Determinants of Structural Change*

Tomasz Święcki†

First Version: June 2013
This Version: January 2017‡

Abstract

In this paper I ask which of the multiple mechanisms suggested in the literature are quantitatively important for understanding the process of structural change. I build a model combining four forces in a common framework: (i) sector-biased technological progress, (ii) nonhomothetic tastes, (iii) international trade and (iv) changing wedges between factor costs across sectors. I calibrate the model using the data for 45 diverse countries over the period 1970-2005 and use counterfactual simulations of the model to systematically assess the relative importance of the four determinants of structural change. I find that sector-biased technological change is overall the most important mechanism and it is essential for understanding the decline of manufacturing labor share and the corresponding growth in services in developed countries. Nonhomothetic preferences are key to accounting for movement of labor out of agriculture, which matters primarily for poorer countries. International trade and changes in relative factor costs across sectors are important for individual countries but their impact on the relocation of labor is less systematic. I also show that a model with homothetic preferences would overstate the importance of agriculture in accounting for differences in aggregate productivity across countries and over time.

JEL Numbers: O11, O47.

Keywords: structural change, structural transformation, nonhomothetic preferences, sectoral productivity.

---

*I am grateful to Gene Grossman, Steve Redding and Esteban Rossi-Hansberg for their encouragement and guidance. I also thank the editors Diego Restuccia and Matthias Doepke, two anonymous referees, as well as Ariel Burstein, Pablo Fajgelbaum, Richard Rogerson and seminar participants at UCLA, CEA 2014 and SED 2014 for helpful comments and suggestions.

†Vancouver School of Economics, University of British Columbia. Email: Tomasz.Swiecki@ubc.ca.

1 Introduction

Structural change is one of the most robust features of economic development. As countries grow richer, we observe secular shifts in their allocation of labor and expenditure across broad sectors of agriculture, manufacturing and services. A number of theoretical explanations of this process have been proposed in the literature.[1] There is little consensus, however, on the relative importance of the suggested mechanisms. The goal of this paper is to assess, quantitatively, how crucial are various forces for understanding the observed patterns of structural change.

To address this issue, I begin by building in Section 2 a quantitative model combining in a unified framework four mechanisms that can drive structural change. The first classic source of structural change is sector-biased technological progress. If productivity growth in a sector is slow relative to other sectors then the relative price of the sluggish sector increases over time. With sectoral outputs being gross complements in consumption, expenditures and labor shift towards sectors with relatively slow productivity growth. The second classic explanation of structural change is based on nonhomothetic preferences. As incomes rise, households spend relatively less on agricultural goods and more on services.

To allow both sector-biased technological change and income effects to operate, I use a flexible specification of consumer preferences. The augmented constant differences of elasticities of substitution (ACDES) preferences that I introduce to the structural change literature nest other commonly used preference specifications. The extra flexibility helps me to better asses the importance of sector-biased productivity growth and of the overall rise in aggregate productivity (income effect) for structural change in a broad sample of countries.[2]

International trade is the third channel affecting the sectoral composition of economies. Matsuyama (2009) formalizes an argument that the the same underlying forces can have quite different implications for structural change in a closed economy and in an interdependent world. For example, whereas fast productivity growth in manufacturing would lead to a decline in the manufacturing labor share in a closed economy, in an open economy manufacturing employment can expand because of specialization according to comparative advantage. This consideration is potentially important given that in recent decades many countries have become substantially integrated with the world economy. I therefore embed my framework in a three-sector general equilibrium model of international trade. I treat agriculture and manufacturing as tradable sectors as in the Ricardian model of Eaton and Kortum (2002) while services are treated as nontradable. In order to better capture the impact of openness on sectoral labor shares, I allow for trade imbalances both at the sectoral and at the aggregate level.

[1]See Herrendorf et al. (2013a) for an overview of the large theoretical and empirical literature on structural transformation.

[2]As explained by Herrendorf et al. (2013b), the relative importance of the two channels might depend on how one defines the commodity space. In their terminology, I use a consumption value added approach so that, e.g., the expenditure share of agriculture reflects the share of value added originating in agriculture in total final expenditure.

The last force influencing structural change is represented by changes in relative labor costs across sectors. It is well known that the breakdown of economic activity at a level of broad sectors looks different when measured in nominal terms (expenditure and value added shares) than in terms of factor allocation (labor shares). Buera and Kaboski (2009) observe that quantitative models therefore need to allow for factor cost differences across sectors in order to be consistent with both nominal and real margins of structural change. In my model, as in most quantitative work on structural change, homogenous labor is the only primary factor of production. Factor cost differentials are therefore summarized by intersectoral labor wedges. An open empirical question is the extent to which changes in wedges over time can account for the relocation of real resources across sectors.

To assess the empirical relevance of the four channels described above, I take the model to the data for 45 countries over the period 1970-2005. I combine data from multiple sources to construct an unbalanced panel featuring countries with diverse levels of economic development. Working with a diverse sample offers two main advantages relative to the prior literature studying the experiences of individual countries. First, it puts more discipline on the calibration of key unknown model parameters. Second, it allows me to systematically assess how the relative importance of different forces behind structural change depends on the country's stage of economic development and other country characteristics, such as the degree of openness to trade or country size.

My baseline calibration, described in Section 3, is designed so that the model exactly accounts for structural change in all countries along two key margins: sectoral shares of employment and value added. I use the general equilibrium predictions of the model for the third margin - sectoral labor productivity growth - to obtain the parameters of ACDES preferences via a GMM procedure.

In Section 4 I use counterfactual simulations of the calibrated model to assess the importance of each of the model's four forces. The counterfactual simulations switch individual channels on and off in order to isolate their impact on structural change. To quantify that impact I introduce the Labor Relocation Index which measures the fraction of the observed labor relocation across sectors that can be accounted for by a specified combination of forces.

I find that the sector-biased technological progress is overall the most important factor. For example, this channel alone can explain 43% of the labor relocation for the median country. At a more disaggregated level, the sector-biased technological progress is particularly important for explaining the net movement of labor from manufacturing to services and is thus crucial for understanding structural change occuring in developed countries. While income effects have on average less power to account for broad shifts in sectoral employment, they remain an important force. In particular, nonhomotheticity of preferences plays a key role in generating the transition of labor out of agriculture and is thus very relevant for countries at earlier stages of economic development. My quantitative exercise therefore shows that the relative importance of the two classic channels depends on how far along a country is in the process of structural change. Both mechanisms are necessary to provide a fully satisfactory account of a complete transition from an agriculture-based to a service-based economy.

International trade and changes in intersectoral wedges play a more idiosyncratic role in accounting for labor relocation. Ignoring either channel would not lead to a systematic bias in predicting the changes in labor allocation over time. Both forces do nevertheless play a significant idiosyncratic role for some countries. For example, trade is relatively important for smaller countries that tend to rely on it more. Moreover, there are strong interactions between trade and wedges in that the latter only matter in an open economy setting.

In Section 5 I investigate the importance of nonhomothetic preferences for modeling structural change from a different perspective. I recalibrate the model restricting preferences to be of the homothetic CES form. I then compare the sectoral productivity patterns derived from the homothetic and the baseline model. Because both models need to match the same expenditure share data, the homothetic model requires larger dispersion of relative prices in the cross section of countries and larger changes in relative prices within countries over time. The CES model achieves this by predicting larger dispersion of labor productivity in agriculture and higher labor productivity growth in that sector than the nonhomothetic model. Consequently, the homothetic model overstates the importance of agriculture in accounting for cross-country differences in aggregate productivity and in accounting for aggregate productivity growth within countries.

Among the growing literature on structural change, three studies are particularly closely related to this paper. Herrendorf et al. (2013b) study the ability of income effects and changes in relative sectoral prices (price effects) to explain shifts in sectoral expenditure shares.[3] They find both mechanisms to be important.[4] There are several important differences between their work and this paper. First, they focus only on the US experience whereas I work with a broad sample of countries. The difference in the sample scope translates to very different methodologies. Estimating preference parameters for a single county requires time series of sectoral expenditure shares and price and quantity indices. In contrast, multi-country estimation also requires knowledge of levels of prices and quantities across countries. Because levels are not as readily available as indices in the data, the multi-country approach solves for the levels using the full general equilibrium structure of the model. Once this is done, the advantage of the general equilibrium approach is that it allows for a richer set of counterfactuals. In particular, it allows me to assess the contribution of different forces to the sectoral labor relocation, whereas Herrendorf et al. (2013b) focus only on how restrictions on preference parameters affect the fit of expenditure shares. My multi-country approach also introduces open economy issues and highlights the patterns that are common across

[3]Price effect is sometimes referred to as a substitution effect in the structural change literature, with a slight abuse of the term from consumer theory. Furthermore, in simple closed economy models changes in relative prices are often driven only by changes in relative productivities so the price effect is synonymous with sector biased productivity growth. In my model relative prices can change also because of international trade and changes in intersectoral wedges.

[4]This is true for their consumption avlue added approach. Of the two commodity spaces they consider the consumption value added is closer to the data used in this paper. The difference is that they separate the consumption and investment parts of value added, whereas this paper abstracts from investment and treats all value added as consumed.

countries, whereas their analysis by construction is restricted to the closed economy for a single country.[5]

The role of international trade in the process of structural transformation is investigated, both theoretically and quantitatively, by Uy et al. (2013). The model I build uses a similar structure on the production and trade side as their paper.[6] However, in their empirical application they concentrate only on trade between Korea and the composite Rest of the World, with key parameters calibrated using only Korean data. Their main finding is that globalization played a big role in Korea's structural change. My results suggest that Korea's experience is not necessarily typical. Moreover, my model-based decompositions indicate that even for Korea sector-biased productivity growth and nonhomothetic preferences are quantitatively more important than trade as drivers of structural change.[7]

The calibration strategy of this paper can be seen as building on Duarte and Restuccia (2010). Like those authors, I use the observed sectoral data together with the general equilibrium structure of the model to back out unobserved levels of sectoral productivity across countries. This paper enhances the calculations of productivity by taking into account an endogenous component of productivity due to trade, an impact of wedges, and by working with a more flexible demand structure. More importantly, whereas Duarte and Restuccia (2010) focus only on the role of sectoral productivity growth in the process of structural transformation in a closed economy, I evaluate the importance of that channel relative to other potential drivers of structural change.

Finally, this paper is related to a literature in macroeconomics featuring factor cost differences across sectors. This literature focuses on the cross-sectional implications of wedges for aggregate productivity differences in a closed economy setting (Restuccia et al. (2008)), open economy setting (Tombe (2015)), or for the welfare gains from trade (Święcki (2017)). In contrast, in this paper I analyze the contribution of various forces to the observed structural change within countries in the time-series.

[5]Boppart (2014) is another recent paper studying the income and price effects as forces behind the structural change in the US. He uses a different form on non-Gorman preferences than employed in this paper. His specification has convenient aggregation properties in a two-sector model that he maps to data on consumption of goods and services. In this paper I consider the classic three-sector split, as agriculture, manufacturing, and services exhibit distinct patterns during the process of structural change.

[6]They abstract from intersectoral wedges and use more restrictive preferences. On the other hand, their model allows for richer input-output linkages across sectors.

[7]Other quantitative studies of structural transformation in an open economy context include Betts et al. (2015), Teignier (2014), Stefanski (2014), and Sposi (2015). The first two papers study how trade affected structural change in individual countries. Stefanski investigates the impact of structural transformation in China and India on the price of oil. Sposi extends the methodology developed in this paper to study transmission of sectoral productivity shocks to composition of sectoral value added in a multi-country setting.

2 Model

The model world consists of N countries. Homogeneous labor is the only primary factor of production and can be employed in one of three sectors: agriculture, manufacturing or services. The sectoral allocation of labor is driven by four forces. (1) Nonhomothetic preferences imply that as countries grow richer expenditure and labor move away from sectors with low income elasticity to sectors with high income elasticity. (2) Sector-biased productivity growth combined with low elasticity of substitution across sectors means that expenditure and labor move towards sectors with sluggish productivity growth. (3) Trade in agriculture and manufacturing leads countries to have higher employment shares in their net-exporting sector. (4) Intersectoral wedges account for the divergence of sectoral employment shares from output shares. All goods are utilized in the period they are produced. I thus abstract from physical investment. Trade need not be balanced each period for individual countries but I abstract from the intertemporal decisions that determine aggregate trade deficits. The model's solution is therefore a sequence of static equilibria. Time subscripts are henceforth omitted except where required for clarity.

2.1 Consumers

Consumers have preferences defined over consumption of aggregate output of agriculture C_A, manufacturing C_M and services C_S. Intratemporal preferences are described by means of an indirect utility function $V(P_A,\, P_M,\, P_S,\, m)$, which gives the maximum level of utility achieved by an individual with nominal expenditure m facing prices $\{P_s\}$:[8]

$$V(P_A,\, P_M,\, P_S,\, m) = \sum_{s \in \{A,\, M,\, S\}} \gamma_s \frac{\left(\frac{m - \sum_{s'} P_{s'} \bar{c}_{s'}}{P_s}\right)^{\alpha_s} - 1}{\alpha_s}, \tag{1}$$

where parameters satisfy the following restrictions:

$$\gamma_s > 0,\, \sum_s \gamma_s = 1,\, \alpha_s \geq -1,\, \bar{c}_s \geq 0.$$

This formulation of preferences augments the constant differences of elasticities of substitution (CDES) preferences by introducing subsistence consumption requirements $\bar{c}_s$. The standard CDES preferences obtained with $\bar{c}_s = 0$ are characterized by Jensen et al. (2011) and Hanoch (1975) among others, and build on the indirect addilog preferences dating back to Houthakker (1960) and beyond. In Appendix A I show that the CDES indirect utility function (1) augmented with subsistence requirements (ACDES) satisfies regularity conditions imposed by the consumer theory.

The demand system associated with (1) generalizes preference structures used in the prior lit-

[8]There is no closed-form solution for the direct utility function $u(C_A\, C_M,\, C_S)$ corresponding to (1) except in some special cases.

erature on structural transformation. Its main advantage over the commonly used functional forms is that it gives non-vanishing roles to the two forces emphasized in the structural transformation literature and thus improves the ability of the model to match the data for a wide range of countries. One tradition attributes the pattern of falling expenditure share of agriculture and rising share of services to income effects, typically modeled by postulating a Stone-Geary utility function as in Kongsamut et al. (2001). The second strand of the literature links changes in sectoral expenditure shares to changes in relative prices. That price effect is usually modeled with the aid of CES preferences with an elasticity of substitution less than one, with Ngai and Pissarides (2007) serving as a recent example. However, as discussed by Buera and Kaboski (2009) and Herrendorf et al. (2013a), models relying on income or relative price channel alone fail to account for important empirical regularities of structural change. More recently, Herrendorf et al. (2013b) and Buera and Kaboski (2009) worked with augmented CES (ACES) preferences that nest both the Stone-Geary and homothetic CES as special cases.[9] However, that specification is still quite restrictive. In particular, the allocation of marginal expenditure across sectors is independent of income level. At low income levels the income effect plays a dominant role but for high enough incomes the demand system essentially behaves like a homothetic CES. This asymmetry becomes problematic in an empirical analysis when the sample contains observations with very different income levels, which is the case in this paper.[10]

In contrast to other preferences commonly used in the literature, the preferences implied by (1) remain nonhomothetic regardless of income level.[11] Denoting the discretionary expenditure by $\tilde{m} = m - \sum_s P_s \bar{c}_s$ to simplify notation, the Marshallian demand for sector s goods is given by:

$$C_s = \bar{c}_s + \frac{\gamma_s \left(\frac{\tilde{m}}{P_s}\right)^{\alpha_s+1}}{\sum_{s'} \gamma_{s'} \left(\frac{\tilde{m}}{P_{s'}}\right)^{\alpha_{s'}}}. \tag{2}$$

The ratio of expenditures on sectors s and s' is asymptotically (for high incomes) given by

$$\frac{\gamma_s \left(\frac{m}{P_s}\right)^{\alpha_s}}{\gamma_{s'} \left(\frac{m}{P_{s'}}\right)^{\alpha_{s'}}},$$

which depends on the level of expenditure as long as $\alpha_s \neq \alpha_{s'}$ and depends on relative prices unless $\alpha_s = \alpha_{s'} = 0$. While standard CDES preferences (with $\bar{c}_s \equiv 0$) already incorporate one form of nonhomotheticity, allowing for subsistence requirements further increases the flexibility of the demand system. The augmented CDES demand in fact nests all the systems mentioned in the preceding paragraphs. With $\alpha_s \equiv \varepsilon - 1$ and $\bar{c}_s \equiv 0$ constant across sectors we obtain the

[9] The utility function in those papers takes the form $U = \left(\sum_s \gamma_s^{-\frac{\varepsilon-1}{\varepsilon}} (C_s - \bar{c}_s)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}$.

[10] Real per capita income in the richest country in my sample (Norway in 2005) is 76 times larger than in the poorest country (China in 1978).

[11] Another indirect utility function which gives rise to non-vanishing income effect is used by Boppart (2014).

standard homothetic CES preferences with elasticity of substitution ε. Taking the limit $\alpha_s \to 0$ in (1) while allowing $\bar{c}_s \neq 0$ we can recover Stone-Geary preferences. Combining $\alpha_s \equiv \varepsilon - 1$ and arbitrary $\bar{c}_s$ yields the demand system consistent with augmented CES. An additional advantage of the ACDES demand system over ACES is that ACDES gives a richer pattern of substitution among goods while still remaining parsimoniously parametrized.[12] As discussed further in Section 3.4, the greater flexibility of the ACDES specification measurably improves the ability of the model to fit the data.

2.2 Production and Trade

In each sector there is a unit measure of intermediate goods indexed by $h \in [0, 1]$. Intermediates in any sector are produced using a constant returns to scale technology combining labor and the aggregate output of that sector. Specifically, the production function for variety h in sector s in country i at time t is:

$$q_{sit}(h) = \kappa_s z_{sit}(h) L_{sit}(h)^{\beta_s} Q_{sit}(h)^{1-\beta_s},$$

where $z_{sit}(h)$ denotes the variety-sector-country-year-specific productivity.[13] Labor shares $0 < \beta_s \leq 1$ are sector-specific but are constant across countries and time.[14]

The nontraded aggregate output Q_{si} of industry s is costlessly assembled from all intermediates of that industry using the CES technology with the elasticity of substitution across varieties σ. The aggregate sectoral output is used both as an input for production of intermediates and to satisfy final demand.

Country i draws productivity $z_{si}(h)$ in variety h from a distribution with cumulative distribution function F_{si}, with draws independent across countries, sectors, varieties and time. Following Eaton and Kortum (2002), the realizations are assumed to come from the Fréchet distribution with $F_{si}(z) = e^{-T_{si} z^{-\theta_s}}$. The parameter T_{si} is related to country i's average efficiency in sector s. The parameter θ_s is an inverse measure of the dispersion of productivity draws and is assumed to be constant across countries and time.

The product market is perfectly competitive. Given prices of intermediates $p_{si}(h)$ prevailing in market i, the price index for the aggregate output is given by $P_{si} = \left[\int_0^1 p_{si}(h)^{1-\sigma} dh\right]^{\frac{1}{1-\sigma}}$. The cost of producing a unit of variety h in sector s and country i is then $c_{si}/z_{si}(h)$, where

$$c_{si} = w_{si}^{\beta_s} P_{si}^{1-\beta_s} \tag{3}$$

is the cost of the input-bundle used by sector s and where w_{si} is the wage in sector s in country i.

[12] For example, ACDES allows pairs of goods to be Allen-complements which is impossible with ACES.

[13] The constant $\kappa_s = \beta_s^{\beta_s}(1-\beta_s)^{(1-\beta_s)}$ is introduced to simplify notation.

[14] The main role for combining labor and aggregate output is to reconcile production data (recorded at value added level) and trade data (recorded at gross output level) in the empirical implementation of the model. I do not pursue richer input-out structures for lack of consistent IO tables series for many countries.

Intermediate goods in agriculture and manufacturing are tradable subject to the standard iceberg transportation costs. Delivering a unit of variety h in sector $s \in \{A, M\}$ from country i to country j requires shipping $\tau_{sji} \geq 1$ units of the good, with $\tau_{sjj} = 1$. Using the results of Eaton and Kortum (2002), the price index in the tradable sectors can be written as[15]

$$P_{sj} = \Gamma_s \left[\sum_i T_{si} \left(c_{si} \tau_{sji} \right)^{-\theta_s} \right]^{-\frac{1}{\theta_s}} , \ s \in \{A, M\} . \tag{4}$$

This structure also gives the following expression for the share of expenditure in sector s in country j going to goods from country i:

$$\pi_{sji} = \frac{T_{si} \left(c_{si} \tau_{sji} \right)^{-\theta_s}}{\sum_m T_{sm} \left(c_{sm} \tau_{sjm} \right)^{-\theta_s}}. \tag{5}$$

In the nontraded service sector the price level is simply given by:

$$P_{Sj} = \Gamma_S^{\frac{1}{\beta_S}} \frac{w_{Sj}}{\left(T_{Sj}^{\frac{1}{\theta_S}} \right)^{\frac{1}{\beta_S}}}. \tag{6}$$

2.3 Intersectoral Wedges

The wage w_{si} appearing in (3) is sector-specific. As noted by Buera and Kaboski (2009), wage differentials are necessary to account for differences in sectoral employment shares and value added shares in the data.[16] Wage differentials across sectors might arise with homogeneous labor for a number of reasons. Wages appearing in (3) are wages faced by producers in sector s. Wage differentials can thus capture the effects of distortions on the domestic labor markets that have the effect of making the effective labor costs diverge across sectors. Examples of such distortions are sector-specific labor taxes or differential market power on the side of workers due to varying degrees of unionization. Wage differentials can also represent compensating differentials when workers have a preference for work in certain sectors. For the purpose of this paper, the exact source of the wage differential is not crucial as long as labor is homogenous.[17] To highlight this agnosticism and to ease notation the wage differentials are therefore summarized by the wedge between wages in agriculture

[15] $\Gamma_s \equiv \Gamma\left(\frac{\theta_s+1-\sigma}{\theta_s}\right)$, where $\Gamma(\cdot)$ is a Gamma function.

