HETEROGENEOUS PATHS OF INDUSTRIALIZATION

By

Federico Huneeus and Richard Rogerson

July 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2253

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Heterogeneous Paths of Industrialization*

Federico Huneeus
Yale University & Central Bank of Chile

Richard Rogerson
Princeton University

July 2020

Abstract

Industrialization experiences differ significantly across countries. We use a benchmark model of structural change to shed light on the sources of this heterogeneity and, in particular, the phenomenon of premature deindustrialization. Our analysis leads to three key findings. First, benchmark models of structural change robustly generate hump-shaped patterns for the evolution of the manufacturing sector. Second, heterogeneous patterns of catch-up in sectoral productivities across countries can generate variation in industrialization experiences similar to those found in the data, including premature deindustrialization. Third, differences in the rate of agricultural productivity growth across economies can account for a large share of the variation in peak manufacturing employment shares.

Keywords: Structural transformation, productivity growth, industrialization.

JEL Codes: E24, O11, O13, O14, O33, O41.

*We thank conference participants at Notre Dame for useful comments. Felipe Leal provided excellent research assistance.

1 Introduction

In his Nobel Prize address, Kuznets emphasized structural transformation–the reallocation of economic activity across broad sectors–as one of the key stylized facts of growth and development. One empirical regularity of structural transformation is that the size of the industrial sector exhibits a hump-shaped pattern, increasing at low levels of development (i.e., the industrialization phase), reaching a peak, and then declining in the later stages of development (i.e., the deindustrialization phase). Recent work by (Rodrik, 2016) documents that many recent developers seem to be experiencing a much lower value for this peak, and that the peak is occurring at a much lower level of development relative to what earlier developers experienced. He coined the term premature deindustrialization to describe this phenomenon.[1]

In this paper we study the industrialization process from the perspective of a simple benchmark model of structural change. We have three key findings. First, we show that the model robustly implies hump-shaped dynamics for the employment share of manufacturing.[2] Second, we show that variation in the profile of sectoral productivity growth rates across countries can generate variation in industrialization patterns that that mimic those found in the data, including the phenomenon of premature deindustrialization. In particular, we show that relatively slow productivity growth in agriculture can give rise to differences comparable to those found in the data. Third, we calibrate our model to match the industrialization process of the US and then use it to study the industrialization experiences of a set of Asian and Latin American economies. Our model accounts for a significant portion of the variation in industrialization paths found in the data and differences in the relative growth rate of agricultural productivity are key to this finding.

The literature on structural change emphasizes the role of sectoral productivities in

[1]This observation was also noted independently by Felipe et al. (2018).

[2]We follow the standard practice in the literature of using the term manufacturing to refer to the broader industrial sector.

shaping the structural transformation process, and consistent with this, relative sectoral productivity growth rates play a central role in our analysis. Following Gollin et al. (2002, 2007), food is a necessity and the agricultural employment share is dictated by productivity of the agricultural sector. Non-agricultural employment is allocated between manufacturing and services and following Boppart (2014), depends both on the overall level of productivity as well as the relative productivity of the two sectors.[3]

The evolution of the manufacturing employment share is determined by the interplay of two forces: productivity growth in agricultural creates a flow of workers into manufacturing, but (for empirically reasonable specifications) productivity growth in the non-agricultural sectors creates a flow of workers out of manufacturing. At low levels of development the first force dominates, while at higher levels of development the second force dominates, thereby giving rise to the hump-shaped pattern for the manufacturing employment share.

Late developing economies are effectively inside the world technology frontier but are moving toward it. It is well established that different countries have moved toward the frontier at significantly different rates. But what is important for our analysis is the fact that this rate varies across sectors within economies. Variation in the rate of convergence across sectors affects the relative magnitudes of the two forces identified in the previous paragraph and therefore affect the path of industrialization. For example, we show that relatively slow growth in agriculture will lead to a lower peak employment share for manufacturing and that this peak will be reached at an earlier point in the development process.

Our paper is intimately related to the recent and growing literature on models of

[3]Our specification thus allows for the allocation of non-agricultural employment between manufacturing and services to be influenced by both income effects as in Kongsamut et al. (2001) and relative price effects as in Baumol (1967) and Ngai and Pissarides (2007). Święcki (2017) and Comin et al. (2015) also allow for both effects. For additional discussion and evidence see Herrendorf et al. (2014).

structural change.[4] Our analysis is most closely linked to those of Duarte and Restuccia (2010) and Świecki (2017). Like us, Duarte and Restuccia (2010) study productivity driven structural transformation in a large set of countries using a benchmark closed economy model of structural change. While our model is somewhat more general because it allows for non-homotheticicities in preferences over manufacturing and services, the key difference between the two analyses is our focus on the industrialization phase and the ability of the model to account for the heterogeneity in industrialization experiences across countries. Świecki (2017) extends Duarte and Restuccia (2010) to a multi-country setting and considers additional driving forces. His analysis focuses on the post 1970 period and again does not focus on the industrialization phase.

Recent papers by Sposi et al. (2020) and Wise (2020) have also sought to isolate factors that might give rise to the premature industrialization phenomenon. These papers both emphasize open economy interactions, and so are complementary to our analysis of forces in a closed economy setting. Our analysis is also related to that in Gollin et al. (2016), who study heterogeneous urbanization experiences and how this relates to industrialization.

An outline of the paper follows. In the next section we present evidence on the heterogeneity in industrialization experiences across a set of Asian and Latin American countries in the post 1950 period, as well as three European economies that also experienced considerable industrialization during this period. While we choose a different representation of the data than Rodrik (2016), our analysis yields a similar characterization. In Section 3 we present a benchmark model of structural change and study the forces shaping industrialization. In Section 4 we calibrate the model to the US industrialization experience, and Section 5 uses the calibrated model to illustrate the ability of the model to capture the quantitative differences in industrialization experiences when sectoral productivity profiles differ. Section 6 connects the model to data for our sample of Asian and Latin American

[4]Important contributions include Kongsamut et al. (2001), Gollin et al. (2002, 2007), Ngai and Pissarides (2007), Buera and Kaboski (2009), Boppart (2014) and Comin et al. (2015).

economies and shows that differences in the growth rate of agricultural productivity across countries can account for a large part of differences in peak manufacturing employment shares across countries. Section 7 discusses extensions and Section 8 concludes.

2 Industrialization Patterns Across Countries

In this section we document patterns of industrialization for a set of Asian and Latin American economies using the Groningen Growth and Development Centre (GGDC) 10-Sector database. By industrialization we refer to the phase of economic development in which the manufacturing sector is growing in terms of its share of the overall labor force. At its core, industrialization reflects the release of labor from agriculture that is then absorbed into non-agricultural activities. To best focus on this dynamic, we study the relationship between the release of labor from agriculture and its absorption into the manufacturing sector. While we focus on an alternative representation of the data, our characterization is very similar to that offered by Rodrik (2016).[5]

2.1 Data

Our selection of countries is dictated by those that are included in the GGDC 10-Sector database. The Asian countries in the dataset are China, India, Indonesia, Japan, South Korea, Malaysia, the Philippines, Taiwan and Thailand. The Latin American countries are Argentina, Bolivia, Brazil, Chile, Colombia, Costa Rica, Mexico, Peru and Venezuela. We also include three countries from Europe that are included in the data base and experienced significant industrialization since 1950–France, Spain and Italy. While the data set generally covers the period from 1950 to 2010, coverage for some countries begins

[5] Rodrik (2016) focused on value added shares and the level of GDP/capita measured using PPP. We focus on employment shares and so do not use any information about relative prices.

after 1950.[6] While the US is included in the GGDC 10-Sector database, the post 1950 data for the US is of limited interest for the simple reason that the post 1950 period does not cover the industrialization phase in which the employment share of manufacturing is increasing. In order to include the US experience as a reference point we will combine data from Carter et al. (2006) for the pre-1930 period with data from the BEA starting in 1929 to cover the US over the period 1880-1980.

We aggregate the ten sectors covered by the GGDC 10-Sector database into three using standard methods. Agriculture is one of the ten sectors in the database, so this does not involve any aggregation. We aggregate four sectors (mining, manufacturing, construction and utilities) to obtain what we will label as manufacturing, and the remaining five sectors (trade, restaurants and hotels, transportation, finance insurance, real estate and business services, government and community, social and personal services) are combined to obtain what we label as services.

For each country we compute time series for the employment shares of the three sectors, which we denote as h_{at}, h_{mt} and h_{st} for agriculture, manufacturing and services, respectively. We define h_{nt}, the nonagricultural employment share, as $1 - h_{at}$. Our empirical analysis focuses on the relationship between h_{mt} and h_{nt}. Because we are interested in trend relationships, we smooth the data by regressing h_{mt} on a fifth order polynomial in h_{nt}, and will use these smoothed profiles in our analysis.[7]

Almost all of the countries in our sample have experienced peaks for their manufacturing employment share. The clear exceptions are China, India, and Thailand. The cases of Bolivia and Colombia are less clear–it appears that they have reached their peaks near the end of the sample period, though absent additional data it is not possible to make a

[6]The notable exceptions are that data for Indonesia, Malaysia and the Philippines does not start until the 1970s.

[7]For each country we only consider the range of values for h_{nt} that are are observed in the data; i.e., we only use our polynomial to smooth the data and do not use it to extrapolate either forward or backward in time.

Figure 1: Paths of Industrialization: Four Emerging Economies

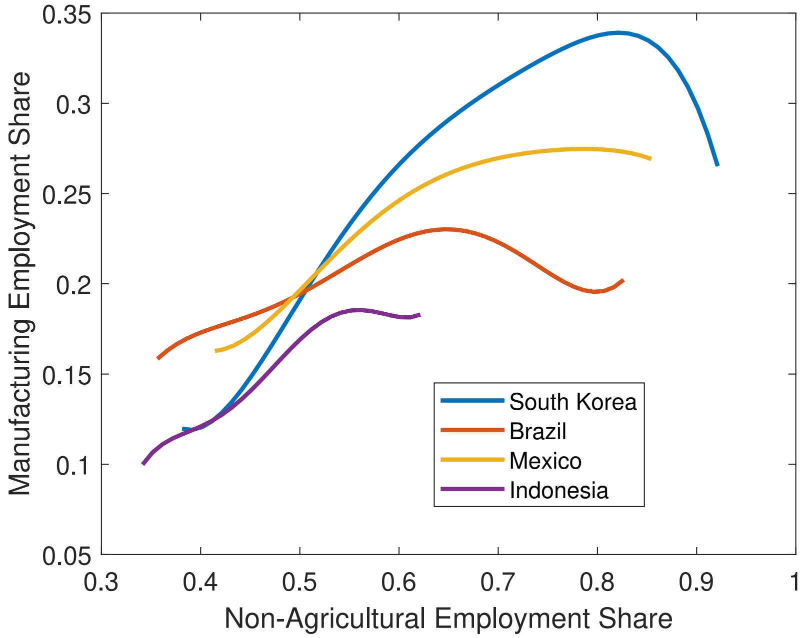

definitive statement. In what follows we treat them as having reached their peaks, but excluding them does not affect the empirical patterns that we document.

