

Structural Change and Global Trade

Logan T. Lewis
Federal Reserve Board

Ryan Monarch
Federal Reserve Board

Michael Sposi
Southern Methodist University
& Federal Reserve Bank of Dallas

Jing Zhang
Federal Reserve Bank of Chicago

August 2019

Abstract

Services, which are less traded than goods, rose from 58 percent of world expenditure in 1970 to 79 percent in 2015. In a trade model featuring nonhomothetic preferences and input-output linkages, we find that such structural change has restrained the growth of world trade to GDP by 15 percentage points over this period. This is about half the magnitude that declining trade costs have boosted openness. Even absent rising protectionism, structural change dampens openness as well as the measured gains from trade through endogenous responses of expenditure shares to the trade regime. In the long run, policies that liberalize services trade complement structural change, boosting openness and welfare. In the face of continued structural change, high-income countries stand to gain more from trade liberalization in services than in goods, and low-income countries, the opposite.

JEL classifications: F41, L16, O41

Keywords: Globalization, Structural Change, International Trade

The views expressed here should not be interpreted as reflecting the views of the Federal Reserve Board of Governors, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Dallas, or any other person associated with the Federal Reserve System. We thank Diego Comin, Kerem Coşar, Lukasz Drozd, Colin Hottman, Kiminori Matsuyama, Marti Mestieri, Sebastian Sotelo, Tomasz Świącki, and Kei-Mu Yi for useful comments. This paper also benefited from audiences at the Chicago Fed, the Federal Reserve Board, the International Monetary Fund, the University of California at Santa Cruz, the University of Pittsburgh, the University of Rochester, the University of Western Ontario, and the University of Wisconsin as well as participants at the 2017 Society for Economic Dynamics Conference, 2017 BNM/IMF Conference on Challenges to Globalization, 2017 EIIT conference, Spring 2017 Midwest Macroeconomics Meetings, Spring 2017 Midwest International Trade Meetings, 2018 AEA Meetings, 2018 SCEIA Meetings, the 2018 Trade Integration and Growth Network (TIGN) conference, the 2018 Econometric Society Summer Asian Meeting, and 2018 NBER SI ITM meeting. Victoria Perez-Zetune provided excellent research assistance.

1 Introduction

Since 1970, world trade has grown faster relative to world income than at any other time in history, with the ratio of global trade to GDP more than doubling from about 20% to nearly 50% by 2015. Understanding this trend of greater “openness” is important, because when trade grows more rapidly than production, economies become more sensitive to changes in trade flows, trade policy and the composition of trade.¹ In addition, openness is closely related to gains from trade. The typical explanation for the rapid growth in the trade to GDP ratio is a decline in trade costs, including tariff reductions, improvements in communications technology, and lower transportation costs. However, there is another factor affecting the long-run movement in openness that has not been studied in the literature: the share of global spending on *services* has risen consistently and substantially. This fact, often referred to as “structural change”, is a well-known foundational component of economic growth and development. Yet because services are traded substantially less than goods are, structural change has dampened the growth in global openness and thus also the potential benefits from trade integration.

We study the joint evolution of international trade flows and structural change and answer two related questions: (i) How much has structural change restricted the growth in the trade to GDP ratio and, hence, the gains from trade? In a similar vein, how does the magnitude of that restriction compare to the effect of loosening of trade barriers over the same period? (ii) What are the policy implications for liberalizing trade in either goods or services going forward in the face of continued structural change? We find that, by 2015, structural change since 1970 has held back openness by 15 percentage points, or about one-third. This magnitude is roughly half as much as declining trade costs have boosted openness over the same time period. The gains from trade estimated in our model with structural change increase only about two-thirds as much since 1970 compared to a model without structural change. With structural change persisting in the future, reducing trade costs in services boosts both openness and the gains from trade, relative to reducing trade costs in goods. Moreover, rich countries stand to benefit relatively more from liberalizing services trade than from liberalizing goods trade, and to poor countries, the opposite.

The key insight is that structural change, as a result of price and income effects, determines what we consume, produce and trade. The data show that the share spent on services increases with real income and with relative services prices, both of which are endogenous to the trade regime. As countries open up to trade, the relative price of services and real income both rise, shifting expenditure to services and attenuating the benefits from more-integrated goods markets.

To fix ideas, we start by calculating a simple reduced-form counterfactual ratio of trade to GDP for a world without structural change. We assume that every country’s expenditure share on each sector is fixed at its 1970 level, while the ratio of trade to expenditure in each sector and country—“sectoral openness”—evolves as in the data. Under these assumptions, by 2015, openness would have been 81 percent, or 34 percentage points higher than in the data. This indicates that shifting consumption towards less traded services substantially suppressed trade growth in the last four decades. At the same time, it is implausible to assume that sectoral openness would have evolved exactly as in the real world in the absence of structural

¹See, for example, Hoekman (2015) and Irwin (2002).

change; the different level of prices and different levels of trade in intermediate inputs implied by a world of fixed expenditure shares should lead to different sectoral trade-to-expenditure ratios. The interactions between these factors call for a general equilibrium model to more accurately measure the effect of structural change on trade.

We build a multi-country, multi-sector, Ricardian trade model that incorporates endogenous structural change and trade patterns over time, similar to Uy, Yi and Zhang (2013) and Sposi (2019). On the production side, labor and intermediates produce a continuum of varieties in each sector. Countries differ in their sectoral input-output linkages, productivity and trade costs, forming the basis for comparative advantage. The evolution of productivity and bilateral trade costs at the sector level influences the patterns of production and trade over time. On the demand side, we adopt nonhomothetic preferences which allow total income and relative prices to shape sectoral expenditure shares, as in Comin, Lashkari and Mestieri (2018).

We calibrate the underlying structural parameters and time-varying processes of the model to relevant observables in 26 countries and a rest-of-world aggregate from 1970–2015. Using data on sectoral expenditures, sectoral prices, and employment levels, we estimate the key preference parameters, namely the elasticity of substitution between goods and services and the income elasticity of demand for both goods and services. Services have a higher income elasticity than goods, generating a positive correlation between the services expenditure share and income, and goods and services are complements. As Comin et al. (2018) show, the higher income elasticity of services is crucial for explaining the trend of higher services expenditure shares seen across countries. Coupling these with input-output coefficients from the World Input-Output Database and bilateral trade data enables us to back out estimates of productivity and trade costs at the sector level from the structural equations of the model.

After calibrating and solving the baseline model, we conduct a counterfactual similar to the reduced-form one. We impose constant expenditure shares across time by setting both the elasticity of substitution and the income elasticity to one. The model differs from the reduced-form calculation in that it allows for the counterfactual expenditure shares to impact prices for goods, services and labor, and trade flows, all of which affect sectoral openness in turn. We show that the model-based counterfactual still implies a substantial increase in the global trade-to-expenditure ratio relative to the data—15 percentage points or 32 percent higher than in reality by 2015. Using the same model, we show that the size of the decrease in openness as a result of structural change is about half the size of the increase in openness that stemmed from lower trade barriers.

The model-based contribution of structural change on world openness (15 percentage points) is smaller than the reduced-form one (34 percentage points). Why is this the case? The primary reason is that “goods openness”—the ratio of goods trade over goods expenditure—in the counterfactual is substantially lower than in the data. When fixing the expenditure shares at the 1970 level, the goods expenditure share rises relative to the data; however, as a result of input-output linkages weakening the overall effect, goods trade does not rise by the same degree.

The model incorporating nonhomothetic preferences requires using *equivalent variation* to measure the welfare gains from trade.² With the trade integration over the past four decades, we estimate that the gains

²The class of additively non-separable nonhomothetic preferences we use is not within the set discussed by Arkolakis, Costinot,

from trade relative to autarky have increased by 6.2 percentage points. This is lower than in a counterfactual world with expenditure shares held fixed from 1970 onwards, where gains from trade increase by 8.6 percentage points. Thus, since structural change holds back goods trade, and goods trade is a major contributor to welfare gains, accounting for structural change implies lower estimated welfare gains. We find that relative price changes, rather than income effects, explain most of the difference between the 6.2 and 8.6 percentage point changes.

Projecting our model into the future, assuming constant trade costs and continued technical progress, demonstrates that openness has peaked and may decline to below 40 percent by 2060. Importantly, the projected downward trend in the trade to GDP ratio is driven by the effects of increased services consumption. At the same time, there is little evidence that the slowdown in international trade growth that started in 2011 is due to new forces; that is, structural change has been a drag on trade growth for decades, and the drag has not been stronger in recent years.

To the extent that structural change reflects the efficient, long-run response of expenditures and production to asymmetric technological progress and aggregate income growth, it would not prove prudent to design policies that restrict the expansion of the service sector. Furthermore, trade policy has become increasingly restrained in its ability to further boost trade in goods, as tariffs are currently low. Modern trade policy could focus on liberalizing trade in services in order to foster the growth in world trade and to stimulate the benefits therein. Indeed, we estimated the projected gains from further reductions in trade barriers in either goods or services in the face of continued structural transformation induced by technical progress. The benefits from liberalizing services trade complements the rising services expenditures, particularly when the services sector becomes substantially more open relative to goods. Moreover, rich countries gain relatively more than poor countries from reducing services trade costs.

This paper contributes to a broad literature on how global trade grows relative to GDP. In an early theoretical contribution, Markusen (1986) includes nonhomothetic preferences in a trade model to be consistent with empirical evidence of a relationship between income and trade volumes. Rose (1991) shows that increases in income and international reserves along with declining tariff rates help explain the differences in trade growth across countries over three decades. Krugman, Cooper and Srinivasan (1995) analyze the growth in world trade since World War II and potential consequences for labor markets. Baier and Bergstrand (2001) find that income growth explains nearly two-thirds of the increase in global trade, with tariffs explaining an additional one-quarter. Imbs and Wacziarg (2003) document a U-shaped pattern of specialization as countries become richer; they first diversify across industries and only later specialize as they grow. Yi (2003) shows how vertical specialization—the splitting of production stages across borders—can amplify gross trade relative to value-added trade and help explain the large increases in trade-to-GDP ratios. Our paper provides an additional reason why the trade-to-GDP ratio is an imperfect measure of true openness, and given our projection exercise, a decline in this ratio does not necessarily reflect a less-open world with increasing protectionism.

A well-established literature documents how international trade and openness affect structural change. Matsuyama (2009) emphasizes that trade can alter patterns of structural change and that using closed-

Donaldson and Rodríguez-Clare (2019).

economy models may be insufficient. Uy et al. (2013) find that rapid productivity growth in South Korea's manufacturing sector contributed to a rise in manufacturing employment share due to improved comparative advantage. In a closed economy, the same productivity growth would have produced a decline in the manufacturing share. Betts, Giri and Verma (2017) explore the effects of South Korea's trade policies on structural change, finding that these policies raised the industrial employment share and hastened industrialization in general. Teignier (2018) finds that international trade in agricultural goods affected structural change in the United Kingdom even more than in South Korea. We show in this paper that structural change may be more consequential for international trade than trade is for explaining structural change in many countries.

More broadly, our findings point to structural change as being an important link between international trade and economic development. McMillan and Rodrik (2011) find that the effect of structural change on growth depends on a country's export pattern, specifically the degree to which a country exports natural resources. Cravino and Sotelo (2019) show that structural change originating from greater manufacturing trade increases the skill premium, particularly in developing countries. Sposi (2019) documents how the input-output structures of advanced economies are systematically different from those of developing economies, which contributes to systematic differences in resource allocations between rich and poor countries. Markusen (2013) shows how including nonhomothetic preferences into a Heckscher-Ohlin model can help explain why we observe less trade than predicted by models without nonhomotheticities.

Some analyses suggest that international trade plays only a small role in explaining structural change. Kehoe, Ruhl and Steinberg (2017) find that relatively faster growth in manufacturing productivity was the primary cause for reduced employment in the goods-producing sector in the United States, with a smaller role for trade deficits. Świącki (2017) also finds differential productivity growth is more important than other mechanisms, including international trade, in explaining structural change. Nonetheless, even if international trade only contributes a small portion to structural change, we show that structural change plays a large role in the growth of world trade.

Nonhomothetic preferences are important in understanding other aspects of international trade as well. Fielor (2011) finds that nonhomothetic preferences can explain why trade grows with income per capita but not population. Caron, Fally and Markusen (2014) documents a positive correlation between the income elasticity and skilled labor intensity across sectors and demonstrates that this is important for understanding international trade patterns. Simonovska (2015) shows that nonhomothetic preferences can match the pattern found in the data that higher-income countries have higher prices of tradable goods. Matsuyama (2015) and Matsuyama (2019) show that nonhomothetic preferences combined with home market effects can lead to high-income countries producing and exporting higher income elasticity goods without assuming they have an exogenous comparative advantage in such goods.

The remainder of the paper proceeds as follows. Section 2 describes the reduced-form counterfactual, and Section 3 sets up the model. Section 4 describes the calibration and solution of the model, and Section 5 presents the quantitative results. Section 6 concludes.

2 Empirics and a Reduced-Form Counterfactual

The ratio of global trade to GDP rose from about 20 percent to 50 percent between 1970 and 2010 before flattening through 2015. How would this trend have differed without the significant shift in expenditures from goods to services over that time? This section presents a direct and simplified answer to the question by holding each country’s expenditure share on goods and services fixed at its 1970 level and tracing out a counterfactual path for the global trade-to-GDP ratio.

2.1 Data

We begin by laying out the key concepts for our exercise and describing how we capture them in the data. First, some definitions: *Expenditure* refers to final demand: consumption, investment, and government spending. *Structural change* refers to changes in the expenditure of goods and services as a share of total expenditure over time. *Openness* is defined as total trade (imports plus exports) as a share of expenditure, with *sectoral openness* defined analogously at the sector (either goods or services) level.

For every country (and for the world as a whole), we can decompose openness in period t as:

$$\frac{Trade_t}{Exp_t} = \frac{Trade_{gt}}{Exp_{gt}} \frac{Exp_{gt}}{Exp_t} + \frac{Trade_{st}}{Exp_{st}} \frac{Exp_{st}}{Exp_t}, \quad (1)$$

where g and s denote goods and services. Clearly, changes in sectoral openness $\frac{Trade_{kt}}{Exp_{kt}}$, and sectoral expenditure shares $\frac{Exp_{kt}}{Exp_t}$, shape the aggregate openness measure over time.

We gather data needed to do the breakdown in equation (1) for the 26 countries and a rest-of-world aggregate over the period 1970–2015.³ In UN nomenclature, we take the goods sector to consist of “agriculture, hunting, forestry, fishing” and “mining, manufacturing, utilities,” while services includes “construction,” “wholesale, retail trade, restaurants, and hotels,” “transport, storage, and communication,” and “other activities”. The trade data is straightforward to assemble. We begin using data for 1995–2011 from the 2013 release of the World-Input-Output Database (WIOD). For the remaining years, we splice country-level goods trade data from the IMF Directions of Trade Statistics Database, and country-level services trade data from the World Development Indicators database.

Next is the construction of sectoral expenditures for each country in the sample. Although this data is provided in the WIOD, only the years 1995 through 2011 are included. In order to generate a longer time series of sectoral expenditure data, we use the aggregate identity that total absorption of each sector (i.e., gross output plus imports minus exports) must go either to final expenditures or intermediate expenditures. Using input-output coefficients, we can calculate what fraction of sectoral absorption went to intermediate usage. The remaining amount corresponds to final sectoral expenditure.

To do this calculation, we take a long time series of data on sectoral value added from the United Nations Main Aggregates Database and convert it into sectoral gross output using gross-output-to-value-added ratios.⁴ After adding imports and subtracting exports, we then use intermediate input coefficients to

³The full list of countries is listed in Appendix A.2.