[16] Abstracting from other factors of production is arguably not the reason for the divergence since factor intensity differences are not large at the level of broad sectors, cf. Valentinyi and Herrendorf (2008). A similar picture emerges if we look at labor compensation rather than value added as discussed further in Appendix E.1.

[17] Homogenous labor is a standard assumption in the quantitative literature on structural change. An important exception is Caselli and Coleman (2001) who focus on the regional convergence in income in the US during the process of structural change. The issue of worker heterogeneity is discussed further in Appendix E.2.

or services, and manufacturing wages, i.e. I call the objects

$$\xi_{Ai} \equiv \frac{w_{Ai}}{w_{Mi}}, \quad \xi_{Si} \equiv \frac{w_{Si}}{w_{Mi}} \tag{7}$$

the wedge in agriculture and the wedge in services, respectively. By construction the wedge in manufacturing is then equal to one, $\xi_{Mi} \equiv 1$.

2.4 Equilibrium

In this subsection I close the model by describing the goods and labor market clearing conditions of the model world economy. Let L_{si} denote employment in sector s in country i and let Y_i denote the GDP of country i, equal to its labor income:

$$Y_i = \sum_{s \in \{A, M, S\}} w_{Mi} \xi_{si} L_{si}.$$

Let D_i be country i's overall trade deficit, with deficits summing to zero at the world level:

$$\sum_i D_i = 0. \tag{8}$$

The budget constraint of agents in country i then dictates that total final demand by consumers in i is given by $X_i^F = Y_i + D_i$.[18] To simplify notation, denote by $\tilde{X}_i^F$ the final demand spending net of subsistence expenditure in country i: $\tilde{X}_i^F = X_i^F - L_i \sum_s P_{si} \bar{c}_s$.

The goods market clearing conditions say that the value of gross output of sector s in country i must be equal to the value of imports by all countries (including i) of goods from i in that sector. Using the fact that value added $w_{Mi} \xi_{si} L_{si}$ constitutes a fraction β_s of gross output, we can write the market clearing conditions as follows: for all $i = 1, ..., N$ and $s \in \{A,\ M,\ S\}$

$$w_{Mi} \xi_{si} L_{si} = \sum_j \pi_{sji} \left\{ (1 - \beta_s)\, w_{Mj} \xi_{sj} L_{sj} + \beta_s L_j P_{sj} \left[\bar{c}_s + \frac{\gamma_s \left(\frac{\tilde{X}_j^F / L_j}{P_s} \right)^{\alpha_s + 1}}{\sum_{s'} \gamma_{s'} \left(\frac{\tilde{X}_j^F / L_j}{P_{s'}} \right)^{\alpha_{s'}}} \right] \right\}. \tag{9}$$

The labor market clearing condition requires that

$$L_{Ai} + L_{Mi} + L_{Si} = L_i, \quad i = 1, ..., N. \tag{10}$$

The world general equilibrium can be characterized by means of the following definition:

[18] All workers within a country are assumed to have the same expenditures. For example, if wedges capture differences in take-home wages across sectors, then agents pool incomes in their extended families whose sectoral employment is representative of the entire economy. Without this assumption aggregation across consumers would be more cumbersome because the indirect utility (1) is not of the Gorman polar form.

Definition 1. Given technology and preference parameters, labor wedges $\{\xi_{Ai},\, \xi_{Si}\}_{i=1}^{N}$, labor endowments $\{L_i\}_{i=1}^{N}$, trade costs $\{\tau_{Aji},\, \tau_{Mji}\}_{i=1,..,2;\, j=1,...,N}$, and trade deficits $\{D_i\}_{i=1}^{N}$ satisfying (8), the world equilibrium can be summarized as a collection of manufacturing wages $\{w_{Mi}\}_{i=1}^{N}$ and labor allocations $\{L_{Ai},\, L_{Mi},\, L_{Si}\}_{i=1}^{N}$ such that (i) goods markets (9) clear and (ii) the labor market clearing condition (10) is satisfied.

3 Data and Calibration

In this section I describe the calibration procedure that maps the theoretical model to the data. The calibration is designed to exactly account for the movement of labor across broad sectors in the data, which is the central feature of the process of structural change. This approach allows me to develop a model-based decomposition of overall labor relocation into contributions of different forces in Section 4.

The calibration strategy involves two main steps. In the first step, I take the observed data on employment, output and trade and use the general equilibrium structure of the model to back out quantities that are not measured directly but are required to simulate the model. Results of this step depend on the assumed values of the preference parameters. In the second step, I use the time-series predictions of the model for sectoral productivity growth from the first step to determine the preference parameters. Before moving on to the details of the calibration methodology, I present a brief description of the data. Detailed discussion of data sources and construction of the variables is relegated to the Data Appendix.

3.1 Data Overview

To capture secular trends associated with structural change, data containing consistent long time series is required. The availability of such data at a sectoral level is rather limited, especially for developing countries. To maximize the size of the sample while maintaining acceptable data quality standards, I combine sectoral data from four sources: EU KLEMS database, GGDC 10-sector database, OECD STAN database and Asian Productivity Organization database. The result is an unbalanced panel of between 26 and 44 countries over the period 1970-2005. These sources provide consistent and comparable series for total employment, gross value added in current prices and value added price deflators. The underlying data is aggregated to three sectors: agriculture, manufacturing and services. All data is smoothed using a Hodrick-Prescott filter with smoothing parameter 25 before it is fed into the calibration.[19]

International trade data comes from the NBER-UN databse (Feenstra et al. (2005)) and the

[19]Smoothing mitigates the problem of implausible short-term productivity jumps predicted by the model due to nominal exchange volatility rate and business-cycle fluctuations. Appendix D shows that the main results of this paper do not depend on smoothing, however, since both the calibration procedure and counterfactual simulations focus on long run changes.

BACI database (Gaulier and Zignago (2010)).[20] I use the SITC to ISIC concordance from WITS (subject to minor adjustments) to map trade flows into agriculture and manufacturing. To reconcile the trade data recorded in gross output terms and production data recorded in value added terms, I calculate the value of gross production by dividing the VA by the share of VA in gross output β_s.[21] Finally, the data on GDP at constant international prices and the level of exchange rates is sourced from version 7.0 of the Penn World Table (Heston et al. (2011)).

3.2 Calculating Sectoral Productivity Levels

In this subsection I show how the data on sectoral employment levels, value added, trade and aggregate productivity is used to determine wages, wedges and sectoral price levels and productivities in the model. These quantities are required for counterfactual simulations that are the basis of decompositions in Section 4. The discussion for now takes the preference parameters $\{\alpha_s, \gamma_s, \bar{c}_s\}$ as fixed. How those parameters are determined will be discussed in Section 3.3.[22]

Since discussions of sectoral labor productivity feature prominently in the following sections, it merits some explanation what exactly is understood by that measure. First, since technologies in the model combine labor and intermediates we can define the "multi-factor" productivity as

$$B_{si} \equiv \Gamma_s^{-1} T_{si}^{1/\theta_s} \pi_{sii}^{-1/\theta_s}, \tag{11}$$

where π_{sii} is the share of domestic varieties in total expenditure on sector s in country i. In a closed economy $\pi_{sii} = 1$ and B_{si} is simply the average efficiency $z_{si}(h)$ across the intermediate goods producers. In tradable sectors of an open economy MFP also captures the selection effect, in that varieties in which country i is not productive enough are not produced domestically but are imported instead.[23] Holding the state of technology in country i fixed, an increased penetration by imports would lead to higher measured multi-factor productivity. The general equilibrium structure of the model allows us to write the MFP as $B_{si} = (w_{si}/P_{si})^{\beta_s}$. Finally, using the fact that value added is a constant share β_s of gross output in industry s we can define labor productivity as

$$A_{si} \equiv B_{si}^{1/\beta_s} = \frac{w_{si}}{P_{si}}. \tag{12}$$

Conditional on sectoral wages there is a one-to-one mapping between sectoral price levels and

[20] Services are treated as nontradable because detailed data on trade in services for a broad range of countries is very limited. For the period under consideration in this paper trade in services accounts for about 20% of world trade. Trade in services also represents only about 5% of value added in services, compared to 15% for agriculture and 70% for manufacturing.

[21] Those shares are calculated as the median share of VA in sector's gross output for the subsample of countries for which I have the required data (EU KLEMS subsample), yielding $\beta_A = 0.50$, $\beta_M = 0.33$, $\beta_S = 0.57$.

[22] The calibration algorithm described in this and next subsection is presented with additional details in Appendix C.

[23] See Finicelli et al. (2013) for an argument why (11) is the correct measure of MFP in the Eaton and Kortum (2002) model.

sectoral labor productivities in the model.

In order to calculate wages, wedges and productivities across countries, the model matches the following data by design:

i) Sectoral employment levels L_{si}

ii) Sectoral nominal value added VA_{si}

iii) Trade flows in agriculture and manufacturing X_{Aji}, X_{Mji}

iv) Aggregate productivity (real GDP per worker) y_i.

Since structural change is most often characterized in terms of changes in sectoral employment and output shares, it is natural to target (i) and (ii) directly. Allowing the full model to exactly match the process of structural change facilitates evaluating the contribution of individual channels to that process in the later sections.

In the data there are often sizable differences between sectoral employment shares and output shares. To simultaneously account for both, any model with labor as the only factor of production requires wage differentials across sectors. Specifically, consistency of the model with both sectoral employment and value added requires that the wedge in sector $s \in \{A, S\}$ is calculated as

$$\xi_{si} = \frac{VA_{si}/L_{si}}{VA_{Mi}/L_{Mi}}. \tag{13}$$

Thus wedges in the model are identified from differences in value added per worker across sectors in the data.

There are two potential concerns with taking the model so literally when calculating the wedges. First, differences in value added per worker can be also capturing differences in factor intensities across sectors. In Appendix E.1 I illustrate that the wedges are similar when based on labor compensation per worker instead and I discuss the broader implications of abstracting from physical capital in a context of structural change. Second, differences in labor costs can be also capturing differences in human capital per worker across sectors. The magnitude of latter differences is debated in the recent literature, given a dearth of reliable data on human capital at a sectoral level in developing countries. I discuss this issue further in Appendix E.2, where I also illustrate the effects of making an extreme assumption that the wedges can be entirely explained by human capital differences across sectors. Finally, it is worth stressing in relation to both concerns that changes in wedges over time rather than their absolute levels are of main interest in this paper since my focus is on structural change within countries.

When countries trade, sectoral value added need not equal sectoral final consumption expenditures.[24] Let E_{sj} denote per worker final consumption expenditure on aggregate output of sector s.

[24]Here sectoral final consumption captures sectoral value added generated in different countries. A distinct issue

It can be verified that matching nominal trade flows (iii) in addition to (i) and (ii) implies via the goods market clearing conditions (9) that manufacturing wages and sectoral expenditures are given by: [25]

$$\begin{aligned} w_{Mi} &= VA_{Mi}/L_{Mi} \\ E_{si} &= \left(VA_{si} + \sum_j X_{sij} - \sum_j X_{sji} \right) /L_i. \end{aligned} \tag{14}$$

With sectoral expenditures determined, I then use the functional form of preferences and data on aggregate productivity to pin down sectoral price levels. To find three sectoral price levels $\left\{ P_{Aj},\, P_{M_j},\, P_{Sj} \right\}$ for each country j, I use $3N$ conditions that prices must satisfy. First, sectoral prices must be such that given those prices consumers optimally choose sectoral expenditures E_{sj} calculated in (14). Formally, sectoral prices $\left\{ P_{Aj},\, P_{M_j},\, P_{Sj} \right\}$ must be consistent with sectoral expenditure share equations:

$$\frac{E_{sj}}{\sum_{s'} E_{s'j}} = \frac{1}{\sum_{s'} E_{s'j}} \left[P_s \bar{c}_s + \left(\sum_{s'} E_{s'j} - \sum_{s'} P_{s'} \bar{c}_{s'} \right) \frac{\gamma_s \left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'} \bar{c}_{s'}}{P_s} \right)^{\alpha_s}}{\sum_{s'} \gamma_{s'} \left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'} \bar{c}_{s'}}{P_{s'}} \right)^{\alpha_{s'}}} \right]. \tag{15}$$

Since expenditure shares sum to one this restriction gives two independent equations for each country. To find three prices of sectoral output for each country I therefore need an additional set of restrictions. I use data on aggregate productivity - target (iv) above - as a source of those additional restrictions. Specifically, I calculate the real GDP per worker in the model using methodology that is analogous to one applied in the development of the PWT.[26] The required restriction is that the model measure of aggregate productivity must match PPP-adjusted GDP from PWT 7.0 divided by total employment L_i.

To summarize, the result of calculations described in this subsection is a set of sectoral wages and prices (and hence sectoral labor productivity levels by (12)) such that the model matches exactly the data on sectoral employment levels, nominal VA, trade flows and aggregate real GDP for all years and all countries in the sample. The model does not match by design the data on sectoral labor productivity growth so these series can be used to determine the remaining parameters of the model.

arises in the presence of cross-sectoral input-output linkages. In that case final expenditures measured in consumer prices could diverge from value added also in a closed economy. The expenditure shares in the paper are therefore conceptually different from final demand shares that can be directly calculated from input-output tables. In terms of the terminology of Herrendorf et al. (2013b), I am using the consumption value added approach.

[25] Nominal variables are rescaled in every year so that manufacturing wage in the US equals one.

[26] See Appendix C.1.

3.3 Calibration of Preference Parameters

The procedure for calculating sectoral productivity in the previous subsection took the preference parameters $\{\alpha_s, \gamma_s, \bar{c}_s\}_{s\in\{A,M,S\}}$ as given. In this subsection I discuss how those parameters are selected.

I begin with some restrictions imposed by consumer theory and normalizations. Ensuring that ACDES preferences are well-behaved requires $\alpha_s \geq -1$, $\gamma_s > 0$, $\sum_s \gamma_s = 1$.[27] As a choice of units I normalize all sectoral prices to one in the US in the reference year 1995. Given $\{\alpha_A, \alpha_M, \alpha_S, \bar{c}_A, \bar{c}_M, \bar{c}_S\}$, the preference weights $\{\gamma_A, \gamma_M, \gamma_S\}$ are then pinned down by U.S. expenditure shares (15) in 1995 for two sectors and a normalization $\gamma_A + \gamma_M + \gamma_S = 1$.

This leaves six preference parameters $\{\alpha_A, \alpha_M, \alpha_S, \bar{c}_A, \bar{c}_M, \bar{c}_S\}$ to be determined. Those parameters are pinned down using the model's prediction for sectoral labor productivity growth over time. Assuming that the difference between productivity growth in the model and the data is the result of measurement error, the preference parameters are chosen to minimize a GMM function of the measurement error.

Specifically, following the procedure described in the previous subsection for any candidate parameter vector $\omega = \{\alpha_A, \alpha_M, \alpha_S, \bar{c}_A, \bar{c}_M, \bar{c}_S\}$ I calculate sectoral labor productivities $A_{sit}(\omega)$ for each year in which country i is in the sample. Let t_l^i and t_f^i denote the last and first year that country i appears in the sample. Then the annualized average log growth of A_{sit} is calculated as $g_{si}(\omega) = \frac{1}{t_l^i - t_f^i} \log\left(\frac{A_{sit_l^i}(\omega)}{A_{sit_f^i}(\omega)}\right)$. Analogous log growth of labor productivity computed from the data is denoted as g_{si}^d. Because sectoral productivity series in the data are calculated using sectoral producer price deflators that are likely to suffer from measurement error, there will be a discrepancy between the model's predictions for sectoral productivity growth and their empirical counterpart. Stated formally

$$g_{si}^d = g_{si}(\omega_0) + \varepsilon_{si}, \quad s \in \{A, M, S\},$$

where ω_0 is the true data-generating value of the parameter vector. The main assumption is that ε_{si} is a mean-zero random measurement error uncorrelated with sectoral employment and expenditure growth. The precise moment conditions are

$$E\left[x_{si}^{(m)}\varepsilon_{si}\right] = 0, \quad s \in \{A, M, S\}, \ m = 1, ..., 3, \tag{16}$$

where the instruments x_s for sector s log productivity growth include a constant, log growth in sector s employment and log growth in expenditure share of sector s (all growth rates on an annualized basis). The sample size n equals the total number of countries appearing in the sample. Let $h_n(\omega)$

[27] See Appendix A.

be the vector of sample analogs of moment conditions (16):

$$h_n(\omega) = \left[\frac{1}{n} \sum_{j=1}^{n} x_{Aj}^{(1)} \left(g_{Aj}^{d} - g_{Aj}(\omega) \right) \ldots \frac{1}{n} \sum_{j=1}^{n} x_{Sj}^{(3)} \left(g_{Sj}^{d} - g_{Sj}(\omega) \right) \right]'.$$

The calibrated parameter vector $\hat{\omega}$ then minimizes the following objective function:[28]

$$\hat{\omega} = \arg\min_{\omega} n \cdot h_n(\omega)' h_n(\omega). \tag{17}$$

3.4 Calibration Results

Having described the calibration procedure, I now discuss its results and performance, focusing on the consumer preferences.

Parameters solving problem (17) are presented in the first column of Table 1. The second column shows the results under the restriction $\overline{c}_M = \overline{c}_S = 0$. The latter is the baseline calibration of the paper. There are two reasons why the restricted model is preferred. The first reason is statistical: looking at the bootstrapped standard errors of $\hat{\overline{c}}_M$ and $\hat{\overline{c}}_S$ shows that these parameters are not precisely calibrated. Moreover, setting subsistence requirement to zero in manufacturing and services has minimal effect on the fit of the model as measured by the value of the GMM objective function.[29] The second reason is that while the additional parameters increase the complexity of the model, $\hat{\overline{c}}_M$ and $\hat{\overline{c}}_S$ at magnitudes calibrated in column 1 have little economic impact. To show this, Table 2 reports what the preference parameters (which by themselves might be hard to interpret) imply for income, price, and substitution elasticities, and for the importance of subsistence in overall consumption (all averaged across countries in the reference year 1995). Comparing the elasticities in the first panel (unrestricted model) and the second panel (baseline model) shows a very similar pattern. This is because the calibrated subsistence consumption as a share of overall consumption is only 0.3% for manufacturing and 2.3% for services. In contrast, subsistence requirement accounts for almost 2/3 of agricultural consumption on average (and for 90% in poorest countries).[30] The baseline calibration thus retains a positive $\overline{c}_A$, as it is economically important, precisely calibrated

[28] As a robustness exercise Appendix D performs a two-step GMM estimation.

[29] It might seem natural to take one more step and conduct a formal hypothesis test of $\overline{c}_M = \overline{c}_S = 0$. Unfortunately, technical difficulties with hypotheses testing arise when some of the estimated parameters are at the boundary of the parameter space, as is the case with here with a binding constraint $\alpha_A \geq -1$. While the GMM estimator remains consistent when the true parameter is on a boundary, its distribution is nonstandard (Andrews (2002)). Neither standard asymptotic formulas nor bootstrap provide consistent estimators of standard errors in this case (Andrews (2000)) While the theoretical econometrics literature suggests some solutions to inference in the presence of binding boundary constraints (Andrews (1999)), they require computations that are impractical in the context of this paper. For this reason, my procedure remains a calibration rather than a full estimation. This is erring on the side of caution: $\overline{c}_A = 0$ and $\overline{c}_S = 0$ would not be rejected using standard tests.

[30] The calibration procedure ensures that consumption exceeds the subsistence requirement for all years and countries.

and substantially improves the ability of the model to fit the data.[31]

Looking closer at the properties of the demand in the baseline calibration in Table 2 shows that both income and relative price effects are clearly at play. All elasticities of substitution significantly less than one hint at a potent relative price effect. Strong nonhomotheticity is illustrated by large differences in income elasticities across sectors. To illustrate the nonhomotheticity further, the first panel of Figure 1 plots the implied expenditure share Engel curves for the US and India. Nonhomotheticity can be seen directly from large changes in expenditure shares induced by empirically relevant changes in expenditure levels. Moreover, as the second panel of Figure 1 illustrates, even the allocation of marginal expenditure across sectors depends on the expenditure level. Comparing the Engel curves for US and India might suggest that income differences alone can go a long way towards explaining differences in sectoral patterns between rich and poor countries. Despite having similar shapes, the Engel curves for the two countries are nevertheless not identical, suggesting that relative price effects also play a role. In the next section I quantify the relative importance of the income and relative price effects, as well as of international trade and intersectoral wedges for structural transformation.

But before doing that, I discuss the performance of the calibration procedure. Figure 2 illustrates how well the calibrated model fits the data by plotting the annualized growth rates of sectoral labor productivity predicted by the model and calculated in the data. The objective function in the calibration procedure puts a lot of weight on minimizing the distance between the two statistics. As can be seen in the figure, the model with the calibrated parameters achieves that goal well. The correlation between the annualized productivity growth in the model and in the data is 0.81, 0.85 and 0.92 for agriculture, manufacturing and services, respectively.

It is worth contrasting my baseline results for ACDES preferences with subsistence consumption in agriculture with results that would be obtained under augmented CES preferences used in recent quantitative studies of structural change (e.g., Buera and Kaboski (2009), Herrendorf et al. (2013b), Uy et al. (2013)). Recall that such preferences are nested by my general ACDES structure as a special case with $\alpha_A = \alpha_M = \alpha_S$. Preference parameters calibrated under this restriction, reported in the third column in Table 1, point to (augmented) Leontief preferences ($\alpha = -1$).[32] Comparing the value of the GMM objective function in columns 2 and 3 shows that the ACES model provides a worse fit to the data than the baseline model. This is a fair comparison since in both cases four preference parameters are calibrated. The baseline model fits better because it is more flexible in two dimensions. First, it allows a richer substitution pattern. While in the Leontief case there is no substitution at all between any goods, the baseline model even allows some pairs of goods to be

[31] Starting from the baseline calibration and imposing $\overline{c}_A = 0$ (i.e. standard CDES preferences) results in significant deterioration of the model fit, with the GMM objective function increasing by three orders of magnitude.