2.2 Patterns

As a first step we illustrate the range of experiences within our sample of countries. Figure 1 shows profiles for four countries that have experienced peak manufacturing employment shares–South Korea, Brazil, Mexico and Indonesia.

We highlight three properties of Figure 1. First, the level of peak employment in manufacturing varies significantly: Indonesia has a peak value below 0.20, whereas South Korea reaches a value of almost 0.35. Second, there is also significant variation in the value of h_{nt} at which the peak is reached, ranging from less than 0.60 for Indonesia to more than 0.80 for South Korea. Third, there is a strong positive correlation between the level of the peak and the value of h_{nt} at which the peak occurs.

While Figure 1 showed that the industrialization process varies substantially across

Figure 2: Paths of Industrialization: Advanced Economies

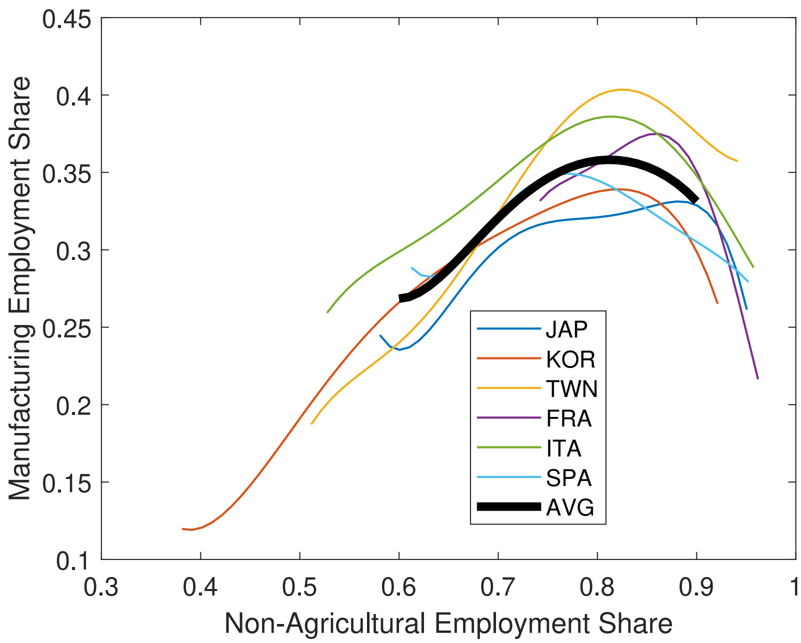

countries, it is also of interest to ask whether there is some sense of a "typical" pattern that current advanced countries have followed. To examine this, Figure 2 shows profiles for the six countries in our sample that currently qualify as advanced: Japan, South Korea, Taiwan, France, Italy and Spain.

While there is still some heterogeneity among the experiences of these countries, the dispersion is quite small relative to what we saw in Figure 1. All six of these countries reach their peak manufacturing employment shares when h_{nt} lies between .80 and .90, and the peak shares range from 0.34 to 0.40. The thick black line in the picture reflects the average for this subset of countries for h_{nt} in the range of 0.60 to 0.90.[8]

As described earlier, we create a profile for the US pattern of industrialization using data from Carter et al. (2006) and the BEA. The early data is decadal, and so cannot really be smoothed, but for the post 1929 period we compute five year moving averages for employment shares. Figure 3 shows the data for the US and the profiles for the six

[8]The average excludes France on account of the fact that the French data do not begin until h_{nt} is already beyond 0.70.

Figure 3: Paths of Industrialization: Advanced Economies and the US

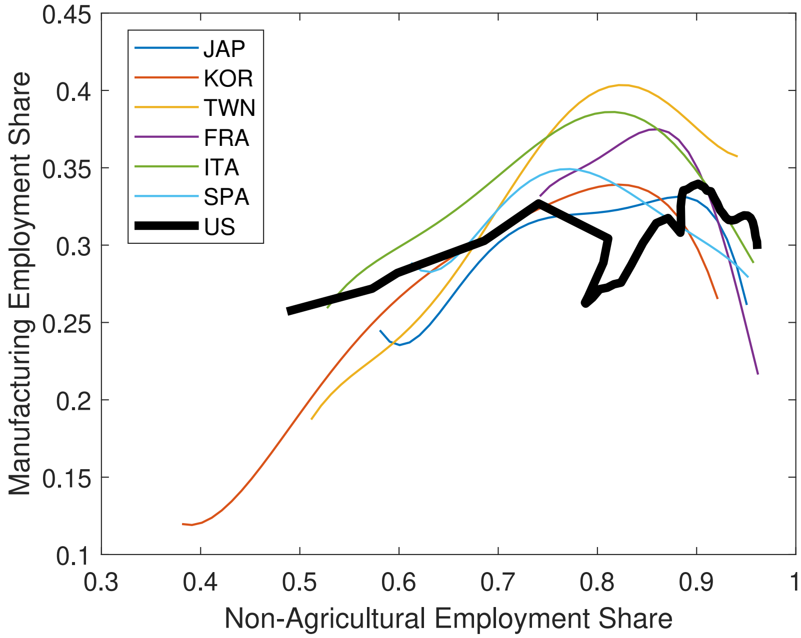

advanced economies shown in Figure 2.

A notable feature of the industrialization process in the US is that it was severely disrupted by the Great Depression, and this effect is readily apparent in the figure. It seems reasonable to infer that the counterfactual profile that would have occurred in the US in the absence of the Great Depression would have been broadly similar to what these countries experienced.

Figure 4 shows the profiles for the three countries that have not yet reached a peak for their manufacturing employment share–China, India, and Thailand.[9] As a reference point, we have included the profile for South Korea on this figure as well.[10]

Given that none of these countries has yet reached their peak employment share for manufacturing it is premature to say anything definitive. But, a casual look at this figure suggests that these three countries are following different paths and that only China seems

[9]Appendix A includes a plot showing industrialization paths for all of the remaining countries in our sample.

[10]We include South Korea since it is the advanced economy that has data coverage for the lowest values of h_{nt}.

Figure 4: Paths of Industrialization: Four Asian Economies

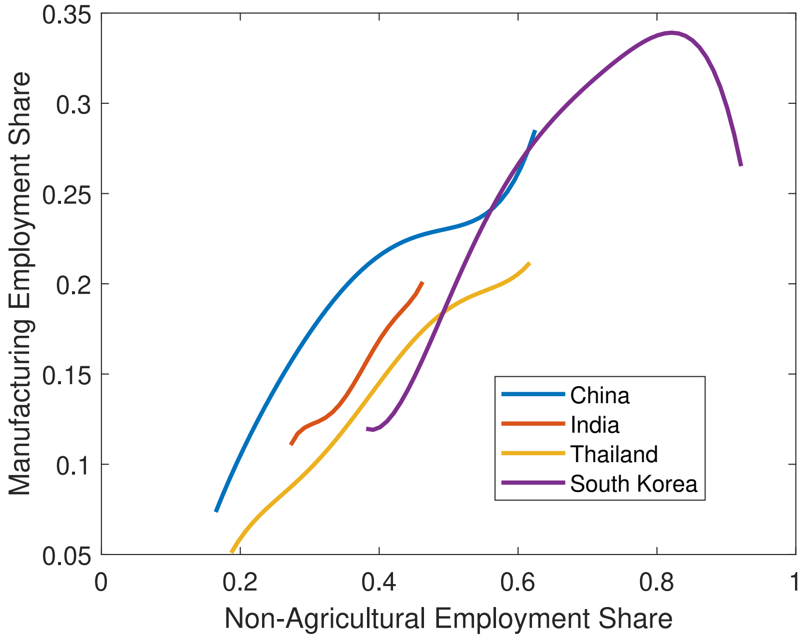

to be exhibiting behavior that is similar to that of South Korea. Comparing with Figure 1, it would appear that Thailand looks to be on a path that is similar to that of Brazil, whereas it is quite difficult to say much about India given its current stage. By way of summary, it appears that the experiences of these three countries will ultimately exhibit a fair bit of heterogeneity.

When summarizing the patterns in Figure 1 we noted a strong positive correlation between the value of h_{nt} at which the peak value of h_{mt} occurred and the value of peak h_{mt}. Denote these two values by h_n^* and h_m^* respectively. We now pursue this pattern further using the full sample of 18 countries who have attained their peak. Table 1 shows the values of h_n^* and h_m^* for each of the 18 countries.

Table 1

Values of h_n^* and h_m^*

Asia			Latin America			Europe		
	h_n^*	h_m^*		h_n^*	h_m^*		h_n^*	h_m^*
IDN	0.56	0.19	ARG	0.82	0.35	FRA	0.86	0.38
JAP	0.88	0.33	BOL	0.78	0.25	ITA	0.82	0.39
KOR	0.82	0.34	BRA	0.65	0.23	SPA	0.77	0.35
MAL	0.83	0.34	CHL	0.73	0.31			
PHL	0.59	0.17	COL	0.81	0.20			
TWN	0.82	0.40	CRI	0.75	0.27			
			MEX	0.79	0.28			
			PER	0.56	0.20			
			Ven	0.83	0.27			

Figure 5 shows a scatterplot for the pairs of h_n^* and h_m^* across countries as well as a fitted linear regression line.

The positive correlation is evident in the picture, and is equal to 0.76. In the remainder of this paper we will try to shed some light on factors that can give rise to this pattern, both qualitatively and quantitatively.

3 A Model of Industrialization

In this section we introduce a simple benchmark model of structural change. While simple, the model captures the key forces that the literature has emphasized as the drivers of structural change. The model's tractability allows us to analytically characterize the industrialization phase.

