HETEROGENEOUS PATHS OF INDUSTRIALIZATION

Federico Huneeus
Richard Rogerson

ABSTRACT

Industrialization experiences differ substantially 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 industrial 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 the majority of the variation in peak industrial employment shares.

Federico Huneeus
1180 Agustinas
Office 352B
Research Department (DIE)
Central Bank of Chile
Santiago, RM 8340454
Chile
fhuneeus@bcentral.cl

Richard Rogerson
Department of Economics
JRR Building Room 292
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

1 Introduction

In his Nobel Prize address, Kuznets emphasized structural transformation–the realloca-
tion 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 devel-
opment (i.e., the industrialization phase), reaching a peak, and then declining in the later
stages of development (i.e., the deindustrialization phase). 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 manu-
facturing.[2] Second, we show that variation in the profile of sectoral productivity growth
rates across countries can generate variation in industrialization patterns that mimic those
found in the data, including the phenomenon of premature deindustrialization. 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. In particular, we show that relatively slow productivity growth in
agriculture can give rise to differences comparable to those found in the data. We also
use our model to assess the welfare consequences of premature deindustrialization caused
by relatively slow growth in agricultural productivity and find the effect to be modest in

---

[1]This observation was also noted independently by Palma (2005, 2014) and 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.

comparison to the overall gaps between rich and poor economies.

The literature on structural change emphasizes the role of sectoral productivities in shaping the structural transformation process, and consistent with this, they play a central role in our theoretical and quantitative analysis. Following Gollin et al. (2002, 2007), food is a necessity and the agricultural employment share is heavily influenced by productivity of the agricultural sector. Non-agricultural employment is allocated between manufacturing and services and in our theoretical analysis, depends on both the overall level of productivity in services as well as the relative productivity of the two sectors.[3] In our calibrated model, we find that only the relative productivity effect is required to match the industrialization phase of the US development path.

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 can vary 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 affects the path of industrialization. For example, we show that relatively slow growth in agriculture will lead to a lower peak employment

---

[3]Thus, our theoretical specification 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). Świecki (2017) and Comin et al. (2015) also allow for both effects. For additional discussion and evidence see Herrendorf et al. (2014).

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 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. 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.[5]

We study a closed economy and abstract from dynamic elements such as capital accumulation. Two recent papers have explored these extensions. Wise (2020) studies a model similar to ours but embeds it in a multi-country setting and allows for trade. Trade implies that recent developers face different relative prices than earlier developers, and he shows that this can account for modest differences in peak manufacturing shares. Sposi et al. (2020) study a model that includes both trade and capital accumulation. Consistent with our results, they find the most important driving force is heterogeneity in sectoral productivity growth patterns, though they also find a significant role for the interaction of sectoral productivity and changes in trade costs.

An outline of the paper follows. In the next section we present evidence on the hetero-

---

[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).

[5]Fujiwara and Matsuyama (2022) study economic forces very similar to the ones that we highlight and theoretically characterize the relationship between productivity catch-up dynamics and premature deindustrialization, but do not use their model to study the quantitative patterns in the data. Our analysis is also related to that in Gollin et al. (2016), who study heterogeneous urbanization experiences and how this relates to industrialization.

geneity in industrialization experiences across a set of Asian and Latin American countries in the post 1950 period, as well as four 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. Section 4 calibrates 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 economies and shows that differences in the growth rate of agricultural productivity across countries can account for the majority 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 (Timmer et al., 2015). 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).[6]

---

[6]Rodrik (2016) studied how employment and value added shares in the "narrow" manufacturing sector varied with the level of GDP/capita measured using PPP. We focus on the broader industrial sector and use only employment shares, and so do not use any information about relative prices.

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 database 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 four countries from Europe that are in the database and experienced significant industrialization since 1950–France, Spain, Denmark and Italy. We exclude West Germany, the United Kingdom, the Netherlands and Sweden, as they are already near the end of the industrialization phase when the data becomes available, and we exclude Hong-Kong and Singapore as they are city-states with little agriculture.[7] While the data set generally covers the period from 1950 to 2010, coverage for some countries begins after 1950.[8] 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 it 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 1880-2000. Finally, we also use international trade data from the World Bank Open Dataset to measure current account surplus as a percent of GDP for the 1950-2010 period for a set of countries. To measure the current account surplus for Taiwan, we combine data from the International Financial Statistics (IFS) from the International Monetary Fund (IMF) and the Federal Reserve Economic Data (FRED).

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

---

[7]Our criterion is to exclude countries for which the agricultural employment share is below 25 percent when data is first available.

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

involve any aggregation. We aggregate four sectors (mining, manufacturing, construction and utilities) to obtain what we 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 non-agricultural 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.[9]

Almost all of the countries in our sample have experienced peaks for their manufacturing employment share. The clear exceptions are China, India, and Thailand that have not yet experienced the peak. 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 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

---

[9]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
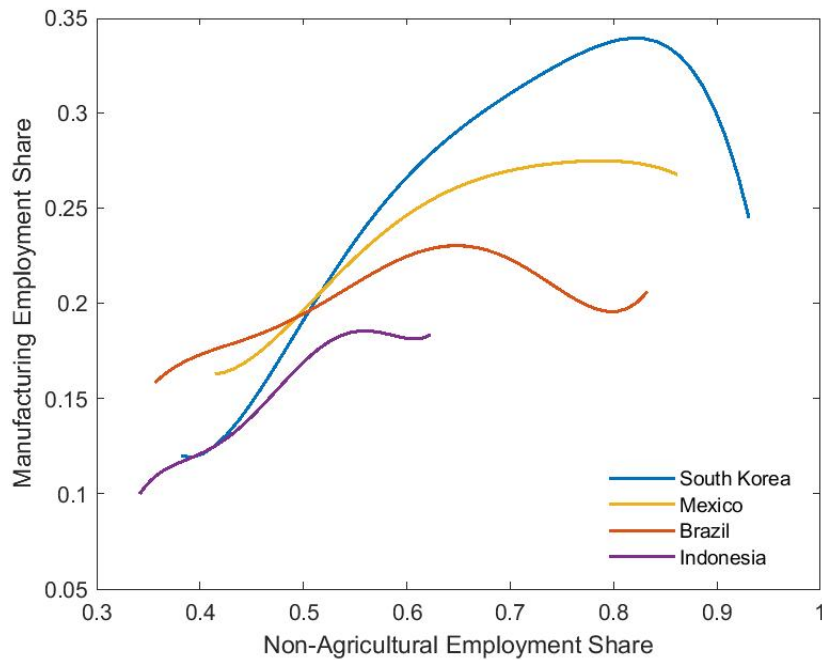
manufacturing varies significantly: Indonesia has a peak value below 0.20, whereas South Korea reaches a value of almost 0.35. Second, there is also significant variation in the value of h_{nt} at which the peak is reached, ranging from less than 0.60 for Indonesia to more than 0.80 for South Korea. 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 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 seven countries in our sample that currently qualify as advanced: Japan, South Korea, Taiwan, France, Denmark, 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 seven of these countries reach their peak manufacturing employment shares when h_{nt} lies between .80 and .90,

Figure 2: Paths of Industrialization: Advanced Economies
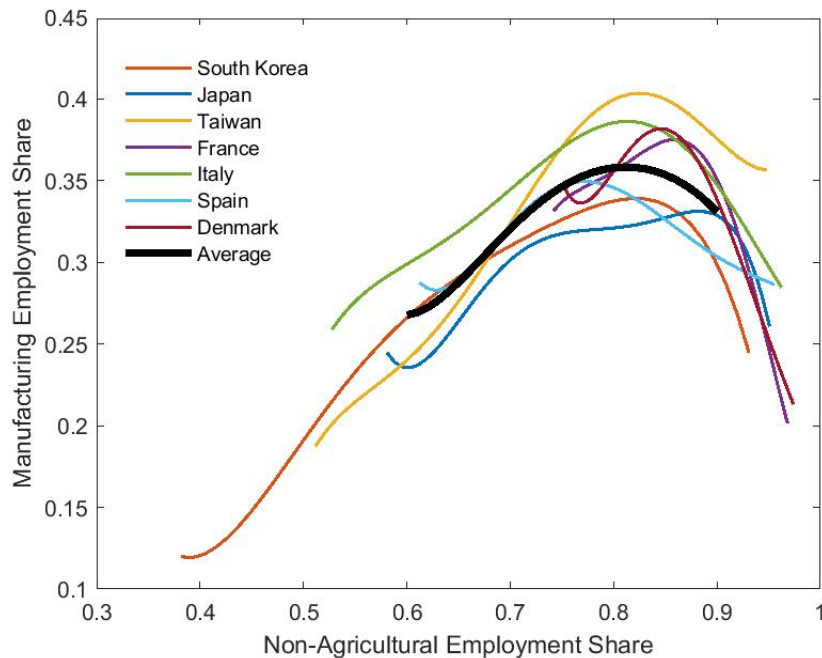
and the peak shares range from 0.34 to 0.40. The thick black line in the picture reflects the average for this subset of countries for h_{nt} in the range of 0.60 to 0.90.[10]

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 seven 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.