⁴Some years require imputing country-level gross-output-to-value-added ratios, where our imputation procedure is based on

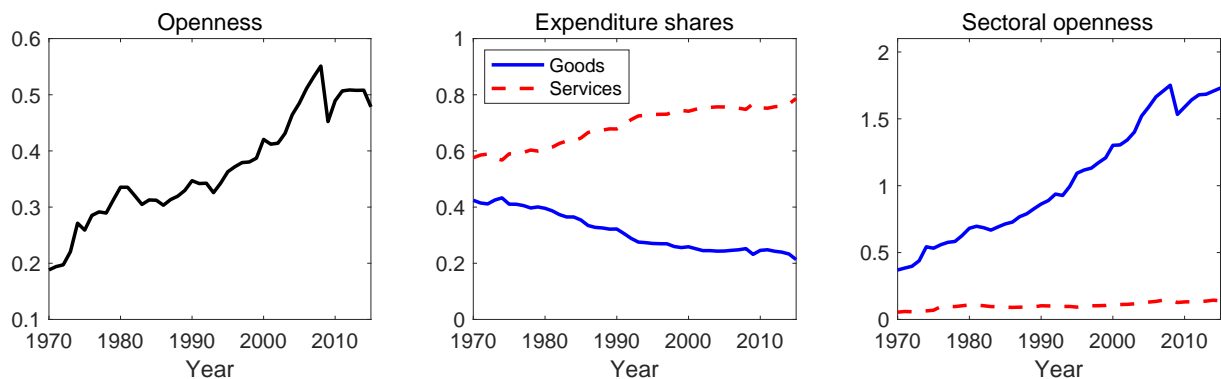
separate intermediate expenditure from final expenditure.⁵ The entire procedure is outlined in Appendix A.2.⁶

2.2 Openness and Structural Change

The long-term trends in openness and expenditure shares are shown in figure 1. The left panel shows the trade to GDP ratio (i.e., “openness”), rising from 19 percent in 1970 and to 55 percent by 2008. Openness grew substantially during much of the period, accelerating in the late 1990s and 2000s. From 2011-2015, the ratio was nearly flat at about 50 percent. The middle panel plots the expenditure shares for goods and services from 1970 to 2015. Clearly, world consumption shifted prominently from goods to services. The services expenditure share increased steadily by a total of 21 percentage points, from 58 percent in 1970 to 79 percent in 2015.

If these two sectors were both traded internationally with similar intensities, the impact of structural change on openness would be small. In the data, however, openness significantly differs between the two sectors. The right panel of figure 1 plots the ratio of sectoral trade to sectoral expenditure over 1970-2015. Clearly, goods are much more open than services; the ratio of trade to expenditure was about 5 percent for services compared to 37 percent for goods in 1970. Over time, trade openness increased for both sectors, while the increase was much more pronounced for goods. By the end of the period, the trade-expenditure ratio was about 14 percent for services and 170 percent for goods.⁷

Figure 1: Openness and structural change



Considering these three figures together presents a puzzle of sorts: How could trade grow so quickly while a relatively less-traded sector gained expenditure share? In fact, trade grew spectacularly *in spite of* the ongoing transition to services in the world economy, meaning structural change prevented even greater

estimating a relationship between value-added-to-gross-output ratios and income per capita, with country and time fixed effects. Details available upon request.

⁵These input-output coefficients also need to be imputed for some country years. Full details are presented in the Appendix, though the results are not sensitive to our choice of imputation.

⁶A stylized depiction of this calculation is in figure A.4 in the Appendix.

⁷The ratio of trade to expenditure can be over 100 percent for two reasons. First, trade refers to the sum of imports and exports. Second, trade is a gross measure including intermediate spending, while expenditure includes only final spending.

increases in trade. This dynamic becomes apparent when calculating the correlation between the growth rates of openness and the services expenditure share. For the world, the correlation is -0.77 meaning that periods of faster openness growth feature a slower-growing service expenditure share. This relationship also exists at the country level.⁸ Next we present reduced-form evidence of how much structural change held back global trade growth.

2.3 A Reduced-Form Counterfactual

To gauge the contribution of structural change to openness, we return to equation (1), but freeze every country's expenditure shares at their 1970 levels. We compute a counterfactual measure of openness as:

$$\frac{\widetilde{Trade}_t}{Exp_t} = \frac{Trade_{gt} Exp_{g0}}{Exp_{gt} Exp_0} + \frac{Trade_{st} Exp_{s0}}{Exp_{st} Exp_0}. \quad (2)$$

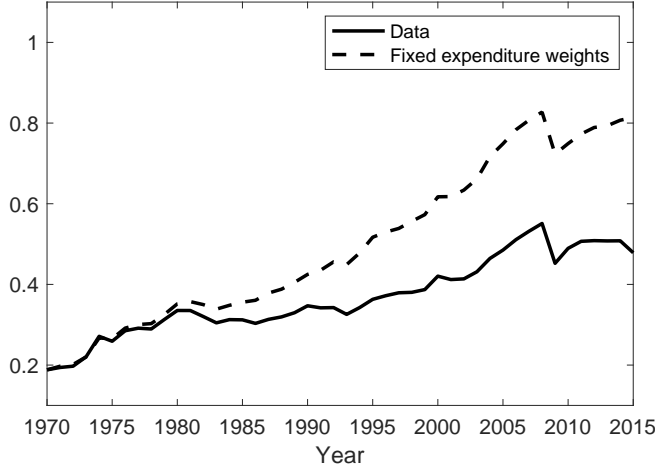
By holding the expenditure shares of sector k fixed at the first period, we shut down the process of structural change in the data. The counterfactual openness measure, $\frac{\widetilde{Trade}_t}{Exp_t}$, is free of structural change but retains the observed sectoral openness.

Figure 2 contrasts the aggregate trade openness measure in the data (the solid line) with the reduced-form counterfactual (the dashed line). The gap between the counterfactual measure and the data widens substantially over the 1990s and early 2000s, indicating that without underlying movements toward less-tradable services, global trade growth would have been far greater. As of 2015, persistent structural change since 1970 had lopped about 34 percentage points off the ratio of trade to expenditure.

Of course, this reduced-form exercise has a major deficiency: Sectoral openness should be *jointly* affected by the same forces that instigated structural change. The dynamics of sectoral productivity and trade barriers not only affect expenditure shares through relative prices and income levels but also affect sectoral openness through comparative advantage and trade flows. Additionally, input-output linkages are critical for identifying how changing expenditure shares feed through into production and trade. Thus, a structural model incorporating these endogenous relationships and featuring intermediate input-output linkages is needed to properly quantify the impact of structural change on international trade.

⁸For detailed results see Appendix A.1. When a country featured higher growth in its service expenditure share, it experienced lower growth in openness, even controlling for its level of income per worker.

Figure 2: Openness: data and empirical counterfactual



3 Model

We consider a multi-country, two-sector, Eaton-Kortum trade model of the global economy with nonhomothetic preferences. There are I countries and the two sectors are goods (g) and services (s). Household preferences have non-unitary income and substitution demand elasticities. In each sector, there is a continuum of varieties, and production uses both labor and intermediate inputs. All varieties are tradable, but trade costs vary across sectors, country-pairs, and over time. Productivities also differ in initial levels and subsequent growth rates across sectors and countries. These time-varying forces drive structural change. We omit the time subscript in this section for brevity.

3.1 Endowments and Preferences

Labor is mobile across sectors within a country, but immobile across countries. Let L_i denote total labor endowment in country i , which varies over time, and L_{ik} denote labor employed in sector k . The factor market clearing condition is given by:

$$L_i = L_{ig} + L_{is}. \quad (3)$$

The household in country i maximizes the level of aggregate consumption, C_i , which is a function of sectoral consumption C_{ig} and C_{is} . Aggregate consumption, which is a stand in for the overall utility level, combines sectoral composite goods according to the implicitly defined function:

$$\sum_{k=g,s} \omega_k^{\frac{1}{\sigma}} \left(\frac{C_i}{L_i} \right)^{\frac{\varepsilon_k(1-\sigma)}{\sigma}} \left(\frac{C_{ik}}{L_i} \right)^{\frac{\sigma-1}{\sigma}} = 1, \quad (4)$$

where for each sector $k \in \{g, s\}$, $\omega_k > 0$ describe the relative weight of each sector in the aggregate consumption bundle. These preferences, known as “Non-Homothetic Constant Elasticity of Substitution”, are

discussed by Hanoch (1975) and are also used by Comin et al. (2018). The elasticity of substitution across sectoral composite goods is σ , and the income elasticity of demand for sector k is ε_k . If $\sigma > 1$, goods and services are substitutes, and if $\sigma < 1$, they are complements. The sector with a greater ε_k is a luxury good, which expands in expenditure shares as the income rises, all else equal.

Hanoch (1975) showed that in order for these preferences to be well-behaved, i.e., monotone and quasi-concave, we require $\varepsilon_k > 0$ and either (i) $0 < \sigma < 1$ or (ii) $\sigma > 1$ must hold.⁹ Given our broad categorization of the two sectors, goods and services are complements empirically, so (i) is the relevant case in our context.¹⁰ As is usual when dealing with non-homothetic preferences, one of the income elasticity parameters needs to be normalized, since only the difference matters for allocations.¹¹ We normalize $\varepsilon_g = 1$. Another attraction of this normalization is that when ε_s is set at one, the preferences collapse into the commonly used homothetic CES preferences in the literature.

Comin et al. (2018) show that this specification of nonhomothetic preferences has two attractive properties for studying long-run structural change. First, the elasticity of the relative demand for the two sectoral composites with respect to aggregate consumption is constant at all levels of consumption. This contrasts with Stone-Geary preferences, where the elasticity of relative demand goes to zero as income or consumption rises—a prediction at odds with the data both at the macro and micro levels. Second, the elasticity of substitution between sectoral composites, given by σ , is constant over income, meaning that there is no functional relationship between income and substitution elasticities.¹²

The representative household maximizes aggregate consumption, C_i , in each period by choosing sectoral consumption levels, C_{ik} , subject to the following budget constraint:

$$\underbrace{P_{ig}C_{ig} + P_{is}C_{is}}_{P_i C_i} + \rho_i w_i L_i = w_i L_i + R L_i, \quad (5)$$

where w_i and P_{ik} denote the wage rate and the price of the sector- k composite good, respectively, and P_i denotes the aggregate consumption price (i.e., the price of one util when expenditures are optimally allocated across goods and services). The household supplies its labor endowment inelastically and spends its labor income on consumption. A fraction ρ_i of income is sent into a global portfolio, and the portfolio disperses R in lump sum equally across countries on a per-worker basis. ρ_i varies over time and R is determined by global portfolio balance in each period. Therefore, each country lends, on net, $\rho_i w_i L_i - R L_i$ to the rest of the world. This aspect enables the model to tractably match aggregate trade imbalances in the data, as in Caliendo, Parro, Rossi-Hansberg and Sarte (2018).

⁹Our notation differs from Hanoch (1975). These conditions are a rewriting of his expression (i) on page 403, with $d = \frac{\sigma-1}{\sigma}$ and $e_i = \varepsilon_k$.

¹⁰For the empirical estimation of σ and ε_k see section 4.1.

¹¹In our notation, C is really utility and the income elasticities are technically elasticities w.r.t. utility levels. Scaling each sector's income elasticity by the same proportion is a monotonic transformation of the utility function and only affects cardinal properties of the level of utility. Comin et al. (2018) employ the same preference structure as we do in a three-sector model and normalize the income elasticity in one sector to 1. Another commonly used class of preferences is "Stone-Geary", which incorporates income effects through sector-specific subsistence requirements. With Stone-Geary preferences, the subsistence term is normalized to zero in at least one sector; see Herrendorf, Rogerson and Valentinyi (2013) and Kongsamut, Rebelo and Xie (2001).

¹²This is a key difference from the preferences used in Fajgelbaum and Khandelwal (2016) and Hottman and Monarch (2018), whose frameworks could be used to ask a similar question to ours.

The first-order conditions imply that the consumption demand of sectoral goods satisfies:

$$C_{ik} = L_i \omega_k \left(\frac{P_{ik}}{P_i} \right)^{-\sigma} \left(\frac{C_i}{L_i} \right)^{(1-\sigma)\varepsilon_k + \sigma}, \quad (6)$$

The sectoral expenditure shares are thus given by:

$$e_{ik} \equiv \frac{P_{ik} C_{ik}}{P_i C_i} = \omega_k \left(\frac{P_{ik}}{P_i} \right)^{1-\sigma} \left(\frac{C_i}{L_i} \right)^{(1-\sigma)(\varepsilon_k - 1)} \Leftrightarrow \frac{e_{is}}{e_{ig}} = \left(\frac{\omega_s}{\omega_g} \right) \left(\frac{P_{is}}{P_{ig}} \right)^{1-\sigma} \left(\frac{C_i}{L_i} \right)^{(1-\sigma)(\varepsilon_s - \varepsilon_g)} \quad (7)$$

and the aggregate price index is given by:

$$P_i = \left[\sum_{k=g,s} \omega_k P_{ik}^{1-\sigma} \left(\frac{C_i}{L_i} \right)^{(1-\sigma)(\varepsilon_k - 1)} \right]^{\frac{1}{1-\sigma}}. \quad (8)$$

Thus, the elasticity of substitution between sectors, σ , and the sectoral elasticity of income, ε_k , govern how relative price and real income per worker shape the sectoral expenditure shares. Specifically, when $\sigma < 1$, a rising sectoral relative price pushes up the expenditure share in that sector, and vice versa. When sectoral income elasticities differ, i.e., $\varepsilon_s - \varepsilon_g > 0$, the service sector's expenditure share also rises with consumption per worker.

3.2 Technology and Market Structure

There is a continuum of varieties, $z \in [0, 1]$, in both the goods (g) and services (s) sectors. The sectoral composite good, Q_{ik} , is an aggregate of the individual varieties $Q_{ik}(z)$:

$$Q_{ik} = \left(\int_0^1 Q_{ik}(z)^{\frac{\eta-1}{\eta}} dz \right)^{\frac{\eta}{\eta-1}},$$

where the elasticity of substitution across varieties within a sector is $\eta > 0$. Each variety z is either produced locally or imported from abroad. The composite sectoral goods are used in domestic final consumption and domestic production as intermediate inputs:

$$Q_{ik} = C_{ik} + \sum_{n=g,s} M_{ink},$$

where M_{ink} is the intermediate input usage of composite good k in the production of sector n .