Figure 5: Correlation of Peak and Timing of Industrialization

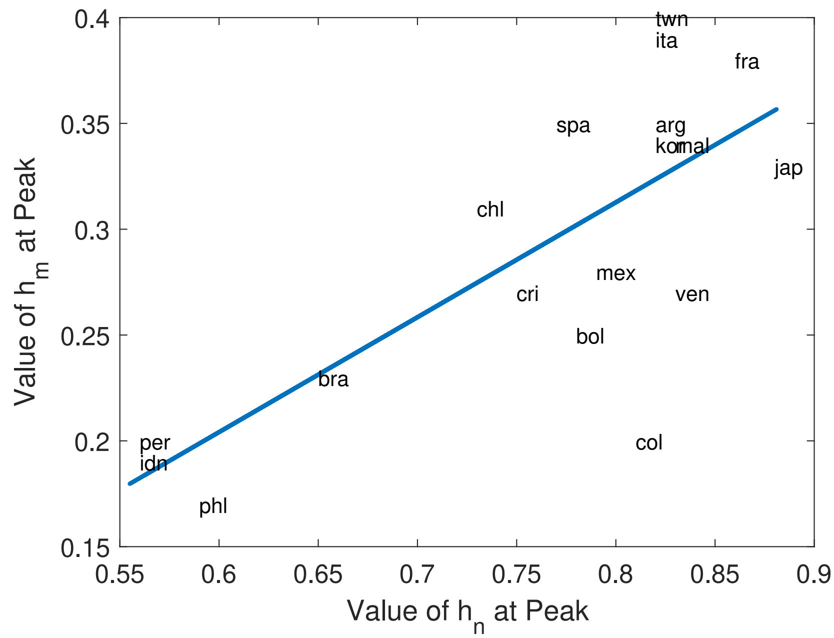

3.1 Model

We consider a simple benchmark model of structural transformation formulated in continuous time. The allocation decisions in the model are all static, with all dynamics generated by exogenous technological change over time. In this subsection we focus on the decisions made at a particular point in time and so suppress time subscripts.

There are three consumption goods in the economy: agriculture, manufacturing and services. Each of the consumption goods is produced using a linear production function with labor as the only input:

$$c_i = A_i h_i,\ i = a, m, s$$

There is a representative household that is endowed with one unit of time and has preferences over the three consumption goods. Preferences reflect a combination of those used by Gollin et al. (2002) and Boppart (2014). In particular, following Gollin et al. (2002) we assume an extreme form of a subsistence constraint in terms of agricultural consumption: households receive no utility from consumption of manufacturing or services if $c_a < \bar{c}_a$, utility from consuming c_a is increasing until c_a reaches $\bar{c}_a$, but they receive no additional utility from consuming additional food beyond $\bar{c}_a$. While somewhat extreme, this assumption serves to generate tractability and facilitate transparency. Moreover, as we show later on, its key prediction for agricultural employment tracks the data quite well for our sample of countries.[11]

Conditional on c_a being at least equal to $\bar{c}_a$ the household will also receive utility $u(c_m, c_s)$ from consuming c_m and c_s. The literature on structural change emphasizes two forces that affect the allocation of expenditure between manufacturing and services and thereby shape structural change between the two sectors: income effects and relative price effects. Boppart (2014) shows that price inelastic generalized linear (PIGL) preferences offer a tractable specification in which these effects operate smoothly along a development path, and we follow him. These preferences have an analytic representation for the indirect utility function but not for the direct utility function $u(c_m, c_s)$. For this reason we will work with the indirect utility function for the utility that derives from consumption of c_m and c_s. Letting E be total expenditure on manufacturing and services, and p_m and p_s be the prices of manufacturing and services, we will assume that indirect utility function $v(E, p_m, p_s)$ is given by:

$$v(E, p_m, p_s) = \frac{1}{\chi}\left(\frac{E}{p_s}\right)^{\chi} - \frac{\alpha}{\varepsilon}\left(\frac{p_m}{p_s}\right)^{\varepsilon} - \frac{1}{\chi} + \frac{\alpha}{\varepsilon},$$

where $\alpha > 0$ and $0 < \chi < \varepsilon < 1$. Boppart (2014) shows that this indirect utility function is well defined given these parameter restrictions.[12]

[11]This finding is also in Üngör (2013).

[12]The parameter χ controls income effects. In the limit as χ goes to zero preferences are homothetic. The parameter ε controls the elasticity of substitution between goods and services, but it should be noted that the elasticity is not equal to ε and in fact is not even constant. In the limiting case of $\chi = 0$, ε tending to one corresponds to Leontief and ε tending to zero corresponds to Cobb-Douglas. See Boppart (2014) for more details.

We focus on the competitive equilibrium allocation for the above economy.[13] We normalize the wage rate to equal unity and let the three prices be denoted by p_i, $i = a, m, s$. Given the linear production functions it follows that in equilibrium we must have:

$$p_i = \frac{1}{A_i},\ i = a, m, s \tag{3.1}$$

It remains to determine the allocation of labor. Because we normalize the wage to unity, total income for the household will also equal unity. The allocation of labor will then be dictated by the demands of the representative household for the three consumption goods given equilibrium prices and total income.

If $A_a \leq \bar{c}_a$ then the household will allocate all of its income to purchasing the agricultural good and the equilibrium allocation of labor is $h_a = 1$ and $h_m = h_s = 0$. In what follows we will focus on the case in which $A_a > \bar{c}_a$. In this case the household will purchase $\bar{c}_a$ units of c_a, implying that the equilibrium allocation of labor to agriculture allocation will be:

$$h_a = \bar{c}_a / A_a \tag{3.2}$$

To solve for the allocation of non-agricultural labor between manufacturing and services we need to consider the optimal allocation of non-agricultural expenditure between manufacturing and services. We begin by using Roy's Identity to uncover the expenditure share for c_m as:

$$\frac{p_m c_m}{E} = \alpha \left(\frac{E}{p_s}\right)^{-\chi} \left(\frac{p_m}{p_s}\right)^{\varepsilon}. \tag{3.3}$$

Note that the share of non-agricultural expenditure devoted to manufacturing depends upon both total expenditure and the relative price of manufacturing to services. Both

[13] The equilibrium allocation will also be the unique Pareto efficient allocation. But because we have specified preferences with an indirect utility function we will solve for allocations by directly considering the equilibrium.

effects display constant elasticities, with χ governing the income effect and ε governing the relative price effect.

Recalling the expressions for equilibrium prices, market clearing implies:

$$p_i c_i = h_i, \quad i = a, m, s \tag{3.4}$$

Since $E = 1 - p_a c_a$, equation (3.4) implies:

$$E = 1 - h_a \tag{3.5}$$

Substituting equations (3.1), (3.4), and (3.5) into equation (3.3) and using $h_n = 1 - h_a$ gives:

$$f_m = \frac{h_m}{h_n} = \alpha \left(h_n A_s\right)^{-\chi} \left(\frac{A_s}{A_m}\right)^{\varepsilon} \tag{3.6}$$

Having solved for the equilibrium values of h_a and h_m, the equilibrium value of h_s is then determined as $1 - h_a - h_m$.

3.2 Hump-Shaped Industrialization Dynamics

We now use the previous expressions to characterize the evolution of sectoral employment shares along a development path when technical change is the sole driving force. To maximize transparency we focus on the case in which there is constant technological progress in each of the three sectors, though possibly at different rates:

$$A_{it} = e^{g_i t}$$

where $g_i > 0$ for $\{i = a, m, s\}$ and we have implicitly normalized all three initial productivities at time zero to unity. It is useful to focus on the empirically relevant part of parameter space, and so consistent with empirical evidence we assume that $\varepsilon < 1$ and that $g = g_m - g_s > 0$.[14]

Two simple properties follow. First, $g_a > 0$ implies that h_{at} will decrease over time. It follows that h_{nt} will be increasing over time. Combined with $g > 0$ and $\varepsilon < 1$, Equation (3.6) then implies that h_{mt}/h_{nt} will be decreasing over time. Since non-agricultural labor is increasing and a decreasing share of it is devoted to manufacturing, it follows that h_{st} will be increasing over time. Importantly, these two properties–a monotonic decline in h_{at} and a monotonic increase in h_{st}–are robust features of the development process.

Next we turn to the evolution of the manufacturing employment share. It is useful to start with the following identity:

$$h_{mt} = h_{nt} \cdot \frac{h_{mt}}{h_{nt}} = h_{nt} \cdot f_{mt}$$

where as before, $h_{nt} = 1 - h_{at}$ and f_{mt} is defined by Equation (3.6).

It follows that:

$$\frac{\dot{h}_{mt}}{h_{mt}} = \frac{\dot{h}_{nt}}{h_{nt}} + \frac{\dot{f}_{mt}}{f_{mt}} \tag{3.7}$$

Equation (3.7) shows the two opposing forces that shape the dynamics of h_{mt}. On the one hand, increases in A_{at} lead to movement of labor out of agriculture, so that the first term on the right hand side is always positive. But, growth in h_{nt} coupled with higher growth in A_{mt} relative to A_{st} leads to a decreasing share of non-agricultural labor in the manufacturing sector, so that the second term is always negative.

Our assumptions on preferences and technology allow us to characterize what happens to the relative magnitude of these effects as technology advances, as summarized in the following proposition.

Proposition 1: Assume $g_i > 0$ for $i = \{a, m, s\}$, $0 < \chi < \varepsilon < 1$ and $g = g_m - g_s > 0$.

[14] See, for example the summary in Herrendorf et al. (2014).

Then,

(i) $\frac{\dot{h}_{nt}}{h_{nt}} > 0$, and decreases monotonically to 0.

(ii) $\frac{\dot{f}_{mt}}{f_{mt}} < 0$, and increases monotonically to $-\chi g_s - \varepsilon g$.

(iii) $\frac{\dot{h}_{mt}}{h_{mt}}$ decreases monotonically and converges to $-\chi g_s - \varepsilon g$.

Proof:

To prove (i) recall that $h_{at} = \bar{c}_a / A_{at}$, so that:

$$\frac{\dot{h}_{nt}}{h_{nt}} = -\frac{\dot{h}_{at}}{1 - h_{at}} = g_a \frac{\bar{c}_a e^{-g_a t}}{1 - \bar{c}_a e^{-g_a t}}$$

The result follows from the fact that the numerator decreases monotonically to 0, while the denominator increases monotonically to 1.

To prove (ii), start with:

$$f_{mt} = \frac{h_{mt}}{h_{nt}} = \alpha \left(h_{nt} A_{st}\right)^{-\chi} \left(\frac{A_{st}}{A_{mt}}\right)^{\varepsilon}.$$

Straightforward calculation implies:

$$\frac{\dot{f}_{mt}}{f_{mt}} = -\chi \frac{\dot{h}_{nt}}{h_{nt}} - \chi g_s - \varepsilon g$$

The second and third terms on the right are both constant and negative, so the result follows from the fact that $\frac{\dot{h}_{nt}}{h_{nt}}$ is positive but decreases monotonically to zero.