---

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

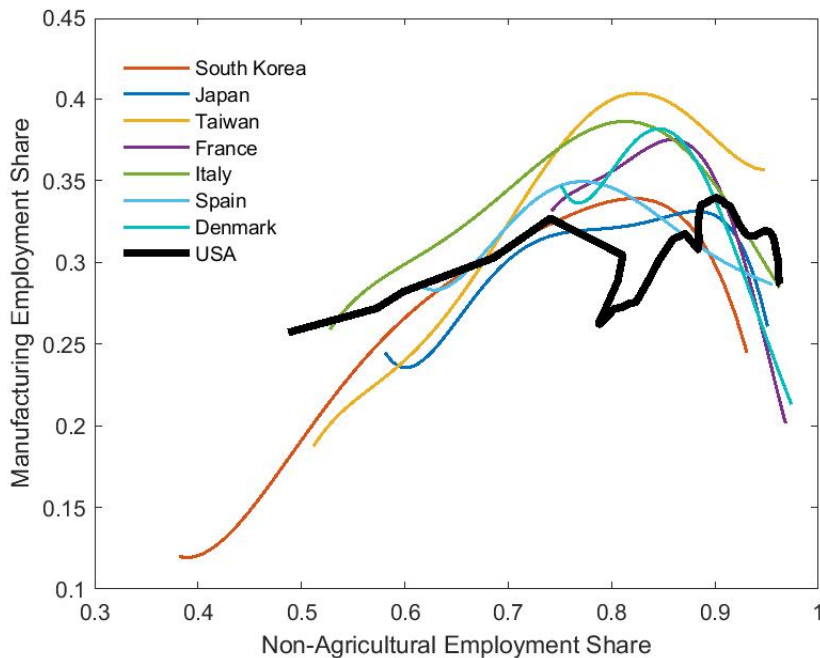Figure 3: Paths of Industrialization: Advanced Economies and the US

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.[11] As a reference point, we have included the profile for South Korea on this figure as well.[12]

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 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.

---

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

[12]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
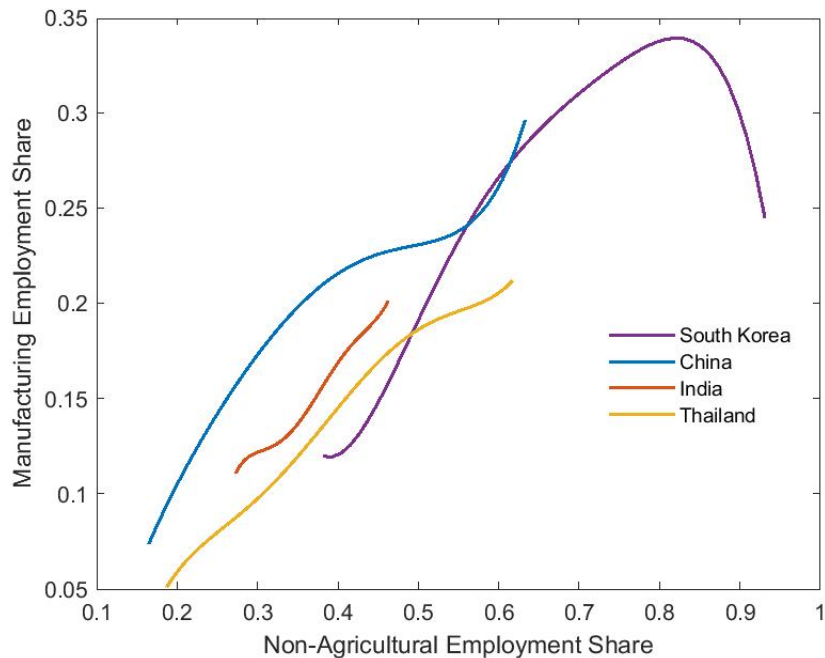
When summarizing the patterns in Figure 1 we noted a strong positive correlation between the value of h_{nt} at which the peak value of h_{mt} occurred and the value of peak h_{mt}. Denote these two values by h_n^* and h_m^* respectively. We now pursue this pattern further using the full sample of 19 countries who have attained their peak. Table 1 shows the values of h_n^* and h_m^* for each of the 19 countries and 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 0.82. 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.

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^*
Indonesia	0.59	0.20	Argentina	0.80	0.35	France	0.89	0.37
Japan	0.85	0.34	Bolivia	0.77	0.25	Italy	0.83	0.39
South Korea	0.84	0.36	Brazil	0.67	0.24	Spain	0.81	0.36
Malaysia	0.85	0.36	Chile	0.76	0.34	Denmark	0.84	0.38
Philippines	0.61	0.18	Colombia	0.75	0.23			
Taiwan	0.85	0.41	Costa Rica	0.77	0.29			
			Mexico	0.84	0.29			
			Peru	0.56	0.22			
			Venezuela	0.87	0.29			

Figure 5: Correlation of Peak and Timing of Industrialization

11
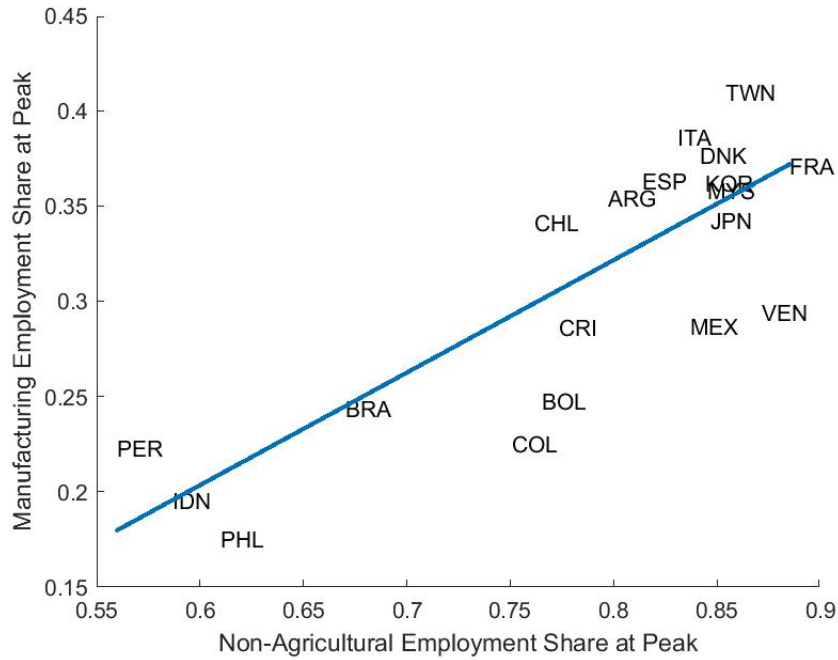
3   A Model of Industrialization

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 provide sufficient conditions for the model to generate the hump-shaped dynamics for h_{mt} that characterize the industrialization phase of development as illustrated in the previous section. We also report some comparative statics results to illustrate how differences in productivity growth rates can generate the qualitative relationship between h_m^* and h_n^* that we documented in the previous section.