Each country possesses technologies for producing all the varieties in both sectors. Production requires labor and intermediate inputs as in Levchenko and Zhang (2016). The production function for variety $z \in [0, 1]$ in sector $k \in \{g, s\}$ of country i is:

$$Y_{ik}(z) = A_{ik}(z) (T_{ik} L_{ik}(z))^{\lambda_{ik}} [\prod_{n=g,s} M_{ikn}^{\gamma_{ikn}}(z)]^{1-\lambda_{ik}}, \quad (9)$$

where λ_{ik} denotes the country-specific value-added share in production, and γ_{ikn} denotes the country-specific

share of intermediate inputs sourced from sector n ; these parameters vary over time to track changes in input-output relationships. $Y_{ik}(z)$ denotes output, $L_{ik}(z)$ denotes labor input, and $M_{ikn}(z)$ denotes sector- n composite goods used as intermediates in the production of the sector k variety z . T_{ik} is the time-varying, exogenous productivity of varieties in sector k and scales value added equally across all varieties. $A_{ik}(z)$ is a variety-specific productivity level that scales gross output, given by the realization of a random variable drawn from the cumulative distribution function $F(a) = Pr[A \leq a]$. Following Eaton and Kortum (2002), we assume that $F(a)$ is a Fréchet distribution: $e^{-a^{-\theta_k}}$. The larger θ_k is, the lower the heterogeneity, or variance, in $A_{ik}(z)$ is.¹³ The parameters governing the distribution of idiosyncratic productivity draws are invariant across countries but different across sectors. We assume that the productivity is drawn each period.¹⁴

Total sectoral labor, input usage, and production in sector k in country i are the aggregates of the variety-level components taken over the set of varieties produced in country i , V_{ik} :

$$L_{ik} = \int_{V_{ik}} L_{ik}(z) dz; \quad M_{ikn} = \int_{V_{ik}} M_{ikn}(z) dz; \quad Y_{ik} = \int_{V_{ik}} Y_{ik}(z) dz.$$

Markets are perfectly competitive; prices are determined by marginal costs of production. The cost of an input bundle in sector k is:

$$v_{ik} = B_{ik} w_i^{\lambda_{ik}} (\prod_{n=g,s} (P_{in})^{\gamma_{ikn}})^{1-\lambda_{ik}},$$

where $B_{ik} = \lambda_{ik}^{-\lambda_{ik}} ((1 - \lambda_{ik}) \prod_{n=g,s} \gamma_{ikn}^{-\gamma_{ikn}})^{\lambda_{ik}-1}$. The cost of an input bundle is the same within a sector, but varies across sectors given different input shares.

3.3 Trade

When varieties are shipped abroad, they incur trade costs, which include tariffs, transportation costs, and other barriers to trade. We model these costs as exogenous iceberg costs, which vary over time to track the pattern of bilateral trade. Specifically, if one unit of variety z is shipped from country j , then $\frac{1}{\tau_{ijm}}$ units arrive in country i . We assume that trade costs within a country are zero, i.e., $\tau_{iig} = \tau_{iis} = 1$. This means that the price at which country j can supply variety z in sector k to country i equals $p_{ijk}(z) = \frac{\tau_{ijk} v_{jk}}{A_{ik}(z) T_{ik}^{\lambda_k}}$. Since buyers will purchase from the cheapest source, the actual price for this variety in country i is $p_{ik}(z) = \min \{p_{ijk}(z)\}_{j=1}^I$.

Under the Fréchet distribution of productivities, Eaton and Kortum (2002) show that the price of composite good $k \in \{g, s\}$ in country i is:

$$P_{ik} = \Gamma_k \left[\sum_{j=1}^I \left(T_{jk}^{-\lambda_{jk}} v_{jk} \tau_{ijk} \right)^{-\theta_k} \right]^{-\frac{1}{\theta_k}}, \quad (10)$$

where the constant $\Gamma_k = \Gamma(1 - \frac{\eta-1}{\theta_k})^{\frac{1}{1-\eta}}$ denotes the Gamma function, and the summation term on the right-

¹³ $A_k(z)$ has geometric mean $e^{\frac{\gamma}{\theta_k}}$ and its log has a standard deviation $\frac{\pi}{\theta_k \sqrt{6}}$, where γ is Euler's constant.

¹⁴ Alternatively, we could assume that the productivity is drawn once in the initial period, and as the T 's change over time, the productivity relative to T remains constant.

hand side summarizes country i 's access to global production technologies in sector k scaled by the relevant unit costs of inputs and trade costs.¹⁵

The share of country i 's expenditure on sector- k goods from country j , π_{ijk} , equals the probability of country i importing sector- k goods from country j , and is given by:

$$\pi_{ijk} = \frac{\left(T_{jk}^{-\lambda_{jk}} v_{jk} \tau_{ijk}\right)^{-\theta_k}}{\sum_{s=1}^I \left(T_{sk}^{-\lambda_{sk}} v_{sk} \tau_{isk}\right)^{-\theta_k}}. \quad (11)$$

Equation (11) shows how a higher average productivity, a lower unit cost of input bundles, and a lower trade cost in country j translates into a greater import share by country i .

3.4 Equilibrium

Combining the goods and factor market clearing conditions and demand equations with the equations for the consumption of the composite good, trade shares, prices, and the global portfolio balance yields a set of conditions that fully characterize the equilibrium of the model. Table 1 collects these conditions. Equations (D1)-(D3) describe the household demand side. (D1) provides the optimal condition for sectoral consumption. (D2) specifies the aggregate price index given the preferences. (D3) is the budget constraint.

Table 1: Equilibrium conditions

D1	$C_{ik} = L_i \omega_k \left(\frac{P_{ik}}{P_i}\right)^{-\sigma} \left(\frac{C_i}{L_i}\right)^{(1-\sigma)\varepsilon_k + \sigma}$	$\forall i, k$
D2	$P_i = \left[\sum_{k=g,s} \omega_k P_{ik}^{1-\sigma} \left(\frac{C_i}{L_i}\right)^{(1-\sigma)(\varepsilon_k-1)}\right]^{\frac{1}{1-\sigma}}$	$\forall i$
D3	$P_i C_i + \rho_i w_i L_i = w_i L_i + R L_i$	$\forall i$
S1	$\pi_{ijk} = \frac{\left(T_{jk}^{-\lambda_{jk}} v_{jk} \tau_{ijk}\right)^{-\theta_k}}{\sum_{\ell=1}^I \left(T_{\ell k}^{-\lambda_{\ell k}} v_{\ell k} \tau_{i\ell k}\right)^{-\theta_k}}$	$\forall i, j, k$
S2	$v_{ik} = B_{ik} w_i^{\lambda_{ik}} \prod_{n \in \{g,s\}} P_{in}^{(1-\lambda_{ik})\gamma_{ikn}}$	$\forall i, k$
S3	$P_{ik} = \Gamma_k \left(\sum_{j=1}^I \left(T_{jk}^{-\lambda_{jk}} v_{jk} \tau_{ijk}\right)^{-\theta}\right)^{-\frac{1}{\theta}}$	$\forall i, k$
S4	$w_i L_{ik} = \lambda_{ik} P_{ik} Y_{ik}$	$\forall i, k$
S5	$P_{in} M_{ikn} = (1 - \lambda_{ik}) \gamma_{ikn} P_{ik} Y_{ik}$	$\forall i, k, n$
S6	$Q_{ik} = C_{ik} + \sum_{n \in \{g,s\}} M_{ink}$	$\forall i, k$
S7	$\sum_{k \in \{g,s\}} P_{ik} Y_{ik} - \sum_{k \in \{g,s\}} P_{ik} Q_{ik} = \rho_i w_i L_i - R L_i$	$\forall i$
G1	$\sum_{i=1}^I \rho_i w_i L_i = R \sum_{i=1}^I L_i$	
G2	$P_{ik} Y_{ik} = \sum_{j=1}^I P_{jk} Q_{jk} \pi_{jik}$	$\forall i, k$

Equations (S1)-(S7) are from the supply side. (S1) gives bilateral import shares in total absorption at

¹⁵We assume $\eta - 1 < \theta_k$ to have a well-defined price index. Under this assumption, the parameter η , which governs the elasticity of substitution across goods within a sector, can be ignored because it appears only in the constant term Γ .

the sector level. (S2) specifies the cost of a unit of the input bundle. (S3) gives sectoral prices. (S4) and (S5) state the optimal value added and intermediate input usages implied by the Cobb-Douglas production function. (S6) links sectoral aggregate absorption with final demand and intermediate input demand. (S7) is the resource constraint at the country level.

Equations (G1)-(G2) are from the global market clearing conditions. Equation (G1) specifies net transfers across countries are zero globally. Equation (G2) links a country's total output in a sector with the sum of demand from all countries. Together, these conditions imply that all labor markets clear.

We define a competitive equilibrium of our model economy with the exogenous time-varying processes for every country: labor endowment $\{L_i\}$, trade costs $\{\tau_{ijg}, \tau_{ijs}\}_{i,j=1}^I$, productivity $\{T_{ig}, T_{is}\}$, and contribution shares to the global portfolio $\{\rho_i\}$; time-varying structural parameters for every country $\{\lambda_{ik}, \gamma_{ikn}\}$; and time-invariant structural parameters $\{\sigma, \varepsilon_k, \omega_k, \theta_k\}_{k=g,s}$ as follows.

Definition 1. A *competitive equilibrium* is a sequence of output and factor prices $\{w_i, P_{ig}, P_{is}, P_i\}_{i=1}^I$, allocations $\{L_{ig}, L_{is}, M_{igg}, M_{igs}, M_{isg}, M_{iss}, Q_{ig}, Q_{is}, Y_{ig}, Y_{is}, C_{ig}, C_{is}, C_i\}_{i=1}^I$, transfers from the global portfolio, R , and trade shares $\{\pi_{ijg}, \pi_{ijs}\}_{i,j=1}^I$, such that each condition in table 1 holds.

3.5 Gains from Trade

In existing multisector trade models that feature homothetic preferences like CES and Cobb-Douglas, the gains from trade are equivalent to changes in real income or aggregate consumption from autarky to trade, and as shown by Costinot and Rodríguez-Clare (2014), there is typically a sufficient statistic for calculating this change. Unfortunately, the structure of the non-homothetic CES preferences used in our model means that changes in real income or consumption are not appropriate for quantitative applications, complicating the welfare calculation. Instead, we rely on the concept of the *equivalent variation* to compare the equilibrium allocation in both trade and autarky.¹⁶ This measure is invariant to any monotonic transformation of the utility function, i.e., the normalization of ε_g .

The calculation begins by considering two equilibrium states characterized by different real incomes and sector prices: the trade equilibrium summarized by (w_i, P_{ig}, P_{is}) and the autarky equilibrium summarized by $(w_i^{AUT}, P_{ig}^{AUT}, P_{is}^{AUT})$. These are the equilibrium outcomes of the baseline model and a counterfactual where the trade costs are set to a prohibitively high level, respectively. The welfare changes between the trade regime and the autarky regime for any country i can be expressed using equivalent variation. Equivalent variation measures the amount of extra income that country i requires in order to obtain the utility level possible in autarky, while facing prices in the trade regime:

$$EV_i = e(P_{ig}, P_{is}, u^{AUT}) - e(P_{ig}, P_{is}, u),$$

where $e(\cdot)$ is the expenditure function, and $u = v(w_i, P_{ig}, P_{is})$ and $u^{AUT} = v(w_i^{AUT}, P_{ig}^{AUT}, P_{is}^{AUT})$ are the utility levels for the indirect utility function $v(\cdot)$ evaluated at the two equilibria. To match the literature, we define

¹⁶Nigai (2016) uses the equivalent variation concept to measure heterogeneous gains from trade across households with non-homothetic preferences.

the gains from trade as the lost income that comes from moving to the lower utility level in autarky, taken as a fraction of initial income:

$$GFT_i = \frac{e(P_{ig}, P_{is}, u) - e(P_{ig}, P_{is}, u^{AUT})}{e(P_{ig}, P_{is}, u)}. \quad (12)$$

To take this definition to our model, we derive the indirect utility function and expenditure functions. For a country i with a per-capita income level w_i and prices (P_{ig}, P_{is}) under the trade regime, equilibrium conditions (D2) and (D3) in Table 1 imply the indirect utility (per capita) function $v(w_i, P_{ig}, P_{is})$ takes the implicit form:

$$\sum_k \omega_k P_{ik}^{1-\sigma} v(w_i, P_{ig}, P_{is})^{(1-\sigma)\varepsilon_k} = w_i^{1-\sigma}. \quad (13)$$

In turn, the expenditure minimization problem yields the expenditure function $e(P_{ig}, P_{is}, u)$ as:

$$e(P_g, P_s, u) = \left(\sum_k \omega_k P_{ik}^{1-\sigma} u^{(1-\sigma)\varepsilon_k} \right)^{\frac{1}{1-\sigma}}. \quad (14)$$

Plugging in the optimal consumption bundle consistent with the model solution as well as a version solved with infinite trade costs will thus permit the calculation of welfare gains.

Importantly, this definition nests the standard method of calculating welfare gains. Since setting $\varepsilon_k = 1$ for all k produces standard CES preferences, gains from trade under equation (12) are:

$$GFT_i = 1 - \left(\frac{\sum_k \omega_k P_{ik}^{1-\sigma}}{\sum_k \omega_k P_{ik}^{AUT 1-\sigma}} \right)^{\frac{1}{1-\sigma}} \frac{w_i^{AUT}}{w_i} = 1 - \frac{w_i^{AUT} / P_i^{AUT}}{w_i / P_i}. \quad (15)$$

Thus, the gains from trade for homothetic preferences are simply the change in real income from trade to autarky.

4 Calibration and Solution

To quantify the role of structural change in global trade flows, we calibrate the exogenous processes and the parameters in the model to match data from 26 countries plus a rest-of-the-world aggregate over 1970–2015. Preference parameters, $(\sigma, \varepsilon_g, \varepsilon_s, \omega_g, \omega_s)$, are estimated using panel data on sectoral prices, expenditure shares, and total expenditure per worker. The trade elasticity, θ_k , is taken from the literature. Trade imbalances, ρ_{it} , and labor endowment, L_{it} , are set to match data on trade deficits and total number of employees. The production coefficients λ_{ikt} and γ_{ikt} are constructed using the input-output data. Processes for sectoral trade costs, τ_{ijkt} , and productivity, T_{ikt} , are constructed to match data on bilateral trade shares and expenditure shares.

We discuss the calibration procedures together with the corresponding data sources in the next three subsections. With the calibrated parameters, we can solve the baseline model fully in levels for each year

$t = 1970, \dots, 2015$. In the following section, we check the model fit by comparing untargeted moments in the model with those in the data.

4.1 Preference parameters

To estimate the preferences parameters $(\sigma, \varepsilon_g, \varepsilon_s, \omega_g, \omega_s)$, we collect data on sectoral prices P_{ik} , sectoral expenditure $P_{ik}C_{ik} = E_{ik}$, aggregate expenditure, $P_iC_i = \sum_k E_{ik}$, and employment levels, L_i . The construction of sectoral expenditure data is discussed in section 2. For sectoral prices, we use the UN Main Aggregates Database and the GGDC Productivity Level Database “2005 Benchmark” from Inklaar and Timmer (2014). Aggregate employment comes from the Penn World Table and corresponds to “number of persons engaged”. Details of data construction and data sources are provided in Appendix A.2.

Before structurally estimating the preference parameters, we illustrate the empirical importance of price and income effects. Taking logs of equation (7), and replacing sectoral expenditure shares e_{ik} with sectoral expenditure levels E_{ik} , gives a simple OLS regression equation:

$$\ln\left(\frac{E_{ist}}{E_{igt}}\right) = \text{constant} + (1 - \sigma^{OLS}) \ln\left(\frac{P_{ist}}{P_{igt}}\right) + (1 - \sigma^{OLS})(\varepsilon_s^{OLS} - \varepsilon_g^{OLS}) \ln\left(\frac{C_{it}}{L_{it}}\right).$$

Holding fixed variation in the last term (aggregate consumption per worker – income effects), the extent that relative expenditure shares move with relative prices helps us identify the price elasticity σ^{OLS} . Holding relative prices fixed, the extent that relative expenditure shares move with aggregate consumption per worker helps us identify the income elasticity ε_s^{OLS} (we maintain the normalization that $\varepsilon_g^{OLS} = 1$). Setting the sector weights ω_{ik} , which are encompassed by the constant term, to be constant across countries and over time allows us to exploit both the cross-section and time-series variation.