To prove (iii), combine the results from (i) and (ii) to get:

$$\frac{\dot{h}_{mt}}{h_{mt}} = (1 - \chi) \frac{\dot{h}_{nt}}{h_{nt}} - \chi g_s - \varepsilon g$$

The result follows from the fact that the first term is monotonically decreasing to zero (recall that $0 < \chi < 1$).■

It follows from Proposition 1 that a sufficient condition for h_{mt} to exhibit hump-shaped dynamics is for $\frac{\dot{h}_{m0}}{h_{m0}}$ to be positive. The following result is essentially a corollary of Proposition 1 and provides a necessary and sufficient condition for this to hold.

Corollary 1: Assume $g_i > 0$ for $i = \{a, m, s\}$, $0 < \chi < \varepsilon < 1$ and $g > 0$. Then $\frac{\dot{h}_{mt}}{h_{mt}} > 0$ if and only if:

$$(1-\chi)g_a \frac{h_{at}}{1-h_{at}} > \chi g_s + \varepsilon g = (\chi - \varepsilon)g_s + \varepsilon g_m \tag{3.8}$$

Proof: Follows immediately from substituting into equation (3.7), noting that $h_{at} = \bar{c}_a e^{-g_a t}$ and using the expressions derived in Proposition 1. ■

Because the inequality in equation (3.8) will always hold as h_{at} tends to one, it follows that the model will always generate a hump-shaped path for h_{mt}. A more interesting issue concerns the point at which the peak occurs. To pursue this, define h_a^* as the value of h_{at} at which equation (3.8) holds with equality:

$$(1-\chi)g_a \frac{h_a^*}{1-h_a^*} = \chi g_s + \varepsilon g = (\chi - \varepsilon)g_s + \varepsilon g_m \tag{3.9}$$

As in the previous section, define $h_n^* = 1 - h_a^*$ as the value of h_{nt} at the peak and h_m^* to be the value of h_{mt} at the peak. We will be interested in how these values are affected by changes in the sectoral productivity growth rates. Accordingly, we can write these values as functions of g_a, g_s and g_m, i.e., $h_n^*(g_a, g_s, g_m)$ and $h_m^*(g_a, g_s, g_m)$. Several comparative statics results of interest follow and are summarized in the next Proposition.

Proposition 2: Assume $g_i > 0$ for $i = \{a, m, s\}$, $0 < \chi < \varepsilon < 1$ and $g = g_m - g_s > 0$. Then,

(i) $h_n^*(g_a, g_s, g_m)$ is increasing in g_a and g_s and decreasing in g_m.

(ii) $h_m^*(g_a, g_s, g_m)$ is increasing in g_a and g_s and decreasing in g_m.

Proof: Part (i) follows immediately from using equation (3.9) to solve for $h_n^* = 1 - h_a^*$:

$$h_n^* = \frac{1}{1 + \frac{(\chi-\varepsilon)g_s+\varepsilon g_m}{(1-\chi)g_a}} \tag{3.10}$$

To show part (ii) we start with the case of an increase in g_a. To show that h_m^* increases it suffices to show that there is some value of t at which the new series for h_{mt} exceeds the original value of h_m^*. To show this, let t' be the value of t at which the new series for h_{nt} equals the original value of h_n^*. Because g_a has increased it follows that if t^* is the time at which the peak occurs for the original value of g_a, then $t' < t^*$. If f_m^* is the value of f_{mt} at the original value of t^*, then $g > 0$ implies that the new value of f_{mt} at t' is greater than f_m^*. It follows that the new value of h_{mt} at t' exceeds the original value of h_m^*.

The other two results in part (ii) are much simpler to derive. The result for g_s follows from the fact that the series for h_{nt} is unaffected, and the series for f_{mt} is everywhere higher. Similarly, the result for g_m follows from the fact that the series for h_{nt} is again unaffected, and the series for f_{mt} is now everywhere lower. ■

Importantly, this proposition shows that variation in sectoral productivity growth rates can qualitatively generate the empirical pattern shown previously in Figure 5. These results are also intuitive: a higher value for g_a serves to increase the flow of workers into manufacturing, while higher values of g_s and lower values of g_m serve to decrease the flow of workers out of manufacturing. In the remainder of the paper we examine the quantitative significance of these effects.

Before proceeding we want to emphasize an important property of the model regarding its implications for the industrialization path plotted in $h_n - h_m$ space. Specifically, this profile is determined by the profile of relative sectoral productivities that the economy experiences and not by the pace at which the economy moves along this profile. More formally, let τ be an indicator for the level of development, and assume that the relationship between sectoral productivities and development is given by $A_{j\tau} = e^{g_i \tau}$. Fixing these profiles, let $\tau(t)$ be a function that describes how quickly a country moves along the development path. The key feature of our model is that the industrialization path in $h_n - h_m$ space is invariant to the function $\tau(t)$.

With this result in mind, consider the comparative static result concerning an increase in the value of g_a. This serves to both increase the pace of overall development and change the profile of relative productivities along the development path. But in view of the previous discussion, the implications for h_n^* and h_m^* are invariant to the pace of development, so the comparative static result should be understood as highlighting changes in the profiles of relative productivity rather than changes in the pace of development.

4 Benchmark Calibration

In this section we present a benchmark calibration of the above model that captures the trend evolution of sectoral employment shares in the US economy during its industrialization period, which we take as 1880-1950. Because our application will focus on the industrialization phase for current developing economies, we want our calibrated model to reflect this phase for the US economy. The changing composition within services might reasonably lead to secular changes in the properties of preferences defined over highly aggregated sectors, and we want our preference parameters to be relevant for the industrialization phase of development. Additionally, the fact that services is increasingly dominated by low productivity growth sectors like education and health care suggests that the gap between manufacturing and services productivity growth is plausibly increasing over time, especially in the post 1970 period.[15]

As noted earlier, our data for the evolution of US sectoral employment shares between

[15]See, for example, Duernecker et al. (2017) for an analysis that disaggregates the service sector.

Figure 6: US Sectoral Employment Shares 1880-1980

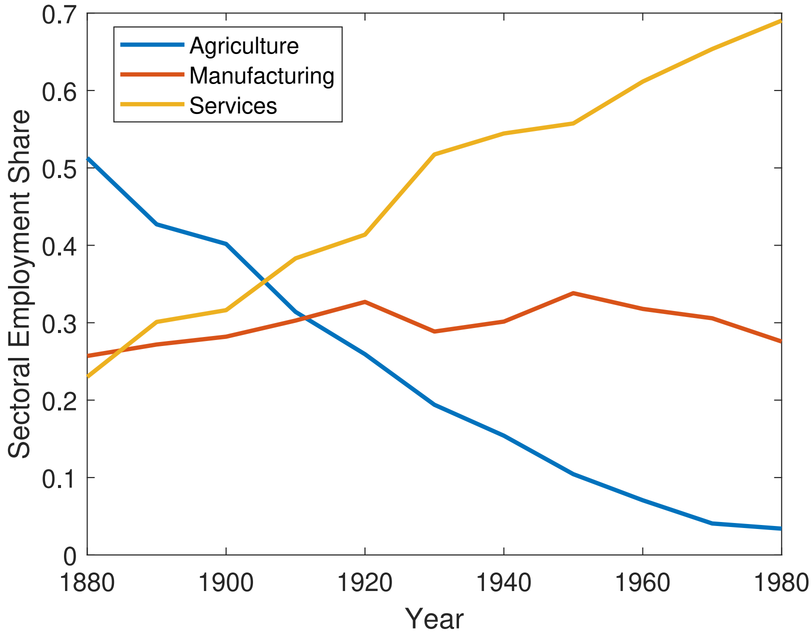

1880 and 1950 comes from Carter et al. (2006) and the BEA. Figure 6 shows the time series for these shares between 1880 and 1980.

The figure displays the monotonic decline in the agricultural employment share as well as the monotonic increase in the services employment share. The trend behavior of the manufacturing employment share reflects a hump-shaped pattern, but as noted earlier, the disrupting effect of the Great Depression on the evolution of the manufacturing employment share is readily apparent. Our calibration procedure will implicitly reflect the evolution that would have occurred if the Great Depression had not occurred.

As of 1880 the agricultural employment share in the US is approximately 0.50. For considering the development paths of current developing countries we will be interested in considering even higher values of h_{at} than witnessed during this period for the US. For this reason our benchmark calibration will consider an economy that begins with $h_a = 0.60$, though parameters will be set so that the model matches the evolution of the US economy as h_{at} decreases from 0.50 to slightly less than 0.10.

Next we assign values for the three (constant) growth rates of sectoral labor productivity. We do not have sectoral productivity data that covers that period from 1880-1950. The growth rate for agricultural productivity, g_a, is set by requiring the model to achieve the observed decrease in h_{at} for the US economy between 1880 and 1950. This implies that $g_a = 1.0239$. For the other two growth rates we use data from the GGDC Ten Sector Data Base for the period 1950-1970 and assume that these are indicative of average productivity growth rates in the preceding 70 year period. While somewhat heroic, this assumption seems somewhat reasonable given the relative constancy of trend aggregate growth over this period. This implies $g_m = 1.0225$ and $g_s = 1.0147$. We normalize all three productivity levels in the initial period to equal unity. Because we start the economy with a fraction 0.60 of employment in the agricultural sector, this implies that $\bar{c}_a = 0.60$.

It remains to pick values for the three remaining preference parameters that define the indirect utility from consumption of manufacturing and services: α, χ, and ε. Although both Boppart (2014) and Herrendorf et al. (2020) estimate parameters for the PIGL specification, we note that neither of them is appropriate for our setting. In both cases they assumed that goods included both manufacturing and agriculture. Importantly, our specification has a very strong income effect for agricultural consumption, whereas in their settings this effect is reflected in the income effect for overall spending on goods. Additionally, whereas Boppart (2014) considered preferences over final expenditure, our preferences should be interpreted as being over value added components of consumption, as in Herrendorf et al. (2020).[16]

Our strategy is to pick values for these three parameters so as to match the industrialization profile for the US as shown in Figure 3. Table 4 displays the calibrated values and Figure 7 shows the fit of the model for the profile of h_{mt} versus h_{nt} as shown in Figure 3.

[16]See also Herrendorf et al. (2014) for additional discussion of this issue and its significance for estimating preference parameters.