3.1   Model

In this section we lay out a continuous time formulation of the simple benchmark model of structural transformation used by Duarte and Restuccia (2010). The allocation decisions in the model are all static, with all dynamics due to 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 consumption good 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. As in Duarte and Restuccia (2010), preferences

12

are given by[13]:

$$U(c_a, c_m, c_s) = \alpha_a \log(c_a - \bar{c}_a) + (1 - \alpha_a) \log \left(\left[\alpha_m c_m^{\frac{\sigma-1}{\sigma}} + (1 - \alpha_m)(c_s + \bar{c}_s)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \right)$$

where $\bar{c}_a$ and $\bar{c}_s$ are positive constants that generate non-homotheticities in preferences, α_a and α_m are preference weights in the interval $(0, 1)$, and $\sigma > 0$ influences the elasticity of substitution between manufacturing and services.[14] These preferences are also the natural extension of those used by Caselli and Coleman II (2001), Lagakos and Waugh (2013), and Adamopoulos and Restuccia (2014) extended to disaggregate the nonagricultural sector.

We focus on the competitive equilibrium allocation for the above economy, which is also the solution to the Social Planner's problem that solves for the efficient allocation of labor across the three sectors. For future reference we note that if we normalize the wage rate to equal unity and let the three prices be denoted by p_i, $i = a, m, s$, the linear production functions imply that equilibrium prices are the inverse of sectoral productivities:

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

The equilibrium allocation of labor is then dictated by the demand 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 focus on the case in which $A_a > \bar{c}_a$. If $\bar{c}_s > 0$ then it is possible that h_s could equal zero, but in what follows we will assume that all three labor allocations are interior, as

---

[13]In the working paper version of this paper (Huneeus and Rogerson, 2020) we assumed that agricultural consumption was always equal to $\bar{c}_a$ and considered preferences with a non-homothetic aggregator for manufacturing and services in the spirit of Boppart (2014). Both the analytic and quantitative results obtained for that specification are very similar to what we present in this paper.

[14]If $\bar{c}_s$ is non-zero then the elasticity of substitution between manufacturing and services is no longer constant and in particular, is not identical to σ.

this will be the relevant case for our quantitative analysis.

Assuming interior solutions, algebraic manipulation of the first order conditions from the Social Planner's problem yields:

$$h_a = \alpha_a + \frac{\bar{c}_a}{A_a} + \alpha_a \left[\frac{\bar{c}_s}{A_s} - \frac{\bar{c}_a}{A_a} \right] \tag{3.2}$$

$$h_m = \frac{(1-\alpha_a)\alpha_m^\sigma}{[\alpha_m^\sigma + (1-\alpha_m)^\sigma (\frac{A_m}{A_s})^{1-\sigma}]} \left[1 - \frac{\bar{c}_a}{A_a} + \frac{\bar{c}_s}{A_s} \right] \tag{3.3}$$

$$h_s = \left(\frac{\alpha_m}{1-\alpha_m} \right)^{-\sigma} \left(\frac{A_m}{A_s} \right)^{1-\sigma} h_m - \frac{\bar{c}_s}{A_s} \tag{3.4}$$

3.2 Industrialization Paths

We now consider the evolution of sectoral employment shares along a development path when technical change is the sole driving force behind development. 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 $\sigma < 1$ and that $g = g_m - g_s > 0$.[15] For some of our results we will need an additional condition that restricts the relative importance of the two sources of non-homotheticities in preferences. With positive and constant growth rates, the time series properties of the expression $\frac{\bar{c}_a}{A_{at}} - \frac{\bar{c}_s}{A_{st}}$ are generally ambiguous; $\frac{\bar{c}_a}{A_{at}}$ is positive, monotone decreasing and convex, while $-\frac{\bar{c}_s}{A_{st}}$ is negative, monotone increasing

---

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

14

and concave. Because both the level and derivative of this term go to zero as t tends to infinity, its properties are only relevant over some interval $[0, T]$. Our additional assumption requires that $\frac{\bar{c}_a}{A_{at}} - \frac{\bar{c}_s}{A_{st}}$ over the horizon considered inherits the properties of $\frac{\bar{c}_a}{A_{at}}$. This is trivially true if $\bar{c}_s = 0$, which turns out to be the case for the calibrated model that we study later in the paper. If $g_a = g_s$, this assumption is equivalent to $\bar{c}_a > \bar{c}_s$. This assumption is also satisfied if $g_s = 0$ and $\bar{c}_a > \bar{c}_s$.

For reference we collect the two previous assumptions as A1 and A2:

Assumption A1: $g_i > 0$ for $i = a, m, s$, $g = g_m - g_s > 0$, and $\sigma < 1$.

Assumption A2: The function $F(t) = \frac{\bar{c}_a}{A_{at}} - \frac{\bar{c}_s}{A_{st}}$ is positive, monotone decreasing and convex for horizon being analyzed.

The next proposition shows that our model can generate three key stylized facts regarding the sectoral employment dynamics observed for current industrialized countries: the share of employment in agriculture declines monotonically, the share of employment in services increases monotonically and the share of employment in manufacturing is hump-shaped.

PROPOSITION 1. *(i) Under A1, $\dot{h}_{at} < 0$ and increases monotonically to 0.*

(ii) Under A1 and A2, $\dot{h}_{st} > 0$

(iii) Under A1 and A2, $\dot{h}_{mt}/h_{mt}$ is monotonically decreasing and converges to $-g(1 - \sigma)$.

Proof: See Appendix B.

Because the Proposition implies that $\frac{\dot{h}_{mt}}{h_{mt}}$ converges to $-g(1 - \sigma)$, it will necessarily become negative at some point. It follows that a sufficient condition for h_{mt} to exhibit hump-shaped dynamics is for $\frac{\dot{h}_{mt}}{h_{mt}}$ to be positive at some point. While our normalization of productivities at time zero need not imply that this is true at $t = 0$, if one works backward from time zero there will be a point $\hat{t}$ at which $h_{a\hat{t}} = 1$, so that h_{mt} will necessarily increase over some interval beginning at $\hat{t}$. In the next section we will show

15

that the model can quantitatively account for the industrialization dynamics of the US for empirically reasonable parameters.

As in the previous section, let h_n^* and h_m^* denote the values for these two employment shares at the point at which h_{mt} reaches its maximum value. We now turn to some comparative statics results for these values. These values are functions of the three productivity growth rates: g_a, g_m and g_s. However, an examination of Equations (3.2), (3.3), and (3.4) suggests that it is useful to instead consider these values as functions of g_a, $g = g_m - g_s$, and g_s, so we will write $\mathrm{h}_n^*(g_a, g, g_s)$ and $h_m^*(g_a, g, g_s)$. We note that in the special case in which $\bar{c}_s = 0$, these functions depend only on g_a and g, so that g_s has no influence independently of g. The next proposition characterizes these functions.

PROPOSITION 2. *(i) Given A1, $h_m^*(g_a, g, g_s)$ is increasing in g_a, decreasing in g and increasing in g_s.*

(ii) Given A1 and A2, $h_n^*(g_a, g, g_s)$ is decreasing in g.

(iii) Given A1 and assuming either $\bar{c}_s = 0$ or $g_s = 0$, $h_n^*(g_a, g, g_s)$ is increasing in g_a.

Proof: See Appendix B.

This proposition tells us how differences in sectoral productivity growth rates can qualitatively generate some of the differences we documented in the previous section. In particular, it provides conditions under which a lower value of g_a will lead to a lower peak value for h_m^*, and for this peak value to occur at a lower value of h_n^*. Similarly, a higher value of $g = g_m - g_s$ will also lead to this qualitative pattern. In the remaining sections we explore the quantitative magnitude 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 in mind, consider the comparative statics concerning an increase in the value of g_a. This serves both to increase the overall pace of development and change the profile of relative productivities along the development path. Becausethe implications for h_n^* and h_m^* are invariant to the pace of development, these results should be understood as reflecting changes in the profile of relative sectoral productivities and not changes in the pace of development.