Data on utility C_{it} (which we have been calling consumption) is not observable. We proxy for this by using data on real domestic absorption, measured at current purchasing power parities.¹⁷ The results are reported in Table 2. We also report the results when either the relative prices or consumption per worker are used alone in the regression (further constraining either $\varepsilon_s^{OLS} = 1$ or $\sigma^{OLS} = 1$). As shown in the third and fourth columns of Table 2, both relative prices and the income effect are important in accounting for variation in expenditure shares, though the marginal importance of the income effect is greater. Moreover, the explanatory power is much higher when including both relative prices and income, compared to including only one at a time. The R^2 is 72% with both variables, while it is only 41% with relative prices alone and 66% with relative income per worker alone. Thus, nonhomothetic preferences are crucial in accounting for the observed expenditure shares.

Notice that including only the relative price on the right-hand side results in a negative price elasticity, which is outside of the range of theoretically plausible estimates. This omitted variable bias stems from the underlying pattern in the data that in all countries, expenditure shifted towards services—the sector which became relatively more expensive. On the other hand, including only aggregate expenditure implies a larger income elasticity compared to including both aggregate expenditures and relative prices, since the relative price of services and consumption per worker both rise with development.

¹⁷This is the variable *cda* in the Penn World Table.

Table 2: Regression results

Variable	Prices & income	Income only	Prices only
σ^{OLS}	0.36*		-0.44
	(0.20)		(0.32)
$\varepsilon_s^{OLS} - \varepsilon_g^{OLS}$	0.47***	0.58***	
	(0.05)	(0.05)	
constant	-4.13***	-5.32***	0.89***
	(0.56)	(0.57)	(0.06)
N	846	846	846
Adjusted R^2	0.72	0.66	0.41

Note: Standard errors clustered at the country level are in parentheses, with *** and * denoting a 99% and 90% significance level, respectively.

Though this exercise demonstrates the importance of both the price and income effect channels, the data shortcut for the unobservable model object C_i is unpalatable. Thus, we structurally estimate the elasticities of both channels. As above, we choose the elasticities that minimize the distance between the observed sectoral expenditures and those implied by the model. However, we rely on the model to infer utility, or real aggregate consumption C_i , since there is no direct empirical counterpart to it. Multiplying the price index in equation (D2) by $\frac{C_i}{L_i}$ on both sides gives equation (19) below, which specifies the expenditure per capita, $\frac{E_{ig}+E_{is}}{L_{it}}$, required to obtain utility C_{it} under sectoral prices P_{ikt} . This is the classic expenditure function, derived above in equation (14). Since the expenditures per capita and sectoral prices are observed, we can infer the real aggregate consumption (utility) implicitly. Formally, the preference parameters are estimated by solving a constrained nonlinear least squares problem:

$$\min_{(\sigma, \varepsilon_s)} \sum_{t=1}^T \sum_{i=1}^I \left(\left(\frac{\omega_s}{\omega_g} \right) \left(\frac{\hat{P}_{ist}}{\hat{P}_{igt}} \right)^{1-\sigma} \left(\frac{C_{it}}{\hat{L}_{it}} \right)^{(1-\sigma)(\varepsilon_s - \varepsilon_g)} - \left(\frac{\hat{E}_{ist}}{\hat{E}_{igt}} \right) \right)^2 \quad (16)$$

$$\text{s.t. } \varepsilon_g = 1 \quad (17)$$

$$\omega_g + \omega_s = 1 \quad (18)$$

$$\frac{\hat{E}_{igt} + \hat{E}_{ist}}{\hat{L}_{it}} = \left(\sum_{k=g,s} \omega_k \hat{P}_{ikt}^{1-\sigma} \left(\frac{C_{it}}{\hat{L}_{it}} \right)^{(1-\sigma)\varepsilon_k} \right)^{\frac{1}{1-\sigma}}, \forall (i,t), \quad (19)$$

where observables are denoted using a “hat”. Equation (17) normalizes the income elasticity for goods to one. As in Comin et al. (2018), this normalization has no influence on equilibrium allocations; instead different values of ε_g reflect monotonic transformations of the same utility function. Equation (18) normalizes the sum of ω_g and ω_s to one without the loss of generality. We have flexibility in setting ω_s by changing the units of relative prices. Specifically, we normalize relative prices such that $\omega_s = 0.58$, the global expenditure share on services in 1970.

We solve the minimization problem as follows. Start by making a guess for (σ, ε_s) . Given this guess, we exploit aggregate expenditure and sectoral price data to impute the aggregate consumption index, C_{it} , for each country in every year using constraint (19), which is a simple nonlinear equation with one unknown.

Given the imputed consumption indexes we then exploit data on sectoral prices and expenditures and use nonlinear least squares on the objective function (16) to obtain updated estimates of (σ, ε_s) . With the updated estimates of the preference parameters in hand, we revisit equation (19) and impute updated consumption indexes and, in turn, obtain new estimates of the preference parameters by minimizing (16). We continue the procedure until converging to a fixed point in the preference parameter space. The result of the estimation delivers $\sigma = 0.16$ and $\varepsilon_s = 1.62$. Implicitly we also obtain estimates of the aggregate consumption index, C_{it} , which has no direct empirical counterpart. This object will be used later on in order to calibrate productivity levels in an internally consistent manner.

Our estimates of the preference elasticities are consistent with the regularity conditions (i) from Section 3.1 required, namely $0 < \sigma < 1$ and $\varepsilon_k > 0$ for all k . While estimation of demand functions may suffer from endogeneity problems, our estimates are remarkably similar to those obtained by Comin et al. (2018) using Consumer Expenditure Survey data at the household level, including instrumenting for household consumption. In addition, their three-sector estimates using cross-country panel data on sectoral employment and value added shares are broadly consistent with our two-sector estimates using cross-country panel data on expenditure shares.¹⁸ Specifically, goods and services are complements, and services have a higher income elasticity than goods. The structural estimates are also within a reasonable range of the reduced form estimates in table 2.

4.2 Parameters directly from the literature and the data

We set the dispersion parameter of productivity draws in the goods sector, θ_g , at 4, following Simonovska and Waugh (2014). There is no reliable estimate of the trade elasticity for services, so we set $\theta_s = 4$ as well. We conduct sensitivity analysis for a smaller θ_s of 2.5 in Appendix A.4; the main results are robust. The elasticity of substitution between varieties in the composite good, η , plays no quantitative role in the model other than satisfying $1 + (1 - \eta)/\theta > 0$; we set this value at 2. The upper panel of Table 3 summarizes these parameter values that are common across countries and constant over time.

The country-specific, time-varying production parameters γ_{iknt} and λ_{ikt} are constructed by condensing multi-sector, input-output tables to a two-sector input-output construct. Specifically, λ_{ikt} is the ratio of value added to total production in sector k , while the γ_{iknt} is the share of sector k intermediate spending that is sourced from sector n . The World Input-Output Database (WIOD) provides country-level data annually from 1995-2011. Prior to 1995, we make use of country-specific input-output tables for 11 countries from multiple sources, covering various years depending on the country, resulting in an unbalanced panel of input-output data from 1970-2011 (see Appendix A.2 for more details). We impute the shares for missing country-year observations by (i) estimating a relationship between each share and GDP per worker in our unbalanced panel and (ii) using the estimates and observed GDP per worker to fill in the missing values.

While these production shares vary significantly across countries, they change only mildly over time. Moreover, there are notable patterns that hold across countries. First, production of services is more value-added intensive than production of goods. The lower panel of table 3 indicates that, on average, 61 percent

¹⁸Comin et al. (2018) estimate a substitution elasticity of 0.50 and $\varepsilon_s - \varepsilon_m$ of 0.21 with data on agriculture, manufacturing, and services value added and employment. Their data goes back to 1947, and includes 37 countries.

of total service production compensates value-added factors, compared to 39 percent in goods. Second, inputs from goods sectors account for 70 of intermediate expenditures by the goods sector. That is, goods production is goods-intensive. Similarly, services production is service intensive: inputs from the service sector account for 64 percent of intermediate expenditures by the service sector. Still, cross-sector linkages are relatively strong: roughly one-third of intermediate inputs in each sector is sourced from the other sector.

Table 3: Parameter values

Preference parameters		
σ	Elasticity of substitution b/w sectors	0.16
ε_g	Elasticity of income in goods	1
ε_s	Elasticity of income in services	1.62
ω_g	Preference weight on goods	0.42
ω_s	Preference weight on services	0.58
Production parameters		
θ_g	Trade elasticity in goods sector	4
θ_s	Trade elasticity in service sector	4
η	Elasticity of substitution b/w varieties in composite good	2
Input-Output parameters (cross-country, cross-time averages)		
λ_g	Ratio of value added to gross output in goods	0.39
λ_s	Ratio of value added to gross output in goods	0.61
γ_{gg}	Good's share in intermediates used by goods sector	0.70
γ_{sg}	Good's share in intermediates used by service sector	0.36

The parameters ρ_{it} are calibrated to match each country's ratio of net exports to GDP. In the model, the ratio of net exports to GDP in country i at time t is $\frac{\rho_{it}w_{it}L_{it}-R_tL_{it}}{w_{it}L_{it}}$. In the calibration, we let $R_t = 0$ and simply set $\rho_{it} = \frac{NX_{it}}{GDP_{it}}$. So long as net exports sum to zero across countries (as it does in our data), the global portfolio is balanced. In the counterfactual analysis, the endogenous term R_t adjusts to ensure that the global portfolio balances period-by-period: $R_t \sum_{i=1}^I L_{it} = \sum_{i=1}^I \rho_{it} w_{it} L_{it}$.

4.3 Technology and Trade Costs

We recover the exogenous productivity terms, T_{ik} , and trade costs, τ_{ijk} , by exploiting structural relationships from our model in order to match data on sectoral expenditures and bilateral trade flows in each country and every year. Our procedure is similar to that of Świącki (2017), but incorporates input-output linkages as in Levchenko and Zhang (2016) and Sposi (2019). By explicitly making use of the observed input-output linkages, our procedure also implies that we simultaneously match sectoral value added.

Two key structural relationships provide identification for productivity and trade costs:

$$T_{ik}^{\lambda_{ik}} = \frac{v_{ik}}{\Gamma_k^{-1} P_{ik} (\pi_{iik})^{-\frac{1}{\theta_k}}}, \quad (20)$$

$$\tau_{ijk} = \left(\frac{\pi_{ijk}}{\pi_{jjk}} \right)^{-\frac{1}{\theta_k}} \left(\frac{P_{ik}}{P_{jk}} \right). \quad (21)$$

Both structural relationships are derived by manipulating equations (10) and (11). Measurement of sectoral productivity takes into account differences between imputed input costs and imputed output prices. Holding fixed the unit costs of inputs, the model assigns a high productivity to a country with a low price, meaning that inputs are converted to output at an efficient rate. It also takes into account the home trade share, which reflects the selection effect in Ricardian trade models.

Measurement of the trade costs takes into account imputed price differences and the bilateral trade shares. Holding fixed the imputed price difference between countries i and j , if country i imports a large share from country j relative to what j sources from itself, the inferred trade barrier is low.

Note that measuring trade costs and productivity require sectoral prices. One option is to use sectoral prices directly. However, we have price data for only about two-thirds of our sample. Moreover, prices are measured with error and noise, particularly for services. Instead, we impute our own measure of sectoral prices by inverting the demand curve and taking into account observed sectoral expenditures and real aggregate consumption per capita. Specifically, given preference parameters, $(\omega_g, \omega_s, \sigma, \varepsilon_g, \varepsilon_s)$, sectoral expenditure shares, e_{ik} , labor endowment, L_i , and aggregate consumption, C_i (imputed in equation (19)), we invert the household's first-order condition (7) to recover model-implied price levels that support the observed expenditure shares.

These imputed prices are not the same as the prices observed in the data. Indeed, since preferences are identical across countries, our model does not have enough degrees of freedom to match data on both sectoral prices and sectoral expenditures simultaneously. We choose to match expenditures since they are of first order interest to our question and expenditures shares have higher quality data. Although the sectoral expenditure data are not directly observed for all countries prior to 1995, our procedure for constructing the missing data provides a little more confidence as it only requires imputing some of the input-output coefficients, which are not very noisy to begin with, and the variation that they do display can be explained well by income levels and country fixed effects. Moreover, we have complete measures of sectoral value added and sectoral net exports. Given an input-output structure and this information, sectoral expenditures are pinned down by an accounting identity.

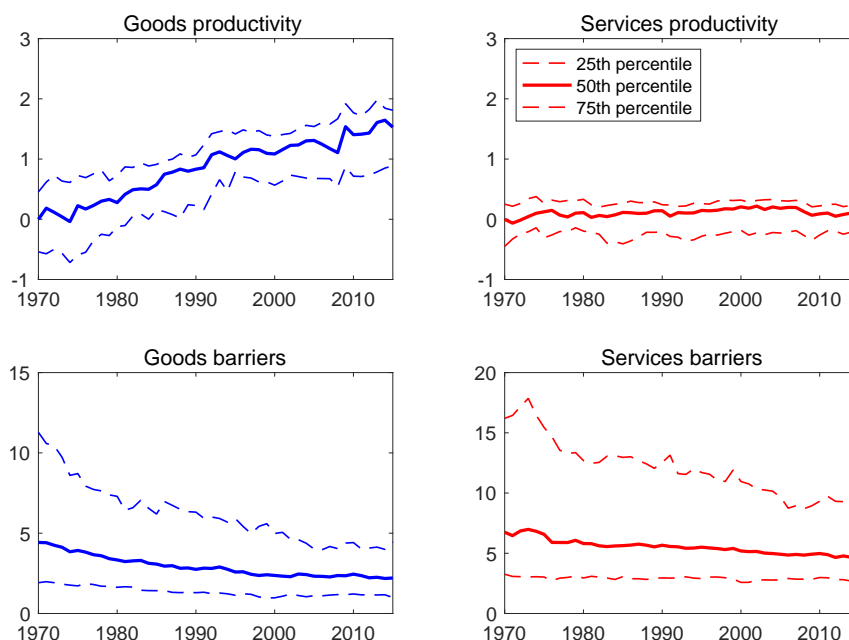
These imputed prices serve as an intermediate ingredient through which productivity and trade costs are recovered, which is an alternative to estimating prices off of trade flows using a gravity specification as in Levchenko and Zhang (2016). The advantage of our procedure is that it guarantees that we match the observed expenditure shares, trade flows, and production.¹⁹

In addition to sectoral prices, we also need data on trade shares π_{ijk} and wages w_i , where wages are used to construct unit costs. With these data series, we use equations (20) and (21) to compute the sectoral productivity and trade costs, which are illustrated in figure 3. The upper panel plots the log of the fundamental productivity levels, T_{ik} , of the median country (solid lines), the 25th percentile country and the 75th percentile country (dashed lines) at the sector level in each year. As shown in the figure, productivity grows faster in goods than in services. Specifically, over the sample period, on average the median fundamental productivity series grows by 3.5% per year in the goods sector and by 0.3% in the services sector. The cross-country productivity dispersion is fairly stable over time in both sectors.

¹⁹The main results of the paper hold if one instead chooses to target sectoral prices rather than sectoral expenditures.

To gauge how reasonable these fundamental productivity series are, we compare the model implied labor productivity with that in the data. In a model with trade selection, the model-implied labor productivity is in general higher than the fundamental productivity. The model-implied sectoral labor productivity is consistent with the data. On average, median labor productivity for goods grows by 4.7% per year in the model, compared to 4% in the data. For services, the median labor productivity grows by 1.7% per year in the model, compared to 1.5% in the data. The calibrated model slightly overshoots in sectoral productivity growth, because the calibration does not target sectoral employment or sectoral prices in the data.