Figure 7: Model Fit to US Data

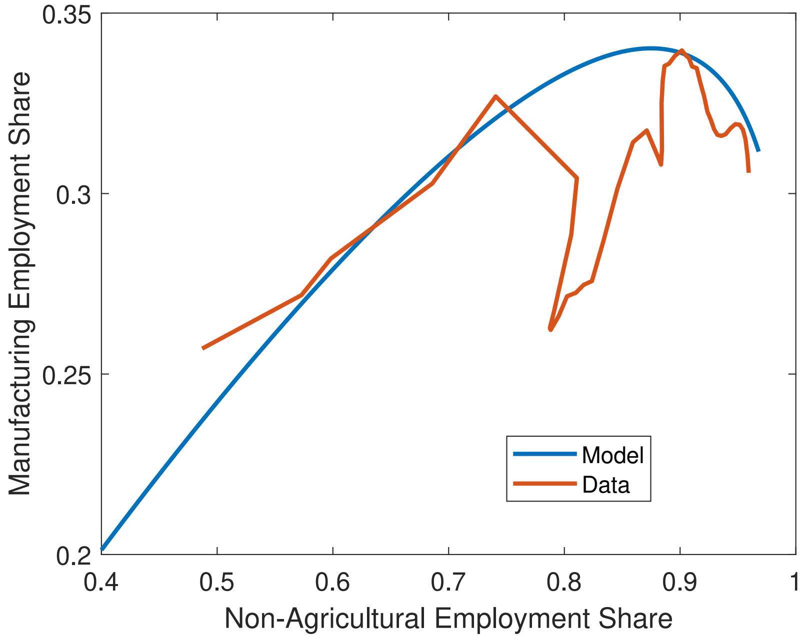

Table 2

Benchmark Calibration

g_a	g_m	g_s	$\bar{c}_a$	ε	χ	α_m
1.0239	1.0225	1.0147	.60	0.30	0.06	0.4762

Recall that our calibrated model assumes that all three productivity growth rates are constant, so that to the extent that trend productivity growth in the US varied over the period 1880-1950, we will see departures of the data from the paths implied by the model. Our calibrated model implicitly assumes that the Great Depression and WWII did not occur.

We make two remarks about the calibrated preference parameters. First, note that the calibrated value of χ is quite small, so that income effects are relatively unimportant for the allocation of expenditure between manufacturing and services. In particular, our value is very small compared to the value estimated in Herrendorf et al. (2020). But as noted earlier, it is critical to emphasize that their estimate implicitly included the income effect operating via agricultural consumption. Additionally, the increasing importance of health and education may result in a drift in this parameter over time.

Second, and related, because we calibrate a relatively low value of χ, our calibrated model is quite close to a model that assumes a homothetic CES specification for preferences over consumption of manufacturing and services. If we adopted this specification then the calibrated elasticity parameter would be close to zero, indicating that preferences over manufacturing and services are approximately Leontief. We note that the results that we present in the next section would be virtually unchanged if we had instead worked with this somewhat simpler specification.

Related to this last point, our calibration procedure does not provide much guidance on disentangling the role of income and substitution effects in the neighborhood of our calibrated values, in that small changes in one can be offset by small changes in the other. However, the value of χ cannot be much larger than 0.10 without a large sacrifice in fit. But consistent with what was previously stated, varying χ within this range has little effect on the results that we present in the next section.

5 Alternative Industrialization Paths

All of the Asian and Latin American economies in our sample lagged behind the US as of 1950, but most of them experienced some degree of catch-up since that time. It is natural to view the calibrated sectoral productivity profile for the US as representing the time paths of the sectoral technology frontiers. An appealing property of our representation of industrialization in $h_m - h_n$ space is that if a late developing country follows the same sectoral productivity profile as the US, though possibly at an accelerated speed, it will produce exactly the same industrialization path as the US.

However, there is no reason that the process of catch-up to the frontier for late developers will necessarily mimic the historical evolution of the frontier; that is, an individual country may converge towards the frontier technology at different rates across sectors. In this case the evolution of its productivity profile may differ from the one experienced historically in the US. In this section we use our calibrated model to learn about the extent to which alternative sectoral productivity profiles can generate the range of industrialization experiences depicted in Figure 5.

Following our theoretical analysis earlier, we will focus on two departures that can give rise to the pattern found in Figure 5. The first is slower growth in agricultural productivity. A large literature has emphasized the relatively large differences in agricultural productivity between rich and poor countries and hence the apparent slow rate of catch-up of agricultural productivity in these countries.[17] Building on the earlier work of Johnston and Mellor (1961) and Johnston and Kilby (1975), Gollin et al. (2002, 2007) emphasize that slow productivity growth in agriculture can delay overall development of the non-agricultural sector. Building on this work, we show that relatively slower catch-up in agricultural productivity can also affect the path of industrialization in a way that quantitatively mimics the findings presented in Section 2.

The second departure is to consider slower productivity growth in services.[18] While this departure can also generate significant differences in the peak employment share for the manufacturing sector, we find that this departure is less able to generate the large differences found in the data. The evidence for this departure is somewhat less strong. Whereas Rodrik (2012) argued that productivity gaps in manufacturing are small suggesting that differences in services must be large, both Hsieh and Klenow (2007) and Herrendorf and Valentinyi (2012) found that differences in manufacturing were relatively large.

[17] See, for example, Restuccia et al. (2008), Caselli (2005) and Gollin et al. (2013).

[18] Because our calibrated model displays preferences over manufacturing and services which are nearly homothetic, the effects of g_s and g_m on the industrialization path are close to mirror images of each other. For this reason we do not report separate results for variation in g_m.

Figure 8: Agricultural Productivity Growth and Peak Industrialization

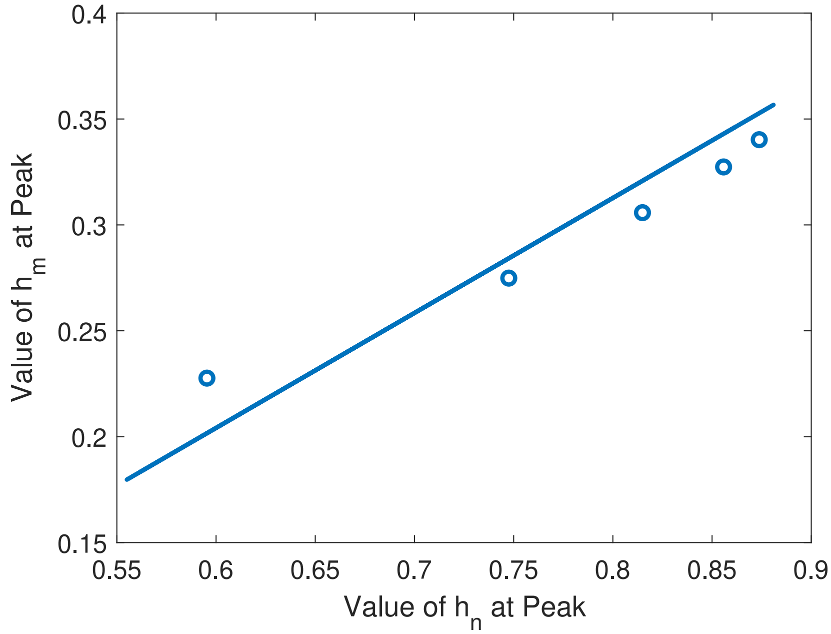

5.1 Slow Catch-up in Agricultural Productivity

To pursue this we consider a set of economies that differ from our benchmark calibrated economy solely in terms of their productivity growth rate in agriculture. Our benchmark economy featured $g_a = 1.0239$. Here we consider four economies with values of g_a that vary from 1.005 to 1.0200 in increments of 0.005. For each economy we simulate outcomes beginning with the same initial conditions as in our benchmark model. In Section 3 we showed that a decrease in g_a will decrease both h_n^* and h_m^*. Our goal here is to examine the quantitative implications of this decrease.

Figure 8 shows the scatter plot that corresponds to the aforementioned counterfactuals, along with the regression line from the scatter plot of Figure 5.

Our calibrated economy lies a little below the regression line from the data, but the model-generated data track the regression line from the data remarkably well. We conclude that differences in agricultural productivity growth are capable of generating differences in industrialization experiences similar to those found in the data.

Figure 9: Agricultural Productivity Growth and Paths of Industrialization

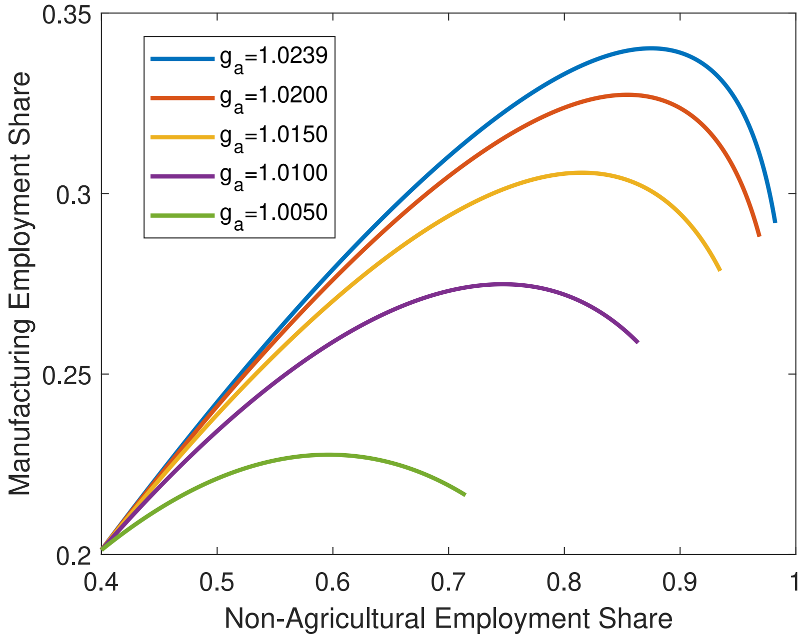

Figure 9 shows the h_{mt} versus h_{nt} profiles for the different values of g_a.

In each case we run the economy forward for 150 years. Note that because the economies differ in their value of g_a they achieve different levels of h_{nt} during the 150 years.[19]

5.2 Slow Catch-up in Services Productivity

In this subsection we repeat the previous exercise but this time considering the possibility of slower catch-up in services.[20] Recall that our calibrated value of g_s was 1.0147. Here we consider five alternative values ranging from 1.0025 to 1.0125 in increments of 0.0025. The results are shown in Figure 10.

We highlight three features of this figure. First, we again see that the points closely

[19]We emphasize again that the resulting path of industrialization as represented in $h_m - h_n$ space is invariant to the speed with which the country travels along the given productivity profile; what matters for this plot is the relative sequence of sectoral productivities and not the overall rate at which a country moves along the profile.

[20]As noted earlier, this exercise is almost identical to considering a higher growth rate for manufacturing productivity.

