micro data as an equilibrium outcome, and we quantify the implications of these departures from the micro data for aggregate trade patterns and prices.

First, 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 (21)) 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 (35) for revealed comparative advantage, and examine how important the assumption of no demand shifts is for understanding patterns of trade. Because we know that the UPI perfectly rationalizes the data, any deviation from the data arising by 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 (21) to incorporate entry and exit.

In Figure 6, we report the results of these comparisons. The top two panels consider exporter price indices, while the bottom two panels examine RCA. In the top-left panel, we compare our common goods exporter price index (the CG-UPI on the horizontal axis) with the Sato-Vartia exporter price index (on the vertical axis), where each observation is an exporter-sector pair. 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 4, which shows that price shifts are highly and 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 our overall exporter price index (the UPI on the horizontal axis) with the Feenstra exporter price index (on the vertical axis), where each observation is again an exporter-sector pair. These two price indices have exactly the same variety correction term, but use different common goods price indices (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 actual changes in RCA (on the horizontal axis) against predicted changes in RCA based on relative exporter Sato-Vartia price indexes (on the vertical 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
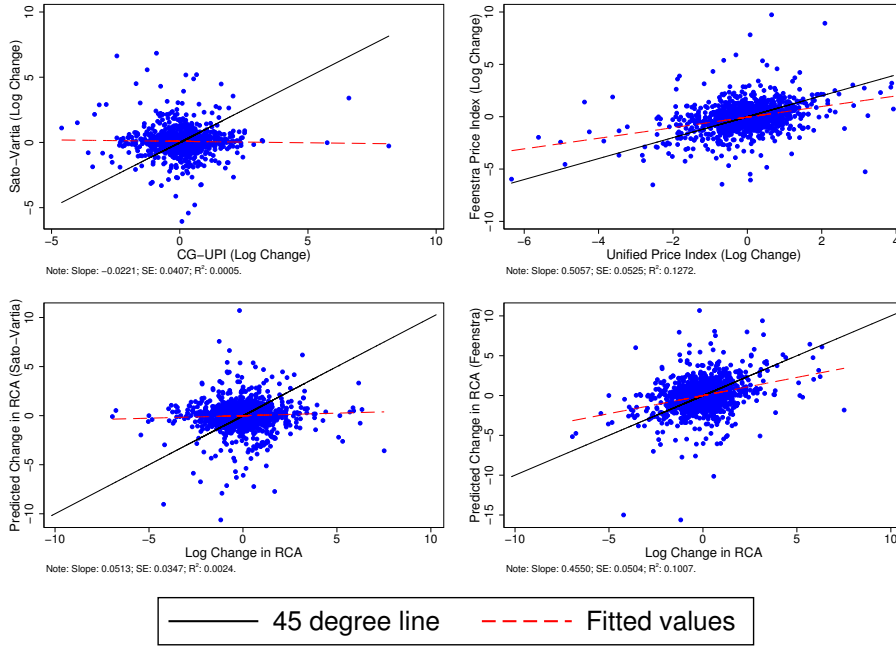
Figure 6: Sector-exporter Price Indexes with Time-Invariant Demand/Quality (Vertical Axis) Versus Time-Varying Demand/Quality (Horizontal Axis)

demand/quality and no entry and 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 indices (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 0.5 and the regression R^2 is about 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. Comparing these empirical findings to existing trade theories, the fact that most existing models are static has meant that relatively little attention has been devoted to changes in the demand/quality for existing goods, but the prominence of this term in the data suggests that it remains an important area for further theoretical research.[22]

Second, an important class of existing trade theories combines the assumption of a constant demand elasticity with the additional restriction of a constant supply elasticity (as for example in the Fréchet and 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

---

[22]These findings of prevalent changes in demand/quality for surviving goods suggest the empirical relevance of quality ladder models of trade and innovation, as in Grossman and Helpman (1991).

[23]A special case of our theoretical framework falls within this class of trade theories, as characterized in Arkolakis, Costinot and Rodriguez-Clare (2012), under the following additional supply-side restrictions: (i) Each industry consists of a continuum of firms and a continuum of products; (ii) Monopolistic competition; (iii) A single factor of production (labor); (iv) Balanced trade; (iii) Constant variable costs; (iv) Firm productivity is Pareto distributed; (v) Firm-product productivity is Pareto distributed; (vi) Firm demand/quality is a power function of firm productivity; (vii) Firm-product demand/quality is a power function of firm-product productivity. Together these assumptions ensure that firm sales (X_{ft}^F), the firm price index (P_{ft}^F), and firm demand/quality (φ_{ft}^F) are all Pareto distributed, as shown for the case of a single sector in Bernard, Redding and Schott (2011).

the observed data for firm sales and our model solutions for the firm price index and firm demand/quality ($\ln V_{ft}^F \in \left\{ \ln \mathbb{X}_{ft}^F, \ln P_{ft}^F, \ln \varphi_{ft}^F \right\}$) with their theoretical predictions under a Pareto distribution.