Figure 3: Calibrated global productivity and trade costs



Note: Productivity plots in logs for each sector are normalized by the 1970 value of the median series of that sector. Trade barrier plots report the net trade cost, $\tau - 1$.

The lower panel of figure 3 plots the net trade barriers, $\tau_{ijk} - 1$, for goods and services over time. Again, the solid line is for the median level, and the dashed lines are for the 25th and the 75th percentiles. As illustrated in the figure, trade costs for both goods and services decline over time, and trade costs in services are generally higher than those in goods. Also, the cross-country dispersion of the trade barriers declines substantially in both sectors over time. Furthermore, the trade barriers decline faster in the goods sector than in the services sector. Over the sample period, the median trade barriers decline by about 50% (from 4.4 to 2.2) for goods, but only by about 30% (from 6.7 to 4.7) for services. The magnitudes of our estimated trade barriers for the goods sector are similar to those in Levchenko and Zhang (2016), and other papers in the literature.

4.4 Model Fit

Our calibration procedure ensures that the model fits data on sectoral value added, sectoral gross output, sectoral absorption, sectoral bilateral trade flows, sectoral expenditures, input-output linkages, and total employment. We now check the fit of the model on observed moments that are not targeted directly by the calibration.

The first two are the sectoral prices, which are illustrated in the left and middle panels of figure 4. Each point corresponds to the price for a country-year with the model value on the y-axis and the data value on the x-axis. Of course, the points are limited by the incomplete coverage of the price data. All prices are taken relative to the U.S. in 2011. The prices fit the data reasonably well; the correlation between the model and the data is 0.91 for goods prices and is 0.96 for services prices.

The second moment we check is the sectoral employment share. In the right panel of figure 4, we plot the services employment shares in the data against those implied by the baseline model. The baseline model succeeds in replicating the employment shares across sectors for all sample countries over time. The correlation between the model and the data is 0.92.

The calibration successfully matches the targeted moments in the data. Moreover, the calibrated model fits well on the above data moments that are not directly targeted by the calibration. Thus, the baseline model closely maps into the relevant data for our analysis and serves as the baseline for the counterfactual analysis in the next section.

Figure 4: Relative prices and services labor shares: model versus data



Note: Prices are in logs, normalized relative to the U.S. in 2011. Employment share depicts the number of workers engaged in services as a share of the entire workforce.

5 Model-based Counterfactuals

This section quantitatively assesses the dampening effect of structural change on global trade volumes in the past four decades by conducting counterfactuals using the calibrated model. We find that the magnitude of this dampening effect from structural change is half as large as the boosting effect of declining trade costs on openness in this period. We also highlight the importance of structural change on the measurement of

the gains from trade. Finally, we project structural change into the future and consider the implications for openness and the policy implications for liberalizing trade in either goods or in services on welfare.

5.1 Global Trade in the Absence of Structural Change

To examine the implications on global trade flows from structural change, we construct a counterfactual in which structural change is absent by restricting expenditure shares to be constant over time. This provides the closest model-based analogue to our reduced-form counterfactual in Section 2. To do so, we assume that the preferences in the counterfactual are given by:

$$C_i = \prod_{k \in \{g,s\}} C_{ik}^{\omega'_{ik}}. \quad (22)$$

With the Cobb-Douglas specification, the income elasticities are one for both sectors and the substitution elasticity is also one across the two sectors. Consequently, expenditure shares across sectors are invariant to economic conditions and constant over time. That is, we choose values for $\omega'_{ik} = e_{ik0}$ so that in 1970, the sectoral expenditure shares are identical to those in the baseline model.

All underlying processes in the counterfactual are identical to those in the baseline. Specifically in the counterfactual, we assume all other parameters and time varying processes for T_{ik} , τ_{ijk} , and L_i are unchanged from the baseline, except that the preference parameters $\{\sigma, \varepsilon_g, \varepsilon_s, \omega_k\}$ in the baseline are set to $\{1, 1, 1, \omega'_{ik}\}$ in the counterfactual experiment. We compute the equilibrium for the counterfactual experiment and analyze how the absence of structural change impacts global trade flows. Our solution procedure is based on Alvarez and Lucas (2007). Start with an initial guess for the vector of wages. Given wages, recover all remaining prices and quantities across countries using optimality conditions and market clearing conditions, excluding the trade balance condition. Then use departures from the trade balance condition to update the wages. Iterate on wages until the trade balance condition holds. Details are available in appendix A.3.

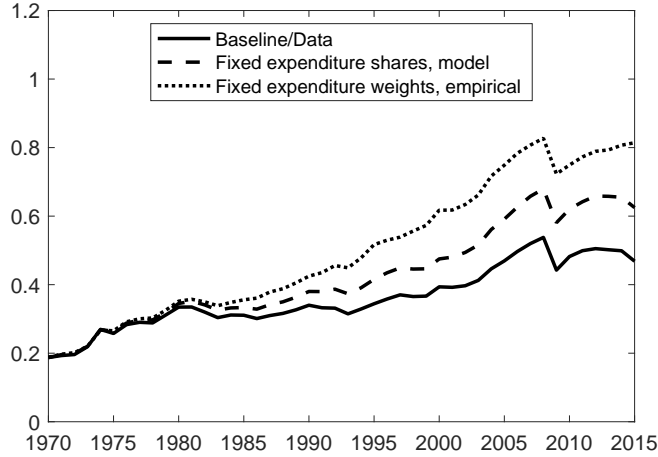
5.1.1 Model counterfactual results

We next present the implications for global trade flows.²⁰ Figure 5 compares openness between the baseline model (solid line), model counterfactual (dashed line), and reduced-form counterfactual (dotted line). In both counterfactuals, global trade would have been much higher had structural change not occurred. By 2015, the reduced-form counterfactual puts openness at 81 percent while the model counterfactual puts it at 62 percent, compared with 47 percent in the data/baseline. The difference between the two counterfactuals peaks in 2015 and is driven by the endogenous changes to sectoral openness generated by the model.²¹

²⁰As shown in figure A.5, the goods share of total expenditure falls from about 42 in 1970 percent to 21 percent in 2015 in both the data and the baseline model. In the counterfactual, the goods expenditure share is held fixed at the 1970 values, country-by-country. The slight rise since 2002 is driven by the increasing weight of China and India in the world economy, both of which have larger expenditure shares in goods compared to the developed countries.

²¹Appendix A.6 shows structural change and the model-based counterfactual for each country in figure A.2 and A.3 respectively, as well as a decomposition of each country's contribution to the aggregate counterfactual in table A.2 for 2015.

Figure 5: Openness: baseline and counterfactuals



5.1.2 Quantitative mechanisms

A key benefit of the general equilibrium structure is its ability to deliver an alternate path for sectoral openness that responds to the same forces that drive structural change. The left two panels of Figure 6 compares sectoral openness in the model counterfactual with the observed sectoral openness. In the model counterfactual, goods openness (the ratio of goods trade to goods expenditure) is about 50 percentage points lower relative to the baseline in 2015, while services openness is about 2 percentage points higher.

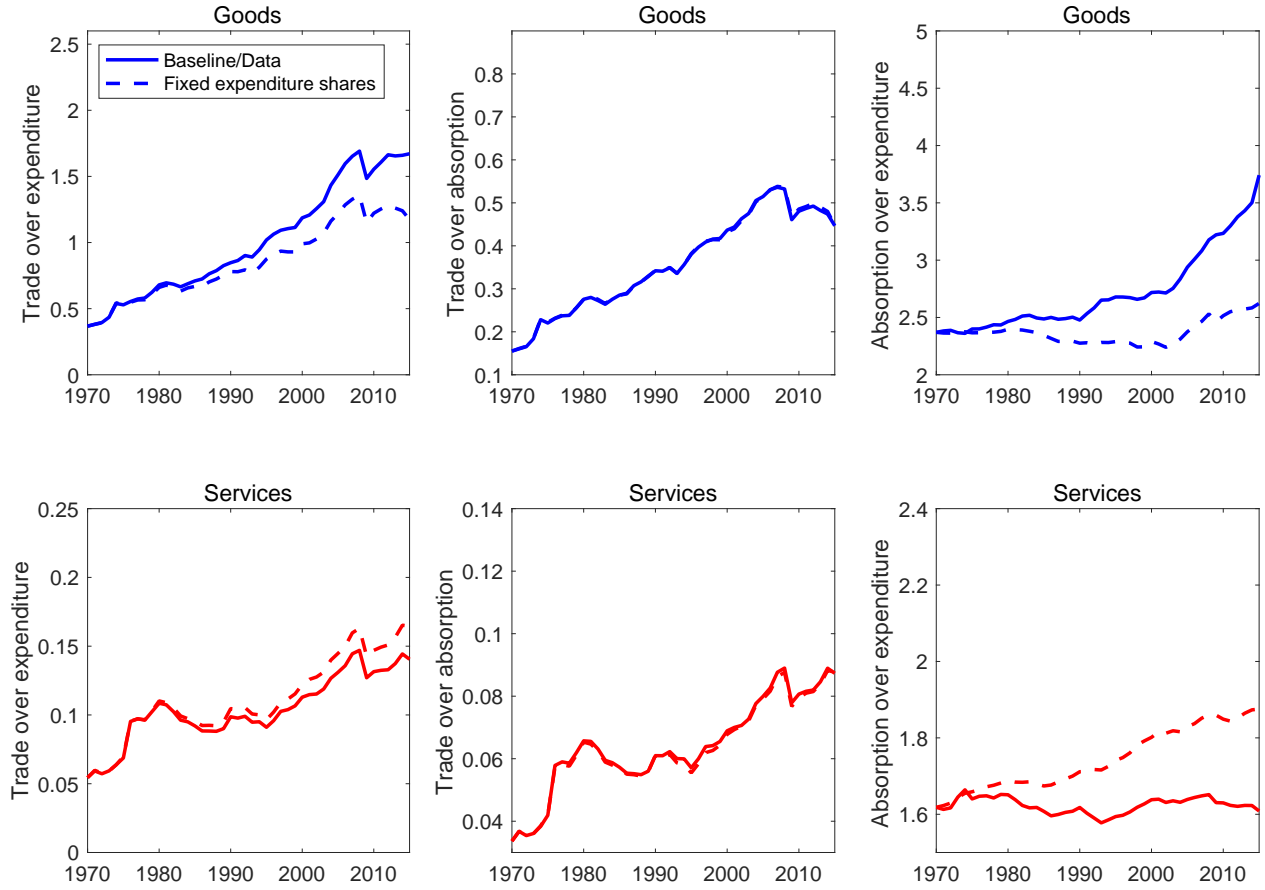
To understand how sectoral openness endogenously responds to changes in expenditure shares, we decompose sectoral trade openness into two terms: (i) the ratio of trade to absorption and (ii) the ratio of absorption to expenditure:

$$\frac{Trade_{kt}}{Exp_{kt}} = \left(\frac{Trade_{kt}}{Abs_{kt}} \right) \times \left(\frac{Abs_{kt}}{Exp_{kt}} \right). \quad (23)$$

These two terms correspond to two potential channels of bias inherent in the reduced-form counterfactual. Through endogenous general equilibrium effects, changing sectoral demand might change the relative wages across countries, and thus the ratio of trade to absorption, which is captured by the first term. In the model, at the country level, the first term resembles $1 - \pi_{ik}$ for each country i and sector k .²² Also, changing the sectoral demand shares might affect the ratio of absorption to expenditure through input-output linkages, as captured by the second term.

²²Sectoral imports over expenditure is exactly equal to $1 - \pi_{ik}$. Sectoral exports differ, but quantitatively they are highly correlated with sectoral imports across countries.

Figure 6: Sectoral openness: baseline and counterfactual with decomposition



We now quantify the bias of each channel. The ratios of trade to absorption in each sector are almost identical in the baseline and in the counterfactual, as shown in the center panels of figure 6. Recall the expression of π_{ik} in equation (S1) in table 1. Since the productivity and the trade cost processes are exogenous and thus unchanged, changing expenditure patterns affect the trade-over-absorption ratios only through their impact on relative wages across countries. It turns out that the general equilibrium effect on relative wages is quantitatively small in the model counterfactual. Thus, the share of each country's absorption that is sourced from abroad in each sector barely changes from the baseline to the counterfactual.

The primary reason why sectoral trade openness in the model counterfactual differs from that in the baseline is due to differences in the ratio of absorption to expenditure, as shown in the right panels of figure 6. Compared with the baseline, the counterfactual ratios of absorption to expenditure in the counterfactual rise by less over time for the goods sector, but rise by more over time for the services sector. Using the expression for sectoral absorption in equation (S6) of table 1, we can write the sectoral ratio of absorption to expenditure as:

$$\frac{Q_{ig}}{C_{ig}} = \frac{C_{ig} + M_{igg} + M_{isg}}{C_{ig}}, \quad \frac{Q_{is}}{C_{is}} = \frac{C_{is} + M_{igs} + M_{iss}}{C_{is}},$$

where sectoral absorption equals final plus intermediate demand for the sectoral composite good. In order

to counterfactually increase consumption of goods, C_{ig} , intermediates must be sourced from both sectors, implying that M_{igg} and M_{igs} rise, since these are directly used to produce the greater demand for goods consumption. At the same time, derived demand for M_{isg} and M_{iss} decline in response to a decline in C_{is} . Consequently, absorption rises by less than expenditure in the goods sector, while absorption declines by less than expenditure in the services sector, implying lower $\frac{Q_{ig}}{C_{ig}}$ and higher $\frac{Q_{is}}{C_{is}}$ in the model counterfactual compared with the baseline.

Going back to figure 6, we conclude that although services trade openness goes up, goods openness decreases sufficiently to imply a lower overall trade openness in the model counterfactual than in the empirical counterfactual. This major bias of the reduced-form counterfactual in predicting global trade openness in the absence of structural change comes from ignoring input-output linkages across sectors. To confirm the importance of cross-sector input-output linkages, we recalibrate the baseline model with no cross-sector input-output linkages ($\gamma_{gg} = \gamma_{ss} = 1$). In this alternative model, we conduct a similar fixed expenditure counterfactual, and find that the absence of structural change has essentially no impact on sectoral trade openness.

5.1.3 Decomposing income versus substitution effects

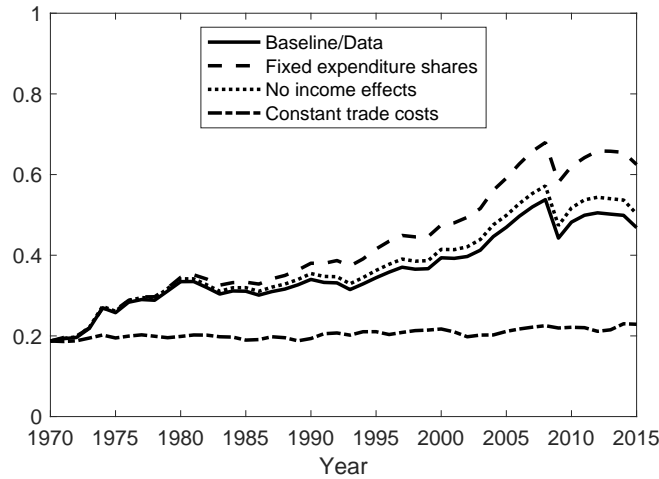
The literature on structural change has established two key mechanisms: income effects and substitution effects. Boppart (2014) provides the first model that incorporates both income and substitution effects to generate structural transformation along a balanced growth path. Herrendorf et al. (2013) demonstrate that when structural change is defined over final expenditures instead of value added, as it is in our paper, then income effects play a nontrivial role relative to substitution effects.