[32] Stone-Geary preferences, also common in the literature (e.g., Caselli and Coleman (2001) and Kongsamut et al. (2001)), are themselves a special case of augmented CES with $\alpha = 0$. The coefficient $\alpha = -1$ is precisely calibrated, thus Stone-Geary specification is strongly rejected in my data.

complements and others to be substitutes (cf. Table 2).[33] Second, nonhomotheticity is also more flexible in the baseline model in that it does not disappear for high incomes. In contrast, in the ACES model the allocation of marginal expenditure is independent of income and preferences effectively become homothetic for sufficiently high incomes. Calibrated to the same data for a diverse set of countries, the ACES model thus generates stronger nonhomotheticity for poor countries and weaker for rich countries.

These differences explain why the worse performance of the ACES model is driven largely by its worse fit for poor but fast growing countries, most notably China. In particular, the ACES model overpredicts by more than 5 p.p. the annualized labor productivity growth in services in China. This is because the model tries to engineer a counterfactual decline in relative price of services in order to reduce the expenditure share of that sector. A force for such reduction is necessary, because with very high income elasticity of services (3.2 compared to 1.6 in the baseline) in China starting poor in 1978, and spectacular subsequent income growth, the income effect alone generates more than sufficient increase in the expenditure share of services in China.

The finding of very little substitutability in consumption across calibrations might appear extreme. It is, however, consistent with comparable estimates in other studies. Both Buera and Kaboski (2009) and Herrendorf et al. (2013b) find that augmented Leontief preferences provide the best fit to the long time series of US data. While Uy et al. (2013) find higher elasticities of substitution based on Korean data, they use a different commodity space. In terms of terminology of Herrendorf et al. (2013b), I use a consumption value added approach while Uy et al. (2013) use a consumption final expenditure approach.[34] Herrendorf et al. (2013b) explain why one should expect more substitutability when sectoral commodities are defined in terms of final goods capturing value added from all sectors via input-output linkages.[35] Nevertheless, the prevalence of corner solutions in the consumption value added approach (Leontief in studies mentioned and $\alpha_A = -1$ in this paper even in the general ACDES case) suggests that even more flexibility might be required to explain the observed patterns in a fully satisfactory manner.

My calibration procedure makes use of both cross-sectional and time-series data for a large number of countries. This breadth offers some benefits relative to an alternative common approach, in which data for a single country is used to calibrate preference parameters that are subsequently imposed in other countries. The outcome of the single-country approach can be sensitive to the choice of the reference country. For example, restricting my procedure to target only the fit for Korea would result in preferences for which consumption in agriculture responds little to both income and

[33] Allowing for Allen complements is also the reason Hanoch (1975) considers the CDES model to be more flexible than the CRES model recently applied to a structural change setting by Comin et al. (2015).

[34] Using a final expenditure approach Boppart (2014) finds an elasticity of substitution between goods and services in the US to be under 0.6.

[35] To give an example, in final expenditure approach a consumer might have a choice between a restaurant meal (services) or food purchased at a farmers market (agriculture), which can be expected to be somewhat substitutable. In value added approach the choice would be between food ingredients produced in agriculture and cooking services, which might not be very substitutable.

relative price changes. However, those preferences would provide poor fit for many countries in the sample. At the same time, the fit for Korea with the baseline parameters is still reasonably good.[36] This asymmetry suggests that it might be difficult to precisely determine the preference parameters by relying on aggregate data for a single country. The calibration strategy used in this paper puts equal weight on information for all countries in a diverse sample, which increases its power to identify unknown parameters assumed to be common across countries.[37] On the other hand, focusing on a single country has the advantage of increased availability of long time-series of high-quality data (such as investment data missing in this paper).

The final exercise of this section is offered to directly illustrate that the calibration strategy can indeed recover the unknown model parameters. For this purpose, I explicitly specify all common and country-specific parameters (of preferences, technology, wedges, trade costs) and simulate the model with those parameters. Then I add random measurement error to the true sectoral labor productivity growth, consistent with assumptions behind the calibration procedure.[38] That procedure is then performed using the simulated series for the same data (on employment, value added, trade flows, real GDP and mismeasured sectoral labor productivity growth) as used in the actual calibration. Table 3 summarizes the results of 200 such simulations. Calibrated preference parameters move within reasonably narrow bands around the true values, with median values almost spot on. These simulations suggest that if the model of this paper is well specified, then the calibration methodology is likely to recover its true parameters accurately.

4 Quantifying Importance of the Four Channels

The model developed in Section 2 incorporates four forces that have been suggested in the literature as important for understanding structural change: (i) sector-biased technological progress, (ii) nonhomothetic tastes, (iii) international trade and (iv) intersectoral wedges. In this section I use model-based counterfactual simulations to quantify the relative importance of these four forces in a unified setting.

While the process of structural change involves many changes in the sectoral constitution of the economy, changes in sectoral labor shares are arguably the most important aspect of the process. It

[36]The value of the GMM objective function (17) would increase by a factor of 10^5 when using the Korea-tailored parameters rather than the calibrated parameters. In contrast, the value of the Korea-specific version of the objective function would increase by a much lower factor of 10^2 when using the baseline parameters rather than Korea-tailored parameters.

[37]Household-level expenditure surveys could also be used as a source of substantial variation in income (across households) and relative prices (across time or across markets). But such micro-level data would only be directly suitable for estimating preferences in the final expenditure approach rather than the value added approach taken in this paper. Mapping the household expenditure data to value added is feasible but requires detailed input-output tables, as in a recent application by Buera et al. (2015).

[38]Parameter values for the simulation (including the measurement error covariance matrix) are chosen to be roughly consistent with values corresponding to the baseline calibration. See Appendix C.3 for details. Note that in the absence of any source of error the calibration procedure can recover the initial parameters exactly.

is therefore natural to assess the importance of various drivers of structural change by quantifying the amount of labor relocation they can account for. To make this exercise transparent, the baseline model is calibrated in the previous section to match the evolution of sectoral employment in the data exactly when all the four forces are operative. That is, the model can explain all the observed labor relocation through a combination of rising incomes and changes in relative productivities across sectors, trading patterns and intersectoral wedges. By switching various channels on and off through counterfactual simulations of the model and checking how the labor relocation is affected we can therefore gauge their importance for structural change.

To measure how much of the labor relocation is captured by a given counterfactual scenario, I introduce the Labor Relocation Index defined as

$$LRI_i \equiv 1 - \frac{\left|\Delta\overline{l_{Ai}^{cf}} - \Delta\overline{l_{Ai}}\right| + \left|\Delta\overline{l_{Mi}^{cf}} - \Delta\overline{l_{Mi}}\right| + \left|\Delta\overline{l_{Si}^{cf}} - \Delta\overline{l_{Si}}\right|}{\left|\Delta\overline{l_{Ai}}\right| + \left|\Delta\overline{l_{Mi}}\right| + \left|\Delta\overline{l_{Si}}\right|}, \tag{18}$$

where $\Delta\overline{l_{si}} = \frac{1}{t_l^i - t_f^i}\left(l_{st_l^i} - l_{st_f^i}\right)$ is the measured annualized change in labor share of sector s in country i between the last and first year the country appears in the sample.[39] Similarly, $\Delta\overline{l_{si}^{cf}}$ measures the change in the counterfactual simulation. The Labor Relocation Index has the following properties:

1. $LRI \leq 1$; $LRI = 1$ when the counterfactual captures the labor reallocation in the data exactly.

2. $LRI = 0$ if the counterfactual predicts no labor reallocation whatsoever.

3. LRI can take negative values if the counterfactual predicts changes in labor shares much larger than observed or going in the wrong direction.

When LRI takes positive values it can be thus interpreted as the fraction of observed changes in labor shares that can be explained by the forces operative in the counterfactual scenario.[40]

There are two natural ways in which we can measure the contribution of a given channel to labor relocation. First, we can ask how much labor relocation we would see if that channel was switched off. For example, if LRI takes low values in the absence of sector-biased productivity growth then it means that this force is important in accounting for structural change. Second, we can ask how much of the labor relocation can be generated by the channel operating on its own. If we only keep nonhomotheticity and switch off three other channels and yet the counterfactual explains a lot of

[39] Changes in labor shares are scaled by the time they are calculated over to make the measure comparable across countries in an unbalanced panel.

[40] Herrendorf et al. (2013b) and Uy et al. (2013) take an approach where neither labor nor expenditure shares are matched exactly in their baseline case. Then they compare the loss of fit of expenditure or labor shares under counterfactual simulations relative to their baseline, which does not have the intuitive decomposition interpretation of the LRI.

labor movement then the nonhomotheticity should be deemed important. Because there are interactions between the channels in the model, the two approaches might present somewhat different picture. Below I present both sets of results as well as the intermediate cases, which allows me to draw some general conclusions on the relative importance of the four determinants of structural change. First I describe in more detail how switching off individual channels is implemented in the counterfactual simulations.[41][42]

Sector-Biased Technological Change

The "Baumol's disease" take on structural change starts with the observation that sectors with low productivity growth see a rise in the relative prices of their output over time. With elasticity of substitution across sectors less than one, this translates into an increase in expenditure share, and consequently also labor share, of the slow-growing sectors over time. At a broad level, we can clearly see the correlation between labor productivity growth and changes in sectoral shares. The average annualized log growth of labor productivity in agriculture, manufacturing and services is 0.037, 0.034, 0.013 [log points]. The corresponding annualized change in labor share is -0.46, -0.16, 0.61 [p.p. p.a.]. The fast-growing agriculture sector sheds labor while slow-growing services expand.

To eliminate the effect of sector-biased technological change, I set the sectoral productivity growth in a country in a way that keeps relative labor productivities approximately constant over time. Specifically, in the first year a country appears in the sample it looks exactly the same in the counterfactual as in the baseline. I then set the growth rate of fundamental technology T_{sit} to $g_{it}^{\theta_s \beta_s}$, where g_{it} is the growth rate of real GDP per worker. This choice ensures that labor productivity in each sector and the aggregate productivity grow at a rate approximately g_{it}, so that the income channel is still operative.[43]

Nonhomotheticity

The calibrated preference parameters imply that the income elasticity is lower than one for agriculture and higher than one for services. To eliminate the effect of this nonhomotheticity in a given country, I replace the common ACDES preferences with Leontief preferences in that country. More

[41] To conduct counterfactual simulations I need to specify values of the parameters governing dispersion of productivity draws in the tradable sectors. I set $\theta_A = \theta_M = 5$ in line with consensus estimates of the trade elasticity (cf. Head and Mayer (2014)), to which these parameters are related. Productivity dispersion in nontradable services does not affect the calculations.

[42] All counterfactuals are simulated country-by-country. E.g., I simulate the evolution of the world equilibrium with sector-biased productivity growth eliminated separately in each country while fundamentals in all other countries evolve as in the baseline calculation. The results are similar if we consider changes to all countries simultaneously, in cases where such counterfactuals are meaningful.

[43] Because fundamental productivity changes relative to the baseline, international trade flows and hence π_{sii} are affected. Therefore the realized labor productivity growth in the tradable sectors is not exactly g_{it}, as it also captures the endogenous part of productivity arising due to specialization in trade (recall from (11) and (12) that labor productivity $A_{sit} = \Gamma_s^{-1} T_{sit}^{(\theta_s \beta_s)^{-1}} \pi_{siit}^{(\theta_s \beta_s)^{-1}}$). This extra effect is conceptually distinct form the differential exogenous growth of fundamental technology, and in practice its magnitude is very small.

precisely, in the first year a country appears in a sample I find sectoral weights in the Leontief utility function such that the model replicates the baseline equilibrium. From then on I recalculate the counterfactual equilibrium with Leontief preferences, given the evolution of fundamental productivity, trade costs and wedges as in the baseline case.

The choice of Leontief specification might appear peculiar as it eliminates substitution possibilities across sectors in addition to nonhomotheticity. There are two reasons for this choice. First, even with baseline ACDES preferences there is very little substitutability as evidenced in Table 1. Second, constraining the calibration of the model to homothetic CES preferences results in fact in the corner case of the Leontief specification, as will be discussed further in Section 5.

International Trade

Recent literature suggests that in a globalizing world domestic factors might not be sufficient to explain the observed patterns of structural change. For example, export-led industrialization might justify a higher manufacturing labor share than would be observed in a closed economy. I shut down this channel by closing the economy to international trade. Formally, I set the counterfactual trade costs to infinity: $\tau'_{sji} = \infty$.[44][45]

Intersectoral Wedges

The allocation of labor across sectors is also affected by domestic factor cost differentials that are captured by the intersectoral wedges. To eliminate the effect of changes in those wedges over time on labor relocation I simply keep the wedges at their initial year level. More explicitly, I set $\xi'_{sjt} = \xi_{sjt_f^j}$ in the counterfactual simulations that eliminate the influence of wedges.

4.1 Shutting Down One Channel at a Time

The first set of counterfactuals illustrates the importance of various drivers of structural change by shutting them down individually. The first panel of Table 4 presents the median Labor Relocation Index for each of the four simulations. Since a lower value of LRI indicates a larger loss of fit

[44] Shutting down trade completely clearly isolates its effect but has a slight disadvantage that the labor allocation in the initial year is different than in the baseline equilibrium (unlike when other channels are eliminated individually). An alternative approach would be to focus only on eliminating changes in trade costs relative to the first year. This approach is only feasible, however, in a balanced panel of countries because when a country disappears from the sample its bilateral trade costs effectively become prohibitive.

[45] While explicitly specifying trade costs in the baseline equilibrium is not necessary to derive any results in this paper, trade costs that rationalize the observed trade flows can be backed out using the model. The average ad-valorem equivalent trade cost implied by the model is 270% (i.e. $\tau = 3.70$) in agriculture and 97% ($\tau = 1.97$) in manufacturing. These magnitudes are consistent with the range estimated in the literature. For example, the headline trade cost in Anderson and van Wincoop (2004) is 170%. Tombe (2015) estimates trade costs separately for agriculture and manufacturing and also finds that trade costs are systematically higher in agriculture (250% on average) than in manufacturing (150% on average). This similarity gives additional credence to the calibration procedure as the magnitude of trade costs was not used as a calibration target.

from eliminating a given channel in this calculation, lower numbers correspond to more important channels. Sector-biased productivity growth is found to be the most important channel by this measure: with productivity growth forced to be uniform across sectors, the remaining channels can account for only 46% of observed labor relocation for a typical country. Changes in intersectoral distortions take a second spot: with labor costs fixed at the initial level the model can explain 65% of labor share changes. The penalty for closing the economies to international trade is a little smaller with a median LRI of 0.71. Income effects turn out to be the least important in this calculation: forcing homothetic preferences results in a median loss of explanatory power of 20%.

The median effects reported in Table 4 naturally mask substantial heterogeneity across countries.[46] Columns (1) to (4) of Table 5 show the corresponding results for individual countries. Since the relative importance of different forces varies across countries, having a broad sample allows me to draw more general conclusions than the prior literature that focused on a couple of channels in a single country. While the median LRI summarizes how important a given channel is for explaining labor relocation associated with the structural change process, we can gain additional understanding about the working of individual channels by looking at more disaggregated numbers.

Figure 3 plots, for each country, the average change in labor share in the three sectors in the counterfactual scenario against the baseline. In the first panel we see the effects of eliminating the sector-biased productivity growth channel. With a few exceptions, the changes in the agricultural labor share are captured relatively well in this scenario. However, the ability to capture the change in labor share of manufacturing and services is compromised. Importantly, there is a systematic pattern to the deviations: without sector-biased technological growth the model underpredicts the movement out of manufacturing and understates the increase in the labor share of services.[47] For example, the low $LRI = 0.54$ in the US (cf. Table 5) reflects the fact that in this scenario the manufacturing labor share declines half as much as in the data (from 0.23 to 0.16 rather than to 0.11), and correspondingly the services share rises just half as much (from 0.74 to 0.81 rather than to 0.88).

The second panel of Figure 3 illustrates analogous results for the case of imposing homothetic preferences. In contrast to the previous scenario, the changes in manufacturing labor shares are captured extremely well. Thus nonhomotheticity appears to have little effect on the share of industry in employment. Instead, the model now has problems capturing the net movement of labor out of agriculture and towards services. The counterfactual consistently understates the decline of agriculture and the rise of services in terms of employment shares.[48] This understatement is par-

[46] In the discussion I focus on medians rather than means since the former are less skewed by instances of large negative LRIs. LRI is not symmetric in that it can take at most the value of one but can take large negative values. Throughout this section the ranking of the results is almost always the same whether we look at means or medians, however.

[47] The average change in the labor share of agriculture, manufacturing and services is -0.46, -0.16, 0.61 [p.p. per year] in the baseline and -0.43, -0.01, 0.44 in the counterfactual without sector-biased technology growth.

[48] The average change in the labor share of agriculture, manufacturing and services is -0.23, -0.16, 0.39 [p.p. per year] in this counterfactual.

ticularly pronounced for countries starting at low levels of development. A look back at the second column of Table 5 confirms that LRIs in this scenario are low for the poor countries, while for the most developed countries eliminating nonhomotheticity has little impact.[49] For example, in the US ($LRI = 0.90$) the decline in agricultural labor share is understated, but it hardly matters because the agriculture labor share is only 0.032 to begin with and the understatement is by just 0.5 p.p. over the entire period 1970-2005. On the other hand, in Thailand ($LRI = 0.44$) without income effects the decline in agricultural share over the same period would be a whooping 20 p.p. lower than in the data (from 0.75 to 0.59 rather than to 0.39), with the corresponding shortfall appearing mostly in services (rise from 0.19 to to only 0.27 rather than to the actual 0.45).

Unlike in the first two cases, shutting down international trade or eliminating changes in wedges does not have a strong systematic bias. On average, changes in labor shares are closer to the data than in the two counterfactuals discussed above. The deviations are simply distributed more symmetrically. Since overall fit measured by median LRI deteriorates noticeably, the deviations can still be quite large. For example, in case of trade Column (3) of Table 5 shows that the fit deteriorates particularly for countries that are relatively open to international trade.[50] In the US (LRI=0.73) the manufacturing labor share would decline by 3.5 p.p. less (8.7 p.p rather than the actual 12.1 p.p.) between 1970 and 2005 if the US did not engage in international trade. To a large extent this reflects a ballooning manufacturing trade deficit in the US over this period (from less than 1% of manufacturing value added to 28%). Not surprisingly, closing trade has little impact on big countries that trade little, such as India ($LRI = 0.93$). On the other hand, India is an example of a country where changes in wedges have significant impact ($LRI = -0.01$). This is because labor costs in agriculture relative to other sectors decline in India over the sample period. When this channel is shut down, Indian agriculture becomes less competitive over time resulting in a large increase in agricultural imports in the counterfactual. Consequently, the decline in agricultural labor share in India is much higher in the constant-wedge counterfactual than in the data (from 0.72 to 0.50 rather than to 0.61). In the US relative labor costs in agriculture also decline, but given the small size of this sector its changing competitiveness has little effect on overall structural change ($LRI = 0.92$), with agricultural labor declining by 1 p.p. more and manufacturing share by 1 p.p. less than in the data.

4.2 One Channel Operative at a Time

An alternative method of quantifying the importance of a given driver of structural transformation is to evaluate how much of the observed labor relocation the driver can explain by itself. The counterfactuals reported in this subsection therefore keep only one force operative while shutting

[49] The correlation between the LRI and the log of real GDP per worker in the initial year is 0.69 and is statistically significant.

[50] The correlation between the LRI and the measure of trade openness defined as (exports + imports)/GDP in the initial year is -0.33 and is statistically significant.

down the remaining three channels. A high value of LRI in this exercise indicates that the only operative channel is important for structural change.

The second panel of Table 4 shows that the sector-biased productivity growth is again the most important force. For a median country that effect alone can explain 43% of the labor relocation across sectors. It also has the largest explanatory power of the four channels for more than half of sample countries, as can be checked in the disaggregated results in columns (5) to (8) of Table 5. The breakdown of labor changes by sector in the first panel of Figure 4 suggests that sector-biased technological progress can capture the average change in the labor share of manufacturing, but does not by itself generate enough net movement out of agriculture and into services on average.

Nonhomothetic preferences are the second most important force in this set of counterfactuals. Rising income coupled with income elasticities varying across sectors can account for 27% of actual labor relocation in a median country. Delving deeper into the sectoral labor relocation patterns illustrated in the second panel of Figure 4, we can see that income effects alone induce little change in the labor share of manufacturing. Consequently, there is not enough net relocation out of manufacturing and towards services, which is the main pattern observed at later stages of development when the share of agriculture is already very low.

Next consider counterfactual simulations with equal fundamental productivity growth in each sector, homothetic preferences and constant wedges. When international trade is the only force generating labor movement across sectors, it can account for 21% of labor relocation in the median country. At a more disaggregated level, the simulation systematically predicts a smaller decline (and often absolute increase) in the share of employment in manufacturing and a smaller rise in the share of services than in the data.

Finally, changes in intersectoral wedges alone do not induce any relocation of labor across sectors at all. To understand why, notice that changes in wedges capture changes in relative factor costs across sectors, which in turn translate to changes in relative prices. But with Leontief preferences, changes in relative prices have no effect on relative real consumption. Moreover, in a closed economy consumption is equal to production, so relative production does not change either. Since relative labor productivity is unchanged by construction, no movement of labor across sectors occurs.

4.3 Generalization: Marginal Effects of Channels

One of the conclusions of the last two subsections is that eliminating changes in wedges has a noticeable impact on evolution of labor shares but changes in wedges by themselves do not affect labor shares at all. This discrepancy illustrates a broader point - the "marginal" effect of a given channel on sectoral labor relocation depends on which other channels were already allowed to operate. Because of the interactions between channels, it is not possible to simply decompose total labor relocation to contributions of individual channels in an additive or multiplicative way. But the pattern of the marginal effects is nevertheless informative about the relative importance of the various channels.