Figure 10: Services Productivity Growth and Peak Industrialization

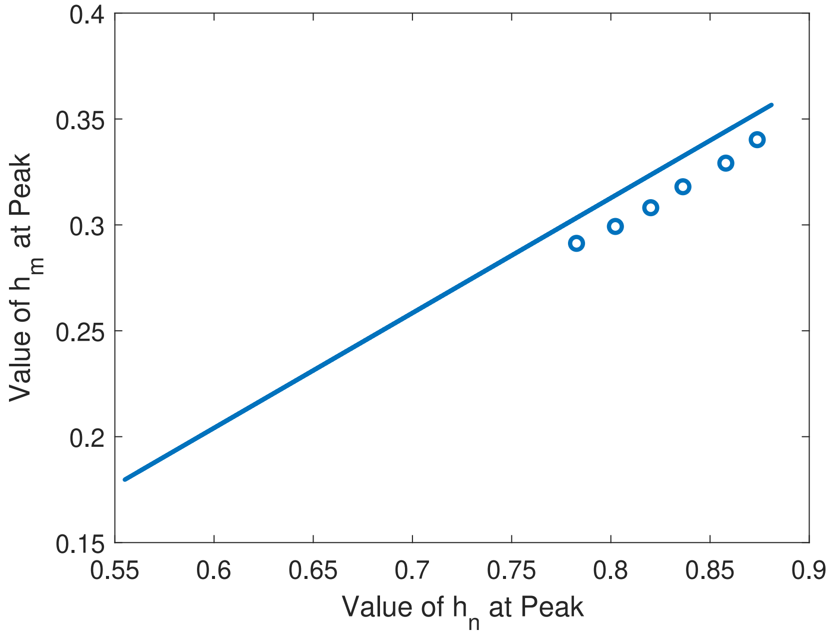

track the regression line from Figure 5. Second, the figure indicates that large differences in g_s can affect the peak manufacturing employment share by as much as five percentage points. And third, comparing Figures 8 and 10, there is somewhat greater scope for differences in g_a to affect these values.

6 Rationalizing the Data

In the previous section we showed that seemingly reasonable cross-country differences in the rate at which sectoral productivities move toward the frontier can generate differences in industrialization paths that mimic those found in the data. In this section we use the model to infer the sectoral productivity profiles that would be required to rationalize the data for each of the 21 countries in our sample and then compare these productivity profiles to those observed in the GGDC Ten Sector Data Base.[21]

[21] An obvious alternative to this two-step procedure would be to simulate the model using empirical productivity profiles. If productivities are measured with error we think our two step procedure is preferable. As we discuss later, there is good reason to think that there is substantial measurement error

6.1 Inferring Productivity Profiles

The first step in our exercise is to use our model calibrated to the US industrialization experience to infer sectoral productivity profiles for each of the countries in our sample using data on employment shares. To do this we assume that preferences are the same across countries, and that the sole source of differences across countries are the time series profiles for sectoral productivities.

Before proceeding it is important to note that there is a basic invertibility issue that one must confront when using the model to recover productivities with data on sectoral employment shares. Specifically, there are three sectoral productivities to determine in each period, but since sectoral employment shares sum to one, employment shares provide only two moments at each point in time. In particular, our model implies the following mapping from the productivity profile (A_a, A_m, A_s) into the employment shares h_a and h_m:

$$h_a = \frac{\bar{c}_a}{A_a} \tag{6.1}$$

$$h_m = \alpha(1-h_a)^{1-\chi} A_s{}^{-\chi} \left(\frac{A_s}{A_m}\right)^{\varepsilon} \tag{6.2}$$

The time series for A_a is uniquely determined by the time series for the employment share h_a. But the second equation shows that the evolution of h_m depends both on the evolution of relative productivity A_s/A_m and the overall level of productivity in the non-agricultural sector.

Our strategy will be to use the data to pin down the level effect and then use the model to infer the profile for the relative productivity of manufacturing and services. That is, we will take the growth rate of A_s as given in the data and use the model to infer series

in the relevant productivity growth rates.

Figure 11: Agricultural Productivity Growth: Model and Data

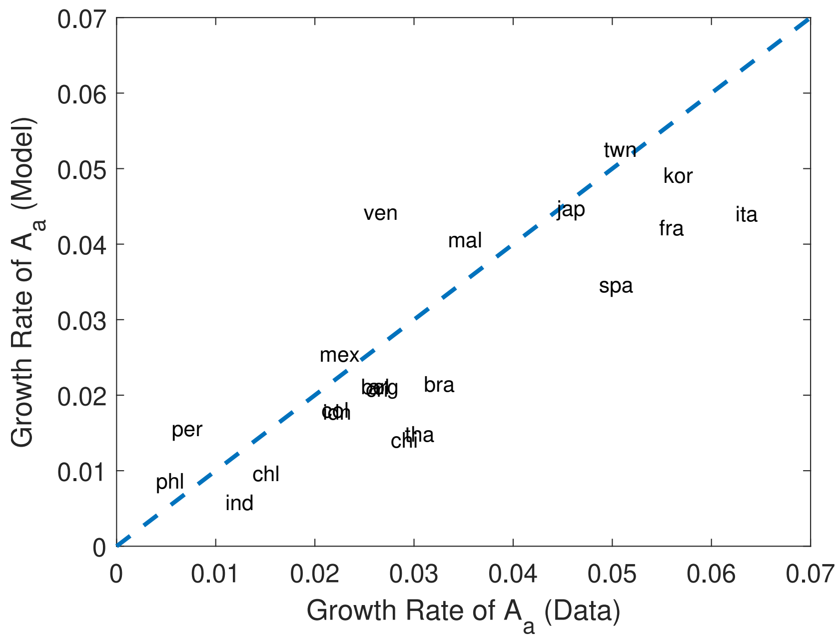

for the growth rate of A_a and the growth rate of (A_s/A_m).[22]

6.2 Results

We carry out the above procedure for each of the 21 countries in our sample up to the point at which they reach their peak employment share in manufacturing.[23] We focus on this period because we calibrated our model to the industrialization phase.

We are particularly interested in the relationship between the productivity growth processes in the data and those implied by this procedure. We begin with the results for growth in agricultural productivity. Figure 11 shows a scatterplot for the values from the data and those inferred from our model based exercise, as well as a 45 degree line.

The figure shows that there is a strong positive correlation between the two, and that

[22] We note that if preferences over manufacturing and services were homothetic then only the ratio of the two non-agricultural productivities would matter for the determination of sectoral employment shares and the invertibility issue would not arise.

[23] We follow each country for a minimum of 25 years to ensure a sufficiently long sample for estimating average productivity growth rates. This only affects a small number of countries.

Figure 12: Growth in Manufacturing-Services Relative Productivity: Model and Data

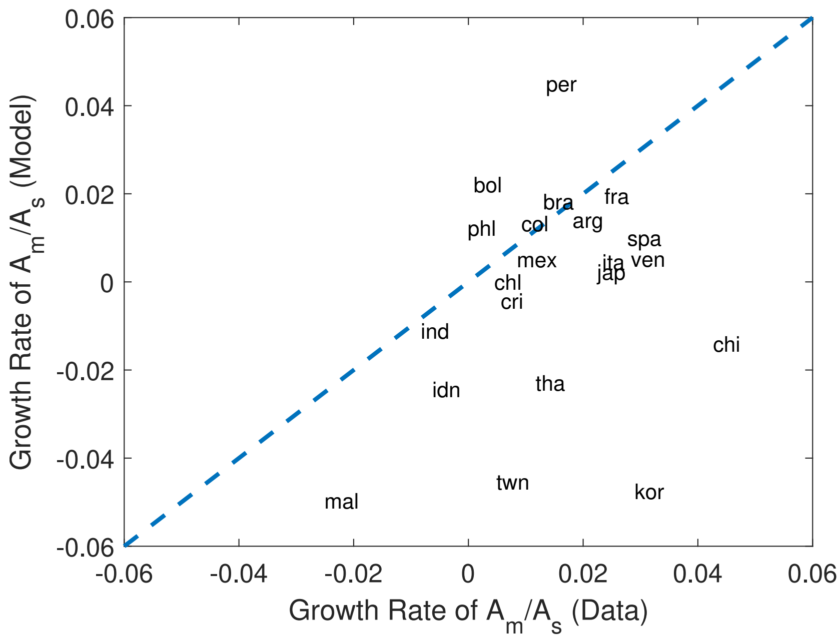

the points tend to track the 45 degree line. The correlation between the two values is 0.83. Being mindful of the fact that there are various issues that result in classical as well as non-classical measurement error, we view this as a very strong correlation. We conclude that the model's relatively stark predictions about the relationship between productivity and employment in the agricultural sector are largely supported by the data.

Next we turn to the results for the growth rate of manufacturing productivity relative to services. Figure 12 shows a scatterplot as well as a 45 degree line.

While many of the points do track the 45 degree line, there are a number of significant outliers relative to the 45 degree line and the overall correlation is only 0.25. The two points in the bottom right corner are China and South Korea. We will revisit these cases in the next section when we discuss extensions, but for now we note if we exclude these two countries the correlation increases quite substantially to 0.56.

Once again, it is important to recognize the potential for measurement error in the productivity series. Importantly, the plot in Figure 12 is comparing the *difference* between two productivity growth rates. Assuming each growth rate is independently measured with error, the difference between the two growth rates will display much greater error. In Appendix B we carry out one exercise to explicitly address the measurement error issue and conclude that measurement of the relative growth rate of manufacturing and services productivity should reasonably be viewed as having a significant amount of error. In view of this the correlation of 0.56 for the subsample without China and South Korea should perhaps be viewed as quite supportive for the model.

6.3 Agricultural Productivity and Industrialization

The previous analysis showed that the benchmark model does a very good job of accounting for the movement of labor out of agriculture in the sense that the model implied values for agricultural productivity growth are closely related to measured values from the GGDC 10-Sector database. The results were a bit more mixed regarding the model's ability to account for the division of non-agricultural labor into manufacturing and services. In this section we show that differences in agricultural productivity profiles play a dominant role in accounting for the observed differences in peak manufacturing employment shares.