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). We start with the empirical distributions. Ordering firms by the value of a given variable V_{ft}^F for $f \in \left\{ 1, \dots, N_{jigt}^F \right\}$ for a given exporter i to importer j in sector g at time t, we observe the empirical quantiles:

$$\mathbb{V}_{ft} = \ln \left(V_{ft}^F \right). \tag{45}$$

We can use these empirical quantiles to estimate the empirical cumulative distribution function:

$$\widehat{\mathcal{F}}_{jigt} \left(V_{ft}^F \right) = \frac{f - b}{N_{jigt}^F + 1 - 2b}, \qquad b = 0.3, \tag{46}$$

where the plot position of $b = 0.3$ can be shown to approximate the median rank of the distribution (see Benard and Boslevenbach 1953). We next turn to the theoretical distributions. Under the assumption that each variable is Pareto distributed, its cumulative distribution function is given by:

$$\mathcal{F}_{jigt} \left(V_{ft}^F \right) = 1 - \left(\frac{\underline{V}_{jigt}^F}{V_{ft}^F} \right)^{a_g^V}, \tag{47}$$

where $\mathcal{F}_{jigt} \left(\cdot \right)$ is the Pareto cumulative distribution function; $\underline{V}_{jigt}^F$ is the lower limit of the support of the distribution for variable V_{ft}^F for exporter i, importer j, sector g and time t; and a_g^V is the Pareto shape parameter for variable V_{ft}^F for sector g. Inverting this cumulative distribution function, and taking logarithms, we obtain the following predicted theoretical quantile for each variable:

$$\ln \left(V_{ft}^F \right) = \ln \underline{V}_{jigt}^F - \frac{1}{a_g^V} \ln \left[1 - \mathcal{F}_{jigt} \left(V_{ft}^F \right) \right]. \tag{48}$$

The QQ estimator regresses the empirical quantile from equation (45) on the theoretical quantile from equation (48), using the empirical estimate of the cumulative distribution function from equation (46) to substitute for $\mathcal{F}_{jigt} \left(\cdot \right)$. We estimate this regression by 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.

In Figure 7, we show the predicted theoretical quantiles (vertical axis) against the empirical quantiles (horizontal axis). We display results for log firm imports (top left), log firm price indexes (top right) and log firm demand/quality (bottom left). In each case, we observe sharp departures from the linear relationship implied by a Pareto distribution, with the actual values below the predicted values in both the lower and upper tails. Estimating the regression in equation (48) separately for observations below and above the median, we find that these departures from linearity are statistically significant at conventional levels.[24]