The third panel of Table 4 shows the median LRIs for counterfactuals in which two channels at a time are switched on. Moving across the panels of this table allows us to trace the marginal effect of different channels. Consider, for example, starting with the model featuring only sector-biased productivity growth and progressively adding extra drivers of structural transformation. Figure 5 illustrates the possible paths to the full model graphically to ease exposition. The marginal benefit of adding nonhomotheticity is substantial - LRI increases from 0.43 to 0.70. Adding trade instead takes LRI to a somewhat lower 0.63. In contrast, adding changes in wedges does absolutely nothing as far as labor shares are concerned for the reason explained earlier. But changes in wedges are not always impotent - with sector-biased productivity growth effect and trade operative, adding changes in wedges increases median LRI form 0.65 to 0.80. In general, wedges are either irrelevant or unimportant in a closed economy setting but adding changes in wedges in an open economy setting has a noticeable positive impact on fitting the labor relocation.[51]

The contribution of trade depends on whether labor wedges are accounted for by the model. For example, starting with a model featuring sector-biased productivity growth and nonhomotheticity but constant wedges, adding trade actually decreases the fit in terms of median LRI from 0.70 to 0.65 because the implied trade patterns imply labor movement in directions not consistent with the data. In all other cases adding trade positively affects the ability of the model to account for labor share changes, however.

While the marginal benefit of switching on trade and changes in wedges is occasionally small or even negative, adding sector-biased productivity is always distinctly beneficial for the fit and adding nonhomotheticity is almost always beneficial.[52] Whether or not the model already features trade or changes in labor wedges, differential productivity growth and income effects are important for generating reallocations seen in the data.

4.4 Taking Stock

Whenever a meaningful pairwise comparison can be made, the median marginal benefit of adding sector-biased productivity growth is larger than that of any other channel. For this reason, I conclude that sector-biased technological change is overall the most important force explaining the relocation of labor across sectors for a broad sample of countries. At a more detailed level, sector-biased productivity growth is key for matching the movement of labor from manufacturing to services and is thus particularly important for countries at more advanced stages of structural transformation. This mechanism is less crucial for understanding the behavior of the agricultural labor share, the evolution of which matters most for countries at earlier stages of development.

[51] In a closed economy with nonhomothetic preferences the impact of changes in wedges on labor allocation is still minimal because the calibrated elasticities of substitution are close to zero so consumption responds little to change in relative prices induced by the change in wedges.

[52] The one exception is that starting from the model with sector bias and trade adding nonhomotheticity has only a minor positive impact.

Nonhomothetic preferences, while on average less important than sector-biased technological progress, remain a significant determinant of structural change. Income effects are particularly important for accounting for the decline of agricultural labor share, especially at early stages of economic development.[53] Nonhomothetic preferences have less power to match the movement of labor between manufacturing and services and are thus less crucial for accounting for developments in rich countries. It is worth emphasizing that this is a quantitative result and not merely a consequence of an assumed modeling framework. This would be the case, e.g., if the model specified Stone-Geary or augmented CES preferences, as is standard in the literature. These preferences converge to being homothetic and therefore income effects necessarily play no role once countries grow sufficiency rich. In contrast, the calibrated demand remains nonhomothetic even for the richest countries. The fact that nonhomotheticity is particularly evident in agriculture for poor countries is certainly related to the subsistence requirement $\bar{c}_A$. But the presence of such term is heavily favored by the data rather than imposed, as similar terms for manufacturing and services are found to be insignificant in Section 3.4.

The conclusion to be drawn from this discussion of the two classic channels is that income effects and sector-biased productivity growth complement each other and are both necessary to quantitatively account for structural change in countries at very different stages of the structural change process.

The role of international trade in accounting for intersectoral labor relocation manifests itself a little bit differently. Ignoring trade would not necessarily lead to a systematic bias in predicting changes in sectoral labor shares over time. Instead, it would lead to large idiosyncratic deviations from the observed patterns for individual countries, especially small open economies and countries in which the sectoral composition of exports or trade balance change significantly over time. Moreover, international trade is a catalyst for the influence of intersectoral wedges over sectoral labor shares. When countries are open to trade, changes in relative factor costs across sectors affect specialization in trade and hence sectoral labor shares. Thus once the model features international trade, changes in intersectoral wedges also become a quantitatively significant driver of structural change.

5 Implications of Homothetic Preferences

In this section I further explore the importance of nonhomothetic preferences for modeling structural change. In the last section I started with a nonhomothetic model that could account for sectoral labor relocation exactly. I then used counterfactual simulations in which the model was restricted to feature homothetic preferences. The loss of ability of the restricted model to account for labor relocation was used as a measure of importance of income effects. In this section I study the

[53] This result is consistent with Caselli and Coleman (2001), who find income effects to be important for transition out of agriculture in the US starting from 1880, when the structure of US economy resembled that of developing countries in my sample.

implications of ignoring nonhomothetic tastes in an alternative way. I recalibrate the restricted homothetic model so that it matches the same data as the full nonhomothetic model. I then compare the implications of the two models for variables not matched exactly. I show that the homothetic and nonhomothetic model provide different answers to some of the questions that quantitative models of structural change are used to address. In particular, the two models have quantitatively different predictions for patterns of sectoral productivity levels across countries and their growth rates over time. This in turn implies differences in results of development and growth accounting exercises, with the homothetic model overstating the importance of agriculture.

Formally, I redo the calibration described in Section 3 under the restrictions that all subsistence requirements are zero ($\bar{c}_s = 0$) and $\alpha_A = \alpha_M = \alpha_S = \alpha$. These restrictions imply that preferences are of the homothetic CES form with elasticity of substitution $\alpha + 1$. The corner solution of Leontief preferences ($\alpha = -1$) provides the best fit of the model among the CES class. This should not be surprising given that a corner solution with no substitutability is found even in the more general augmented CES model in Section 3.4. The intuition for the Leontief case comes from an observation that for a number of countries in the sample the relative real consumption of services increases over time just as the relative price of services increases. Since this behavior is not consistent with homothetic preferences, the best the CES calibration can do is to choose the lowest possible level of substitutability across sectors. The CES restriction is costly in terms of the model fit measured by the value of the objective function: from the value of 3.22×10^{-7} for ACDES (cf. Table 1) it goes up to 1.17×10^{-2} for CES.

5.1 Cross-Sectional Results

In this subsection I compare the cross-sectional patterns of sectoral labor productivities generated by the two calibrations. Comparing productivity levels at a sectoral level across countries directly is difficult since the required data on producer price based purchasing power parities at a sector level exists only for a handful of developed countries.[54] Structural sectoral models like the one developed in this paper are therefore often used to back out sectoral productivities from easily measured data, with Duarte and Restuccia (2010) serving as a recent example. Below I illustrate that the form of preferences used in this class of models has important implications for calculated productivities.

Figure 6 plots the labor productivity from the CES calibration relative to the number from the baseline ACDES calibration for each sector in 1995.[55] A robust pattern emerges: the ratio of labor productivities from the two models is strongly increasing in aggregate productivity for agriculture, strongly decreasing for services and modestly rising for manufacturing. Since the homothetic model predicts lower agricultural productivity in poor countries than the baseline model, and poor

[54]Arguably the best effort in that direction is the GGDC Productivity Level Database covering 30 OECD countries in 1997 (Inklaar and Timmer (2008)).

[55]I report the cross-sectional results for the benchmark year 1995 but the patterns I highlight are robust throughout the sample period.

countries have low levels of agricultural productivity to begin with, the CES model implies larger dispersion of A_A across countries than the baseline model. The pattern is the opposite for the service sector. In numbers, the coefficients of variation for A_A, A_M, A_S in the CES/ACDES cases are: 1.42/1.31, 0.74/0.71 and 0.52/0.57. The intuition behind those differences is simple. Both models need to match the same expenditure shares. For example, to account for high expenditure share of agriculture in poor countries the homothetic model needs a higher relative price of agriculture than the ACDES model since the latter can rely partially on a low income elasticity of agriculture. With the same wedges in both models, higher P_A for poor countries in the CES model translates into lower A_A, resulting in a higher dispersion of A_A than in the baseline.[56]

Which model gives better estimates of sectoral labor productivity levels? It is difficult to give an irrefutable answer this question since reliable independent measures of sectoral productivity for developing countries are hard to come by. However, the baseline numbers are arguably more plausible. In Appendix F I compare my results to existing estimates of sectoral productivity levels and I also calculate productivities using alternative methodologies. While these alternative calculations have their problems, I generally find that their results are closer to ACDES estimates than to CES estimates. These observations suggest that the model relying on homothetic preferences might be understating agricultural productivity levels in poor countries. Low agricultural productivity surfaces in development accounting exercises, discussed next.

Development Accounting

One of the key reasons why deriving sectoral productivity levels is useful is that it allows us to better understand the proximate causes of large cross-country differences in aggregate real output per worker. If income differences can be in large part attributed to differences in sectoral composition of economies then it is another argument why studying the process of structural change is important. I now use the calibrated models for a development accounting exercise. Specifically, I decompose the gap between aggregate productivity in country i and in the US as follows:

$$\frac{y_i - y_{US}}{y_{US}} = \underbrace{\sum_s \frac{(y_{si} - y_{sUS})\,\bar{l}_s}{y_{US}}}_{\text{within-sector effect}} + \underbrace{\sum_s \frac{\overline{y}_s\,(l_{si} - l_{sUS})}{y_{US}}}_{\text{labor-share effect}}, \tag{19}$$

where l_{si} is the labor share of sector s, y_i is the aggregate labor productivity, y_{si} sectoral labor productivity and for any variable x we have $\overline{x}_s = \frac{1}{2}(x_{si} + x_{sUS})$. In this exercise labor productivity is measured in international prices, i.e. $y_{si} = p_s A_{si}$, where p_s is the Geary-Khamis international

[56] The high dispersion in the CES is driven by lack of ability to rely on income effects rather than because of the Leontief corner solution. Performing a similar calculation for the nonhomothetic ACES model would yield productivity dispersion much closer to the baseline ACDES than to CES, even though ACES model would still offer no substitutability (as calibrated in Section 3.4).

price calculated during the calibration of the model.[57] The within-sector effect gives the share of aggregate productivity gap that can be attributed to differences in sectoral labor productivity between country i and the US. The labor-share effect shows how much of the aggregate productivity gap can be explained by differences in labor shares in country i and in the US.

Figure 7 shows the mean contribution of the within-effect (separately for each sector) and of the overall labor-share effect in the nonhomothetic and homothetic versions of the model. The most striking difference is that the CES model attributes a much larger fraction of the aggregate productivity gap to the productivity gap within agriculture: 52.1% compared to 21.8% in the ACDES case. In contrast, the contribution of the within-manufacturing gap is similar between the models and modest in both cases: 23.3% for CES and 19.3% for ACDES. Within-services productivity gaps are important in absolute terms but again quite similar across the models: 62.7% for CES and 67.5% for ACDES. Since the contributions sum to 100%, the homothetic model attributes a larger negative effect to the labor-share effect: -38.2% vs. -8.6% in the nonhomothetic case. These differences across models consistently appear at a country level, as can be verified in Table 6 with disaggregated results.

Why is the importance of within-agriculture productivity differences so much larger in the homothetic model? Because labor shares are the same in both models (and equal to values observed in the data), the discrepancy comes only from $(y_{Ai} - y_{AUS})$ taking more extreme values in the CES case. These more extreme values are related to larger dispersion of A_{Ai} across countries in the CES model discussed earlier. Since $y_{Ai} - y_{AUS} = p_A (A_{Ai} - A_{AUS})$, larger gaps in $(A_{Ai} - A_{AUS})$ in the CES case directly translate to larger productivity gaps measured in international prices. This direct effect is not very strong, however, given that the decomposition is additive rather than multiplicative.[58] The main effect occurs through the level of international prices in two models: p_A is twice as high in the CES case. The reason is simple: the CES model requires much higher prices of agriculture in poor countries to account for their large expenditure shares of agriculture. Since Geary-Khamis international prices are a particular weighted average of domestic sectoral price levels, higher values of P_{Ai} in poor countries result in higher p_A.[59] In contrast, the international prices of manufacturing and services are similar in the two models.

High p_A in the CES case also explains the large negative contribution of the labor-share effect. When sectoral output is evaluated at international prices, labor in the US is most productive in agriculture and least productive in services. The differences in y_{sUS} across sectors are modest in the nonhomothetic calibration and much larger in the homothetic case. Since poor countries have high employment shares in agriculture and low employment shares in services, a hypothetical relocation

[57] The decomposition is exact since y_{it} calibrated to the data on real GDP per worker is calculated in the model as $y_{it} = \sum_s p_s A_{sit} l_{sit} = \sum_s y_{sit} l_{sit}$.

[58] Take actual values of $A_{AUS}^{ACDES} = A_{AUS}^{CES} = 0.57$ for the US and $A_{AIND}^{ACDES} = 0.007$, $A_{AIND}^{CES} = 0.0008$ for India. Then in absolute value $A_{AIND}^{CES} - A_{AUS}^{CES}$ is very close to $A_{AIND}^{ACDES} - A_{AUS}^{ACDES}$ even though A_{AIND}^{ACDES} is an order of magnitude larger than A_{AIND}^{CES}.

[59] Comparing price levels across models is meaningful because nominal wages and real GDP per worker y_i take the same values in the two models.

of labor shares in poor countries to U.S. levels therefore widens their aggregate productivity gap with respect to the US when sectoral productivities are fixed at the average level in the two countries. Differences in sectoral labor shares therefore on average mitigate the income gap with the US, and much more strongly so in the homothetic model.

An important caveat about simple shift-share calculations merits a brief discussion at this point. The finding that labor in the US is most productive in agriculture when measured in international prices does not imply that welfare in the US would be higher if more more workers in the US toiled on farms. In Święcki (2017) I take a stronger stand on the underlying sources of intersectoral wedges that allows me to evaluate welfare levels. Under the assumption that wedges represent labor distortions, the US would be in fact better off with fewer workers in agriculture. This suggests that branding intersectoral labor relocation as good or bad based on shift-share analysis as in recent work of McMillan and Rodrik (2011) can be misleading.

To summarize, this subsection illustrates that sectoral productivity levels obtained from models of structural change can be sensitive to the treatment of nonhomotheticity. Given my calibration strategy, imposing CES preferences results in significantly larger dispersion of agricultural productivity. To the extent that the predictions of the homothetic model are less plausible, the CES model likely overstates the importance of agriculture for understanding cross-country income differences.

5.2 Time-Series Results

Differences between the baseline model and the homothetic model also appear in patterns of sectoral productivity growth over time within countries. In this subsection I describe these differences and illustrate their importance through growth accounting decompositions.

Table 7 shows the average log growth rates of sectoral labor productivity as calculated in the data and as predicted by the baseline and CES models. Recall that average sectoral productivity growth is one of the moments used in calibrating the preference parameters. Since the CES model is a restricted version of the baseline model, it is not surprising that it does not do as well as the non-homothetic model in matching the moments. Most notably, the homothetic calibration overstates labor productivity growth in agriculture on average, delivering 0.052 log points of growth per year across countries compared to 0.037 in the data and ACDES calibration. But it is important that the overstatement is not merely on average but is systematic. Figure 8 plots the labor productivity growth in the CES model against the ACDES model for each country and each sector. The CES model generates noticeably higher growth in agriculture in all but two cases. The explanation for the divergent predictions for agriculture is simple. Both models need to account for the same evolution of expenditure shares over time. To explain the falling share of agriculture, the CES model requires a larger fall in the relative price of agriculture than the ACDES model, since in the latter the low income elasticity of agriculture explains much of the expenditure shift. To achieve a faster decline of relative price of agriculture, the homothetic model requires faster productivity growth in that sector. Similar reasoning explains why the homothetic model generates systematically slightly lower productivity growth in services.

Growth Accounting

The discussion above suggests that a model of structural change with homothetic preferences has limited success in matching the evolution of sectoral productivities within countries, conditional on exactly matching the evolution of expenditure and employment. What if we were not aware of this limitation and nevertheless used the calibrated homothetic model to measure the proximate sources of aggregate growth? We would likely overstate the importance of agriculture and understate the importance of services for aggregate productivity growth.

To see why, decompose the aggregate productivity growth in country i between the first and last year as:

$$\frac{y_{it_l} - y_{it_f}}{y_{it_f}} = \underbrace{\sum_s \frac{\Delta y_{sit} \bar{l}_{sit}}{y_{it_f}}}_{\text{within-sector effect}} + \underbrace{\sum_s \frac{\Delta l_{sit} \overline{y}_{sit_f}}{y_{it_f}}}_{\text{labor-share effect}}, \tag{20}$$

where $\Delta x_{sit} = x_{sit_l} - x_{sit_f}$ and $\overline{x}_{sit} = \frac{1}{2}\left(x_{sit_l} + x_{sit_f}\right)$. The first part of the decomposition gives the contribution of the within-sector productivity growth. The second term measures the increase in aggregate productivity accounted for by the relocation of labor across sectors.

Figure 9 shows contributions of within-sector growth and labor relocation averaged across countries.[60] In the baseline model productivity growth in agriculture explains on average only 11.0% of aggregate growth. At 19.8%, the contribution of agriculture is much larger in the homothetic model. Productivity growth in manufacturing is equally important in both models, accounting for roughly 37% of total growth. In contrast to agriculture, the growth in services has a noticeably larger contribution in the baseline model (40.2%) than in the CES case (32.9%). Changing employment shares explain a modest 11-12% of aggregate productivity growth in both models.

Since the baseline model is better at matching within-country sectoral productivity growth rates, the above numbers indicate that the homothetic model tends to attribute a larger share of aggregate growth to growth in agriculture and lower share to growth in services than they deserve.

6 Conclusions

This paper seeks to assess how important, quantitatively, are the various mechanisms that have been postulated in the literature as driving the process of structural change. To answer this question,

[60]Four countries (Bolivia, Brazil, Germany and Mexico) are excluded from the calculation of means because they had very little growth in output per worker over the sample period. As a result of dividing by a very small denominator the decomposition in those countries generates extreme results. The overall picture is similar if we look at medians for the entire sample instead. The detailed decomposition for all countries is reported in Table 8.

I build a quantitative model featuring (i) sector-biased technological change, (ii) nonhomothetic preferences, (iii) international trade and (iv) wedges between factor costs across sectors.

Taking the model to a large sample of diverse countries allows me to robustly calibrate the key model parameters and to illustrate how the relative importance of different mechanisms depends on the stage of economic development and other country characteristics such as country size or the level of integration with world markets.

My calibration strategy involves exactly matching the two key margins of structural change: the evolution of sectoral employment and expenditure. Predictions of the model for the third margin - sectoral productivity growth - are then used to pin down the parameters of a flexible specification of consumer preferences. I use counterfactual simulations of the calibrated model to highlight the relative strength of the four channels in accounting for sectoral labor relocation.

I find that sector-biased technological change is the single most important driver of structural change. It is particularly indispensable for explaining the net movement of labor from manufacturing to services at later stages of structural change. Nonhomothetic preferences, on the other hand, are vital for accounting for the relocation of labor out of agriculture at earlier stages of development. The importance of international trade is best evaluated on a case by case basis. Trade contributes to structural change particularly in smaller countries that become more globalized. Trade also interacts with intersectoral wedges. Once countries are allowed to trade, changes in relative factor costs across sectors explain a sizable portion of labor relocation.

In the second part of the paper I explore the importance of nonhomothetic preferences in more detail. I show that a model calibrated with homothetic preferences overstates productivity differences in agriculture across countries and overstates productivity growth in agriculture over time. Consequently, ignoring nonhomotheticities results in overstating the influence of developments in agriculture on aggregate productivity.

This paper takes an accounting approach to structural change. As such it abstracts from one interesting consideration: deeper dynamic links between the underlying forces behind structural change. For example, openness to international trade might affect the speed of technological progress within a country, which I take as an exogenous process. Studying such interactions in a dynamic setting would be a fruitful avenue for future research.

References

Acemoglu, D. and V. Guerrieri (2008): "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 116, 467–498.

Alvarez-Cuadrado, F., N. V. Long, and M. Poschke (2014): "Capital-Labor Substitution, Structural Change and the Labor Income Share," Mimeo, McGill University.

Anderson, J. and E. van Wincoop (2004): "Trade Costs," *Journal of Economic Literature*, 42, 691–751.

Andrews, D. (1999): "Estimation When a Parameter is on a Boundary," *Econometrica*, 67, pp. 1341–1383.

——— (2000): "Inconsistency of the Bootstrap when a Parameter Is on the Boundary of the Parameter Space," *Econometrica*, 68, 399–406.

——— (2002): "Generalized Method of Moments Estimation When a Parameter Is on a Boundary," *Journal of Business & Economic Statistics*, 20, 530–44.

APO (2010): "Asian Productivity Organization Database v. 2010.03a," Accessed online at http://www.apo-tokyo.org/.

Betts, C., R. Giri, and R. Verma (2015): "Trade, Reform, and Structural Transformation in South Korea," Mimeo, University of Southern California.

Boppart, T. (2014): "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences," *Econometrica*, 82, 2167–2196.

Buera, F. J. and J. P. Kaboski (2009): "Can Traditional Theories of Structural Change Fit The Data?" *Journal of the European Economic Association*, 7, 469–477.

Buera, F. J., J. P. Kaboski, and R. Rogerson (2015): "Skill Biased Structural Change," Working Paper 21165, National Bureau of Economic Research.

Caselli, F. and J. Coleman (2001): "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 109, 584–616.

Comin, D., D. Lashkari, and M. Mestieri (2015): "Structural Change with Long-run Income and Price Effects," Mimeo, Northwestern University.

Diewert, W. E. (1999): *International and Interarea Comparisons of Income, Output, and Prices*, University of Chicago Press, chap. Axiomatic and Economic Approaches to International Comparisons, 13–107.

Duarte, M. and D. Restuccia (2010): "The Role of the Structural Transformation in Aggregate Productivity," *Quarterly Journal of Economics*, 125, 129–173.

Eaton, J. and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.

Feenstra, R., R. Lipsey, H. Deng, A. Ma, and H. Mo (2005): "World trade flows: 1962-2000," NBER Working Paper 11040, National Bureau of Economic Research.

Finicelli, A., P. Pagano, and M. Sbracia (2013): "Ricardian selection," *Journal of International Economics*, 89, 96–109.

Gaulier, G. and S. Zignago (2010): "BACI: International trade database at the product-level The 1994-2007 version," CEPII Working Paper 2010-23, CEPII.