4   Benchmark Calibration

In this section we calibrate the above model to capture 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 same phase for the US economy. In particular, the changing composition within the services sector that is especially important in the later stages of development 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

Figure 6: US Sectoral Employment Shares 1880-1980
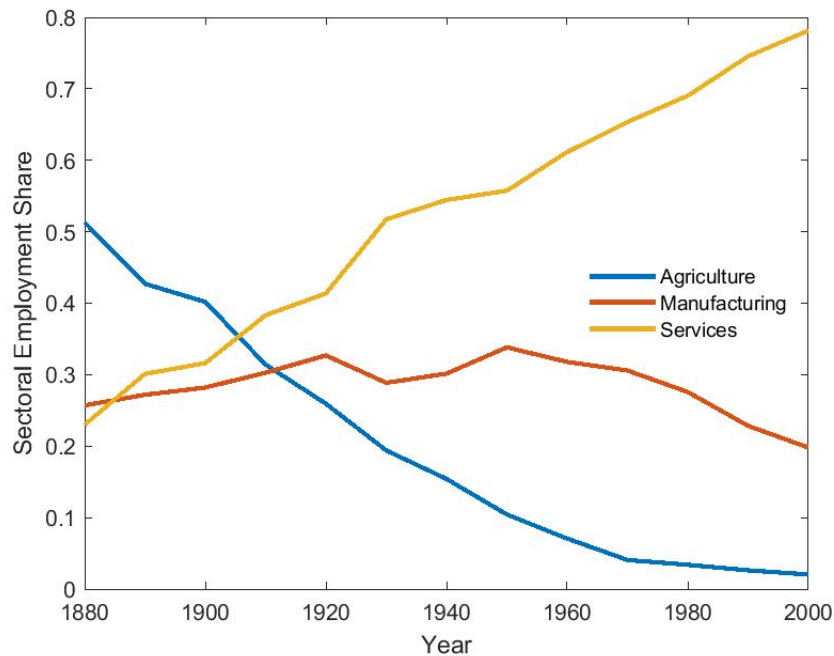
post 1970 period.[16]

As noted earlier, our data for the evolution of US sectoral employment shares between 1880 and 1950 comes from Carter et al. (2006) and the BEA. For the post 1950 period we use data from the GGDC 10-Sector Database. Figure 6 shows the time series for these shares between 1880 and 2000.

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 disruptive 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.

For our quantitative work we adopt a discrete time version of the model and set the

---

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

Table 2: Benchmark Calibration

| g_a	g_m	g_s	$\bar{c}_a$	$\bar{c}_s$	α_a	α_m	σ
1.0239	1.0225	1.0147	0.49	0	0.02	0.42	0.35

period length equal to one year. Relative to our continuous time formulation the only impact is to change the representation of productivity growth: we now write $A_{it+1} = g_i A_{it}$ for each sector i.

We first describe our procedure for determining the sectoral productivity series taking the values for the three preference parameters α_a, $\bar{c}_a$ and $\bar{c}_s$ as given. We normalize all three productivity levels in the initial period to equal unity, which is essentially a choice of units. Unfortunately, we do not have sectoral productivity data that covers the period from 1880-1950. For g_m and g_s we use data from the GGDC 10-Sector Database for real valued added per worker 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 $g_s = 1.0147$. Given this value for g_s, and the assumed values for α_a, $\bar{c}_a$ and $\bar{c}_s$, 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 (see Equation (3.2)). This implies that $g_a = 1.0239$.

Next we describe the calibration of the preference parameters. The model predicts that as time tends to infinity, h_{at} converges to α_a. Therefore, α_a captures the long-run employment share in agriculture, which in the data is around 2%. Thus, we set $\alpha_a = 0.02$.

This leaves four parameter values to assign: σ, $\bar{c}_a$, α_m and $\bar{c}_s$. Our strategy is to pick values for these four parameters so as to match the industrialization profile for the US as shown in Figure 3. Table 2 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.

Figure 7: Model Fit to US Data
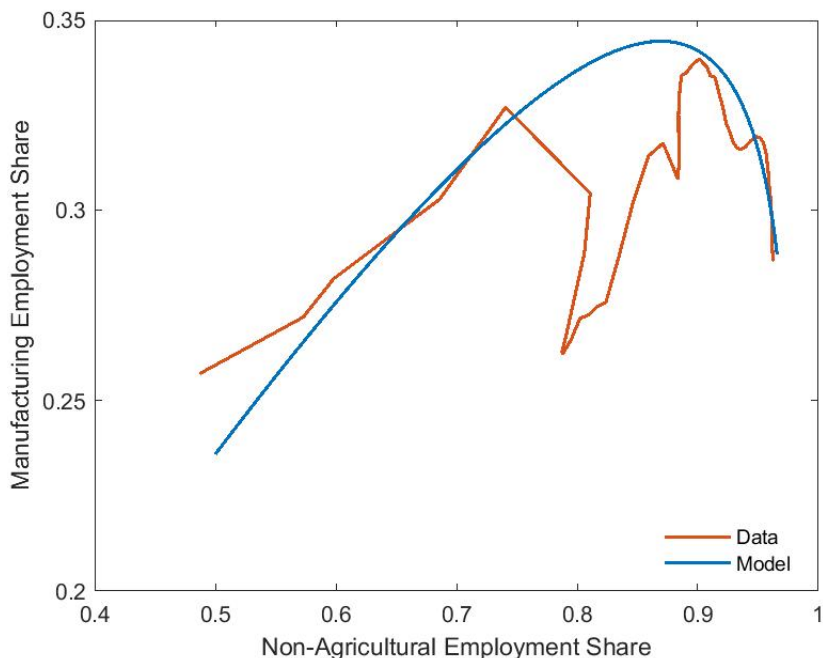
Allowing for the fact that we abstract from the effects of the Great Depression, the fit is reasonably close overall. Recall that our calibrated model assumes that all three productivity growth rates are constant. To the extent that trend productivity growth in the US varied somewhat over the period 1880-1950, we will see some departures of the data from the paths implied by the model.

We make three remarks about the calibrated preference parameters. First, consistent with the existing literature, we find $\sigma < 1$, implying that manufacturing and services are complements in the preference of the representative household. Second, we find $\bar{c}_a$ to be large which implies that income effects are very important for agricultural consumption. consistent with previous literature Herrendorf et al. (2020).

Third, we find that allowing for non-homotheticities in the aggregator for manufacturing and services has little effect on the model's ability to fit the data and so set $\bar{c}_s = 0$.[17]

---

[17]This finding is consistent with the results in our earlier working paper (Huneeus and Rogerson, 2020). In that paper we assumed a non-homothetic aggregator over manufacturing and services as in

The implication is that all of the income effects in our calibrated model of the industrialization phase are driven by agriculture. As noted earlier, the increasing importance of health and education in the later stages of development may give rise to changes in the nature of income effects. Notably, Duarte and Restuccia (2010) found a significant role for income effects on services for the US in the post 1950 period.[18]

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

_____

Boppart (2014) but also found very small income effects.

[18]We note that the preferences studied in Alder et al. (2022) flexibly allow for income effects of a particular consumption category to be on different sides of unity at different stages of development.

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.[19] 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 a larger gap between productivity growth in manufacturing and services, driven by slower productivity growth in services.[20] This departure can also generate significant differences in the peak employment share for the manufacturing sector, though the evidence for this departure is somewhat less strong. Whereas Rodrik (2013) 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.