We use our model to evaluate the relative importance of each effect in shaping global trade flows. In our model counterfactual, we set $\varepsilon_s = 1$ so that preferences are homothetic, i.e., income elasticity of demand in each sector equals 1.²³ By comparing global trade openness implied by this experiment with that of the counterfactual with both effects shut off, we can see to what extent the income effect drives our results. Alternatively, the comparison will illustrate the power of the substitution effect alone.

Figure 7 plots the world ratio of trade to expenditure implied by our model counterfactual without the income effect, depicted with the dotted line. For comparison, we also plot trade openness in the data with the solid line and the one implied by our model counterfactual without the income and substitution effect with the dashed line. As can be seen in the figure, the model that shuts down the income effect leads to a ratio of trade to expenditure that is only about 3 percentage points higher than in the data, or about one-fourth of the difference between the data and the fixed-expenditure-shares counterfactual. Thus, the substitution effect exerts a greater influence than the income effect in terms of attenuating the growth in international trade.

²³We adjust the preferences weights, ω_{ik} , so that in 1970 the sectoral expenditures are identical to those in the baseline model.

Figure 7: Openness: no-income-effect and constant trade cost counterfactuals



5.2 Global Trade in the Absence of Declining Trade Costs

Arguably, declining trade costs are the most common factor attributed to the rise in global openness. Indeed, the past few decades have witnessed drastic reductions in both shipping costs and tariffs. To examine the role of declining trade barriers, consider a counterfactual in the model where trade barriers are held at their 1970 levels. The resulting trade openness is illustrated by the relatively flat line in figure 7. In this world, the global ratio of trade to expenditure grows by only about 4 percentage points instead of about 30 over the sample period. That is, declining trade costs since 1970 have added about 26 percentage points to the ratio of trade to expenditures by 2015. Of course, trade costs in the baseline model are calculated as the residuals required to account for changes in trade not driven by technology or demand. As such, they incorporate a wide variety of economic forces, including tariff reductions, improvements in shipping technology, or even compositional changes in demand at a finer level of disaggregation than our goods and services distinction.

The constant-trade-cost counterfactual demonstrates the quantitative significance of structural change on global trade openness. As shown in figure 7, structural change has held back trade by roughly half the magnitude that reductions in trade costs have boosted trade over the past four decades. However, structural change has not received enough attention in the trade literature when accounting for the dynamics of openness. The following section shows that incorporating structural change also impacts one of the central themes in international trade: the measurement of the gains from trade.

5.3 Importance of Structural Change for the Gains from Trade

In this subsection, we describe how accounting for structural change affects the measurement of the welfare gains from trade using equation (12). Table 4 summarizes the gains from trade from our full model. The left panel shows the distribution of the gains from trade in the first and last years. The gains from trade increase for the median country from about 4% in 1970 to over 10% in 2015 as trade integration rises over time. Countries differ substantially in their gains from trade. Countries at the 75th percentile gain about 5 percentage points more than those at the 25th percentile in 1970. This dispersion in the gains from trade

across countries is persistent over time. The difference between the 75th and 25th percentile is about 6 percentage points in 2015.

Table 4: Gains from trade

(a) Baseline model by percentile			(b) Alternative models at 50th percentile		
	1970	2015		1970	2015
25th percentile	2.2%	7.0%	Baseline	4.1%	10.3%
50th percentile	4.1%	10.3%	No income effects	4.1%	11.0%
75th percentile	6.8%	13.1%	Fixed expenditure shares	4.1%	12.7%

To highlight the role of structural change, the right panel of table 4 contrasts the gains from trade in the baseline with those from the no-income-effects and fixed-expenditure-shares counterfactuals. In the first counterfactual, we adopt homothetic CES preferences to shut down the income effects, while in the second, we use Cobb-Douglas preferences to shut down both the income and price effects. In 1970, the median gains from trade are practically the same in all three models since we set the initial sectoral expenditure shares to be the same across specifications. However, over time, the gains differ due to the presence (or lack) of structural change.

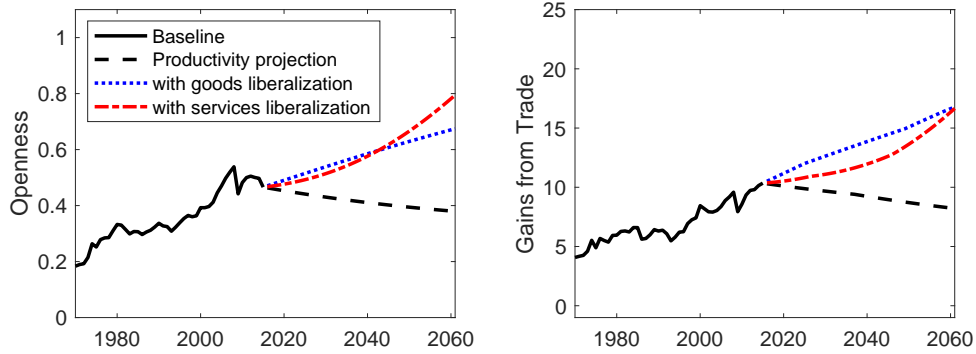
In our baseline model, the median gains from trade increase from 1970 to 2015 by about 6.2 percentage points. By contrast, in the fixed-expenditure-shares counterfactual, the gains increase by about 8.6 percentage points. Since trade integration occurs more prominently in the goods sector and structural change over time shifts expenditure away from goods, the measured gains from trade with structural change are 2.4 percentage points lower in 2015 than those without structural change. In addition, the income and price effects are both important in explaining the difference in the measured gains, while the price effect accounts for roughly two-thirds of the difference. Allowing for price effects alone (the no-income-effects counterfactual) accounts for 1.7 percentage points of the total 2.4-percentage-point difference between the baseline and the fixed-expenditure-shares counterfactual.

5.4 Projections

The recent slowdown in the growth of international trade has prompted careful consideration of the forces that might be restraining trade or no longer boosting it (IMF 2016b, Lewis and Monarch 2016). While structural change has not been a stronger drag on trade growth recently than it was in preceding decades, world trade as a share of total expenditure is likely to fall in the future absent additional trade cost reductions. We show this possibility quantitatively through the lens of our model. Specifically, we extrapolate our sample of countries for another 46 years, holding trade costs fixed at their 2015 value and letting goods and services productivity grow at their respective world average rates observed between 1970-2015.²⁴ Without additional factors boosting trade, our model implies that the ratio of trade to expenditure would fall from 47 percent in 2015 to 38 percent in 2061, shown as the productivity projection in the left panel of figure 8.

²⁴Goods fundamental productivity grows 3.5 percent and services grows 0.3 percent annually.

Figure 8: Projection of trade openness and gains from trade



This quantitative example highlights the importance of paying attention to the role of the prevalent process of structural change when considering trade flows. Without incorporating structural change into the model, the downward pattern in the ratio of trade to GDP from figure 8 would be attributed to rising trade costs. However, we find such a result even without any change in trade costs, as the effects of increased services consumption in a world without rapid trade growth materially affects the trajectory of global trade openness. In other words, it is perfectly within reason to imagine a decline in the ratio of trade to GDP, or even a decline in total trade flows, without any increased trade barriers. All that would be necessary is the combination of ongoing changes in services consumption along the lines of that seen in the past four decades with the continuation of current levels of trade barriers.

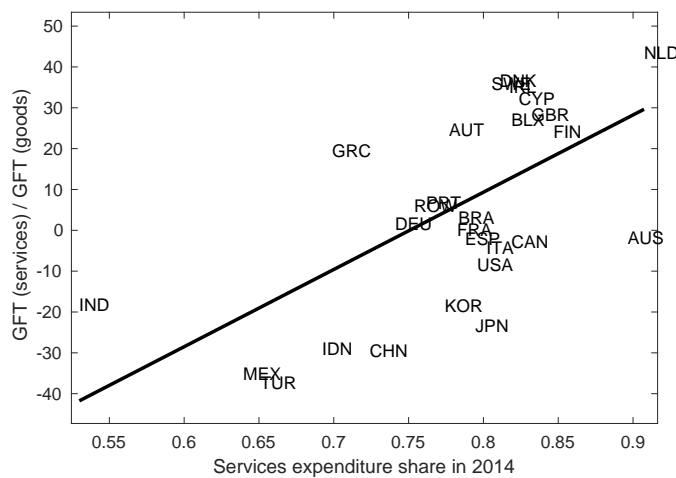
Trade policy in the presence of structural change We next evaluate the implications of different future trade policy scenarios on global openness and welfare in an environment of ongoing structural change. One scenario is continued reductions in trade costs in the goods sector over the next 46 years, occurring at the same rate as in the previous 46 years. The median goods trade cost has been declining by 1.5% per year, i.e., it is about halved from 1970 to 2015. The other scenario is instead reductions in services trade costs at the same rate of 1.5% per year, over the next 46 years. In both experiments, we keep the productivity growth process the same as in the productivity projection. We plot the implications of the two experiments also in figure 8. Openness rises in both cases with the declining trade costs boosting trade. However, in the long run, the boost in openness is larger for the declining services trade costs: the global trade to GDP ratio is predicted to be 79 percent with declining services trade costs and is 67 percent with declining goods trade costs. The reason is that initially services is less open than goods even with declining services trade costs, so as structural change shifts expenditures from goods to services aggregate openness does not increase very rapidly. However, as services become more open over time, when coupled with a shift in expenditures toward services, aggregate openness increases at an increasing rate.

The sector-biased trade policies have analogous implications for the overall gains from trade as they do for openness. The gains from trade under declining services trade costs are lower than those under declining goods trade costs initially, but would overtake as the time goes beyond the projection period.

We next illustrate that countries of varied income levels benefit differently from these two trade policy experiments: services liberalization versus goods liberalization. We plot the percentage difference in the

gains under services liberalization relative to the gains under goods liberalization against the 2014 services expenditure shares in figure 9. Clearly, differences in expenditure shares correlate with differences in gains from trade across countries. Indeed, all else equal, higher income countries with larger expenditure shares in services tend to benefit more from liberalizing services trade because services account for a systematically larger share of their consumption basket. Similarly, developing economies with low services expenditure shares will have greater gains from trade than advanced economies with trade liberalization in the goods sector.

Figure 9: Gains from trade comparison



Note: Vertical axis is the percentage difference in the gains from trade under services liberalization relative to that under goods liberalization. Gains are measured using equivalent variation in equation (12). The horizontal axis is the services expenditure share in 2014 (we use 2014 here since Ireland's sectoral expenditures in 2015 are tainted by unusual balance-of-payments adjustments).

6 Conclusion

We show that structural change, whereby an increasing share of world income is allocated away from goods toward services, has exerted a significant drag on global trade growth over the past four decades. In the absence of structural change, defined as fixing expenditure shares in goods and services at their 1970 level, the global ratio of trade to GDP would be 15 percentage points higher, or one-third higher, than in the data. This is about the half the magnitude that declining trade costs have contributed to the increase in global openness over the same period.

We quantify the effect of structural change on global trade using a general-equilibrium model incorporating comparative advantage, nonhomothetic preferences, and an input-output structure. The model highlights that sectoral openness is endogenous and that holding expenditures fixed at their 1970s levels would have resulted in lower goods openness through the presence of input-output linkages. On the other hand, had structural change not occurred, aggregate openness would have been higher, as goods openness is much greater than services openness. The model also implies that income effects alone account for about one-

third of the effect structural change has had on trade volumes and on the gains from trade. Indeed, in the face of prolonged structural change over the next few decades, reductions in trade costs in services would eventually yield greater gains than reductions in trade costs in goods, particularly for advanced economies.

Though structural change has been a significant drag on global trade growth over recent decades, it has not been a stronger drag since the global financial crisis. Instead, the recent slowdown in trade can be attributed to an absence of factors that have historically caused trade to rise relative to expenditure. Indeed, our paper demonstrates how unusual the 1990s and 2000s were: Even as the share of services in expenditure rose, international trade flows expanded, as input-output linkages proliferated across country borders. For the same reasons, however, our results indicate that world trade as a fraction of GDP may have peaked, and similar patterns of structural change projected into the future foreshadow declines in this measure of openness.

References

- Alvarez, Fernando and Robert E. Lucas**, “General equilibrium analysis of the Eaton–Kortum model of international trade,” *Journal of Monetary Economics*, September 2007, 54 (6), 1726–1768.
- Arkolakis, Costas, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare**, “The Elusive Pro-Competitive Effects of Trade,” *The Review of Economic Studies*, January 2019, 86 (1), 46–80.
- Baier, Scott L. and Jeffrey H. Bergstrand**, “The growth of world trade: tariffs, transport costs, and income similarity,” *Journal of international Economics*, 2001, 53 (1), 1–27.
- Betts, Caroline, Rahul Giri, and Rubina Verma**, “Trade, Reform, and Structural Transformation in South Korea,” *IMF Economic Review*, November 2017, 65 (4), 745–791.
- Boppart, Timo**, “Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences,” *Econometrica*, 2014, 82 (6), 2167–2196.
- Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte**, “The Impact of Regional and Sectoral Productivity Changes on the U.S. Economy,” *The Review of Economic Studies*, October 2018, 85 (4), 2042–2096.
- Caron, Justin, Thibault Fally, and James R. Markusen**, “International Trade Puzzles: A Solution Linking Production and Preferences *,” *The Quarterly Journal of Economics*, August 2014, 129 (3), 1501–1552.
- Comin, Diego, Danial Lashkari, and Martí Mestieri**, “Structural Change with Long-run Income and Price Effects,” 2018.
- Costinot, Arnaud and Andrés Rodríguez-Clare**, “Trade Theory with Numbers: Quantifying the Consequences of Globalization,” *Handbook of International Economics*, January 2014, 4, 197–261.
- Cravino, Javier and Sebastian Sotelo**, “Trade-Induced Structural Change and the Skill Premium,” *American Economic Journal: Macroeconomics*, July 2019, 11 (3), 289–326.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, September 2002, 70 (5), 1741–1779.
- Fajgelbaum, Pablo D. and Amit K. Khandelwal**, “Measuring the Unequal Gains from Trade,” *The Quarterly Journal of Economics*, August 2016, 131 (3), 1113–1180.
- Fieler, Ana Cecília**, “Nonhomotheticity and Bilateral Trade: Evidence and a Quantitative Explanation,” *Econometrica*, July 2011, 79 (4), 1069–1101.
- Hanoch, Giora**, “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 1975, 43 (3), 395–419.