Gollin, D., D. Lagakos, and M. E. Waugh (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129, 939–993.

Hanoch, G. (1975): "Production and Demand Models with Direct or Indirect Implicit Additivity," *Econometrica*, 43, 395–419.

Head, K. and T. Mayer (2014): "Chapter 3 - Gravity Equations: Workhorse, Toolkit, and Cookbook," in *Handbook of International Economics*, ed. by E. H. Gita Gopinath and K. Rogoff, Elsevier, vol. 4 of *Handbook of International Economics*, 131 – 195.

Herrendorf, B., C. Herrington, and A. Valentinyi (2015): "Sectoral Technology and Structural Transformation," *American Economic Journal: Macroeconomics*, 7, 104–33.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2013a): *Handbook of Economic Growth*, Elsevier, vol. 2, chap. Growth and Structural Transformation, 855–941.

——— (2013b): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103, 2752–89.

Herrendorf, B. and T. Schoellman (2015): "Wages, Human Capital, and Structural Transformation," Mimeo, Arizona State University.

Heston, A., R. Summers, and B. Aten (2011): "Penn World Table Version 7.0," Tech. rep., Center for International Comparisons at the University of Pennsylvania.

Houthakker, H. (1960): "Additive preferences," *Econometrica*, 28, 244–257.

Inklaar, R. and M. Timmer (2008): "GGDC productivity level database: International comparisons of output, inputs and productivity at the industry level," Research Memorandum GD-104, Groningen Growth and Development Centre, University of Groningen.

Jensen, B., P. Boer, J. Daal, and P. S. Jensen (2011): "Global restrictions on the parameters of the CDES indirect utility function," *Journal of Economics*, 102, 217–235.

Ketz, P. (2015): "Testing Overidentifying Restrictions when the True Parameter Vector Is near or at the Boundary of the Parameter Space," Mimeo, Paris School of Economics.

Kongsamut, P., S. Rebelo, and D. Xie (2001): "Beyond Balanced Growth," *Review of Economic Studies*, 68, 869–882.

Mas-Colell, A., M. Whinston, and J. Green (1995): *Microeconomic Theory*, Oxford University Press.

Matsuyama, K. (2009): "Structural Change in an Interdependent World: A Global View of Manufacturing Decline," *Journal of the European Economic Association*, 7, 478–486.

McMillan, M. S. and D. Rodrik (2011): *Making globalization socially sustainable*, International Labour Organization (ILO) and World Trade Organization (WTO), chap. Globalization, Structural Change and Productivity Growth.

Ngai, L. and C. Pissarides (2007): "Structural Change in a Multisector Model of Growth," *The American Economic Review*, 97, 429–443.

OECD (2011): "STructural ANalysis Database, ISIC Rev. 3 version," Available online at http://www.oecd.org/sti/ind/stanstructuralanalysisdatabase.htm.

O'Mahony, M. and M. P. Timmer (2009): "Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database," *Economic Journal*, 119, F374–F403.

Restuccia, D., D. T. Yang, and X. Zhu (2008): "Agriculture and aggregate productivity: A quantitative cross-country analysis," *Journal of Monetary Economics*, 55, 234–250.

Santos Silva, J. and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

Sposi, M. (2015): "Evolving Comparative Advantage, Sectoral Linkages, and Structural Change," Mimeo, Federal Reserve Bank of Dallas.

Stefanski, R. (2014): "Structural Transformation and the Oil Price," *Review of Economic Dynamics*, 17, 484–504.

Święcki, T. (2017): "Intersectoral distortions and the welfare gains from trade," *Journal of International Economics*, 104, 138 – 156.

Teignier, M. (2014): "The role of trade in structural transformation," Mimeo, Universitat de Barcelona.

Timmer, M. and G. de Vries (2009): "Structural change and growth accelerations in Asia and Latin America: a new sectoral data set," *Cliometrica*, 3, 165–190.

Tombe, T. (2015): "The Missing Food Problem: Trade, Agriculture, and International Productivity Differences," *American Economic Journal: Macroeconomics*, 7, 226–258.

Uy, T., K. Yi, and J. Zhang (2013): "Structural Change in an Open Economy," *Journal of Monetary Economics*, 60, 667–682.

Valentinyi, A. and B. Herrendorf (2008): "Measuring Factor Income Shares at the Sector Level," *Review of Economic Dynamics*, 11, 820–835.

Table 1: Calibrated Preference Parameters

Model	(1) ACDES General	(2) ACDES Baseline	(3) ACES
α_A	-1.00	-1.00	-1.00
	(0.12)	(0.29)	(0.06)
α_M	-0.99	-0.89	
	(0.12)	(0.16)	
α_S	-0.75	-0.67	
	(0.22)	(0.25)	
$\overline{c}_A$	4.10	3.89	4.17
	(1.21)	(1.53)	(1.11)
$\overline{c}_M$	0.12		1.34
	(1.84)		(1.07)
$\overline{c}_S$	2.40		-4.27
	(3.76)		(3.50)
GMM Obj.	2.31×10^{-7}	3.22×10^{-7}	3.35×10^{-5}
N	45	45	45

Notes: ACDES: Augmented CDES model; Baseline: ACDES model with restrictions $\overline{c}_M = \overline{c}_S = 0$; ACES: Augmented CES model, i.e. ACDES model with restrictions $\alpha_A = \alpha_M = \alpha_S$. GMM Obj: value of the GMM objective function (17); N : number of observations (countries). Standard errors from nonparametric residual bootstrap in parentheses. In the bootstrap randomly drawn residuals $\hat{\varepsilon}_s$ from the original calibration are added to predicted values $g_{si}(\hat{\omega})$ of the dependent variable (sectoral labor productivity growth) to generate artificial data; the calibration procedure is then applied to the generated data. Number of bootstrap simulations: 1000 for the Baseline model and 200 for other models. Subsistence requirements $\overline{c}_s$ and their standard errors multiplied by a factor 10^3 to increase readability.

Table 2: Demand Properties

Model	(1) General ACDES			(2) Baseline ACDES		
Demand elasticities						
Income elasticity	η_A	η_M	η_S	η_A	η_M	η_S
	0.31	0.91	1.14	0.30	0.91	1.14
Own-price elasticity	ϵ_{AA}	ϵ_{MM}	ϵ_{SS}	ϵ_{AA}	ϵ_{MM}	ϵ_{SS}
	-0.02	-0.28	-0.78	-0.02	-0.34	-0.81
Elast. of substitution	σ_{AM}	σ_{AS}	σ_{MS}	σ_{AM}	σ_{AS}	σ_{MS}
	-0.06	0.02	0.09	-0.06	0.02	0.19
Subsistence consumption						
Subs. expenditure share		0.076			0.060	
Subs. consumption share	$\overline{c}_A/C_A$	$\overline{c}_M/C_M$	$\overline{c}_S/C_S$	$\overline{c}_A/C_A$	$\overline{c}_M/C_M$	$\overline{c}_S/C_S$
	0.644	0.003	0.023	0.610	0	0

Notes: Income elasticity: $\eta_s = \frac{\partial \log x_s(p,m)}{\partial \log m}$; Own-price elasticity: $\epsilon_{ss} = \frac{\partial \log x_s(p,m)}{\partial \log p_s}$; Allen-Uzawa elasticity of substitution: $\sigma_{ij} = \frac{1}{e_j} \frac{\partial \log h_i(p,U)}{\partial \log p_j}$, where $x_s(p,m)$ is Marshallian demand and $h_s(p,U)$ is Hicksian demand for sector s and e_s is the expenditure share of sector s. Subsistence expenditure share: $\sum_s p_s\overline{c}_s / \sum_s p_s C_s$. Table reports mean elasticities and subsistence shares across countries computed for 1995.

Table 3: Calibration Robustness: Monte Carlo Simulation

Parameter	α_A	α_M	α_S	$\bar{c}_A$
Model simulation: True value	-1.00	-0.89	-0.67	3.98×10^{-3}
Monte Carlo: 25th percentile	-1.00	-1.00	-0.80	3.39×10^{-3}
Monte Carlo: Median	-1.00	-0.94	-0.67	4.35×10^{-3}
Monte Carlo: 75th percentile	-0.80	-0.75	-0.45	5.09×10^{-3}
Monte Carlo: Mean	-0.87	-0.86	-0.63	4.29×10^{-3}

Notes: Distribution of parameter values obtained from the calibration procedure applied 200 times to simulated data (with known true data generating process) subject to random measurement error.

Table 4: Labor Relocation Index

Operative channels	N+T+W (no P)	P+T+W (no N)	P+N+W (no T)	P+N+T (no W)
Median *LRI*	0.46	0.80	0.71	0.65

Operative channels	P	N	T	W
Median *LRI*	0.43	0.27	0.21	0.00

Operative channels	P+N	P+T	P+W	N+T	N+W	T+W
Median *LRI*	0.70	0.65	0.43	0.32	0.33	0.22

Notes: Labor Relocation Index computed according to formula (18). Table shows the median *LRI* computed across 45 countries in the sample. Channels driving structural change: P - sector-biased productivity growth, N - nonhomothetic preferences, T - international trade, W - changes in intersectoral wedges.

Table 5: Labor Relocation Index by Country

Country	(1) no P	(2) no N	(3) no T	(4) no W	(5) P	(6) N	(7) T	(8) W
Argentina	0.31	0.77	0.90	0.11	0.87	0.19	-0.17	0.00
Australia	0.22	0.93	0.88	0.86	0.83	0.18	0.17	0.00
Austria	0.77	0.76	0.68	0.61	0.39	0.42	0.45	0.00
Bangladesh	0.61	0.17	0.62	-0.24	-0.29	0.02	0.01	0.00
Belgium	0.98	0.85	0.24	0.91	0.04	0.23	0.78	0.00
Bolivia	-0.56	0.88	0.48	-0.47	0.56	0.27	0.21	0.00
Brazil	0.18	0.96	0.34	0.27	0.43	0.05	-0.34	0.00
Canada	0.40	0.95	0.93	0.73	0.96	0.18	0.02	0.00
Chile	0.39	0.85	0.62	0.70	0.66	0.32	0.03	0.00
China	-0.29	-0.73	0.87	0.65	-0.86	-0.34	-0.68	0.00
Colombia	0.51	0.76	0.91	0.45	0.65	0.38	0.23	0.00
Czech Rep.	0.02	0.89	0.48	0.80	0.69	0.63	-0.35	0.00
Denmark	0.34	0.94	0.78	0.86	0.86	0.26	0.04	0.00
Finland	0.56	0.80	0.83	0.91	0.76	0.47	0.11	0.00
France	0.55	0.91	0.79	0.97	0.68	0.26	0.37	0.00
Germany	-0.20	0.87	0.37	0.84	0.29	0.12	-0.40	0.00
West Germany	0.39	0.96	0.83	0.96	0.89	0.25	0.21	0.00
Greece	0.66	0.81	0.55	0.53	0.43	0.27	0.26	0.00
Hungary	0.46	0.70	-0.29	0.80	-0.76	0.36	0.54	0.00
India	-0.26	-0.58	0.93	-0.01	-0.61	-1.03	0.21	0.00
Indonesia	0.66	0.18	0.67	0.59	-0.09	0.81	0.49	0.00
Ireland	-0.70	0.95	-0.04	0.30	-0.09	0.39	-0.37	0.00
Italy	0.53	0.90	0.92	0.81	0.82	0.31	0.34	0.00
Japan	0.45	0.86	0.88	0.99	0.88	0.43	0.08	0.00
Korea	0.75	0.65	0.84	0.87	0.59	0.72	0.26	0.00
Malaysia	0.41	0.55	0.15	0.27	-0.11	0.34	0.75	0.00
Mexico	0.61	0.98	0.70	0.81	0.71	0.10	0.37	0.00
Netherlands	0.52	0.99	0.85	0.82	0.80	0.17	0.27	0.00
Norway	-0.08	0.95	0.03	0.30	0.13	0.34	-0.90	0.00
Pakistan	0.77	0.04	0.71	0.01	-0.18	0.58	0.12	0.00
Peru	-8.17	-1.53	-0.03	-5.85	-1.52	-3.15	-0.25	0.00
Philippines	0.92	0.61	-0.19	0.31	-0.13	0.29	0.76	0.00
Poland	-1.10	0.16	0.71	-1.85	0.16	-0.22	-0.69	0.00
Portugal	0.36	0.12	0.08	0.33	-0.54	0.70	0.40	0.00
Slovakia	0.07	0.66	0.51	0.36	0.71	0.46	-0.35	0.00
Spain	0.64	0.94	0.71	0.38	0.63	0.30	0.50	0.00
Sri Lanka	0.41	0.05	0.72	0.23	-0.27	0.44	-0.13	0.00
Sweden	0.57	0.88	0.92	0.76	0.76	0.27	0.12	0.00
Switzerland	0.70	0.72	-0.19	0.25	-0.50	0.17	0.29	0.00
Taiwan	0.82	0.58	0.88	0.92	0.32	0.95	0.32	0.00
Thailand	0.73	0.44	0.84	0.76	0.24	0.92	0.26	0.00
UK	0.46	0.94	0.72	0.94	0.70	0.17	0.27	0.00
US	0.54	0.90	0.73	0.92	0.63	0.19	0.34	0.00
Venezuela	0.48	0.77	0.98	0.67	0.73	-0.46	0.24	0.00
Viet Nam	0.85	0.11	0.89	-0.23	-0.16	0.37	0.48	0.00

Notes: Labor Relocation Index computed according to formula (18). Channels driving structural change: P - sector-biased productivity growth, N - nonhomothetic preferences, T - international trade, W - changes in intersectoral wedges.

Table 6: Development Accounting Decomposition by Country

	ACDES				CES			
Country	within A	within M	within S	labor sh.	within A	within M	within S	labor sh.
Argentina	12.3	21.3	72.5	-6.1	32.0	24.2	64.6	-20.7
Australia	-7.1	60.2	98.4	-51.5	16.6	85.8	169.2	-171.6
Austria	137.8	79.1	-125.1	8.2	307.2	76.5	-126.7	-157.0
Bangladesh	44.6	15.9	56.6	-17.1	100.7	17.2	50.6	-68.5
Belgium	-25.3	-38.7	151.9	12.1	-52.5	-27.5	145.7	34.3
Bolivia	39.2	19.3	54.1	-12.7	89.0	20.7	48.1	-57.9
Brazil	29.3	17.2	61.3	-7.8	67.0	19.8	53.7	-40.5
Canada	27.2	-36.1	128.5	-19.6	69.7	-21.3	117.1	-65.4
Chile	14.9	24.6	68.9	-8.4	38.3	27.1	61.4	-26.8
China	48.8	20.7	49.7	-19.2	110.0	22.1	43.8	-75.9
Colombia	29.9	19.1	60.0	-9.0	68.6	21.1	53.1	-42.8
Czech Rep.	10.5	44.2	49.3	-4.1	25.0	46.7	43.5	-15.2
Denmark	-9.4	55.8	78.2	-24.6	-2.3	69.6	92.9	-60.2
Finland	39.1	15.6	63.8	-18.5	89.7	22.4	48.8	-60.9
France	-128.8	-5.4	187.1	47.1	-284.7	-16.4	205.7	195.4
Germany	31.5	97.2	-4.0	-24.7	71.8	100.1	-19.1	-52.8
Greece	71.1	57.3	-34.5	6.1	159.2	58.7	-36.2	-81.7
Hungary	13.7	37.4	55.2	-6.3	33.0	40.5	48.4	-21.8
India	58.5	17.2	45.3	-21.1	131.7	18.7	38.7	-89.1
Indonesia	39.2	19.0	56.2	-14.3	88.8	20.7	49.7	-59.3
Ireland	68.4	-22.2	89.1	-35.3	158.1	-12.3	74.5	-120.4
Italy	-262.7	-308.4	581.0	90.2	-583.9	-305.3	583.2	406.0
Japan	34.7	37.6	40.9	-13.2	78.8	41.9	28.4	-49.0
Korea	22.5	34.7	49.2	-6.3	51.5	37.2	42.5	-31.1
Malaysia	23.5	34.1	50.4	-8.1	54.7	36.2	44.9	-35.8
Mexico	28.2	24.5	55.1	-7.7	64.4	27.6	47.2	-39.2
Netherlands	7.6	-29.7	147.2	-25.2	35.3	-5.2	157.1	-87.2
Norway	-9.9	56.4	42.3	11.3	-22.6	56.7	21.0	44.9
Pakistan	45.8	15.1	53.6	-14.6	103.4	17.2	46.9	-67.5
Peru	36.9	16.8	56.4	-10.2	83.7	19.1	49.4	-52.3
Philippines	39.4	16.2	58.0	-13.6	89.1	17.6	51.9	-58.7
Poland	24.9	34.8	46.7	-6.4	57.0	37.3	40.4	-34.7
Portugal	27.4	37.5	39.3	-4.3	62.2	40.3	32.3	-34.8
Slovakia	13.2	36.1	57.8	-7.2	31.1	39.6	50.0	-20.6
Spain	85.4	102.9	-75.6	-12.7	190.9	100.4	-72.0	-119.4
Sri Lanka	39.4	20.6	54.5	-14.5	89.3	22.7	47.6	-59.6
Sweden	19.9	4.4	93.4	-17.8	47.2	14.2	71.3	-32.8
Switzerland	45.4	32.8	44.6	-22.8	103.0	39.5	21.7	-64.3
Taiwan	26.7	46.9	34.6	-8.2	61.2	50.4	27.9	-39.5
Thailand	49.4	19.1	45.5	-14.0	112.3	20.7	39.1	-72.1
UK	11.6	44.0	50.1	-5.7	27.3	41.7	40.4	-9.5
Venezuela	20.1	17.8	68.6	-6.6	46.3	20.3	60.8	-27.5
Viet Nam	63.4	15.4	44.8	-23.6	142.8	16.5	37.8	-97.1

Notes: Table shows decomposition (19) of the aggregate productivity gap relative to the US in 1995. Contributions of within-sector productivity differences and overall labor-share effect expressed as a percentage of aggregate productivity gap.

Table 7: Average Labor Productivity Growth

Sector:	Agriculture	Manufacturing	Services
ACDES	0.037	0.034	0.013
CES	0.052	0.035	0.010
Data	0.037	0.034	0.013

Notes: Table reports the annualized average log growth of labor productivity, averaged across countries, in the data, in the baseline calibration and in the calibration imposing CES preferences.

Table 8: Growth Accounting Decomposition by Country

	ACDES				CES			
Country	within A	within M	within S	labor sh.	within A	within M	within S	labor sh.
Argentina	45.2	47.9	29.2	-22.3	56.4	51.3	14.7	-22.5
Australia	8.7	49.9	49.0	-7.5	25.5	53.7	41.9	-21.2
Austria	3.6	28.9	38.4	29.2	10.6	33.2	34.2	21.9
Bangladesh	9.8	53.7	14.3	22.1	5.5	38.3	11.8	44.4
Belgium	4.6	33.7	57.6	4.1	14.3	40.5	52.2	-7.1
Bolivia	88.0	-98.3	-20.5	130.7	52.6	-84.6	-27.8	159.8
Brazil	149.6	194.1	-427.2	183.5	133.0	200.3	-440.5	207.2
Canada	16.1	65.5	33.9	-15.5	37.1	68.4	26.4	-31.9
Chile	68.6	37.1	11.1	-16.7	77.7	38.3	1.2	-17.1
China	7.3	37.0	37.3	18.4	6.7	31.1	30.1	32.0
Colombia	16.5	37.8	2.4	43.3	17.8	34.2	-3.5	51.5
Czech Rep.	6.5	40.8	47.0	5.7	12.5	41.4	44.1	2.0
Denmark	42.2	41.6	34.3	-18.1	85.4	39.4	20.8	-45.7
Finland	7.4	37.0	44.5	11.1	18.1	40.1	39.3	2.5
France	9.6	42.6	45.1	2.8	25.4	44.8	40.8	-11.0
Germany	7.7	140.6	-41.1	-7.3	16.8	134.6	-31.3	-20.1
West Germany	6.5	47.0	24.8	21.6	14.0	50.0	23.1	12.9
Greece	7.7	10.8	15.6	65.9	21.7	13.5	13.2	51.5
Hungary	11.7	28.9	52.5	6.9	23.2	31.0	45.7	0.1
India	7.3	23.8	40.1	28.8	7.0	16.6	35.9	40.4
Indonesia	11.5	16.6	34.2	37.8	9.4	14.0	23.5	53.0
Ireland	6.4	78.7	8.0	6.9	13.1	75.0	10.0	2.0
Italy	9.7	35.2	31.8	23.3	24.0	38.2	28.5	9.2
Japan	5.8	44.4	39.1	10.6	13.7	45.9	36.4	4.0
Korea	9.1	26.4	35.8	28.7	17.8	28.3	30.7	23.2
Malaysia	9.5	9.5	55.9	25.1	17.9	11.0	43.7	27.4
Mexico	25.1	78.3	-215.0	211.6	35.5	67.4	-208.7	205.8
Netherlands	20.4	61.6	38.7	-20.8	46.0	58.9	28.7	-33.6
Norway	17.3	67.0	28.2	-12.5	41.8	72.0	18.8	-32.6
Pakistan	7.1	32.0	28.1	32.9	7.8	26.8	15.3	50.1
Peru	1.4	54.2	60.2	-15.8	4.8	48.9	59.7	-13.4
Philippines	-10.4	49.5	-6.9	67.8	3.4	29.7	-53.2	120.2
Poland	1.4	29.8	66.5	2.2	5.7	31.1	61.1	2.1
Portugal	2.3	16.8	65.2	15.7	7.0	19.2	59.2	14.6
Slovakia	8.3	31.7	54.8	5.2	15.7	33.8	48.9	1.6
Spain	17.0	22.0	35.6	25.4	40.4	23.6	31.6	4.4
Sri Lanka	9.0	18.4	45.4	27.2	8.9	13.9	44.0	33.2
Sweden	5.5	56.4	43.7	-5.5	14.0	59.5	39.6	-13.1
Switzerland	-0.4	36.6	70.7	-7.0	2.4	46.1	63.6	-12.2
Taiwan	6.2	23.9	49.2	20.7	12.7	25.4	44.7	17.2
Thailand	12.5	16.4	19.8	51.3	11.7	16.0	9.6	62.7
UK	3.4	51.7	51.1	-6.2	9.7	57.2	47.4	-14.3
US	5.6	43.2	62.5	-11.3	17.2	48.8	55.8	-21.8
Venezuela	-6.4	-6.8	126.5	-13.4	-3.1	0.8	115.7	-13.5
Viet Nam	18.6	21.6	28.6	31.2	10.1	16.8	12.6	60.4

Notes: Table shows decomposition 20 of aggregate productivity growth to contributions (in percentage of total growth) of within-sector productivity growth and labor-share relocation.