To do this we carry out the following exercise for each of the 18 countries in our sample that reach a peak employment share in manufacturing. First, we assume that initial productivity levels for each country are such that model implied employment shares perfectly match the observed employment shares in the first period for which data is available. Second, for each country we take productivity growth in agriculture as measured from the GGDC. Specifically, for each country we assume that productivity growth in agriculture is constant, equal to its average value during the industrialization phase. Third, we assume that the growth rates of productivity in both manufacturing and services are those that

Figure 13: Agricultural Productivity Growth and Peak Manufacturing Employment: Model and Data

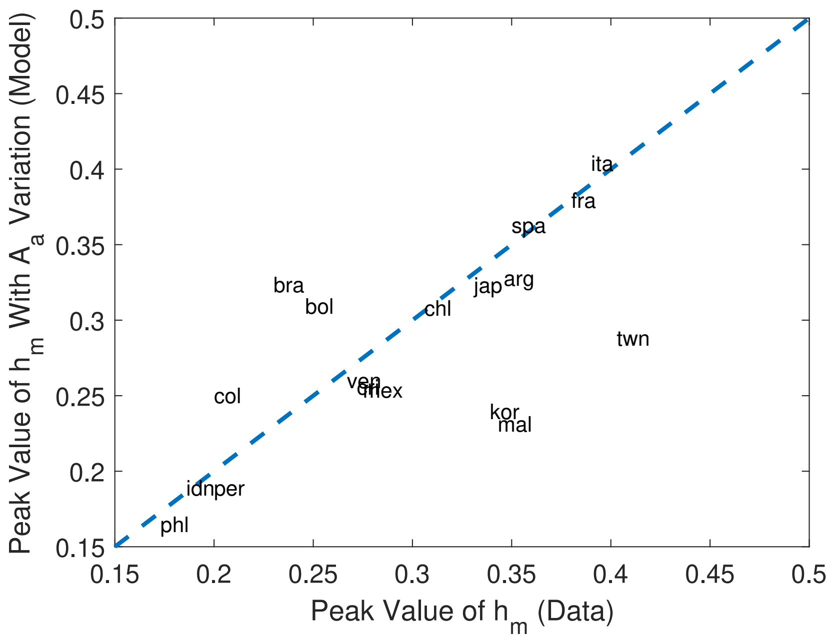

we calibrated for the US economy.[24] Note that this exercise differs from the counterfactuals reported in Section 5 in which we varied g_a because the current exercise assumes differences in both initial conditions and the growth rate of agricultural productivity.

We simulate data for each of the 18 economies and find the peak employment share for the manufacturing sector. Figure 13 plots the values from this exercise against the values reported in Section 2 as well as a 45 degree line to facilitate comparison.

The figure shows that the specification in which observed differences in agricultural productivity growth are the only source of difference across countries does an excellent job of accounting for the observed variation in peak levels of h_m for most of the countries in our sample. Twelve of the countries lie very close to the 45 degree line, indicating that the differences in agricultural productivity are essentially sufficient to account for the large differences in peak values for h_m.

[24]Because our preferences over manufacturing and services are not that far from being homothetic, the key assumption in this exercise is that the gap between these two productivity growth rates is the same as in the US.

There are six countries for which the gap between the model predicted value and the actual value exceeds 5 percentage points. Three of these countries lie below the 45 degree line and are all from Asia: South Korea, Malaysia and Taiwan. For these countries the model requires significant differences in the growth rate of A_m/A_s relative to the US to replicate the evolution of employment shares. It is noteworthy that this is not the case for Japan. The other three countries lie above the 45 degree line and are all from Latin America: Bolivia, Brazil and Colombia.

Importantly, Figure 13 shows that differences in the evolution of agricultural productivity alone can account for gaps in peak manufacturing employment shares that range from less than 0.20 to almost as high as 0.40.

7 Discussion

We view the previous results as supportive of the view that an important part of the heterogeneity in paths of industrialization among our sample of Asian and Latin American economies can be rationalized within the context of a simple benchmark model of structural change, with the differences across countries driven by differences in sectoral productivity dynamics. This finding is consistent with the growing literature on structural change that stresses productivity dynamics as central to understanding the stylized facts of structural transformation.

But our analysis also suggests a role for additional factors in some countries that are particularly relevant for the division of non-agricultural labor between manufacturing and services. With this in mind, in this section we discuss factors that our model abstracts from which we believe may play an important role and which future research should seek to incorporate into the analysis.

Our model is static and views economies as closed. This raises the issue of how trade might matter, and in particular the role of dynamic trade imbalances. Several remarks are in order regarding the potential importance of these channels. First, consider the case in which trade is statically balanced, so that for each country, imports are equal to exports period by period. If trade occurs entirely within the manufacturing sector, the associated specialization would manifest itself as productivity increases and so would be picked up by our analysis. More generally, if trade occurs in other sectors but is balanced within each sector then the same comment would apply.

Trade that is balanced across time but not across sectors would affect our analysis; if a country imports food and exports manufacturing goods this would necessarily affect the employment shares that are the focus of our analysis. However, from an empirical perspective, net trade flows in agriculture tend to be relatively small for most countries, and trade in services has been much less important than trade in manufacturing, though the amount of trade within services continues to rise and is becoming more important in some countries. From a theoretical perspective there is an important question as to why countries that are relatively unproductive in agriculture do not simply import food from abroad (Tombe, 2015). But from an empirical perspective the assumption of no net trade flow in agriculture is not strongly counterfactual.

Next consider the case in which trade is not balanced period by period. Of particular relevance is the possibility that a country chooses to have a trade surplus and that the source of this surplus is exports of manufacturing goods. In this case current consumption is no longer the same as current production. If we take the amount of labor used to produce net exports of manufacturing as given, our model determines the optimal allocation of the remaining labor. This would imply a larger overall share for manufacturing employment.

This mechanism may be important for understanding the dynamics of some of the Asian economies in our sample. As noted earlier, China and South Korea were both notable outliers in terms of model predictions for the growth of A_m relative to A_s. In

Figure 14: Trade Surplus and Non-Agricultural Employment Shares: Asia

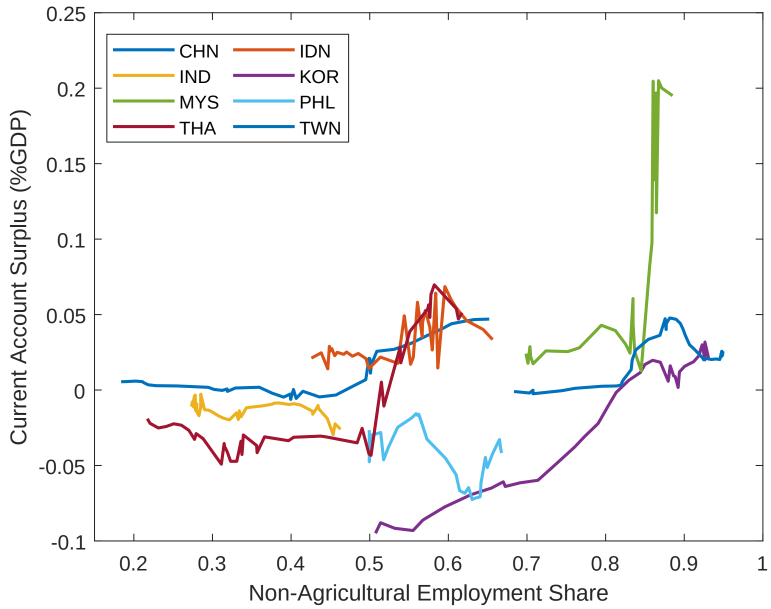

the data both countries had a relatively high value for this ratio, which in our model would imply a counterfactually high reallocation of labor from manufacturing to services. Dynamic trade imbalances may well play an important role for these countries. Figure 14 shows the trade surplus as a percentage of GDP versus the non-agricultural employment share for the Asian countries in our sample. It shows that both China and South Korea exhibit a significant increase in the trade surplus along the industrialization path.[25]

Dynamics may also matter for another reason. Recent work by Garcia-Santana et al. (2019) notes that the investment sector draws much more heavily from the manufacturing sector than does consumption sector. In the standard one sector growth model, a one time increase in TFP will generate a period of high investment as part of the transition dynamics. More generally, this raises the possibility that countries experiencing growth miracles that were associated with periods of relatively high investment may experience high peak employment shares in manufacturing that are at least in part driven by in-

[25]Malaysia and Taiwan also experience an increase but it is effectively at the end of the industrialization phase. Trade data is taken from the World Bank and the OECD national accounts database.

vestment rather than consumption. This is also more likely to be relevant for some of the Asian economies. Consistent with this, our analysis found that it was a set of high growth Asian economies that had higher peak employment shares than predicted solely by agricultural productivity dynamics.

Lastly, our analysis has abstracted from distortions that may impact sectoral labor allocations. Some of these might relate to distortions of consumption versus savings and so relate to the previous discussion. But they may also impact the composition of consumption. Many activities within services are either carried out by the government or subsidized by the government. Differences in government policies may therefore also play a role.

8 Conclusion

Countries exhibit significant heterogeneity in their paths of industrialization. In particular, industrialization paths of many recent developers differ from that of earlier developers (Rodrik, 2016). We have studied a benchmark model of structural change in order to assess the extent to which it can shed light on the sources of the heterogeneous industrialization experiences found in the data.

Our analysis led to three key findings. First, benchmark models of structural change naturally generate hump-shaped patterns for evolution of the manufacturing sector. Second, heterogeneous patterns of catch-up in sectoral productivities across countries naturally give rise to heterogeneous patterns of industrialization similar to those found in the data. Third, differences in the rate of agricultural productivity growth across economies can account for the majority of the variation in peak manufacturing employment shares.

The key message from our analysis is that simple benchmark models of structural provide the foundation for analyzing heterogeneous industrialization experiences across countries. An important next step is to extend the simple benchmark model used here to explore the role of additional factors beyond differences in sectoral productivity growth. We think it will be valuable will be to include trade and capital accumulation and to focus on dynamic implications. It will also be of interest to extend the analysis here to additional countries.