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

---

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

[20]Because our calibrated model displays preferences over manufacturing and services which are homothetic, the effects of g_s and g_m on the industrialization path are 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
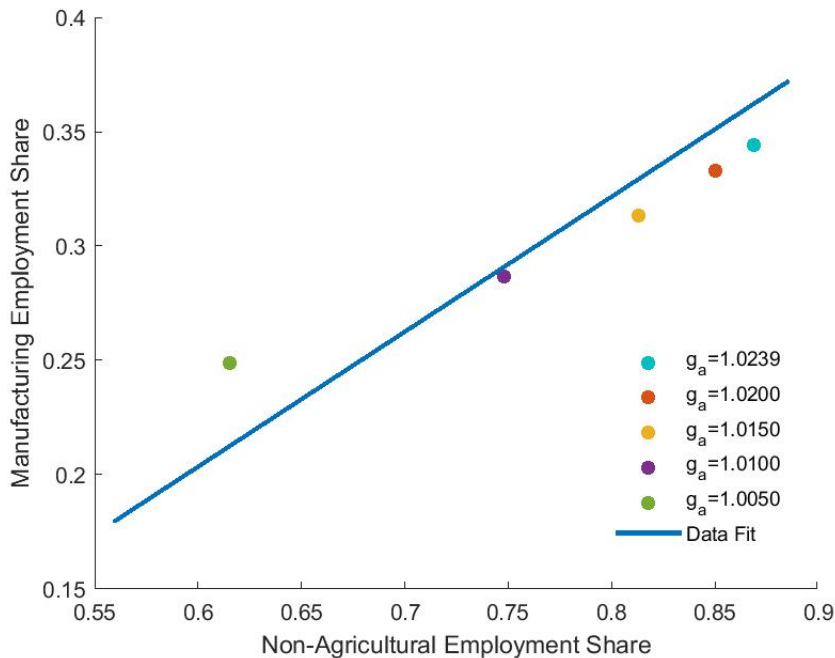
the quantitative implications of this decrease.

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

The point for our calibrated economy (i.e., $g_a = 1.0239$ 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 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 values of g_a they achieve different levels of h_{nt} during the 150 years.

Figure 9: Agricultural Productivity Growth and Paths of Industrialization
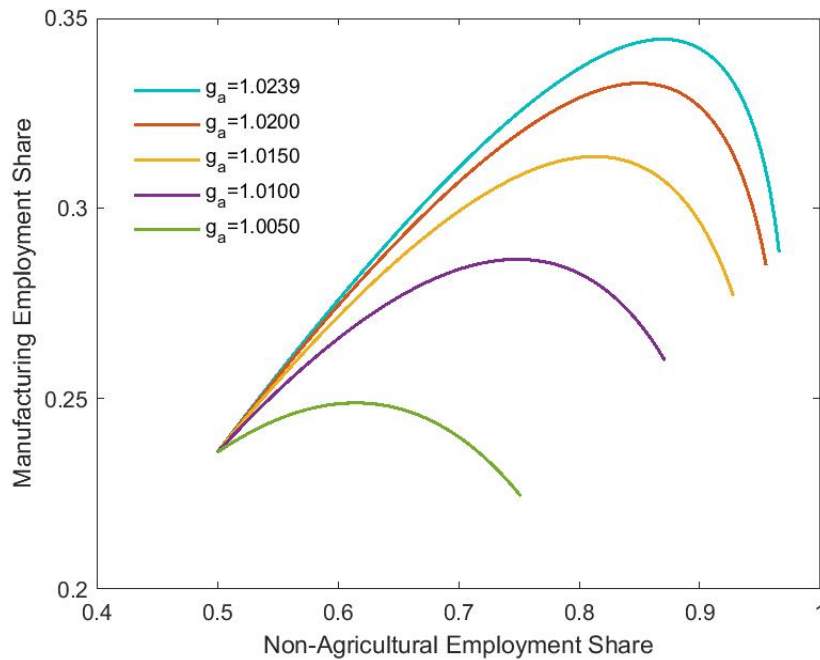

5.2 Slow Catch-up in Services Productivity

In this subsection we repeat the previous exercise but this time considering the possibility of slower catch-up in services.[21] 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 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 slower growth in g_a to affect these values, for the simple reason that the value of g_a exceeds the value of g_s in our benchmark calibration. However, a one percentage point gap in the

---

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

Figure 10: Services Productivity Growth and Peak Industrialization
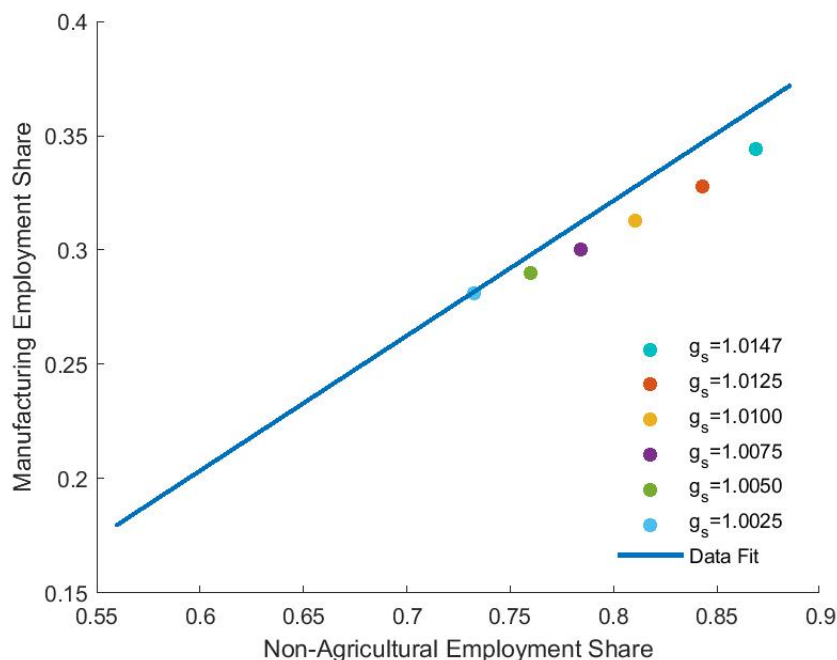
growth rate relative to the US has roughly comparable effects in both cases.

5.3 Welfare Implications

The previous calculations show that our model can generate the heterogeneity in industrialization patterns found in the data. It is also of interest to assess the welfare implications associated with this heterogeneity. In this subsection we report the results of an exercise to examine this.

We carry out the following calculation for each value of g_a depicted in Figure 9. First, we compute the present discounted utility for the representative household, assuming a discount factor of 0.94. Second, we compute the common scale factor applied to all sectoral productivity levels in the benchmark specification (i.e., $g_a = 1.0239$) that would yield the same lifetime utility as in the specification with a lower value for g_a. We choose to do our welfare comparison with a proportionate adjustment of productivity levels rather than

consumption levels because of the non-homotheticity in preferences.

The resulting productivity scale factors for the four values of g_a (1.02, 1.015, 1.01, 1.005) are .98, .95, .91, and .87 respectively. In terms of welfare, a decrease in the peak manufacturing employment share of roughly 11 percentage points driven by a lower value of g_a is equivalent to a 13 percentage point drop in the level of productivity.[22] While this effect is significant, these productivity level effects are somewhat small relative to the overall productivity gaps between the rich and poor countries in our sample.

5.4   Discussion

The previous examples in this section considered cases in which either g_a or g_s was smaller than in the benchmark calibration, so that in each case we were considering economies that were experiencing slower development. Note that because a uniform change in both g_m and g_s has no effect in our calibrated model, our results for lower values of g_a would also apply to a situation in which all three growth rates are uniformly lower. Our results indicate that such a slow developing economy can exhibit premature deindustrialization relative to the benchmark economy.

However, it is also important to note that our framework does not imply that the patterns associated with premature deindustrialization are necessarily associated with slow development. We highlight three reasons why. First, as we noted earlier, our results depend on the evolution of the profile of sectoral productivities and not the speed with which a country moves along the profile. Second, one of the earlier examples considered decreases in g_s. But our calibrated model implies that g_s affects the evolution in $h_n - h_m$ space only through its effect on $g = g_m - g_s$. So we would have obtained a similar evolution in $h_n - h_m$ space if we had instead assumed that the increase in g was due to a higher value of g_m, which would imply a faster rate of overall development in the economy. Third, and

---

[22]While not reported, the welfare effects are similar if the lower peak is caused by a lower value of g_s.

related, because our model implies that g_m and g_s matter only through their difference, uniform changes in the levels of these two growth rates do not affect the industrialization dynamics in $h_n - h_m$ space.

In the next section we will connect our model with the data and show that much of the heterogeneity in industrialization experiences is accounted for by variation in observed values of g_a holding g constant.