---

[24]A similar analysis can be undertaken under the assumption of a Fréchet distribution. In this specification, we again find a similar pattern of statistically significant departures from the predicted linear relationship between the theoretical and empirical quantiles.
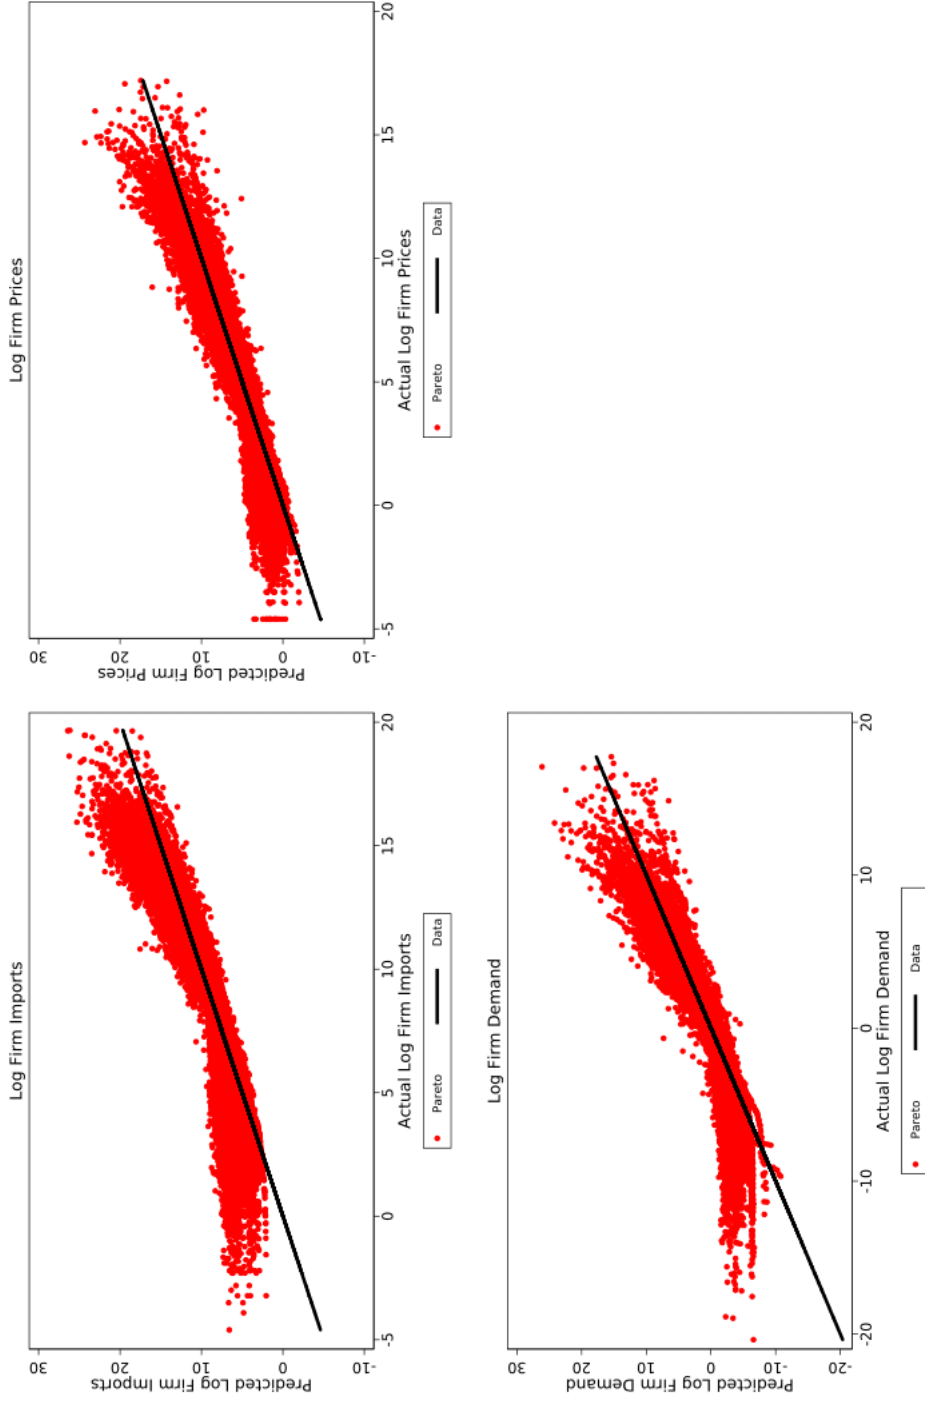
Figure 7: Theoretical and Empirical Quantiles (Pareto Distribution)

We now compare the empirical quantiles to those predicted by a log normal distribution. In particular, we suppose that observed firm sales and our model solutions for the firm price index and firm demand/quality ($V_{ft}^F \in \left\{ \mathbb{X}_{ft}^F, P_{ft}^F, \varphi_{ft}^F \right\}$) are log normally distributed:

$$\ln \left(V_{ft}^F \right) \sim \mathcal{N} \left(\kappa_{jigt}^V, \left(\chi_g^V \right)^2 \right), \tag{49}$$

where κ_{jigt}^V is the mean for variable V_{ft}^F for exporter i in importer j and sector g at time t and χ_g^V is the standard deviation for variable V_{ft}^F for sector g. It follows that the standardized value of the log of each variable is drawn from a standard normal distribution:

$$\mathcal{F}_{jigt} \left(V_{ft}^F \right) = \Phi \left(\frac{\ln \left(V_{ft}^F \right) - \kappa_{jigt}^V}{\chi_g^V} \right), \tag{50}$$

where $\Phi \left(\cdot \right)$ is the standard normal cumulative distribution function. Inverting this cumulative distribution function, we obtain the following predictions for the theoretical quantiles of each variable:

$$\frac{\ln \left(V_{ft}^F \right) - \kappa_{jigt}^V}{\chi_g^V} = \Phi^{-1} \left(\mathcal{F}_{jigt} \left(V_{ft}^F \right) \right), \tag{51}$$

which can be re-expressed as:

$$\ln \left(V_{ft}^F \right) = \kappa_{jigt}^V + \chi_g^V \Phi^{-1} \left(\mathcal{F}_{jigt} \left(V_{ft}^F \right) \right). \tag{52}$$