References

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: The anatomy of urban crisis. The American Economic Review 57(3), 415–426. 3

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. Econometrica 82(6), 2167–2196. 1, 4, 3.1, 12, 4

Buera, F. J. and J. P. Kaboski (2009, 05). Can Traditional Theories of Structural Change Fit the Data? Journal of the European Economic Association 7(2-3), 469–477. 4

Carter, S. B., S. S. Gartner, M. R. Haines, A. L. Olmstead, R. Sutch, and G. W. (Eds.) (2006). Historical Statistics of the United States, Earliest Times to the Present: Millennial Edition. Cambridge University Press, New York. 2.1, 2.2, 4

Caselli, F. (2005). Accounting for cross-country income differences. In: Aghion, Philippe, Durlauf, Steven (Eds.), Handbook of Economic Growth, vol. 1A. North Holland, Amsterdam and New York, pp. 679–742 (Chapter 9). 17

Comin, D. A., D. Lashkari, and M. Mestieri (2015, September). Structural change with long-run income and price effects. Working Paper 21595, National Bureau of Economic Research. 3, 4

Duarte, M. and D. Restuccia (2010, 02). The Role of the Structural Transformation in Aggregate Productivity*. The Quarterly Journal of Economics 125(1), 129–173. 1

Duernecker, G., B. Herrendorf, and A. Valentinyi (2017). Structural change within the service sector and the future of baumol's disease. Discussion Paper 12467, CEPR. 15

Felipe, J., A. Mehta, and C. Rhee (2018, 02). Manufacturing matters. . . but it's the jobs that count. Cambridge Journal of Economics 43(1), 139–168. 1

Garcia-Santana, M., J. Pijoan-Mas, and L. Villacorta (2019). Investment demand and structural change. CEPR Working Paper 11636. 7

Gollin, D., R. Jedwab, and D. Vollrath (2016). Urbanization with and without industrialization. Journal of Economic Growth 21, 35–70. 1

Gollin, D., D. Lagakos, and M. E. Waugh (2013, 12). The Agricultural Productivity Gap *. The Quarterly Journal of Economics 129(2), 939–993. 17

Gollin, D., S. Parente, and R. Rogerson (2002). The role of agriculture in development. The American Economic Review 92(2), 160–164. 1, 4, 3.1, 5

Gollin, D., S. L. Parente, and R. Rogerson (2007). The food problem and the evolution of international income levels. Journal of Monetary Economics 54(4), 1230 – 1255. 1, 4, 5

Herrendorf, B., R. Rogerson, and Valentinyi (2020). Structural Change in Investment and Consumption—A Unified Analysis. The Review of Economic Studies. 4, 4

Herrendorf, B., R. Rogerson, and Ákos Valentinyi (2014). Chapter 6 - growth and structural transformation. In P. Aghion and S. N. Durlauf (Eds.), Handbook of Economic Growth, Volume 2 of Handbook of Economic Growth, pp. 855 – 941. Elsevier. 3, 14, 16

Herrendorf, B. and Valentinyi (2012, 04). Which Sectors Make Poor Countries so Unproductive? Journal of the European Economic Association 10(2), 323–341. 5

Hsieh, C.-T. and P. J. Klenow (2007, June). Relative prices and relative prosperity. American Economic Review 97(3), 562–585. 5

Johnston, B. F. and P. Kilby (1975). Agriculture and Structural Transformation: Economic Strategies in Late-Developing Countries. Oxford University Press. 5

Johnston, B. F. and J. W. Mellor (1961). The role of agriculture in economic development. The American Economic Review 51, 566–593. 5

Kongsamut, P., S. Rebelo, and D. Xie (2001, 10). Beyond Balanced Growth. The Review of Economic Studies 68(4), 869–882. 3, 4

Ngai, L. R. and C. A. Pissarides (2007, March). Structural change in a multisector model of growth. American Economic Review 97(1), 429–443. 3, 4

Restuccia, D., D. T. Yang, and X. Zhu (2008). Agriculture and aggregate productivity: A quantitative cross-country analysis. Journal of Monetary Economics 55(2), 234 – 250. 17

Rodrik, D. (2012, 11). Unconditional Convergence in Manufacturing *. The Quarterly Journal of Economics 128(1), 165–204. 5

Rodrik, D. (2016). Premature deindustrialization. Journal of Economic Growth 21, 1–33. 1, 2, 5, 8

Sposi, M., K. Yi, and J. Zhang (2020). Structural change and deindustrialization. Working paper, University of Houston. 1

Tombe, T. (2015, July). The missing food problem: Trade, agriculture, and international productivity differences. American Economic Journal: Macroeconomics 7(3), 226–58. 7

Wise, A. (2020). Global dynamics of structural change. Working paper, Princeton University. 1

Üngör, M. (2013). De-agriculturalization as a result of productivity growth in agriculture. Economics Letters 119(2), 141 – 145. 11

Świecki, T. (2017). Determinants of structural change. Review of Economic Dynamics 24, 95 – 131. 3, 1

Appendix A Industrialization Across Countries

In this appendix we document industrialization (and deindustrialization) paths across countries using the same structure of Figure 1 but for all the other countries in our sample.

Figure A.1 shows similar patterns to the ones in Figure 1. In particular, four patterns stand out. First, the level of peak employment in manufacturing varies significantly: Venezuela has a peak value below 0.20, whereas Philippines reaches a value of almost 0.35. Second, there is also significant variation in the value of h_{nt} at which the peak is reached, ranging from around 0.60 for Venezuela to more than 0.80 for Argentina. Third, there is a strong positive correlation between the level of the peak and the value of h_{nt} at which the peak occurs. Finally, compared to Figure 2, Figure A.1 highlights that there is significantly more heterogeneity in the paths of industrialization in Asia and Latin America than there is in advanced economies.

Figure A.1: Industrialization Paths in Latin America and Asia

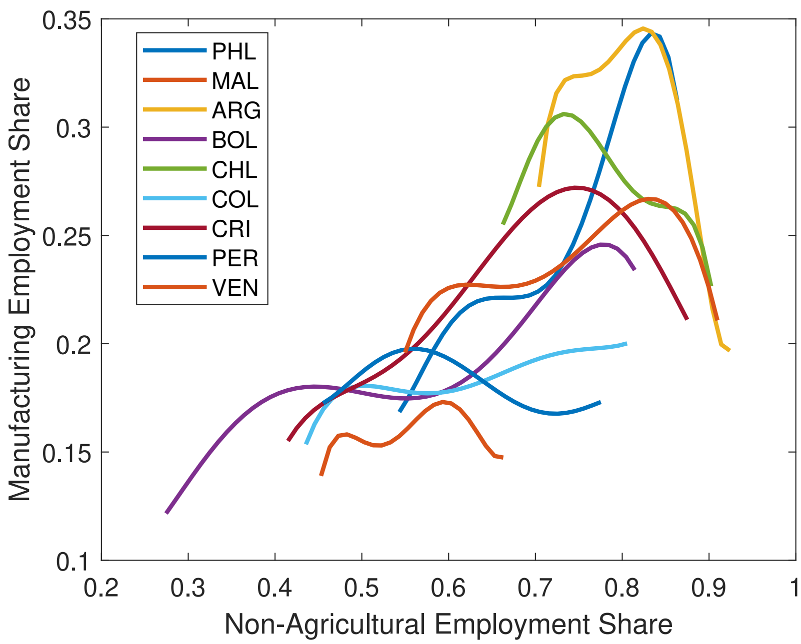

Appendix B Measurement of Growth in A_m/A_s

In this appendix we propose an alternative measure for the growth rate of A_m/A_s and use it to support our claim that the growth of A_m/A_s likely contains a substantial amount of error. In the text we directly computed productivity growth using data on real valued added and employment. In this section we follow the dual approach. In the context of our model, there is a one-to-one mapping between growth in A_m/A_s and the growth in p_s/p_m. Because the GGDC 10-Sector database includes data on nominal value added as well as real value added measured in 2005 prices, we can use the data to infer the growth rates of sectoral price indices and use this as an alternative measure of the relative growth rate of productivity in manufacturing relative to services.[26]

There is one limitation to applying this method to our sample of countries. Specifically, because so many of the Latin American economies in our sample experienced one or more periods of hyperinflation, the nominal value added numbers are either not available for the early years or are rounded to zero. The only exceptions to this are Colombia, Costa Rica and Mexico, so for this exercise we restrict our attention to the nine Asian countries and these three Latin American countries.[27]

The mean gap between the primal and dual measures of A_m/A_s is 0.0094, with a range of -0.0145 to 0.0389 and a mean absolute gap of 0.0139.[28] Figure B.1 plots the dual measure versus the primal measure for these twelve countries.

The differences between the two series are large and support the view that measured differences in sectoral productivity growth rates should reasonably be viewed as noisy

[26]Note that this approach can only be used to compute measures of relative sectoral productivity growth. In particular, it cannot be used to provide an alternative measure of growth in A_a. But since growth in A_m/A_s is a key value the approach offers very relevant information.

[27]The nominal value added series for the three European countries in our sample only start in 1970 in this database, so they are also not included.

[28]When we do this same calculation for the US data over the period that we used for our calibration (1950-1970) the difference between the primal and dual measures is -0.084, indicating that relative prices changed more than relative labor productivity.

Figure B.1: Measurement Discrepancy in Growth of Manufacturing-Services Relative Productivity

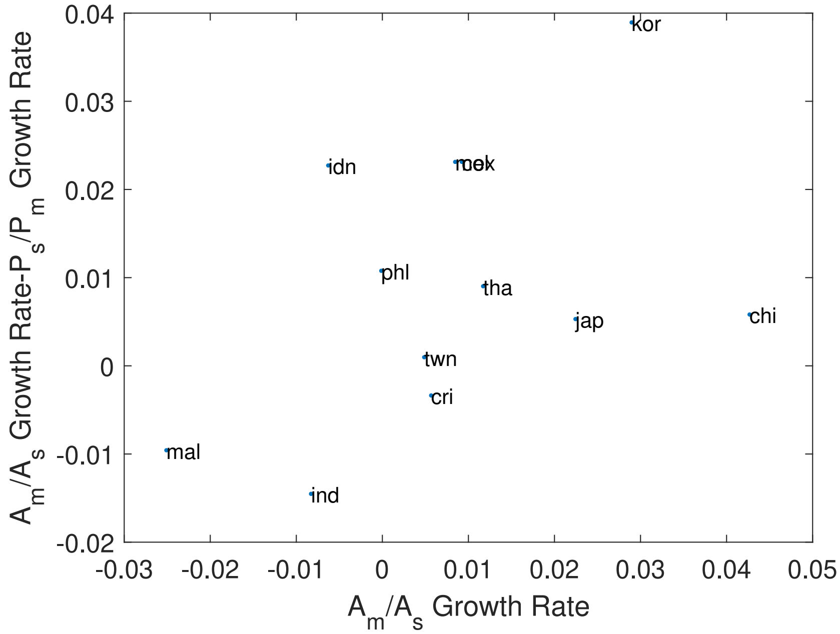

measures of actual differences. The primal approach yields an estimate that is roughly one percentage point larger than the dual approach. It is perhaps of interest to see how Figure B.1 looks if we split the difference and adjust our estimate of the differential growth rate down by one half of one percent for all countries. The results of this are in Figure B.2.

The only point we want to make with regard to this figure is that with this adjustment there is a larger mass of points that are clustered around the 45 degree line, so that a simple uniform adjustment in relative productivity growth seems to improve the fit of the model.

Figure B.2: Growth Rates for Manufacturing-Services Relative Productivity: Model and Adjusted Data