6 Rationalizing the Data

In the previous section we showed that seemingly plausible 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 growth profiles that would be required to rationalize the data for each of the 19 countries in our sample that experience the hump shape in manufacturing employment share, and then compare these productivity growth profiles to those observed in the GGDC 10-Sector Database.[23]

6.1 Inferring Productivity Growth Profiles

The first step in our exercise is to use our model calibrated to the US industrialization experience to infer sectoral productivity growth 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.

---

[23]An obvious alternative to this two-step procedure would be to simulate the model using empirical productivity growth profiles. If productivities are measured with error we think our two step procedure is preferable. In particular, in our model it is the gap between productivity growth rates in manufacturing and services that is key, and the gap between two productivity growth rates might be especially subject to measurement error.

Because sectoral employment shares sum to one, only two of them are independent, implying that only two sectoral productivity growth rates can be inferred. But since our calibrated model features $\bar{c}_s = 0$, the two growth rates g_m and g_s matter only via their influence on $g = g_m - g_s$. For this reason we can use data on changes in employment shares to recover both g_a and g. Specifically, our calibrated model implies the following mapping from the productivity profile (A_{at}, A_{mt}, A_{st}) into the employment shares h_{at} and h_{mt}:

$$h_{at} = \frac{\bar{c}_a}{A_{at}} + \alpha_a \left(1 - \frac{\bar{c}_a}{A_{at}}\right) \tag{6.1}$$

$$h_{mt} = h_{nt} \frac{\alpha_m^{\sigma}}{\alpha_m^{\sigma} + (1 - \alpha_m)^{\sigma} \left(\frac{A_{mt}}{A_{st}}\right)^{1-\sigma}} \tag{6.2}$$

The time series for A_{at} is uniquely determined by the time series for the employment share h_{at} and the time series of (A_{mt}/A_{st}) is uniquely determined by the time series of the employment share ratio h_{mt}/h_{nt}.

6.2   Results

We carry out the above procedure for each of the 19 countries in our sample up to the point at which they reach their peak employment share in manufacturing. We focus on this period because we calibrated our model to the industrialization phase.

We are particularly interested in the relationship between the average productivity growth rates in the data and those implied by this procedure. Because our model contains labor as the only input, we focus on real value added per worker as our measure of productivity in the data. 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

Figure 11: Agricultural Productivity Growth: Model and Data
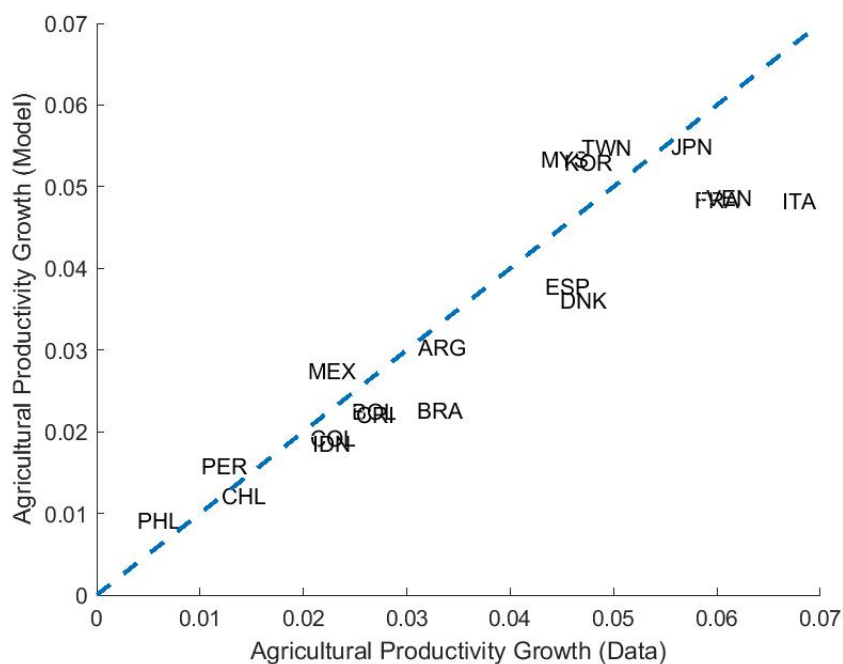
the points tend to track the 45 degree line. The correlation between the two values is 0.91. Being mindful of the fact that there are various issues that result in classical as well as non-classical measurement error, we view this as a very strong correlation. We conclude that the model's relatively 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 some significant outliers relative to the 45 degree line and the overall correlation is only 0.20. The two points in the bottom right corner are Taiwan and South Korea. We will revisit these cases in the next section when we discuss extensions, but for now we note that if we exclude these two countries the correlation increases quite substantially to 0.58.

Importantly, the plot in Figure 12 is comparing the *difference* between two productivity

Figure 12: Growth in Manufacturing-Services Relative Productivity: Model and Data
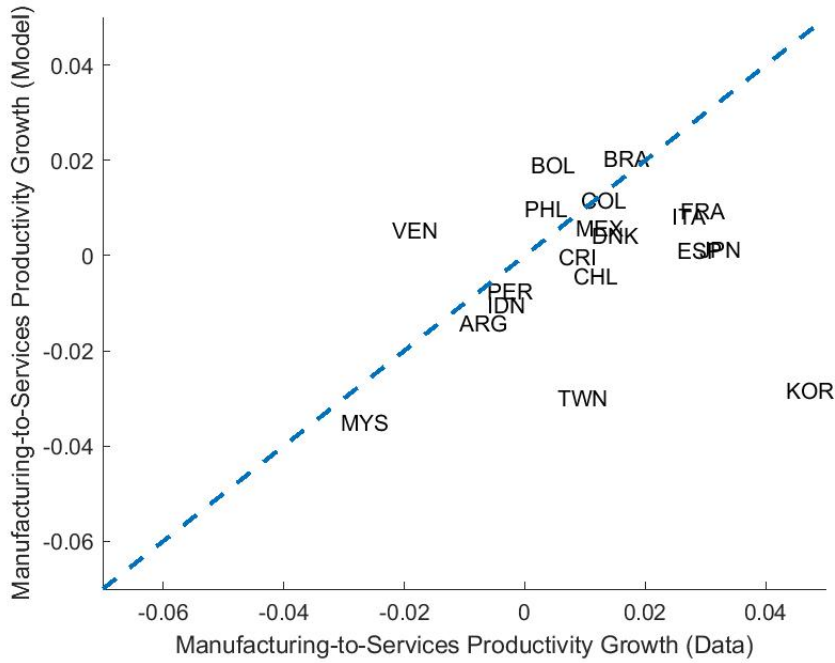
growth rates. Assuming each growth rate is independently measured with error, the difference between the two growth rates will display much greater error. Given this we view the correlation of 0.58 for the subsample without Taiwan and South Korea as also quite supportive of the model.

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
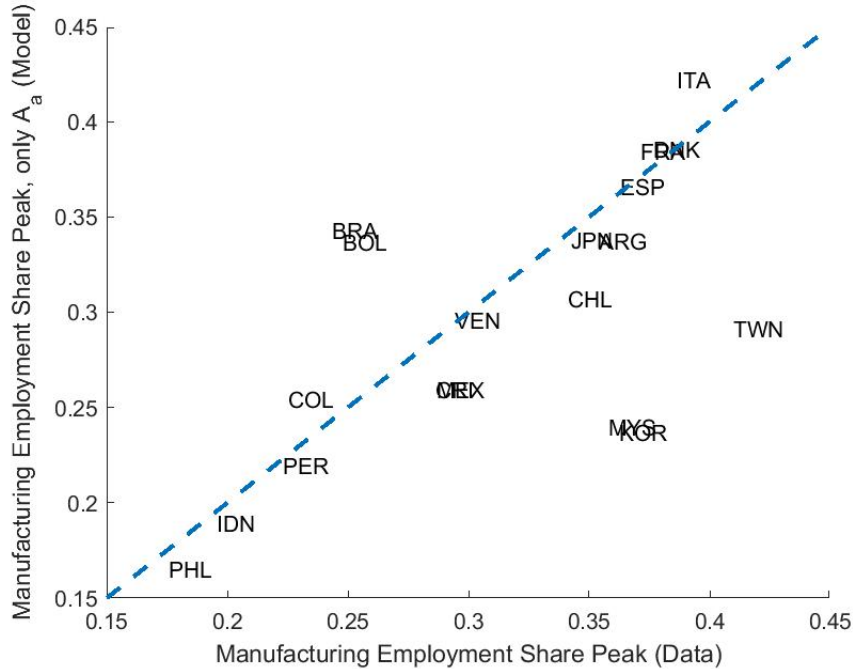
shares.