Figure 1: Calibrated Demand Curves

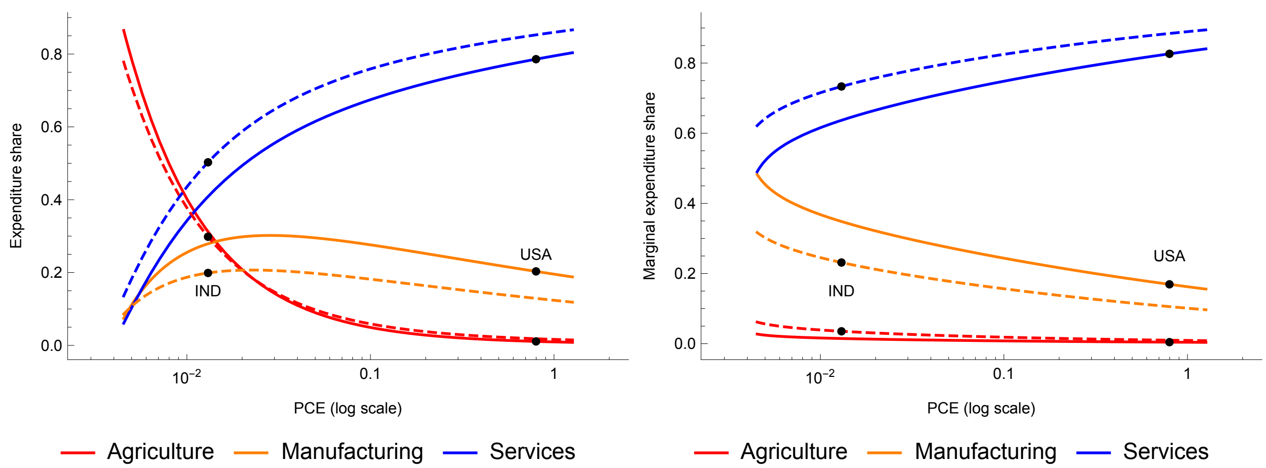

Notes: PCE is per capita expenditure, with scale running from 20% below the lowest value to 20% above the highest value of PCE across countries in 1995. Solid (dashed) lines evaluated holding sectoral pries fixed at a level calculated for the US (India) in 1995. Heavy dots illustrate the actual position of US in India that year.

Figure 2: Sectoral Labor Productivity Growth

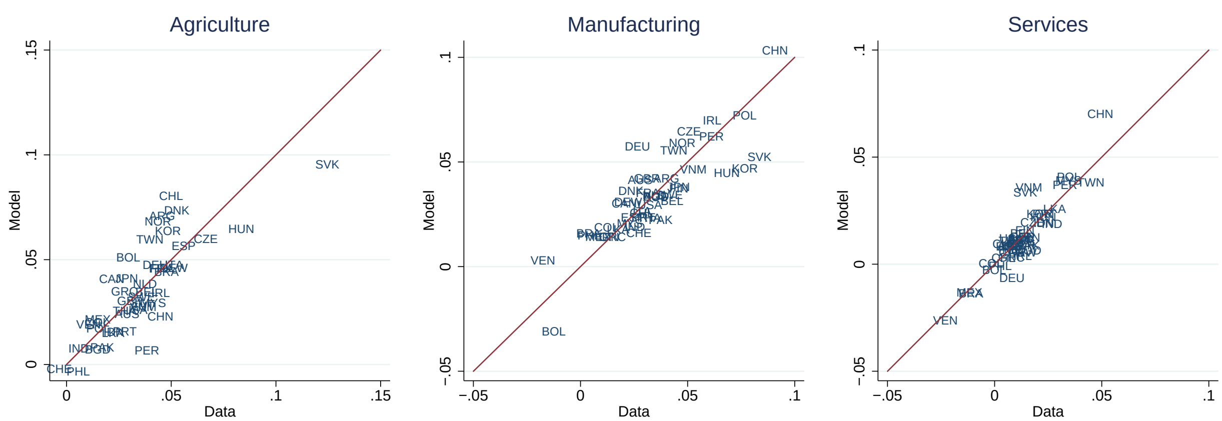

Notes: Annualized average log growth rate of labor productivity A_s for country i computed as $\frac{1}{t_l^i - t_f^i} \log\left(A_{sit_l^i}/A_{sit_f^i}\right)$, where t_l^i and t_f^i is the last and first year that country i appears in the sample.

Figure 3: Labor Relocation with One Channel Switched Off

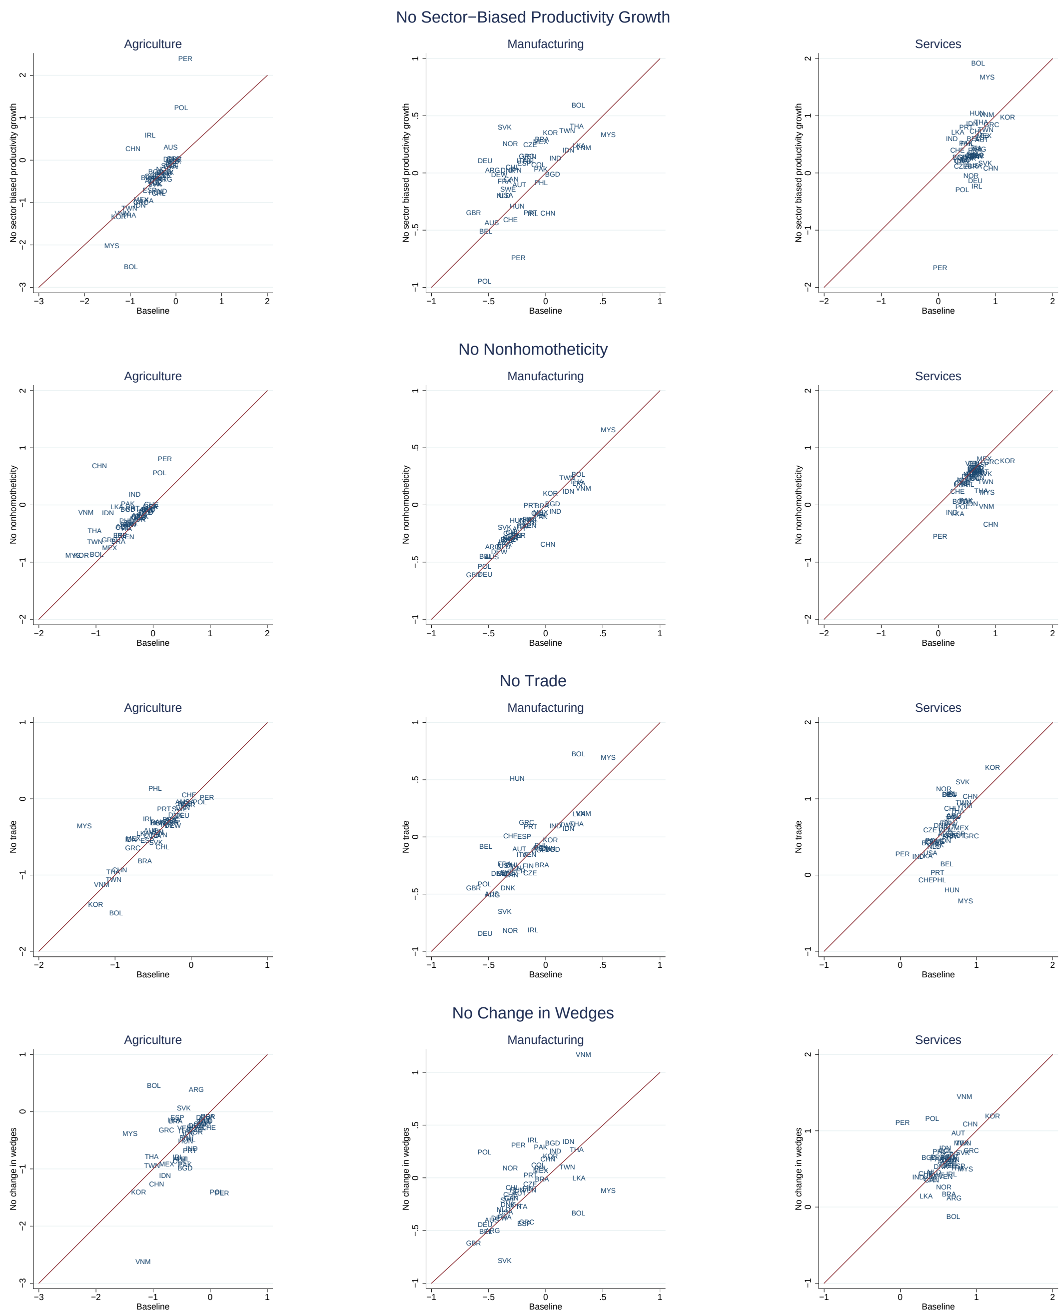

Notes: Graphs show the average annual change in a sector's labor share (expressed in percentage points per year) between the last year and the first year a country is in the counterfactual scenario against the baseline.

Figure 4: Labor Relocation with One Channel Switched On

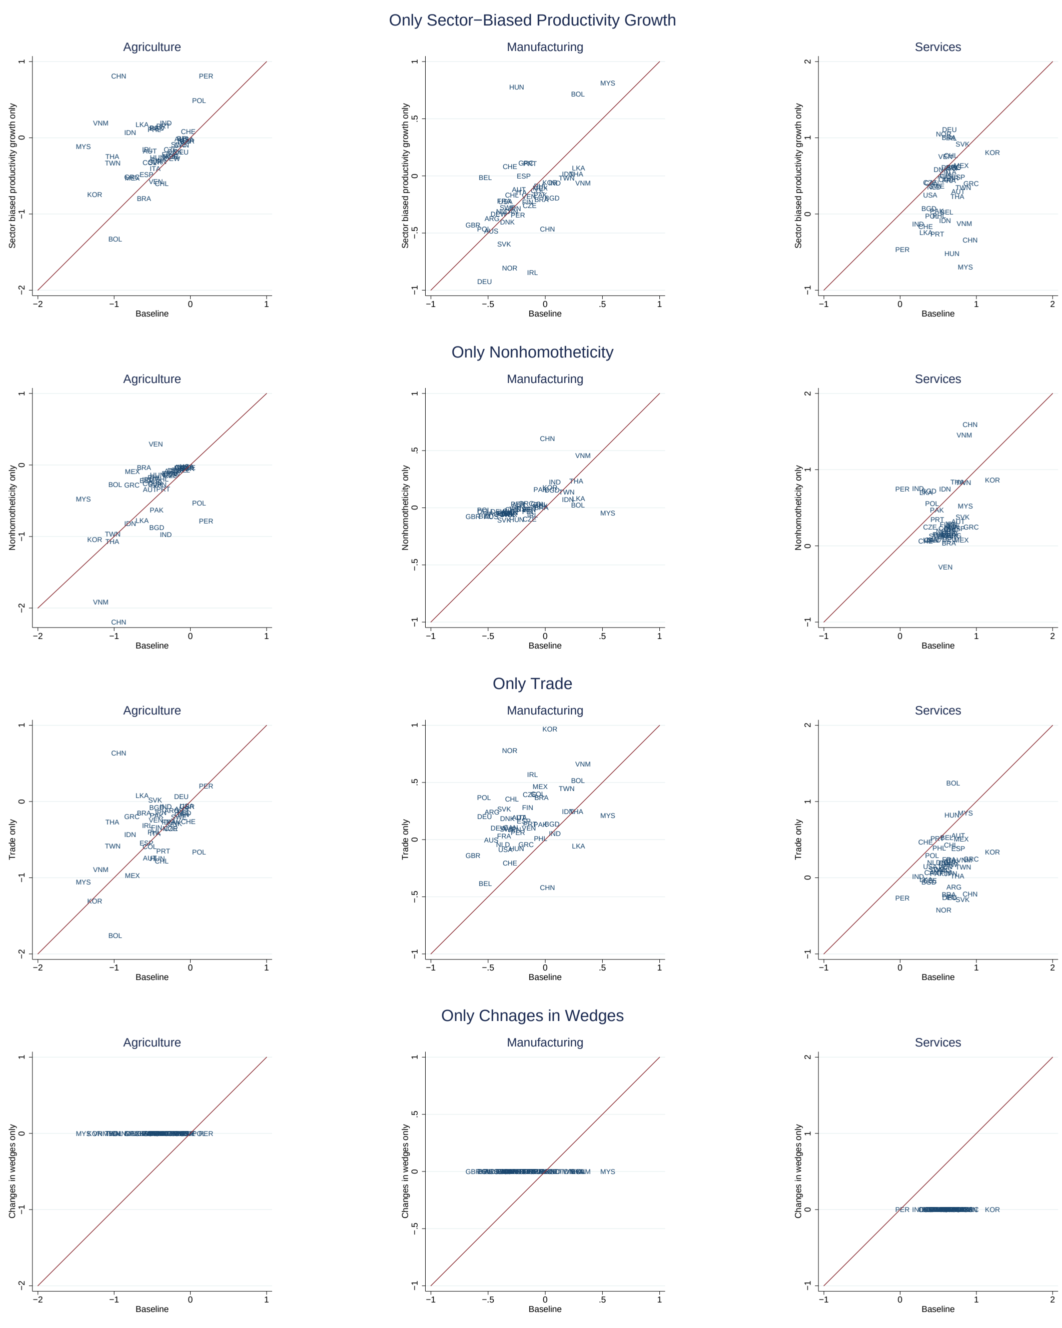

Notes: Graphs show the average annual change in a sector's labor share (expressed in percentage points per year) between the last year and the first year a country is in the counterfactual scenario against the baseline.

Figure 5: Marginal Effects of Adding New Channels

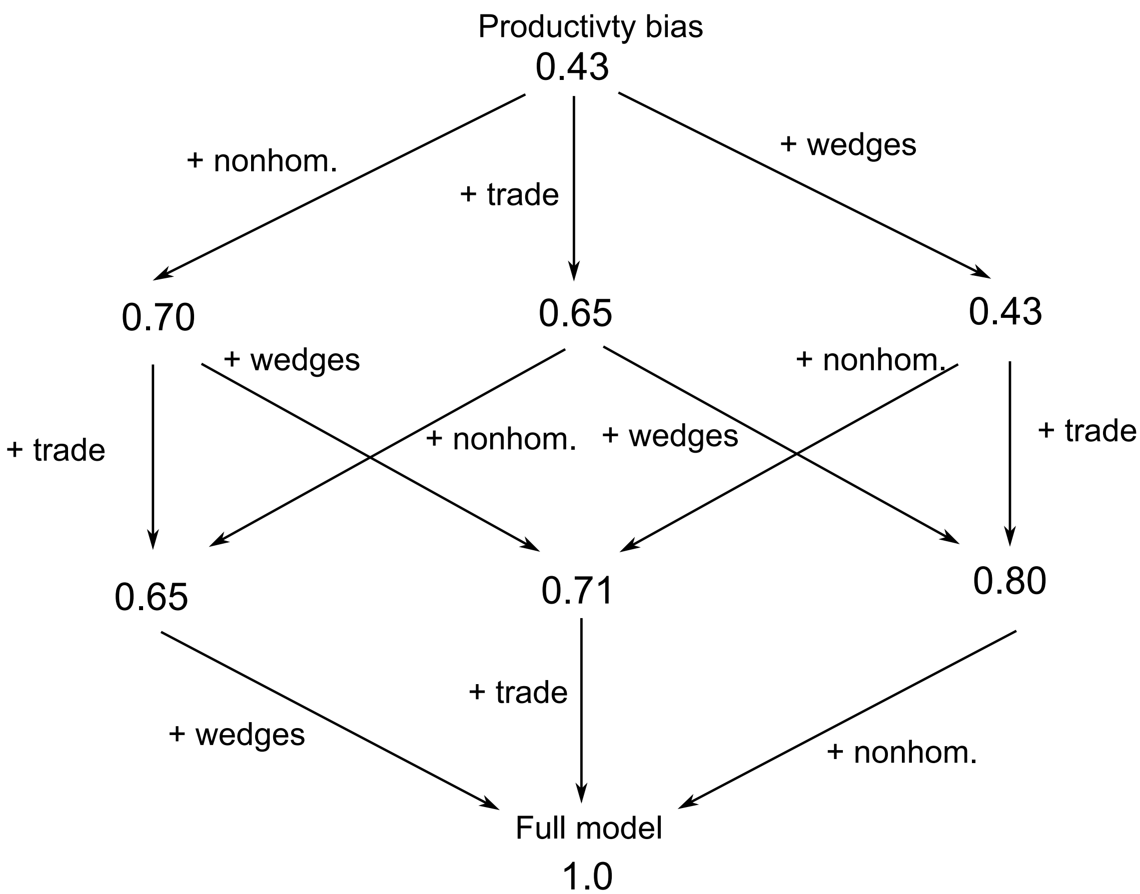

Notes: Figure shows the median Labor Relocation Index for calculations adding progressively more drivers of structural change.

Figure 6: Labor Productivity in CES and Baseline Calibrations

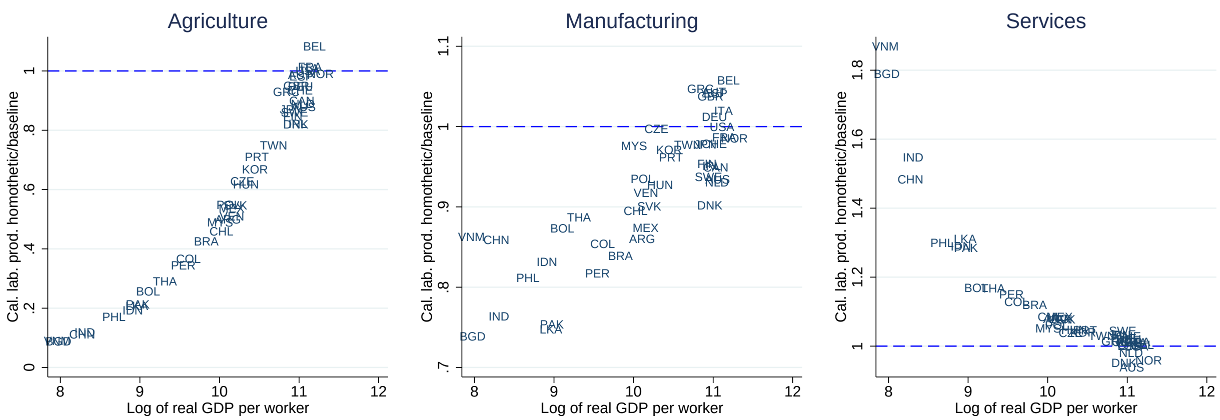

Notes: Figure shows the labor productivity in the CES calibration relative to the baseline ACDES calibration in 1995.

Figure 7: Development Accounting Decomposition

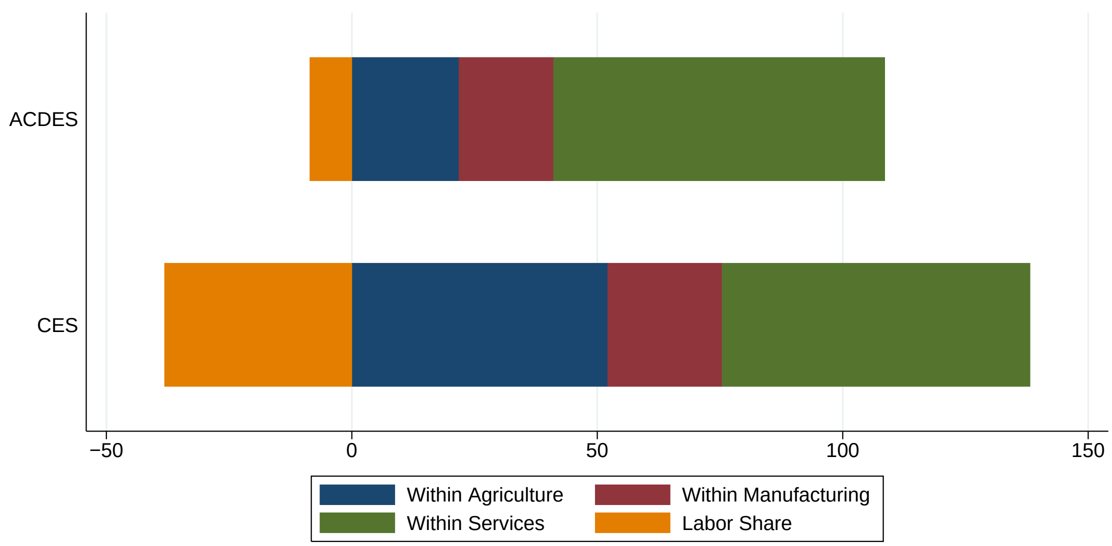

Notes: Decomposition of the aggregate productivity gap relative to the US computed according to formula (19). Figure shows the mean contributions [as a percentage of aggregate productivity gap] across countries in 1995 for the baseline and homothetic CES calibrations.