- Herrendorf, Berthold, Richard Rogerson, and Ákos Valentinyi**, “Two Perspectives on Preferences and Structural Transformation,” *American Economic Review*, December 2013, 103 (7), 2752–2789.
- Hoekman, Bernard**, “Trade and growth – end of an era?,” in “The Global Trade Slowdown: A New Normal?,” VoxEU.org, June 2015, pp. 3–19.
- Hottman, Colin J. and Ryan Monarch**, “Estimating Unequal Gains across U.S. Consumers with Supplier Trade Data,” *International Finance Discussion Paper*, January 2018, 2018 (1220), 1–51.
- Imbs, Jean and Romain Wacziarg**, “Stages of Diversification,” *American Economic Review*, March 2003, 93 (1), 63–86.
- IMF, “Direction of Trade Statistics,” 2016.
- , “Global Trade: What’s Behind the Slowdown?,” in “World Economic Outlook,” Washington, D.C: International Monetary Fund, October 2016, pp. 63–119.
- Inklaar, Robert and Marcel P. Timmer**, “The Relative Price of Services,” *Review of Income and Wealth*, December 2014, 60 (4), 727–746.
- Irwin, Douglas A.**, “Long-run trends in world trade and income,” *World Trade Review*, 2002, 1 (1), 89–100.
- Kehoe, Timothy J., Kim J. Ruhl, and Joseph B. Steinberg**, “Global Imbalances and Structural Change in the United States,” *Journal of Political Economy*, December 2017.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie**, “Beyond Balanced Growth,” *The Review of Economic Studies*, 2001, 68 (4), 869–882.
- Krugman, Paul, Richard N. Cooper, and T. N. Srinivasan**, “Growing World Trade: Causes and Consequences,” *Brookings Papers on Economic Activity*, 1995, 1995 (1), 327–377.
- Levchenko, Andrei A. and Jing Zhang**, “The evolution of comparative advantage: Measurement and welfare implications,” *Journal of Monetary Economics*, April 2016, 78, 96–111.
- Lewis, Logan and Ryan Monarch**, “Causes of the Global Trade Slowdown,” November 2016.
- Markusen, James R.**, “Explaining the Volume of Trade: An Eclectic Approach,” *The American Economic Review*, 1986, 76 (5), 1002–1011.
- , “Putting per-capita income back into trade theory,” *Journal of International Economics*, July 2013, 90 (2), 255–265.
- Matsuyama, Kiminori**, “Structural Change in an Interdependent World: A Global View of Manufacturing Decline,” *Journal of the European Economic Association*, April 2009, 7 (2-3), 478–486.
- , “The home market effect and patterns of trade between rich and poor countries,” 2015.

- _____, “Engel’s Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade,” *Econometrica*, 2019, 87 (2), 497–528.
- McMillan, Margaret S. and Dani Rodrik**, “Globalization, Structural Change and Productivity Growth,” Working Paper 17143, National Bureau of Economic Research June 2011.
- Nigai, Sergey**, “On Measuring the Welfare Gains from Trade Under Consumer Heterogeneity,” *The Economic Journal*, 2016, 126 (593), 1193–1237.
- Rose, Andrew K.**, “Why Has Trade Grown Faster Than Income?,” *The Canadian Journal of Economics / Revue canadienne d’Economie*, 1991, 24 (2), 417–427.
- Simonovska, Ina**, “Income Differences and Prices of Tradables: Insights from an Online Retailer,” *The Review of Economic Studies*, October 2015, 82 (4), 1612–1656.
- _____ **and Michael E. Waugh**, “The elasticity of trade: Estimates and evidence,” *Journal of International Economics*, January 2014, 92 (1), 34–50.
- Sposi, Michael**, “Evolving comparative advantage, sectoral linkages, and structural change,” *Journal of Monetary Economics*, May 2019, 103, 75–87.
- Teignier, Marc**, “The role of trade in structural transformation,” *Journal of Development Economics*, January 2018, 130, 45–65.
- Timmer, Marcel P., Erik Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries**, “An Illustrated User Guide to the World Input-Output Database: the Case of Global Automotive Production: User Guide to World Input-Output Database,” *Review of International Economics*, August 2015, 23 (3), 575–605.
- UN**, “UN National Accounts Main Aggregates Database,” 2017.
- Uy, Timothy, Kei-Mu Yi, and Jing Zhang**, “Structural change in an open economy,” *Journal of Monetary Economics*, September 2013, 60 (6), 667–682.
- Yi, Kei-Mu**, “Can Vertical Specialization Explain the Growth of World Trade?,” *The Journal of Political Economy*, February 2003, 111 (1), 52–102.
- Świącki, Tomasz**, “Determinants of structural change,” *Review of Economic Dynamics*, March 2017, 24, 95–131.

A Appendix

A.1 Correlation of openness and service share

Table A.1 shows the results of regressing the country-level growth rate of openness on the country-level growth rate of the service share for our sample. We find strong evidence of a negative correlation; when a country featured higher growth in its service expenditure share, it experienced lower growth in openness, even accounting for its level of income per worker. In the next subsection, we present reduced-form evidence of how much structural change held back global trade growth.

Table A.1: Country-level openness and service share

<i>Dependent variable: openness growth</i>		
Services share growth	−0.379*** (0.070)	−0.404*** (0.069)
Log GDP per-capita		−0.036*** (0.008)
Country FE	YES	YES
N	1215	1215
R^2	0.05	0.07

Note: Robust standard errors are in parentheses, with *** denoting a 99% significance level.

A.2 Data Appendix

This section describes the data used to construct the empirical counterfactual in Section 2 and to estimate the model in Section 4. These data cover 1970–2015 for 27 countries/regions: Australia, Austria, Belgium-Luxembourg, Brazil, Canada, China, Cyprus, Denmark, Finland, France, Germany, Greece, India, Indonesia, Ireland, Italy, Japan, South Korea, Mexico, Netherlands, Portugal, Spain, Sweden, Turkey, United Kingdom, and United States, plus a “Rest of World”. The empirical counterfactual requires time series of 1) total exports and imports of goods and services and 2) final expenditure in goods and services. The model estimation requires these series as well as 3) bilateral goods and services trade data; 4) input-output coefficients; 5) value added to gross output ratios; 6) sectoral price indices; and 7) the aggregate wage bill, and 8) total employment.

Our strategy is to work with the World Input-Output Database (WIOD) from 1995-2011, which is described in Timmer, Dietzenbacher, Los, Stehrer and de Vries (2015), then build the rest of the time sample out from those numbers using splicing techniques with other longer-running datasets. This ensures that the WIOD-based input-output coefficients generate sensible expenditure shares during WIOD years—otherwise, the input-output coefficients would be applied to trade data that may not match the underlying WIOD data used to generate those coefficients.

Labor endowment by country We take total employment data in the Penn World Table as our measure of L_i . These data correspond to the number of workers engaged in market activity. Since these data only go

through 2014, we create a splicing factor with WDI total employment data in 2015 in order to estimate the model through 2015.

Wage by country Dividing aggregate value added in current US\$ by the labor endowment gives the imputed wage w_i .

Total exports and imports by country For each of the 27 groupings above, we take total goods and services exports and imports from the WIOD from 1995-2011. Then, for all other years (i.e. 1970-1994 and 2012-2015), we splice with other data. The splicing procedure divides the average of three years of the WIOD data by the average of three years of a longer dataset to generate a splicing factor, then applying that splicing factor to the longer dataset in non-WIOD years. The averages are calculated from 1995-1997 for all years before 1995, and from 2009-2011 for all years after 2011. For goods trade, we splice the WIOD trade data with total trade from the IMF Direction of Trade Statistics (IMF 2016a) database from 1970 to 2015. For services, we use aggregate services exports and imports data from the World Development Indicators (WDI) from 1970 to 2015 as the comparison. If WDI data on services are not available, we supplement in growth rates where necessary with OECD services data.

Bilateral goods and services trade by sector and country As with total goods trade, when not taken directly from the WIOD, goods trade between two regions in our sample is generated by splicing importer-reported bilateral goods trade data in the IMF DOTS database with WIOD data, using the same three-year splicing factors as above. Bilateral services data are sparse, so instead of splicing, we simply apply average bilateral shares (e.g., U.S. imports from Canada as a share of U.S. imports from the world) over three year periods to the total services trade data calculated as above. Again, for all years prior to 1995, we use average bilateral shares from 1995-1997, and for all years after 2011, we use average bilateral shares from 2009-2011.

Value added by sector and country For value added data, we rely on the UN Main Aggregates Database (UN 2017). We take nominal goods value added in a country to be the combination of expenditure in “Agriculture, hunting, forestry, fishing” and “Mining, Manufacturing, Utilities”, while services value added is expenditure in “Construction”, “Wholesale, retail trade, restaurants and hotels”, “Transport, storage, and communication”, and “Other Activities”.²⁵ Years: 1970-2015.

Input-output coefficients and value added to gross output ratios To construct γ_{ikn} , the country-specific share of intermediate inputs that sector k sources from sector n , we use the numbers directly from WIOD for 1995-2011. The value added to gross output ratio in sector k , λ_{ik} is also a straightforward manipulation of data in the WIOD for 1995-2011. Since the WIOD covers only 1995-2011, we impute γ_{ikn} and λ_{ik} for the remaining years by (i) supplementing the WIOD with incomplete country-specific input-output tables to

²⁵Results are qualitatively similar defining construction as a goods category, but given the lack of direct trade in construction, categorizing it as a service will make goods sectoral openness higher and services sectoral openness lower. Both the model-based counterfactual and especially the reduced-form counterfactual would be smaller in magnitude relative to the data.

create a longer, yet unbalanced, panel and (ii) imputing all missing country-year observations based on an estimated relationship between the shares and real income per worker.

The specific countries and years that we add are the following: Japan (1970, 1973-1995, from the Japan Industrial Productivity Database); South Korea (1970, 1975, 1980, 1985, 1986, 1987, 1988, 1990, 1993, 1995, from the Bank of Korea); United States (1972, 1977, 1982, 1987, 1992, 1997, from the Bureau of Economic Analysis);. We also appeal to data from the OECD: Australia (1968, 1974, 1986, 1989); Canada (1971, 1976, 1981, 1986, 1990); Denmark (1972, 1977, 1980, 1985, 1990); France (1972, 1977, 1980, 1985, 1990); (West) Germany (1978, 1986, 1988, 1990); Italy (1985); Netherlands (1972, 1977, 1981, 1986); United Kingdom (1968, 1979, 1984, 1990).

For each country we linearly interpolate between the observed years and then splice the shares series to that from the WIOD data in 1995. If these data do not overlap with 1995, we linearly extrapolate forward until 1995, using data from the latest two years, and then splice. This gives us an unbalanced panel from which we can observe variation across countries and over time prior to 1995. This helps provide discipline on the estimated shares that we impute prior to 1995 for countries that have no data.

Using our unbalanced panel we first estimate:

$$\log\left(\frac{\lambda_{ikt}}{1-\lambda_{ikt}}\right) = \alpha + \beta \log(y_{it}) + \varepsilon_{it}, \quad (\text{A.1})$$

and then impute the ratio, $\log\left(\frac{\lambda_{ikt}}{1-\lambda_{ikt}}\right)$, for missing country-years using the observed GDP per worker, y_{it} and the estimates of α and β . The imputed value for λ_{ikt} is ensured to be between 0 and 1. We follow a similar procedure to impute values for γ_{igst} and γ_{isst} , and then set $\gamma_{iggt} = 1 - \gamma_{igst}$ and $\gamma_{isgt} = 1 - \gamma_{isst}$.

Sectoral expenditure We construct sectoral expenditures consistent with the other underlying data and model structure. First, combining (S5)-(S7) yields the following expression:

$$P_{ik}C_{ik} = P_{ik}Q_{ik} - \sum_{n=\{g,s\}} (1-\lambda_{in}) \gamma_{ikn} (P_{in}Q_{in} + NX_{in}), \quad (\text{A.2})$$

where NX_{ik} is net exports in country i sector k , and $P_{ik}Q_{ik}$ is total absorption. From equilibrium condition S4, we also know total absorption of the composite good can be written as:

$$P_{ik}Q_{ik} + NX_{ik} = \frac{w_i L_{ik}}{\lambda_{ik}}. \quad (\text{A.3})$$

Using data on sectoral value added, $w_i L_{ik}$, along with sectoral net exports, NX_{ik} , and the production share, λ_{ik} , we can calculate total expenditure, $P_{ik}C_{ik}$, via equations (A.2) and (A.3).²⁶ Once the sectoral expenditures are generated, the expenditure shares e_{ik} are straightforward to compute. (Note that $w_i L_{ik}$ is the value added in sector k , so had we used value added data from WIOD directly for the years 1995-2011, then we would impute exactly the expenditures observed in WIOD as well for those years.)

²⁶Equations (A.2) and (A.3) exactly summarize how we constructed sectoral expenditure for the empirical results in Section 2, and is detailed in words in Section 2.1 and figure A.4.

Sectoral prices In order to estimate the preference parameters, we need gross-output sectoral prices. First, we take nominal and real value added (indexed to 2005) data in goods and services from the UN Main Aggregates Database. We generate sectoral price indices for each country-year as the ratio of nominal to real value added. We then splice these indices to the sectoral value added price levels in PPP terms from the GGDC Productivity Level Database “2005 Benchmark” (Inklaar and Timmer 2014) to make the country-level price indices comparable across countries. Finally, we “gross up” the value added prices using the equation for the value added deflator in Appendix C4 of Sposi (2019). Note that these prices are only used in our estimating equation for the preference parameters; the imputed sectoral prices in the calibration of the productivities and trade costs are separate model-specific objects.

A.3 Solution Algorithm

This appendix details the solution algorithm for each period of the model economy. Equations that we refer to are listed in table 1. For each time period:

- Guess the vector of wages, w_i , across countries.
- Compute the sectoral unit costs v_{ik} and the sectoral prices P_{ik} using conditions (S2) and (S3) jointly.
- Compute the sectoral bilateral trade shares π_{ijk} using condition (S1).
- Compute the per-worker transfers from the global portfolio R using condition (G1).
- Compute the aggregate price levels P_i and aggregate consumption indices C_i using conditions (D2) and (D3) simultaneously.
- Compute sectoral consumption C_{ik} using condition (D1).
- Compute sectoral labor demand L_{ik} using condition (S4).
- Compute sectoral intermediate input demand M_{ikn} using condition (S5).
- Compute sectoral gross absorption Q_{ik} using condition (S6).
- Compute sectoral gross production Y_{ik} using condition (S7).
- Define excess demand as net exports minus net contributions to the global portfolio:

$$Z_i^w = \frac{\sum_{k \in \{g,s\}} (P_{ik}Y_{ik} - P_{ik}Q_{ik}) - (\rho_i w_i L_i - R L_i)}{w_i}.$$

Condition (G2) requires that $Z_i^w = 0$, for all i , in equilibrium. If this is different from zero in at least some country, then update the wage vector as follows:

$$w'_i = w_i \left(1 + \kappa \frac{Z_i^w}{L_i} \right),$$

where w'_i is the updated guess of wages and κ is chosen to be sufficiently small so that $w'_i > 0$. Use the updated wage vector and repeat every step to get a new value for excess demand. Continue this procedure until the excess demand is sufficiently close to zero in every country simultaneously. Note that Walras' Law ensures that the labor market clears in each country.

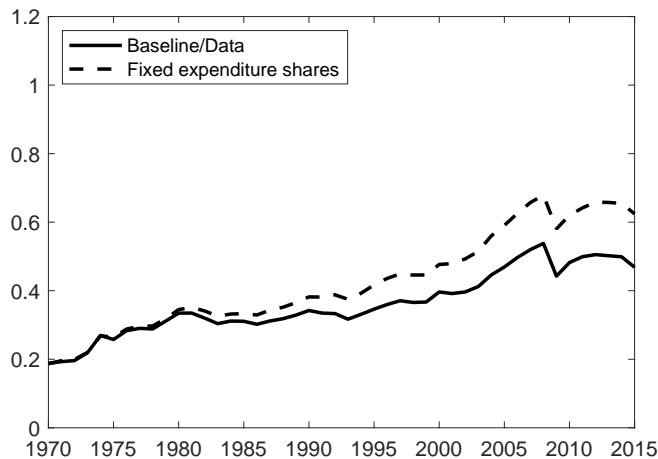
A.4 Robustness

A.5 Sector-specific trade elasticity

Throughout the paper we assumed that $\theta = 4$ in both goods and services sectors. The trade elasticity for goods is widely accepted in the literature. However, there is no widely accepted value for the trade elasticity for services, largely because many studies ignore trade in services. We now consider a robustness exercise by setting $\theta_s = 2.5$, while keeping $\theta_g = 4$. For this experiment, we need to recalibrate T_{is} and τ_{ijs} as in our baseline calibration, to target observed sectoral trade and sectoral expenditure. The remaining parameters are unchanged relative to our baseline calibration.