To do this we carry out the following exercise for each of the 19 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 difference in productivity growth rates in manufacturing and services is the same as those that we calibrated for the US economy. 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 19 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 a good job of accounting for the observed variation in peak levels of h_m for most of the countries in our sample. Fourteen 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.

There are five 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

31

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

replicate the evolution of employment shares. It is interesting to note that this is not the case for all of the so called Asian Tigers, as Japan is very close to the 45 degree line. The other two countries lie above the 45 degree line and are both from Latin America: Bolivia and Brazil.

Importantly, Figure 13 shows that differences in the evolution of agricultural productivity alone can generate differences in peak manufacturing employment shares that range from less than 0.20 to almost as high as 0.40. Furthermore, the correlation between the peak manufacturing employment share from the data and the one from the model with only differences in agricultural productivity growth is 0.64.

6.4 Changes in Value Added Shares and Relative Prices

Our quantitative analysis has focused on changes in employment shares and changes in productivity. In the competitive equilibrium of our model, value added shares are equal to employment shares and relative prices are the inverse of relative productivities. It is of interest to examine the extent to which these relationships hold in the data. Because our main exercise takes initial levels as given and focuses on the subsequent changes in employment shares and productivities, we focus here on the extent to which these relationships hold in the data in changes.[24] That is, we examine the relationship between long run changes in employment shares and value added shares, and between long run changes in relative productivity and inverse relative prices.

To do this we use the data on sectoral value added in current domestic prices contained in the GGDC 10-Sector Database. We drop Brazil and Peru because of data issues, leaving us with 16 of our initial 18 countries from Asia and Latin America.[25]

We begin with a comparison of changes in value added shares and changes in employment shares. Because the three shares sum to one we focus on just two of them, agriculture and manufacturing. For each sector we regress the change in the log of the employment share between initial and final years on a constant and the change in the log of the value added share between initial and final years. We are particularly interested in whether the coefficient on the change in the log of the value added share differs from one. When we run this regression for manufacturing the coefficient is .99 with a standard error of .27.

When we run the same regression for agriculture we obtain a coefficient of .35 with

---

[24]More formally, to the extent that there are level differences between employment and value added shares, or inverse relative prices and relative productivities, this could be captured by introducing wedges into our competitive equilibrium. These wedges would have little impact on the responses of employment shares and relative prices to changes in productivity.

[25]For Brazil the issue is that the data only begins in 1990. For Peru the issue is that value added from services is negative for many years. We note that value added data for Bolivia begins in 1958 and for Venezuela begins in 1960.

Figure 14: Agricultural Employment Share versus Value-Added Shares

A. Argentina   **B.** Venezuela
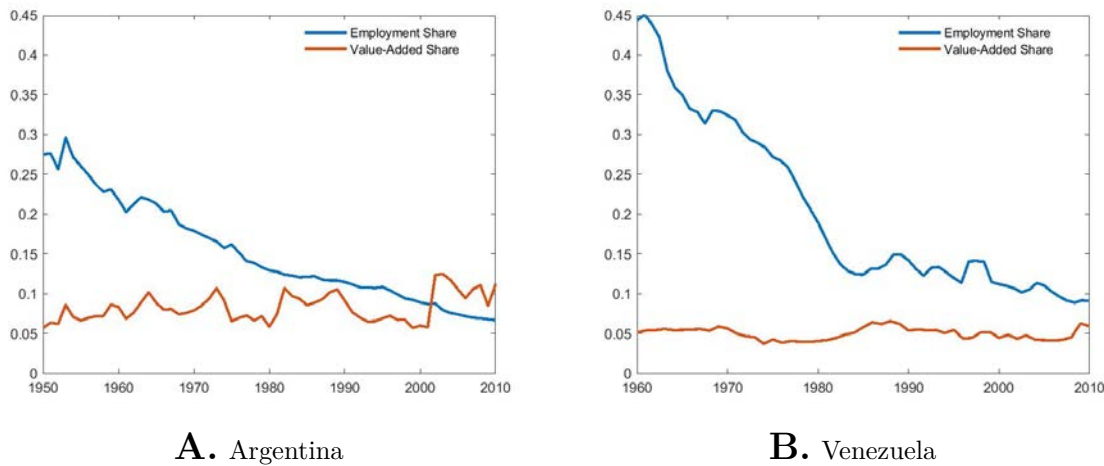
a standard error of .16. However, it turns out that this result is heavily influenced by two outliers: Argentina and Venezuela. We provide more detail on these two cases below, but first note that if we run the same regression for agriculture excluding these two countries we obtain a coefficient of .88 with a standard error of .15. Modulo excluding the two outliers, this suggests that the patterns for changes in employment shares and value added shares are quantitatively similar.

The two panels of Figure 14 show the evolution of employment and value added shares in agriculture for both Argentina and Venezuela. The evolution of the employment share in each country follows the standard declining pattern found for virtually all countries. In sharp contrast, the value added shares do not exhibit any tendency to decline and in Argentina appear to even show a modest upward trend.[26] This behavior is contrary to what is found in virtually all countries. One conjecture is that these countries engaged in a pattern of price controls and producer subsidies that served to keep the measured value added share artificially low. We do not pursue this anomalous behavior further in this

---

[26]Another noteworthy feature of the data for Argentina is the very dramatic movements in the value added share series relative to the employment share series. There are several periods where the value added share changes by a factor of two.

paper, but emphasize that none of the patterns we highlighted earlier in this section rely on observations for these two countries.

Next we consider the patterns for changes in relative prices and relative productivities. The two comparisons we focus on are between agriculture and manufacturing and between services and manufacturing. In each case we regress the change in log relative productivity between initial and final years on a constant and the negative of the change in log relative prices between initial and final years. Once again, we are particularly interested in whether the coefficient on the log change in relative prices is significantly different from one. When we run the regression for services on the entire sample of 18 countries we obtain a coefficient of .79 with a standard error of .20. For agriculture we obtain a coefficient of .72 with a standard error of .16. For the same reasons as earlier, the result for agriculture is affected by the anomalous behavior of Argentina and Venezuela. If we drop them the coefficient for agriculture increases to .91 with a standard error of .10. We again conclude that the patterns for changes in relative productivity and inverse relative prices are quantitatively similar.

7   Other Channels: Trade and Capital Accumulation

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. In this section we discuss factors that our model abstracts from which we believe may play an important role in some cases and which future research should seek to incorporate into the analysis: trade and capital accumulation.[27]

Our analysis focuses on a static closed economy. This raises the issue of how trade might matter, and in particular the role of dynamic trade imbalances. The amount of global trade within services continues to rise and is becoming more important in some countries. But because much of our data refers to the period prior to 2000 and is for countries in the early stages of development, we do not think trade in services is a quantitatively important consideration. For this reason we focus our comments on the potential effects of trade in agriculture and manufacturing.

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. The same comment would apply to trade that is balanced within agriculture.