Figure 8: Labor Productivity Growth in CES and Baseline Calibrations

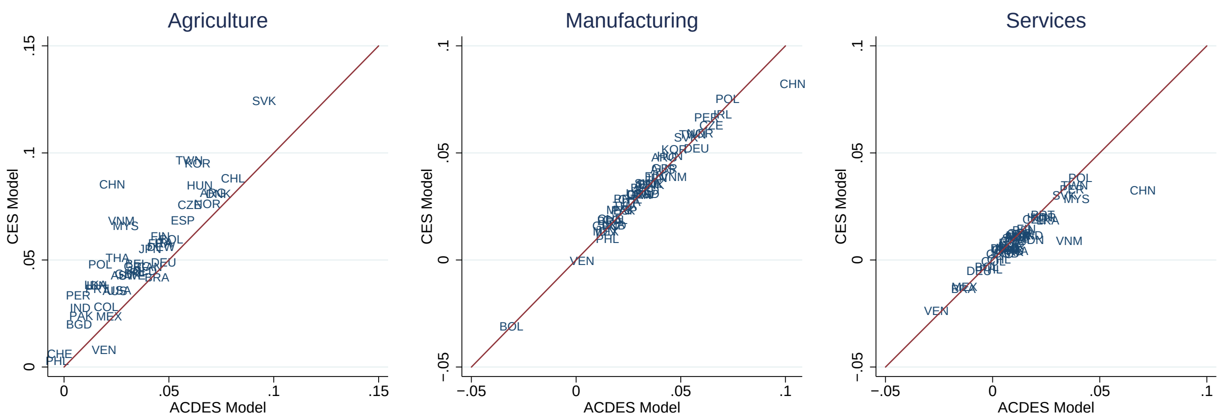

Notes: Figure shows the annualized average log growth of labor productivity in the CES calibration relative to the baseline ACDES calibration.

Figure 9: Growth Accounting Decomposition

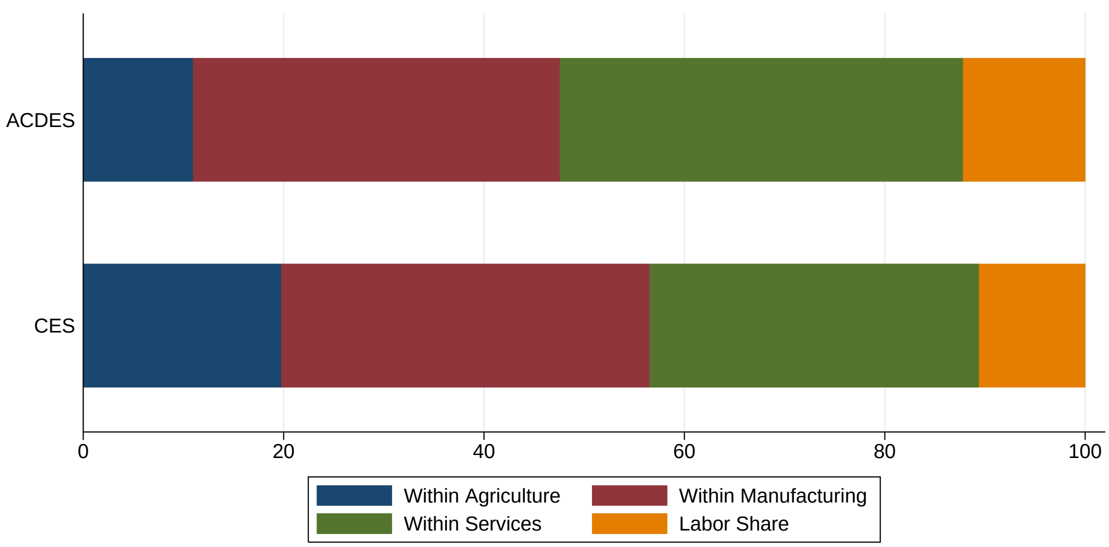

Notes: Figure shows the decomposition of aggregate productivity growth to contributions (in percentage of total growth) of within-sector productivity growth and labor-share relocation. The figure reports mean contributions across all countries except Bolivia, Brazil, Germany and Mexico. Those four countries generate extreme values due to low aggregate growth (small denominator) . Results are similar if medians of the full sample are used instead.

A Proof of Regularity of Augmented CDES

In this Appendix I show that the ACDES indirect utility function (1) satisfies the standard regularity conditions, i.e. that it satisfies the constraints implied by utility maximization subject to a budget constraint. A regular indirect utility function $V(P, m)$ is (c.f. Proposition 3.D.3 in Mas-Colell et al. (1995)):

(i) Homogeneous of degree zero.

(ii) Continuous in P and m.

(iii) Strictly increasing in m and nonincreasing in all prices P_s.

(iv) Qausiconvex.

Proposition. *The augmented CDES indirect utility function $V(P, m)$ given by*

$$V(P, m) = \sum_s \gamma_s \frac{\left(\frac{m - \sum_{s'} P_{s'} \bar{c}_{s'}}{P_s}\right)^{\alpha_s} - 1}{\alpha_s}, \tag{1}$$

with the following parameter restrictions:

$$\gamma_s > 0, \sum_s \gamma_s = 1, \alpha_s \geq -1, \bar{c}_s \geq 0,$$

where the equality $\alpha_s = 1$ can hold for at most one good s, satisfies the regularity conditions (i)-(iv) on a convex set $X = \{(P, m) \ : \ P > 0 \wedge m > 0 \wedge m - \sum_{s'} P_{s'} \bar{c}_{s'} > 0\}$.

Proof. To simplify notation define $\tilde{m} \equiv m - \sum_{s'} P_{s'} \bar{c}_{s'} > 0$. Inspecting (1) it is easy to see that homogeneity of degree zero (i) and continuity (ii) are satisfied. To establish (iii), first observe that

$$\frac{\partial V(P, m)}{\partial m} = \sum_s \gamma_s \left(\frac{\tilde{m}}{P_s}\right)^{\alpha_s - 1} P_s^{-1} > 0,$$

so $V(P, m)$ is strictly increasing in m. Next, observe that

$$\frac{\partial V(P, m)}{\partial P_s} = -\gamma_s \left(\frac{\tilde{m}}{P_s}\right)^{\alpha_s - 1} \frac{\tilde{m} + P_s \bar{c}_s}{P_s^2} - \sum_{k \neq s} \gamma_s \left(\frac{\tilde{m}}{P_k}\right)^{\alpha_k - 1} \frac{\bar{c}_s}{P_k} < 0,$$

so $V(P, m)$ is strictly decreasing in all prices.

Finally, I show that V is strictly quasiconvex (thus satisfying (iv)), i.e. that for any $x, y \in X$ and $\lambda \in (0,1)$, $V(\lambda x + (1-\lambda) y) < \max\{V(x), V(y)\}$, where it is now convenient to use vector notation so that $x = [p_1^x \ p_2^x \dots p_S^x \ m]^T$. The proof relies on two observations: (a) the CDES indirect utility function $\tilde{V}(x)$ obtained by setting $\bar{c}_s = 0$ for all s is strictly quasiconvex on $\tilde{X} = \mathbb{R}_+^{S+1}$ under the given parameter restrictions, as shown by Jensen et al. (2011); (b) the augmented CDES can be expressed as $V(x) = \tilde{V}(Ax)$, where

$$A = \boldsymbol{I}_{S+1} + \begin{bmatrix} 0 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 0 & 0 \\ -\bar{c}_1 & \cdots & -\bar{c}_S & 0 \end{bmatrix};$$

furthermore, if $x \in X$ then $\tilde{x} = Ax \in \tilde{X}$. Combining these observations, we have for any $x, y \in X$ and $\lambda \in (0,1)$

$$\begin{aligned} V(\lambda x + (1-\lambda) y) &= \tilde{V}(\lambda Ax + (1-\lambda) Ay) \\ &= \tilde{V}(\lambda \tilde{x} + (1-\lambda) \tilde{y}) \\ &< \max\left\{\tilde{V}(\tilde{x}), \tilde{V}(\tilde{y})\right\} \\ &= \max\{V(x), V(y)\}. \end{aligned}$$

This establishes that V is strictly quasiconvex and completes the proof. □

As a final remark, if $\bar{c}_s < 0$ for any s, then the regularity conditions are satisfied for sufficiently high m (guaranteeing that the demand is positive for each good).

B Data Appendix

In this Appendix I document the sources of the data and describe the construction of variables used in my quantitative analysis.

B.1 Aggregate Data

I calculate the PPP-adjusted GDP as a product of real GDP per capita ($rgdpch$) and population (POP) taken from version 7.0 of the Penn World Table (Heston et al. (2011)). I HP-filter the resulting series with smoothing parameter 25 (falling in the 6.25-100 range standard in the literature for annual data) and divide by HP-filtered employment (see below) to obtain the smoothed real GDP per worker. PWT 7.0 is also used as a source for the level of nominal exchange rate ($XRAT$).

B.2 Sectoral Output, Employment and Price Data

To conduct the analysis of structural transformation at a sectoral level I construct an unbalanced panel of between 26 and 44 countries over the period 1970-2005. I assemble data from four sources: EU KLEMS database [O'Mahony and Timmer (2009)], GGDC 10-sector database [Timmer and de Vries (2009)], OECD STAN database [OECD (2011)] and Asian Productivity Organization database [APO (2010)]. Table B.1 presents the sample coverage and the primary source of information for each country. These sources provide information at a higher level of disaggregation than used in this study. I therefore aggregate the data by constructing a three sector classification: agriculture (comprising ISIC Rev. 3 sectors 01-05: agriculture, hunting, forestry and fishing), tradable industry (comprising ISIC sectors 10-37: mining and quarrying and manufacturing industries) and nontradables (comprising all other activities). In the paper I refer to those sectors as agriculture, manufacturing and services. To eliminate the effects of cyclical fluctuation, which are

Table B.1: Sample Coverage

Country	Coverage Period	Primary Source	Country	Coverage Period	Primary Source
Argentina	1991-2005	GGDC 10-sector	Japan	1970-2005	EU KLEMS
Australia	1970-2005	EU KLEMS	Korea	1973-2005	EU KLEMS
Austria	1970-2005	EU KLEMS	Malaysia	1975-1997	GGDC 10-sector
Bangladesh	1985-2004	APO	Mexico	1970-2005	GGDC 10-sector
Belgium	1970-2005	EU KLEMS	Netherlands	1970-2005	EU KLEMS
Bolivia	1986-2003	GGDC 10-sector	Norway	1970-2005	STAN
Brazil	1995-2005	GGDC 10-sector	Pakistan	1970-2005	APO
Canada	1970-2005	STAN	Peru	1991-2005	GGDC 10-sector
Chile	1979-2005	GGDC 10-sector	Philippines	1971-1997	GGDC 10-sector
China	1978-2005	APO	Poland	1995-2005	EU KLEMS
Colombia	1970-2005	GGDC 10-sector	Portugal	1970-2005	EU KLEMS
Czech Republic	1995-2005	EU KLEMS	Slovakia	1995-2005	EU KLEMS
Denmark	1970-2005	EU KLEMS	Spain	1970-2005	EU KLEMS
Finland	1970-2005	EU KLEMS	Sri Lanka	1971-2005	APO
France	1970-2005	EU KLEMS	Sweden	1970-2005	EU KLEMS
Germany	1991-2005	EU KLEMS	Switzerland	1991-2005	STAN
West Germany	1970-1990	GGDC 10-sector	Taiwan	1970-1997	GGDC 10-sector
Greece	1970-2005	EU KLEMS	Thailand	1970-2005	GGDC 10-sector
Hungary	1992-2005	EU KLEMS	United Kingdom	1970-2005	EU KLEMS
India	1970-2004	GGDC 10-sector	USA	1970-2005	EU KLEMS
Indonesia	1973-2005	GGDC 10-sector	Venezuela	1970-2003	GGDC 10-sector
Ireland	1970-1999	EU KLEMS	Vietnam	1991-2005	APO
Italy	1970-2005	EU KLEMS			

beyond the scope of this paper, I smooth the time-series of interest using the Hodrick-Prescott filter with smoothing parameter 25. The following paragraphs present more detailed description of the construction of individual variables.

The measure of sectoral employment I use is Total Employment (Number of Persons Engaged). This broad concept of labor input is the only measure consistently available for a large set of countries in all four databases. To obtain the smoothed series I simply filter the time series with sectoral employment separately for each sector and country.

To construct the sectoral value added in U.S. dollars I proceed in several steps. I begin by summing up all sectoral VA in current local prices to calculate the nominal GDP and apply the nominal exchange rate to obtain the GDP in U.S. dollars. Then I HP-filter the resulting series. Next I use the raw sectoral VA numbers to compute the VA shares and smooth those shares with the HP-filter. The smoothed sectoral VA in U.S. dollars is then computed by applying the smoothed VA series to the smoothed GDP series. This calculation guarantees that aggregating smoothed VA across sectors yields the smoothed GDP. I find that this procedure is more robust than smoothing individual sectoral series separately as it filters the annual-frequency movements in nominal exchange rate in a consistent way across all sectors.

Calculations of labor productivity require data on VA in constant (or chained) prices to compute the quantity index of sectoral VA. I thus begin by using the price deflators for VA to convert the raw nominal VA for disaggregated industries to VA in constant prices. Summing across industries within a sector yields VA in constant prices at a sector level. To smooth the series I proceed similarly as for nominal VA - I first smooth separately the GDP in constant prices and sectoral shares of that GDP and then combine the smoothed GDP with smoothed shares to obtain smoothed constant-price VA levels for agriculture, manufacturing and services. Finally, I divide the smoothed constant-price VA series by the smoothed employment series to obtain series of quantity of VA per worker in each sector. I use those series, normalized to one in reference year 1995 in each sector and each country as the empirical measure of sectoral labor productivity growth.

To calculate the evolution of sectoral relative prices I start by calculating a smoothed price deflator for each sector by dividing smoothed sectoral VA in U.S. dollars by the quantity index of sectoral VA described in the previous paragraph. Next I divide the deflator for agriculture and services by the price deflator for manufacturing. Finally, I normalize the two indices to one in 1995 in each country.

For a couple of countries additional steps are required to calculate consistent time series over the relevant sample period. The data for Japan comes from GGDC 10-sector database for 1970-72 and from EU KLEMS for 1973-2005. To link the data from both sources I essentially combine the growth rates over 1970-73 from the GGDC 10-sector database with levels from EU KLEMS database in 1973. The case of Germany is a little more complicated in that I use data for West Germany (from GGDC 10-sector database) for 1970-1990 and for unified Germany (from EU KLEMS) starting in 1991. To make the levels of variables comparable between the two entities when needed I exploit the fact that for 1991 data is available both for the unified Germany and the hypothetical West Germany.

B.3 International Trade Data

In order to compute bilateral trade flows in agriculture and manufacturing over the sample period I combine data from two datasets: the NBER-UN dataset [Feenstra et al. (2005)] and the BACI database prepared by researchers at CEPII [Gaulier and Zignago (2010)].

The NBER-UN dataset records bilateral trade flows at a 4-digit level according to SITC rev.2 classification. To map these disaggregated flows into two tradable sectors of the paper, agriculture and manufacturing, I develop a required concordance. As a starting point I use the SITC rev.2 5-digit to ISIC rev.2 4-digit concordance available from World Integrated Trade Solutions (WITS) project of the World Bank. On the production side I classify all industries with ISIC 4-digit code below 2000 as agriculture and the rest as tradable industry (called manufacturing in the paper). In the next step I adjust the mapping from trade classification to sector classification for a limited number of products which mostly involves moving some categories of meat, milled grains, and vegetable oils and their byproducts to agriculture. The rationale for this somewhat subjective adjustment is that industry classification is based on the final producer of a good with disregard of the share of value added in the last production stage. Since I use data on sectoral VA in my analysis I believe it is more appropriate if trade flows are assigned to sectors based on the VA content of the product and not the identity of the final processing industry. As measures of VA content at a product level are not readily available I had to use my judgment to conservatively reclassify some product categories. For example, WITS assigns both product 0113 ("Meat of swine, fresh, chilled or frozen") and product 0121 ("Bacon, ham & other dried, salted, smoked meat/ swine") to manufacturing industry 3111 ("Slaughtering, preparing and preserving meat"). I reclassify the first product as agriculture while keeping the processed meat assigned to manufacturing. Finally, in a very small number of cases I change the classification at 5-digit SITC level so that all SITC 4-digit code that appear in NBER-UN dataset can be unambiguously classified as agriculture or manufacturing.

The version of BACI dataset used in this paper provides bilateral trade flows by 6-digit HS92 product categories. To map these flows into agriculture and manufacturing in a way consistent with the treatment of NBER-UN data I first use the HS92 6-digit to SITC rev.2 5-digit concordance from WITS and then assign the SITC products in the same way as for the NBER-UN case.

Within the time span of the sample NBER-UN covers years between 1970-1995 while BACI data is available for 1995-2005. Since there are small differences in corresponding bilateral flows recorded by the two sources in overlapping years I compute a weighted average when both numbers are available.[61] In order to avoid discrete jumps in the data due to changing methodology, the weight on BACI flows is gradually increasing between 1995 and 2000.

The bilateral trade flows measured in U.S. dollars are then smoothed to reduce the effect of cyclical fluctuations and nominal exchange rate movements and thus to be more easily comparable with the data on smoothed VA in U.S. dollars described in the preceding subsection. Specifically, I apply the HP filter with smoothing parameter 25 separately to each available time series $\{X_{sjit}\}$ of imports in industry s by country j from country i. Using the filtered series I then compute total imports by country j and total exports by country i as $IMP_{sjt} = \sum_{i \neq j, i=1}^{N_t} X_{sjit}$ and $EXP_{sit} = \sum_{j \neq i, j=1}^{N_t} X_{sjit}$. Because the country coverage varies by year also the set of countries over which total imports and exports are calculated changes over time. This is necessary to make sure trade in the model world is balanced.

Finally, smoothed trade flows and smoothed VA in U.S. dollars VA_{sj} are used to calculate bilateral trade shares as:

$$\pi_{sji} = \frac{X_{sji}}{VA_{sj}\beta_s^{-1} + IMP_{sj} - EXP_{sj}},$$

[61] The two measures are very highly correlated with correlation coefficient above 0.99. R^2 from the regression of log NBER-UN flow on log BACI-flow is 0.97.

where β_s is a median share of value added in gross output in the subsample of countries for which data on both value added and gross output is available (EU KLEMS subsample). Imports from home are computed as $X_{sjj} = VA_{sj}\beta_s^{-1} - EXP_{sj}$ which ensures that the import shares sum to one for each country.[62]

Trade flows and VA series, smoothed and expressed in U.S. dollars, are also used to compute the overall trade deficit of a country relative to its nominal GDP through the formula:

$$\delta_{jt} = \frac{IMP_{Ajt} - EXP_{Ajt} + IMP_{Mjt} - EXP_{Mjt}}{VA_{Aj} + VA_{Mj} + VA_{Sj}}.$$

In less than 2% of country-sector-year observation aggregate trade flows derived by following the procedures described above are too large relative to the scale of domestic industry to be consistent with the Eaton and Kortum structure. Those cases (Belgium, Netherlands, Denmark, Taiwan and Slovakia) are primarily small countries with high levels of reexports and processing trade that the model does not account for. To deal with most of those cases I use time trends of bilateral flows to extrapolate to the problematic years. In two particularly stark cases (agricultural trade of Belgium and the Netherlands) I go further and restrict bilateral trade flows in agriculture involving those countries in a way that stabilizes their trade/output ratio at a level compatible with the model.

C Calibration Details

In this Appendix I provide details of the algorithm used to calibrate the model.

C.1 Calculating International Prices

Given sectoral wages, employment levels and prices in the reference year I find the model international prices through the following procedure. I first calculate the quantity of sectoral output as $q_{sit_R} = w_{sit_R}L_{sit_R}/P_{sit_R}$. The Geary-Khamis price of good s is then

$$p_s = \sum_{i=1}^{N} \frac{q_{sit_R}}{\sum_{j=1}^{N} q_{sjt_R}} \frac{P_{sit_R}}{p_i}, \tag{C.1}$$

where p_i is the PPP price level in country i defined as

$$p_i = \frac{\sum_s P_{sit_R} q_{sit_R}}{\sum_s p_s q_{sit_R}}. \tag{C.2}$$

Equations (C.1)-(C.2) need to be solved simultaneously for PPP price levels p_i and international prices p_s. In practice I use the matrix representation of the problem described in Diewert (1999). Aggregate real income per worker of country i relative to the US in the reference year can than be computed as

$$\frac{(\sum_s p_s q_{sit_R}/L_{it_R})}{(\sum_s p_s q_{sUSt_R}/L_{USt_R})}.$$

[62] Coefficients β_s are assumed to be constant. In the data they show a slight downward trend: between 1970 and 2005 the median VA/GO ration in the EUKLEMS subsample declined from 0.52 to 0.45 in agriculture, 0.34 to 0.30 in manufacturing, and 0.56 to 0.54 in services. This means the the calculations above slightly overestimate the degree to which trade intensity (π_{sjj}) increased over time.

Similarly, the growth of aggregate productivity between the reference year t_R and year t in country i can be calculated as

$$\frac{\left(\sum_s p_s q_{sit}/L_{it}\right)}{\left(\sum_s p_s q_{sit_R}/L_{it_R}\right)}.$$

C.2 Calibration of Preference Parameters

The calibrated parameter vector $\hat{\omega} = \left\{\hat{\alpha}_A,\, \hat{\alpha}_M,\, \hat{\alpha}_S,\, \hat{\bar{c}}_A,\, \hat{\bar{c}}_M,\, \hat{\bar{c}}_S\right\}$ minimizes the GMM objective function $J(\omega)$:

$$\hat{\omega} = \arg\min_{\omega} J(\omega)$$

Below I describe how the function $J(\omega)$ is evaluated. Given a set of parameters $\{\alpha_A,\, \alpha_M,\, \alpha_S,\, \bar{c}_A\, \bar{c}_M\, \bar{c}_S\}$:

1. Find normalized preference weights parameters $\{\gamma_A,\, \gamma_M,\, \gamma_S\}$ such that U.S. expenditures in the reference year are consistent with household optimization given normalization $P_{sUSt_R} = 1$, i.e. find $\{\gamma_A,\, \gamma_M,\, \gamma_S\}$ satisfying:

$$\begin{aligned}
\frac{E_{AUSt_R}}{\sum_{s'} E_{s'USt_R}} - \frac{1}{\sum_{s'} E_{s'USt_R}}\left[\bar{c}_A + \left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)\frac{\gamma_A\left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)^{\alpha_A}}{\sum_{s'} \gamma_{s'}\left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)^{\alpha_{s'}}}\right] &= 0\\
\frac{E_{MUSt_R}}{\sum_{s'} E_{s'USt_R}} - \frac{1}{\sum_{s'} E_{s'USt_R}}\left[\bar{c}_M + \left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)\frac{\gamma_M\left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)^{\alpha_M}}{\sum_{s'} \gamma_{s'}\left(\sum_{s'} E_{s'USt_R} - \sum_{s'} \bar{c}_{s'}\right)^{\alpha_{s'}}}\right] &= 0\,.\\
\gamma_A + \gamma_M + \gamma_S - 1 &= 0
\end{aligned}$$

   Note that expenditures are computed as in (14) and do not depend on ω.