We compute the equilibrium for two versions of the recalibrated model: (i) a baseline case with endogenous structural change, and (ii) a counterfactual with fixed expenditures (Cobb-Douglas preferences with expenditure shares fixed at 1970 levels). Figure A.1 plots the results. Holding fixed expenditure shares in 1970 does result in greater increases in openness, just as in our baseline calibration. In other words, structural change dampens growth in openness. By 2015, global openness is 32 percent higher in the model with fixed expenditure shares compared to the model with endogenous structural change (0.62 compared to 0.47). This effect is almost identical to the effect obtained in our baseline calibration.

Figure A.1: Openness: Baseline and counterfactual under $\theta_s = 2.5$



A.6 Country Results

In this appendix, we break down structural change and the structural model-based counterfactual for each country in our sample and highlight their contribution to the aggregate counterfactual. Figure A.2 shows the goods and services expenditure shares for each country and the rest of world aggregate. In all countries, the

expenditure share of goods is falling, though for some countries, including Greece, Mexico, and Sweden, the shift is more gradual.

Figure A.3 shows the baseline model solution and the model-based counterfactual result holding expenditure shares fixed for each country. The trade to expenditure ratio in the counterfactual is higher for every country, though by starkly different amounts. The counterfactual tends to be more consistent in percent, rather than percentage point, terms across countries. For example, Belgium-Luxembourg starts out with a high degree of openness, and the counterfactual is about 25 percent higher than the baseline. The same is roughly true for other countries, like India and Japan, with a far lower degree of openness.

For some countries, however, the counterfactual level of openness is not much greater. This tends to relate directly to the degree to which the countries are experiencing structural change: Greece, Mexico, and Sweden all have fairly modest increases in their openness in the model-based counterfactual, which echoes their modest structural change from figure A.2.

Table A.2 shows the contribution to the aggregate fixed expenditure counterfactual depicted in figure 5 for the year 2015, the last year of the sample. The first column provides the expenditure share of each country in the world aggregate, while the second is its trade share (exports plus imports in each country as a share of world trade). The third column represents the percentage point contribution of each country to the difference between the model-based counterfactual and the baseline, which sums to 0.156. The final column shows the equivalent percent contribution. The table makes clear that the contribution to the aggregate counterfactual largely follows the country's trade share, not its expenditure share. For example, with the United States being relatively closed, with an expenditure share about twice its trade share, the contribution of the U.S. to the aggregate counterfactual is close to the trade share. By contrast, China has a similar trade share and a smaller expenditure share and contributes the most of any single country to the aggregate counterfactual.

Figure A.2: Sectoral expenditure shares by country

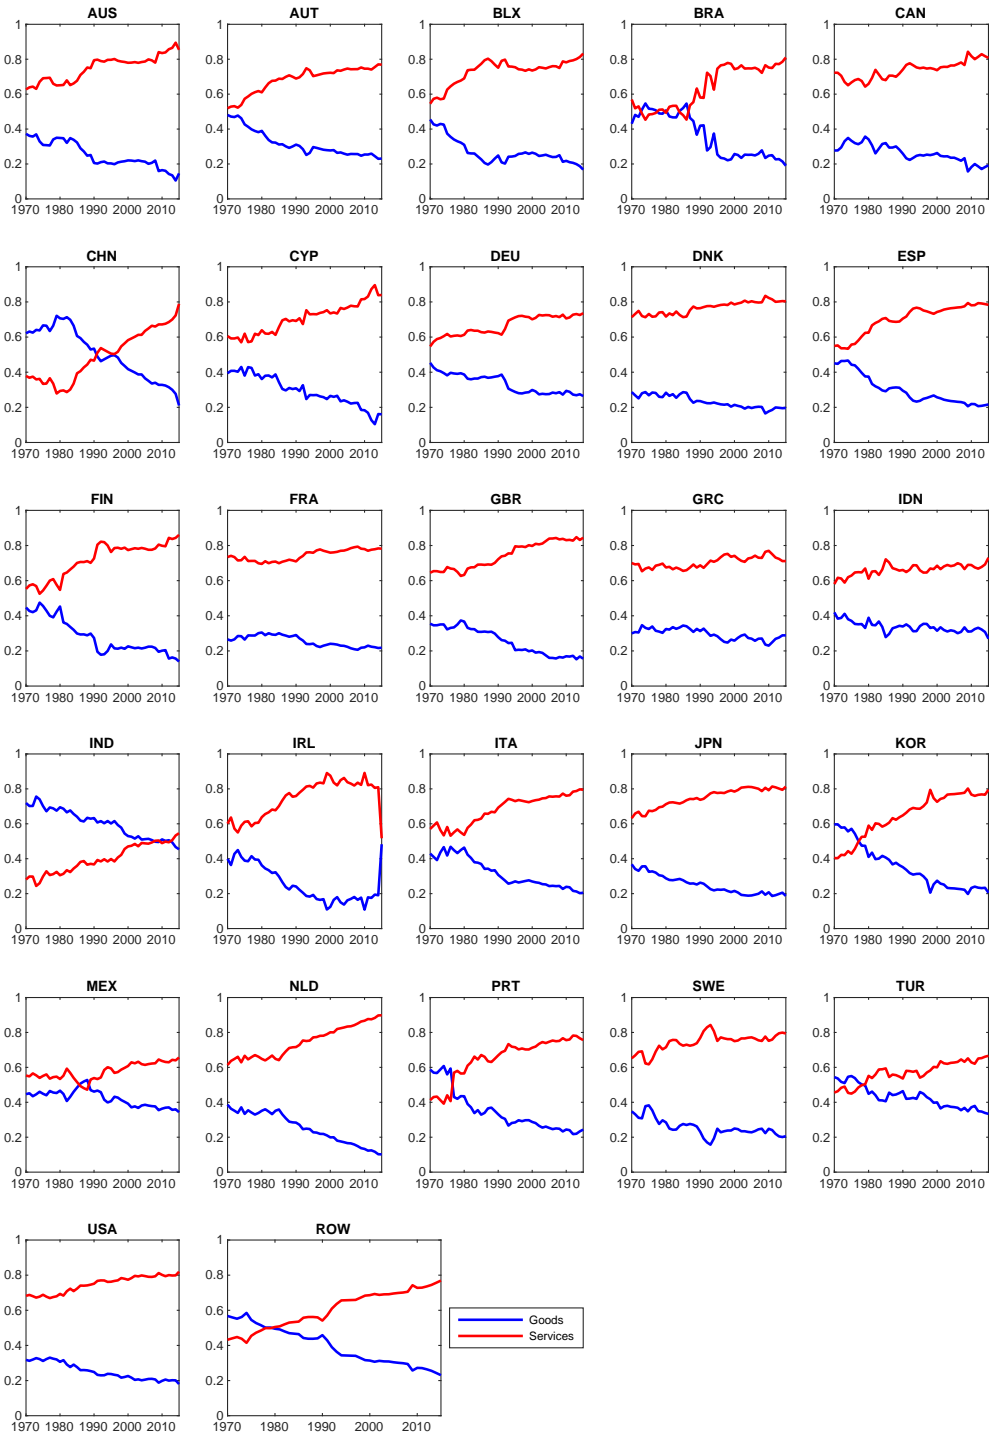


Figure A.3: Trade to expenditure ratio by country

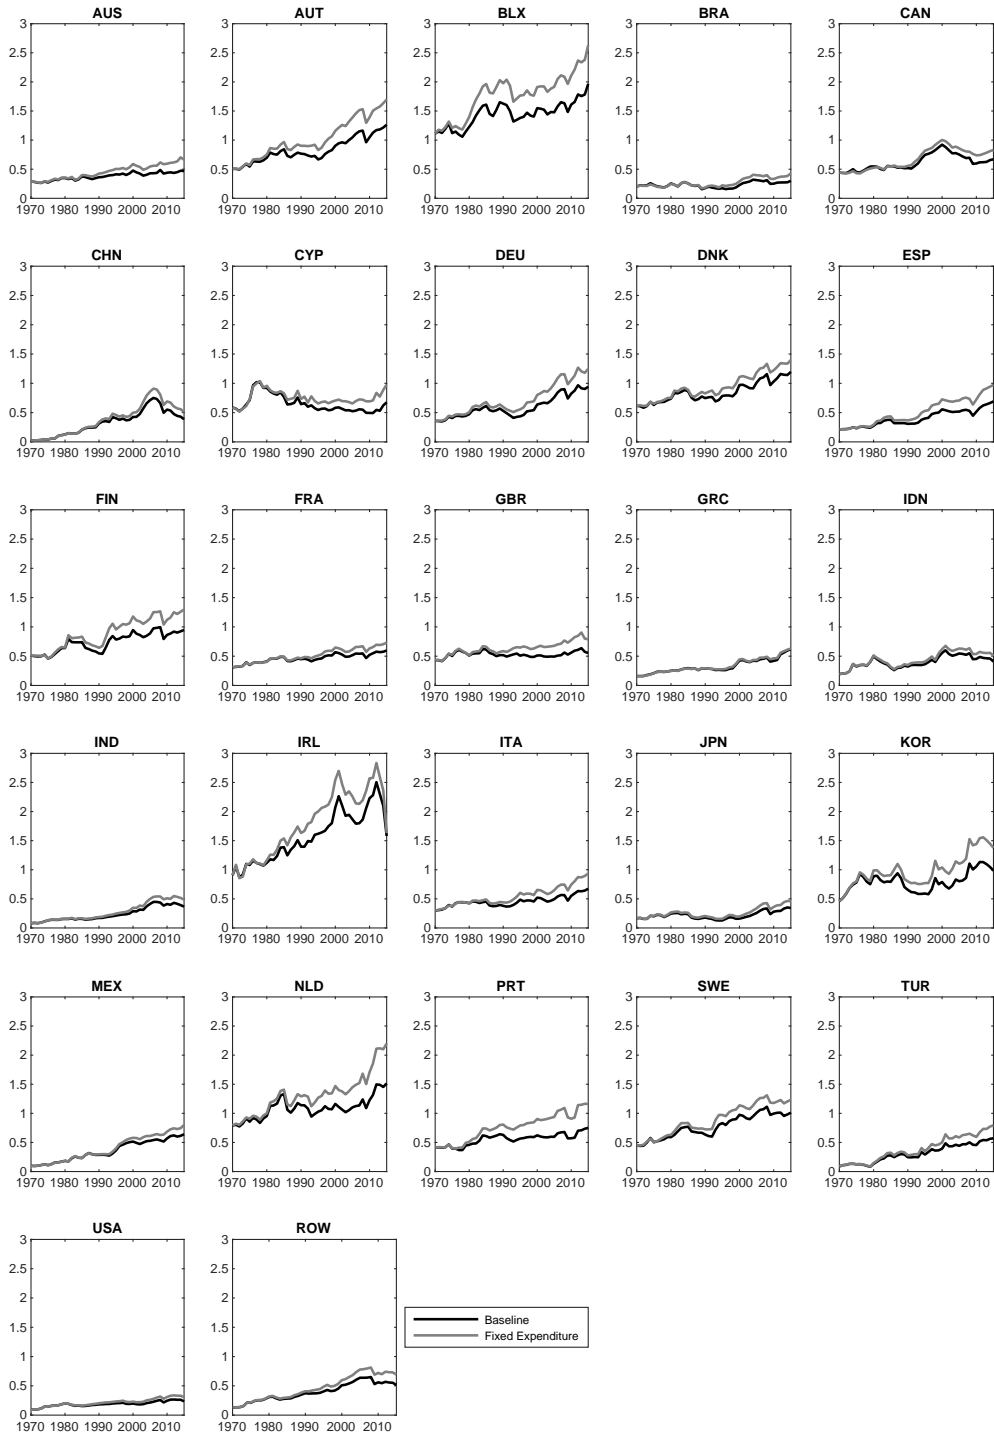
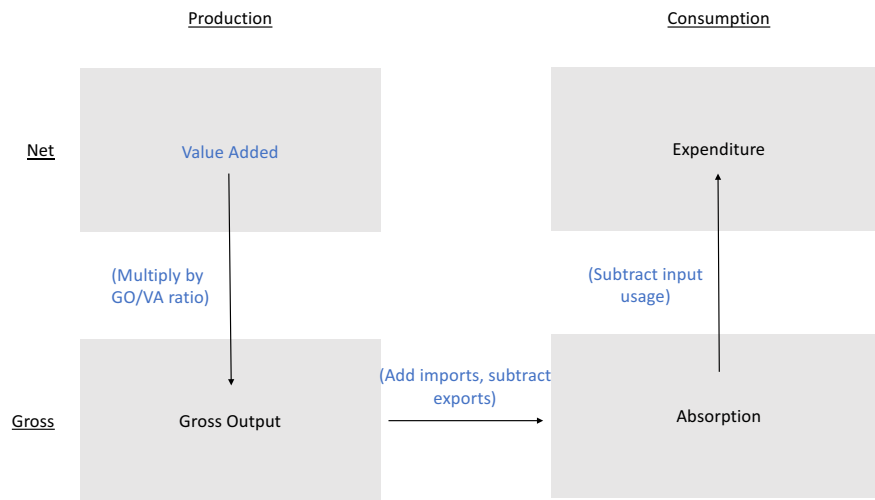


Table A.2: Contributions to fixed expenditure counterfactual in 2015

Country	Expenditure Share	Trade Share	Ppt. Contribution	Pct. Contribution
Australia	1.6%	1.6%	0.003	1.9%
Austria	0.4%	1.2%	0.002	1.2%
Belgium-Luxembourg	0.6%	2.4%	0.003	2.0%
Brazil	2.2%	1.4%	0.003	1.8%
Canada	2.1%	3.0%	0.004	2.3%
China	15.2%	12.5%	0.021	13.4%
Cyprus	0.0%	0.0%	0.000	0.0%
Germany	3.6%	7.2%	0.012	7.4%
Denmark	0.3%	0.9%	0.001	0.4%
Spain	1.5%	2.2%	0.004	2.7%
Finland	0.3%	0.5%	0.001	0.6%
France	3.0%	3.9%	0.005	3.2%
United Kingdom	3.6%	4.2%	0.007	4.7%
Greece	0.3%	0.3%	0.000	0.0%
Indonesia	1.2%	1.0%	0.001	0.9%
India	2.8%	2.2%	0.003	2.2%
Ireland	0.3%	1.1%	0.000	0.1%
Italy	2.2%	3.1%	0.006	3.7%
Japan	6.1%	4.4%	0.007	4.6%
Korea	1.6%	3.4%	0.007	4.2%
Mexico	1.6%	2.1%	0.003	1.8%
Netherlands	0.7%	2.4%	0.004	2.8%
Portugal	0.3%	0.4%	0.001	0.6%
Sweden	0.6%	1.2%	0.001	0.9%
Turkey	1.0%	1.2%	0.002	1.4%
United States	26.1%	13.1%	0.017	10.9%
Rest of World	21.0%	23.1%	0.038	24.2%
Total	100.0%	100.0%	0.156	100.0%

A.7 Additional Figures

Figure A.4: Deriving sectoral expenditures from sectoral value added



Note: Categories in blue represent publicly available data, while categories in black represent imputed moments.

Figure A.5: Goods expenditure shares: baseline and counterfactual