Following a similar approach as for the Pareto distribution above, the QQ estimator estimates equation (52) using the empirical quantile ($\ln V_{ft}^F \in \left\{ \ln \mathbb{X}_{ft}^F, \ln P_{ft}^F, \ln \varphi_{ft}^F \right\}$) on the left-hand side and substituting the empirical estimate of the cumulative distribution function for $\mathcal{F}_{jigt} \left(\cdot \right)$ on the right-hand side. Again we estimate this regression separately across foreign firms for each sector g (allowing the slope coefficient χ_g^V to vary across sectors) and including fixed effects for each exporter-year-sector (allowing the intercept κ_{jigt}^V to vary across exporters, sectors and time).

In Figure 8, we show the predicted log normal theoretical quantiles (vertical axis) against the empirical quantiles (horizontal axis). Again we display results for log firm imports (top left), log firm price indexes (top right) and log firm demand/quality (bottom left). In each case, we find that the relationship between the theoretical and empirical quantiles is closer to linearity for a log-normal distribution than for a Pareto distribution, which is consistent with Fernandes et al. (2015). Nonetheless, we observe substantial departures from the theoretical predictions of a log-normal distribution, and we reject the null hypothesis of normality at conventional levels of significance for the majority of sectors using a Shapiro-Wilk test.

As a concluding point, we examine the implications of these departures from a Pareto and log-normal distributions for understanding aggregate 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 the 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,
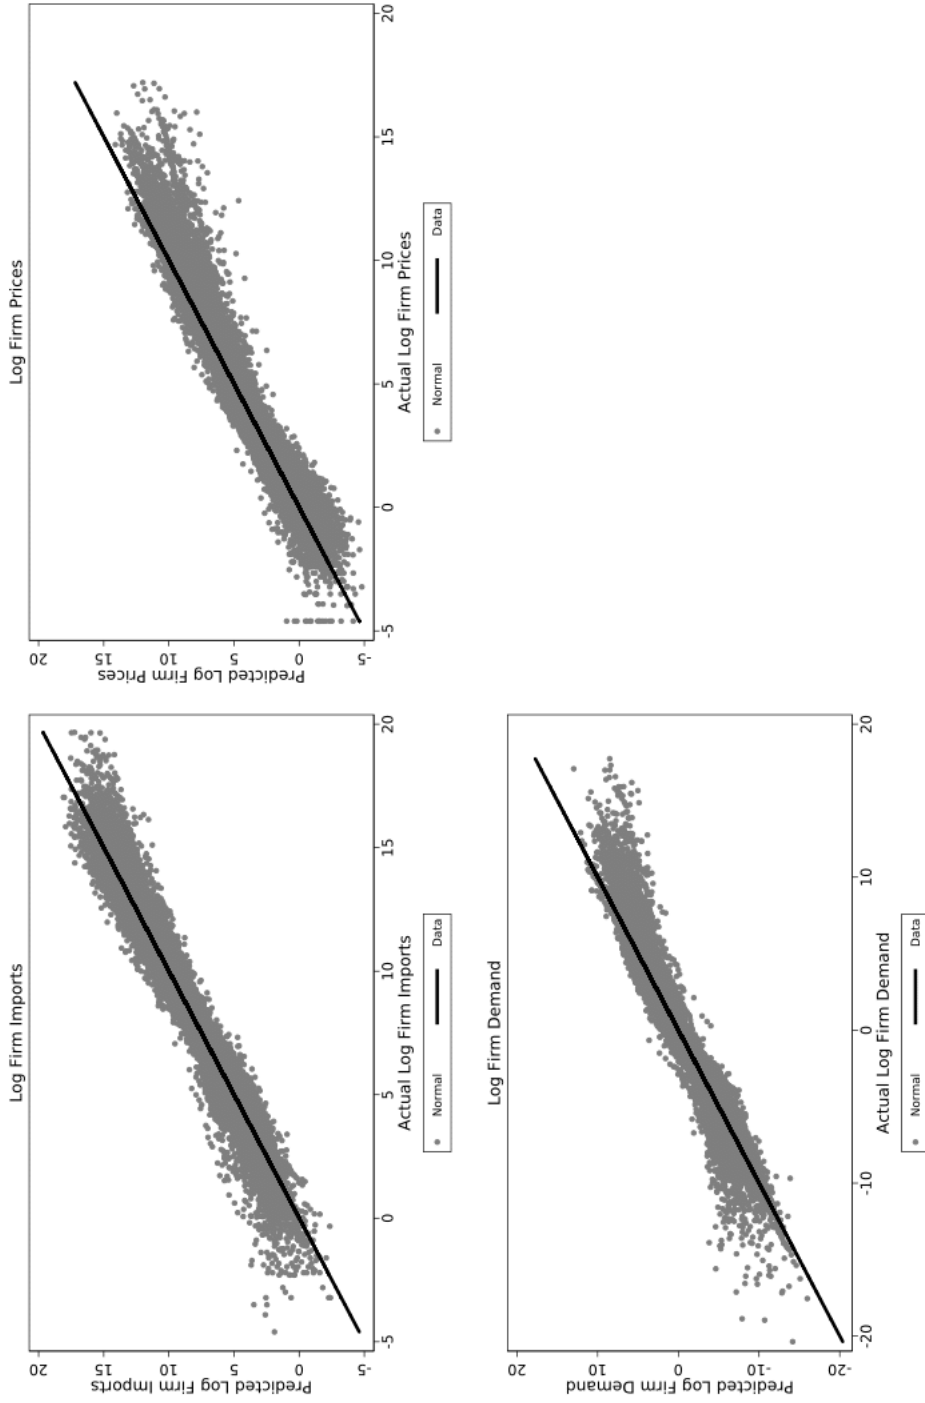
Figure 8: Theoretical and Empirical Quantiles (Log Normal Distribution)