2. In the reference year solve for $\{P_{Aj},\, P_{Mj},\, P_{Sj}\}$ the system of equations

$$\begin{aligned}
\frac{E_{Aj}}{\sum_{s'} E_{s'j}} - \frac{1}{\sum_{s'} E_{s'j}}\left[P_A\bar{c}_A + \left(\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}\right)\frac{\gamma_A\left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}}{P_A}\right)^{\alpha_A}}{\sum_{s'} \gamma_{s'}\left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}}{P_{s'}}\right)^{\alpha_{s'}}}\right] &= 0\\
\frac{E_{Mj}}{\sum_{s'} E_{s'j}} - \frac{1}{\sum_{s'} E_{s'j}}\left[P_M\bar{c}_M + \left(\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}\right)\frac{\gamma_M\left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}}{P_M}\right)^{\alpha_M}}{\sum_{s'} \gamma_{s'}\left(\frac{\sum_{s'} E_{s'j} - \sum_{s'} P_{s'}\bar{c}_{s'}}{P_{s'}}\right)^{\alpha_{s'}}}\right] &= 0 \quad , \; j = 1, ..., N\\
\frac{\sum_s p_s q_{sj}/L_j}{\sum_s p_s q_{sUS}/L_{US}} - \frac{y_j}{y_{US}} &= 0
\end{aligned}$$

   where the procedure for calculating Geary-Khamis international prices p_s is described in Section C.1 and where y_j denotes real GDP per capita in the data. In non-reference years replace the last equation in the system above with an equation matching real income per worker in constant international prices

$$\frac{\sum_s p_s q_{sit}/L_{it}}{\sum_s p_s q_{sit_R}/L_{it_R}} - \frac{y_{jt}}{y_{jt_R}} = 0.$$

3. Given wages from (14) and prices from the previous step calculate labor productivities as $A_{sit} = w_{sit}/P_{sit}$. Let t_l^i and t_f^i denote the last and first year that country i appears in the sample. Calculate annualized average log growth of A_{sit} as $g_{si}(\omega) = \frac{1}{t_l^i - t_f^i} \log\left(\frac{A_{sit_l^i}(\omega)}{A_{sit_f^i}(\omega)}\right), \quad s \in \{A, M, S\}$.

4. Using time series described in Appendix B calculate annualized average log growth g_{si}^d of labor productivity in the data. Also create instruments x_s for sector s log productivity growth: a constant, log growth in sector s employment and log growth in expenditure share of sector s (all growth rates on an annualized basis).

5. Compute a vector of sample moments

$$h_n(\omega) = \left[\frac{1}{n}\sum_{j=1}^{n} x_{Aj}^{(1)}\left(g_{Aj}^d - g_{Aj}(\omega)\right) \ldots \frac{1}{n}\sum_{j=1}^{n} x_{Sj}^{(3)}\left(g_{Sj}^d - g_{Sj}(\omega)\right)\right]',$$

where $n = N^c$ is the sample size and N^c is the total number of countries appearing in the sample.

6. Given weighting matrix $\mathbf{W}$ evaluate the GMM objective function as

$$J(\omega) = n \cdot h_n(\omega)' \mathbf{W} h_n(\omega).$$

I use the identity matrix as the weighting matrix while implementing the algorithm.

C.3 Monte Carlo Simulation

The algorithm below describes details of the simulation discussed at the end of Section 3.4. To ease computational burden, the simulation uses only half of the countries in the main dataset (21 countries present in the sample throughout 1970-2005).

1. Set model parameters: preference parameters at (approximately) calibrated values; common technology parameters (β_s, θ_s) at values assumed in the baseline calibration; labor endowments L_i and labor wedges ξ_{si} at values used in baseline calibration; technology levels T_{si} in the reference year at calibrated levels, with random growth rates chosen so that the average sectoral labor productivity growth is similar as in the data; bilateral trade costs τ_{sji} in the reference year at values consistent with the baseline calibration, with annual declines of 0.4% in agriculture and 0.6% in manufacturing; trade deficits D_i at zero.

2. Simulate the model with assumed parameters and export as observables the same data as used by the calibration procedure.

3. Add measurement error to the simulated data on sectoral labor productivity growth. Draw the measurement error independently across countries from a Normal distribution with a covariance matrix (across sectors) based on the covariance of residuals from the baseline calibration.

4. Apply the calibration procedure to the simulated data and record the value of calibrated preference parameters.

5. Repeat steps 3-4 200 times.

D Additional Robustness Checks

In this Appendix I report additional robustness checks regarding the calibration of preference parameters and their impact on counterfactual-based decompositions.

Two-Step GMM

The GMM procedure gives equal weight to all 9 moment conditions used in the calibration of preference parameters. Column 2 in the first panel of Table D.1 shows parameters calibrated instead by a two-step GMM in which the results from the first step (the baseline calibration) are used to construct the weighting matrix. Parameter estimates are close to the baseline case (for convenience repeated in column 1), but suggest even slightly less substitutability. The second panel of Table D.1 summarizes the key substantive results on the pattern of the Labor Relocation Index obtained for this specification (for brevity focusing on counterfactuals with one channel operative). The ranking of the four forces is the same as in the baseline and the numerical magnitude of LRI's is similar.

The two-step estimation can be used to perform a test of overidentifying restrictions as there are 9 moment conditions and 4 parameters. The value of the J-Statistic (GMM Objective function) is 6.24, below the 11.07 value of 0.95 percentile of χ^2_5 distribution. Thus the null hypothesis of model validity is not rejected. Because some of the estimated parameters are on the boundary of the parameter space, the J-test might have a size different than the nominal 95% confidence level. However, recent research by Ketz (2015) shows that when a parameter is at or close to the boundary the standard J-test overrejects the null, strengthening the case for model validity.

Data Smoothing

Time series data used in the paper is smoothed using a Hodrick-Prescott filter with smoothing parameter 25, primarily to mitigate jumps due to nominal exchange rate volatility and business cycle fluctuations. A potential concern is that the HP filter can be inaccurate at endpoints, and endpoints are used both to calculate long-run sectoral productivity growth for the calibration and changes in labor shares for the calculation of LRI in the counterfactuals. To address this concern, column 3 in the first panel of Table D.1 reports preference parameters calibrated ignoring the first and last two years for each country. Differences in estimates relative to baseline are well within what might be expected from normal data variation (typically within half a bootstrapped standard error from the baseline).

The third row in the second panel of Table D.1 summarizes the corresponding LRI's, also calculated by scrapping the first and last two years for each country. Here also the results as similar as in the baseline. Thus HP-filtering the data does not seem to have any systematic impact on the results of interest.

E Physical Capital and Worker Heterogeneity

In this Appendix I discuss consequences of omitting physical capital and worker heterogeneity for the interpretation of results in this paper.

Table D.1: Additional Results

(A) Calibrated Preference Parameters

	(1)	(2)	(3)
Model	ACDES Baseline	Two Step GMM	Trimmed Sample
α_A	-1.00	-1.00	-0.86
α_M	-0.89	-1.00	-1.00
α_S	-0.67	-0.85	-0.78
$\bar{c}_A$	3.89	4.15	4.57

(B) Labor Relocation Index

Operative channels	P	N	T	W
Baseline	0.43	0.27	0.21	0.00
Two step GMM	0.39	0.21	0.17	0.00
Trimmed Sample	0.39	0.28	0.15	0.00

Notes: Panel A: Col. (1): ACDES model with restrictions $\bar{c}_M = \bar{c}_S = 0$ (Baseline); Col. (2): Two-step GMM procedure; Col. (3): Estimation ignoring the first and last two years for each country. Subsistence requirements $\bar{c}_s$ multiplied by a factor 10^3 to increase readability. Panel B: Table shows the median LRI computed across 45 countries in the sample. Channels driving structural change: P - sector-biased productivity growth, N - nonhomothetic preferences, T - international trade, W - changes in intersectoral wedges.

E.1 Physical Capital

In this paper labor is the only primary factor of production. This choice is largely driven by a lack of reliable data on the evolution of capital stocks at a sectoral level in developing countries. Below I explain how incorporating capital would affect the measurement and interpretation of some variables, and argue that it would be unlikely to affect the main results of the paper.

An obvious issue affected by treatment of capital is the measurement of wedges. When value added is generated only by labor, the intersectoral labor wedges can be computed as relative value added per worker across sectors. In the presence of capital, the appropriate measure is relative labor compensation per worker. The true wedge should thus be calculated as:

$$\tilde{\xi}_s = \frac{\eta_s VA_s/L_s}{\eta_M VA_M/L_M} = \frac{\eta_s}{\eta_M}\xi_s,$$

where η_s is the share of labor compensation (i.e. $w_s L_s$) in value added in sector s and where ξ_s is the simple wedge based on value added per worker. The simple wedge ξ_s should be corrected for the difference in labor intensity across sectors. However, there is suggestive evidence that such an adjustment would not have a major effect on reducing the calculated wedges. First, Valentinyi and Herrendorf (2008) carefully calculate factor income shares in the US, finding very similar labor shares in manufacturing and in services. Labor share in agriculture is, if anything, lower than in manufacturing, which would tend to increase the magnitude of the wedge between agriculture and manufacturing (since $\xi_A < 1$ almost universally in the data). Second, I have calculated the adjusted wedges for a subset of countries in my sample with required data available (EUKLEMS subsample excluding observations with reported labor compensation in excess of value added). The adjusted and simple wedges are similar: pooled across 540 available country-year observations, the geometric means of $\tilde{\xi}_A$ and ξ_A are 0.57 and 0.55, and the correlation between the two measures is 0.73. For services, the corresponding geometric means are 0.95 and 0.93, and the correlation is 0.51. These similarities reflect the fact that the average labor share of value added in all sectors is nearly identical in this sample at 0.66-0.69.

The second issue requiring clarification is the interpretation of productivity growth. The key empirical measure used is the growth of labor productivity, measured as the growth of real value added per worker. In the model this growth is accounted for by an exogenous growth in fundamental technology (scaled by presence of intermediates and adjusted for the selection effect of international trade). In the presence of capital, the growth in labor productivity also reflects an increase in capital per worker. In my framework such capital deepening would be attributed to technological progress.

Such an interpretation of productivity growth could be problematic if capital deepening was an important independent driver of structural change. This is the case in Acemoglu and Guerrieri (2008). In their model capital deepening occurs at different rates across sectors due to differences in sectoral labor intensities. In such setting, labor productivity could grow at different rates across sectors due to differential growth in capital-labor ratios rather than differences in exogenous technological progress. Recent quantitative work suggests, however, that this mechanism is not likely to be very important. Herrendorf et al. (2015) estimate sectoral production functions for the US and conclude that the postwar U.S. experience can be well explained by Cobb-Douglas production functions with common factor shares and different rates of technological progress in agriculture, manufacturing and services. In a similar spirit, Alvarez-Cuadrado et al. (2014) find that the amount of structural transformation between manufacturing and services in the US can be explained primarily by differences in rates of technological progress in the two sectors, with capital deepening playing only a very minor role. To the extent that the U.S. experience is representative of other countries, these results suggest that the main findings of the paper on the relative importance of different channels would not be affected by considering capital accumulation more explicitly.

E.2 Worker Heterogeneity

Throughout this paper I interpret intersectoral wedges as representing different wages for equivalent labor faced by producers in different sectors. The assumption of homogenous labor allows me to connect closely to existing theories of structural change and to focus the analysis on sectoral labor relocation and sectoral labor productivity patterns, which are the main objects studied in the literature. For completeness I now briefly explore an alternative interpretation of the wedges.

While the existence of gaps in labor productivity (nominal VA per worker) and in wages across broad sectors is well established, the interpretation of these gaps is debated. Gollin et al. (2014) find that a significant fraction of the productivity gap between agriculture and nonagriculture remains after controlling for a host of measurement issues (hours worked, education, quality of education). Using census data for a smaller set of countries, Herrendorf and Schoellman (2015) argue instead that the gaps in sectoral wages can be largely explained by human capital differences across sectors once one allows for sector-specific returns to education. I now take the latter view to the extreme and attribute wage differentials entirely to differences in human capital per worker across sectors. The measure of sectoral productivity I now use, denoted $\tilde{A}_s$, is the real output per manufacturing-equivalent worker, i.e. a worker with the same level of human capital as a manufacturing worker. There is a simple relationship between $\tilde{A}_s$ and the measure of labor productivity A_s used earlier: $\tilde{A}_s = A_s/\xi_s$. Conditional on values of preference parameters, finding the calibrated $\tilde{A}_s$ requires only a simple transformation of previously calculated productivity measures. Two observations clarify how previous findings might be affected.

First, worker heterogeneity would complicate the calculation of the counterfactual exercises in Section 4. The reason is that simulating the model would require additional modeling of how workers switch sectors. However, the assumption of labor homogeneity does not significantly affect the accounting exercises from Section 5. Conditional on a model (i.e. values of preference parameters), the cross-sectional patterns of $\tilde{A}_s$ and A_s are similar. Because ξ_A is positively correlated with A_A, dispersion in $\tilde{A}_A$ is somewhat lower than dispersion in A_A. But the overall ranking is robust: productivity dispersion is largest in agriculture and smallest in services. For example, the coefficient of variation in the baseline ACDES calibration is 1.31/0.71/0.57 for $A_A/A_M/A_S$ and 1.12/0.71/0.62 for $\tilde{A}_A/\tilde{A}_M/\tilde{A}_S$. Furthermore, comparing productivity across models is not affected by the choice of productivity measure, since $\tilde{A}_{si}^{CES}/\tilde{A}_{si}^{ACDES} = A_{si}^{CES}/A_{si}^{ACDES}$. Thus the assumption of labor homogeneity is not important for analyzing the implications of imposing homothetic preferences in a model of structural change.

F Alternative Calculations of Sectoral Labor Productivity Levels

In Section 5 I show that the baseline calibration of the model and the calibration imposing homothetic preferences deliver quantitatively different patterns of sectoral labor productivity levels. Specifically, the CES model implies larger dispersion of labor productivity in agriculture and somewhat smaller dispersion in services. A natural question is: which model - baseline or homothetic - delivers more reasonable productivity estimates? In this Appendix I shed some light on this question by presenting alternative calculations of sectoral labor productivity. These calculations are meant to be only illustrative since the methods I use have their clear limitations. With this caveat in mind, the baseline numbers are generally closer than the homothetic counterparts to the alternative calculations. This suggests that labor productivity levels derived from the baseline model are likely to be more reliable.

Three Sectors: Productivity Level Database

Sectoral labor productivity as understood in this paper is supposed to measure quantity of value added generated per worker. Calculating employment at the level of broad sectors is relatively straightforward. The main difficulty lies in aggregating the quantity of value added and then comparing it across countries. The preferred approach would be to use detailed producer prices of comparable goods across countries to construct producer price based PPPs and use them to convert nominal output in each country into comparable quantity index. Unfortunately, sector-level producer price PPPs are not easily available for a broad range of countries. The most comprehensive source of such data is probably the GGDC Productivity Level Database (Inklaar and Timmer (2008)). I therefore use the information on value added and producer price based VA PPPs in 1997 from the PLD to calculate measures of labor productivity in agriculture, manufacturing and services that are comparable across countries.

The first panel of Table F.1 shows the dispersion of labor productivity in the baseline calibration, in the calibration with the CES restriction and in PLD-based calculation. Both models predict larger dispersion in agriculture than PLD while delivering very similar magnitudes for manufacturing and services. Unfortunately, this exercise does not help to differentiate the ACDES and CES estimates. The reason is that the overlap between my base sample and PLD is restricted to only 23 countries (out of 44 in the base sample). Moreover, the remaining countries are mainly the richer ones.

Because the largest differences between CES and ACDES exist for the poor countries, they are not captured in the PLD subsample.

Agriculture: Restuccia et al. (2008)

Comparing productivity in agriculture is somewhat easier than in other sectors since the bulk of agricultural output consists of commodities. Restuccia et al. (2008) calculate real output in agriculture using the FAO data on factory-gate prices of commodities in 1985. Then they divide their real VA in agriculture by employment in agriculture, arriving at their measure GDPaLa which I use directly in the following comparison. On the sample of 32 countries in 1985 (my full sample in that year except for China and Taiwan) I find that the dispersion of GDPaLa is a little lower than dispersion in my baseline agricultural labor productivity. Since CES model predicts even higher dispersion, it is less consistent with the results of Restuccia et al. (2008).

Three Sectors: International Comparison Program

Data on consumer prices across countries is available for a much wider set of countries thanks to the International Comparison Project. It is therefore common to use PPPs derived from consumer prices to compare real output across countries. Because retail prices include value added from multiple sectors, this approach becomes problematic when the object of interest is sectoral rather than aggregate productivity. Nevertheless, I use information from the 1996 Benchmark PWT to calculate real sectoral output in two alternative ways. In the first approach, I assign the disaggregated expenditure categories from Table Vb to one of the three broad sectors. Then I add up expenditures in international dollars within each sector. Thus defined sectoral expenditure in international prices divided by sectoral employment becomes the sectoral productivity measure denoted ICP1 in the third panel of Table F.1. In the second approach, I use my baseline measure of nominal value added and divide it by sectoral PPP to construct the measure ICP2. Sectoral PPPs are obtained by aggregating PPPs for detailed categories from the 1996 Benchmark PWT Table IIIb.

Once again, agricultural productivity dispersion (for both approaches) is lower than predicted by the baseline model, which in turn is lower than predicted by the CES model. For manufacturing and services there are smaller discrepancies among all measures.

Agriculture and Manufacturing: Gravity

A common approach in the international trade literature is to infer measures of productivity in tradable sectors from structural gravity equations. Below I replicate a common empirical strategy with the aid of bilateral trade cost variables from CEPII. Equation (5) implies that:

$$\ln\left(\frac{\pi_{sji}}{\pi_{sjj}}\right) = \ln\left(T_{si}\left(w_{si}^{\beta_s}P_{si}^{1-\beta_s}\right)^{-\theta_s}\right) - \ln\left(T_{sj}\left(w_{sj}^{\beta_s}P_{sj}^{1-\beta_s}\right)^{-\theta_s}\right) - \theta_s \ln \tau_{sji}.$$

Suppose that bilateral trade flows are modeled as

$$\ln \tau_{sji} = d_{sd} + b_s + l_s + im_{sj} + \nu_{sji},$$

where d_d denotes a distance bin, b is an indicator for a common border, l for a common language and im is the importer-specific component of trade costs. Combining the last two expression yields

Table F.1: Alternative Measures of Sectoral Labor Productivity

	ACDES	CES	PLD
A	0.97	0.99	0.62
M	0.47	0.46	0.47
S	0.23	0.22	0.18
23 countries in 1997			

	ACDES	CES	GDPaLa
A	1.20	1.31	1.07
32 countries in 1985			

	ACDES	CES	ICP1	ICP2
A	1.23	1.32	0.97	0.81
M	0.67	0.70	0.60	0.65
S	0.54	0.50	0.48	0.43
39 countries in 1996				

	ACDES	CES	Gravity
A	1.31	1.42	0.67
M	0.71	0.74	0.42
44 countries in 1995			

Notes: Table reports the coefficient of variation of sectoral labor productivity computed over the subsample indicated in the bottom row of each panel. A, M and S correspond, respectively, to agriculture, manufacturing and services.

$$\ln\left(\frac{\pi_{sji}}{\pi_{sjj}}\right) = \underbrace{S_{si}}_{\text{exporter f.e}} - \underbrace{[S_{sj} + \theta_s im_{sj}]}_{\text{importer f.e.}} - \theta_s d_{sd} - \theta_s b_s - \theta_s l_s - \theta_s \nu_{sji}, \tag{F.1}$$

where $S_{si} = \ln\left(T_{si}\left(w_{si}^{\beta_s} P_{si}^{1-\beta_s}\right)^{-\theta_s}\right)$. I estimate equation (F.1) using the popular method of Santos Silva and Tenreyro (2006). Let $\hat{f}_{Esi}$ denote the estimated exported fixed effect relative to the U.S. (omitted category), i.e. $\hat{f}_{Esi} = \widehat{S_{si} - S_{sUS}}$. Using (4) and (5) it can be shown that sectoral productivity in tradable sector s relative to the U.S. level can be calculated as

$$\widehat{\frac{A_{si}}{A_{sUS}}} = \left(e^{\hat{f}_{Esi}}\right)^{\frac{1}{\theta_s}} \left(\frac{w_{si}}{w_{sUS}}\right)^{\beta_s} \left(\frac{\sum_j e^{\hat{f}_{sj}}\left(\hat{\tau}_{sij}\right)^{-\theta_s}}{\sum_j e^{\hat{f}_{sj}}\left(\hat{\tau}_{sUSj}\right)^{-\theta_s}}\right)^{-\frac{1}{\theta_s}(1-\beta_s)} \left(\frac{\hat{\pi}_{sii}}{\hat{\pi}_{sUSUS}}\right)^{-\frac{1}{\theta_s}},$$

where $\hat{\tau}_{sij}$ and $\hat{\pi}_{sii}$ are the predicted trade costs and trade intensities. To evaluate the expression above I measure wages as nominal value added per worker, just as in the main body of the paper.

The last panel of Table F.1 shows that the dispersion of sectoral labor productivity inferred from the gravity approach is much lower in both sectors than in either my baseline calibration or in the CES calibration (which gives the highest dispersion in agriculture). Given that the gravity approach predicts consistently lower productivity differences than any other method, this might, however, suggest that it is not the most robust way of determining productivity levels.