log firm-price indexes and log firm-demand/quality, and using our estimated elasticities of substitution to construct the predicted log common-goods unified price index for each exporter and sector:

$$\widehat{\ln \mathbb{P}_{jigt}^E} = \mathbb{E}_{jigt}^F \left[\widehat{\ln P_{ft}^F} \right] - \mathbb{E}_{jigt}^F \left[\widehat{\ln \varphi_{ft}^F} \right] + \frac{1}{\sigma_g^F - 1} \mathbb{E}_{jigt}^F \left[\widehat{\ln \mathbb{S}_{ft}^{EF}} \right], \tag{53}$$

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 imports for importing country j, exporting country $i \neq j$, sector g and time t). A notable feature of this equation is that if we remove the hats, the equation is simply 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 logs 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. Since the inclusion of the fixed effects in equations (48) and (52) implies that both of the estimated distributions will be centered on the correct means of the logs of each variable, these distributional assumptions do not matter for our accounting of RCA. This further illustrates the point that the demand system and its parameters alone provide a framework for quantifying the contributions of different mechanisms to observed trade patterns, and once these are pinned down, other supply-side factors do not matter.

6   Conclusions

Existing trade research is largely divided into micro and macro approaches. Macro models impose distributional assumptions that can be at odds with what is observed in the micro data, while micro studies lack a clear template for aggregation to the macro level. In contrast, we develop a quantitative framework analogous to growth accounting approach that exactly rationalizes observed disaggregated trade data by firm, product, source and destination as an equilibrium of the model, while permitting exact log-linear decompositions of aggregate trade and prices into the contributions of different mechanisms proposed in existing trade theories.

Our approach nests most existing macro trade models, because we do not impose functional form restrictions on the supply-side. Our structural estimates of the elasticities of substitution between goods imply imperfect substitutability across products within firms and across firms within sectors. In such a differentiated goods environment, we show that measuring the cost of sourcing goods across countries and sectors requires taking a stand on the structure of demand. In general, the theoretically-correct measures of import prices depend not only on product prices, but also on demand/quality, the number of varieties and the dispersion of expenditure shares across varieties (which depends on the dispersion of demand-adjusted prices). We show that these non-conventional terms are only weakly correlated with conventional price measures. Therefore, empirical findings of elasticity puzzles, in which conventional measures of price shocks have weak effects on real economic variables, may in part reflect the failure to control for these non-conventional terms.

Our approach nests most existing macro trade models, because we do not impose functional form restrictions on the supply-side. Our structural estimates of the elasticities of substitution between goods imply imperfect substitutability across products within firms and across firms within sectors. In such a differentiated goods environment, we show that measuring the cost of sourcing goods across countries and sectors requires taking a stand on the structure of demand. In general, the theoretically-correct measures of import prices depend not only on product prices, but also on demand/quality, the number of varieties and the dispersion of expenditure shares across varieties (which depends on the dispersion of demand-adjusted prices). We show that these non-conventional terms are only weakly correlated with conventional price measures. Therefore, empirical findings of elasticity puzzles, in which conventional measures of price shocks have weak effects on real economic variables, may in part reflect the failure to control for these non-conventional terms.

We use the structure of our model to derive a measure of comparative advantage that allows for many countries, goods and factors of production as well as different market structure assumptions and is similar to existing measures of revealed comparative advantage (RCA). Whereas traditional trade theories emphasize

relative prices as the determinant of comparative advantage, we find that they account for relatively little of the observed variation in patterns of trade. Instead, relative demand/quality is more important, as emphasized in the literature on the Linder hypothesis. However, most of the observed cross-section and time-series variation in trade patterns is accounted for by firm entry and exit and heterogeneity across firms. These same forces also dominate changes in countries' aggregate shares of imports. For example, most of the increase in China's aggregate import market penetration over the sample period is explained by demand/quality upgrading, firm entry and an increase in the dispersion of firm characteristics.

Comparing our framework to special cases that impose additional theoretical restrictions, 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 can still only account for a small proportion of changes in comparative advantage over time. We show that standard supply-side distributional assumptions are rejected by the micro data, but do not matter for understanding revealed comparative advantage across countries and sectors, as long as these distributional assumptions are centered on the correct means for the logs of prices, demand/quality, and expenditure shares.

References

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

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

[3] 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.

[4] Antràs, Pol (2015) *Global Production: Firms, Contracts, and Trade Structure*, Princeton: Princeton University Press.

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

[6] Arkolakis, Costas, Marc-Andreas Muendler and Sharat Ganapati (2014) "The Extensive Margin of Exporting Products: A Firm-Level Analysis," *NBER Working Paper*, 16641.

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

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

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

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

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

[12] Benard, A and E. C. Bos-Levenbach (1953) "The Plotting of Observations on Probability Paper," *Statistica Neerlandica*, 7, 163-73.

[13] 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.

[14] Bernard, Andrew B., and J. Bradford Jensen (1999) "Exceptional Export Performance: Cause, Effect or Both?" *Journal of International Economics*, 47(1), 1-25.

[15] 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.

[16] 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.

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

[18] Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott (2017) "Global Firms," *Journal of Economic Literature*, forthcoming.

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

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

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

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

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

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

[25] 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.

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

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

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

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

[30] 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.

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

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

[33] Fajgelbaum, Pablo, Gene Grossman and Elhanan Helpman (2011) "Income Distribution, Product Quality, and International Trade," *Journal of Political Economy*, 119, 721-765.

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

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

[36] 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.

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

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

[39] Fieler, Ana Cecilia, Marcela Eslava and Daniel Xu (2016) "Trade, Quality Upgrading, and Input Linkages: A Theory with Evidence from Colombia," Duke University, mimeograph.

[40] 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.

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

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

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

[44] 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.

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

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

[47] Helpman, Elhanan (2006) "Trade, FDI and the Organization of Firms," *Journal of Economic Literature*, 44(3), 589-630.

[48] Helpman, Elhanan and Paul Krugman (1985) *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*, Cambridge MA: MIT Press.

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

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

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

[52] 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.

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

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

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

[56] Krugman, P., and A. J. Venables (1995) "Globalization and Inequality of Nations," *Quarterly Journal of Economics*, 60, 857-80.

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

[58] Linder, S. (1961) *An Essay on Trade and Transformation*, Stockholm: Almqvist and Wiksell.

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

[60] Mayer, Thierry and Marc Melitz and Gianmarco I. P. Ottaviano (2013) "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104(2), 495-536.

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

[62] 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.

[63] Melitz, Marc J. and Daniel Trefler (2015) "Gains from Trade when Firms Matter," *Journal of Economic Perspectives*, 26(2), 91-118.

[64] 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.

[65] Nocke, Volker and Stephen R. Yeaple (2014) "Globalization and Multiproduct Firms," *International Economic Review*, 55(4), 993-1018.

[66] Proudman, James and Stephen J. Redding (2000) "Evolving Patterns of International Trade," *Review of International Economics*, 8(3), 373-96.

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

[68] 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.

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

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

[71] Sutton, John (1998) *Technology and Market Structure: Theory and History*, Cambridge MA: MIT Press.

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