Trade that is statically balanced in aggregate but not necessarily within each of agriculture and manufacturing 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, trade flows in agriculture tend to be relatively small for most countries, especially those that are in the early stages of development. Tombe (2015) documents that for poor countries, over 90 percent of agricultural consumption is sourced domestically. From a theoretical perspective there is an important question as to why countries that are relatively unproductive in agriculture

---

[27]We note that recent work in this vein has already moved in this direction. As noted in the introduction, Wise (2020) considers the role of trade and Sposi et al. (2020) considers both trade and capital accumulation.

do not simply import food from abroad, but from an empirical perspective the assumption of no net trade flow in agriculture is strongly supported by the data.[28]

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 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 15 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.[29]

Dynamics may also matter for another reason. Recent work by García-Santana et al. (2021) documents that final investment draws much more heavily from the manufacturing sector than does final consumption. 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. This raises the possibility that countries experiencing growth miracles featuring a period of relatively high investment may experience high peak employment shares in

---

[28]Tombe (2015) emphasizes the longer trade delays at the border in poor countries and the asymmetric effect that this has on perishable items. More generally, the transport and distribution of perishables requires a level of infrastructure that is not present in many poor countries.

[29]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.

Figure 15: Trade Surplus and Non-Agricultural Employment Shares: Asia
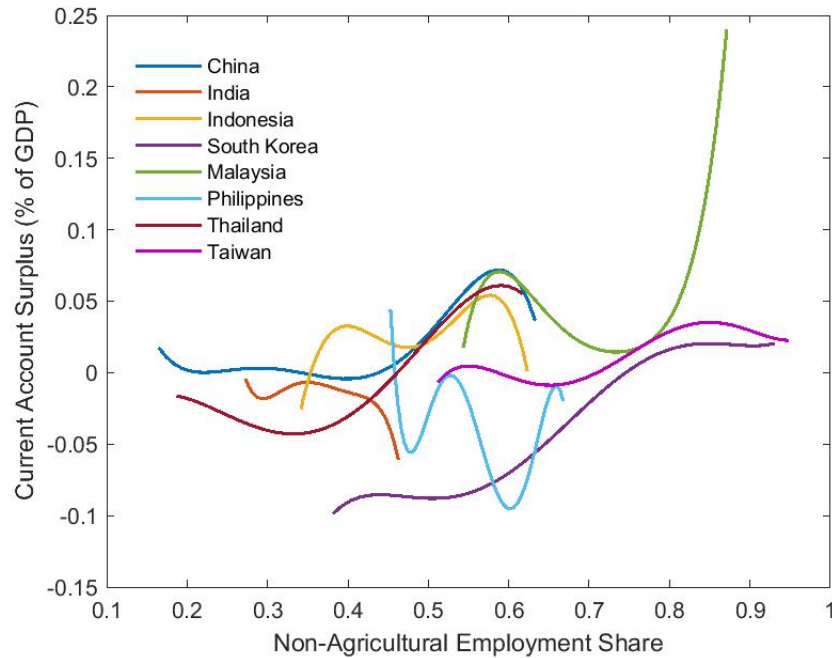
manufacturing, driven at least in part by investment rather than consumption. This is potentially more relevant for some of the Asian economies. Consistent with this, our analysis found that it was a set of high growth Asian economies that had higher peak employment shares than predicted solely by agricultural productivity dynamics.

Lastly, our analysis has abstracted from distortions that may impact sectoral labor allocations. Some of these might reflect distortions to the consumption/saving decisions and so relate to the previous discussion. But they may also impact the composition of consumption. For example, because 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 to include trade and capital accumulation, to focus on dynamic implications and to include the role of frictions, especially in the service sector. It will also be of interest to extend the analysis here to additional countries.

References

Adamopoulos, T. and D. Restuccia (2014, June). The size distribution of farms and international productivity differences. American Economic Review 104(6), 1667–97.

Alder, S., T. Boppart, and A. Müller (2022). A theory of structural change that can fit the data. American Economic Journal: Macroeconomics, Forthcoming.

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

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.

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.

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.

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).

Caselli, F. and W. J. Coleman II (2001). The u.s. structural transformation and regional convergence: A reinterpretation. Journal of Political Economy 109(3), 584–616.

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.

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.

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.

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.

Fujiwara, I. and K. Matsuyama (2022). A technology-gap model of premature deindustrialization. CEPR Discussion Paper 15530.

García-Santana, M., J. Pijoan-Mas, and L. Villacorta (2021). Investment demand and structural change. Econometrica 89(6), 2751–2785.

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

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

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

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.

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

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.

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

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

Huneeus, F. and R. Rogerson (2020, July). Heterogeneous paths of industrialization. Working Paper 27580, National Bureau of Economic Research.

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

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

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

Lagakos, D. and M. E. Waugh (2013, April). Selection, agriculture, and cross-country productivity differences. American Economic Review 103(2), 948–80.

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

Palma, J. G. (2005). Four sources of de-industrialisation and a new concept of the Dutch Disease. In: Jose Antonio Ocampo (Eds.), Beyond Reforms: structural dynamic and macroeconomic vulnerability, Stanford University Press and the World Bank.

Palma, J. G. (2014). De-industrialisation, 'premature' de-industrialisation and the dutch-disease. Revista NECAT 5(3), 7–23.

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.

Rodrik, D. (2013, 11). Unconditional convergence in manufacturing. The Quarterly Journal of Economics 128(1), 165–204.

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

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

Timmer, M. P., G. J. de Vries, and K. de Vries (2015). Patterns of structural change in developing countries. In J. Weiss and M. Tribe (Eds.), Routledge Handbook of Industry and Development, pp. 65 – 83. Routledge.

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

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

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

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: Philippines has a peak value below 0.20, whereas Argentina 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 Philippines 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
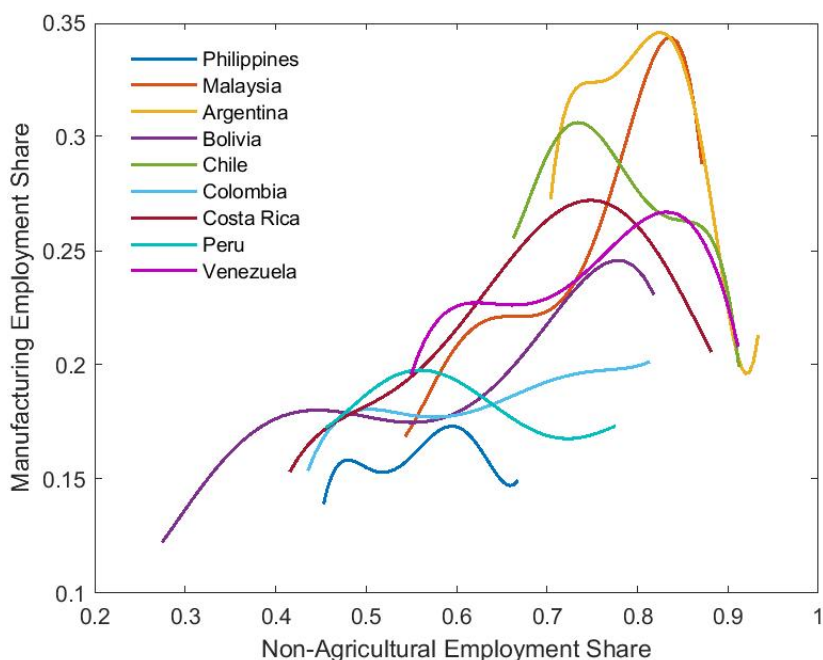
Appendix B   Proofs

B.1   *Proof of Proposition 1*

(i) Note that Equation (3.2) can be rewritten as:

$$h_{at} = \alpha_a + (1 - \alpha_a) \frac{\bar{c}_a}{A_{at}} + \alpha_a \frac{\bar{c}_s}{A_{st}}$$

Constant growth rates imply:

$$\dot{h}_{at} = -[(1 - \alpha_a) \bar{c}_a g_a e^{-g_a t} + \alpha_a \bar{c}_a g_s e^{-g_s t}]$$

The result follows directly.

(ii) To show this, combine Equations (3.3) and (3.4) to get:

$$h_{st} = \frac{(1 - \alpha_a) \left(\frac{A_{mt}}{A_{st}} \right)^{1-\sigma}}{\left[\frac{\alpha_m^\sigma}{(1-\alpha_m)^\sigma} + \left(\frac{A_{mt}}{A_{st}} \right)^{1-\sigma} \right]} \left[1 - \frac{\bar{c}_a}{A_{at}} + \frac{\bar{c}_s}{A_{st}} \right]$$

The result follows from the fact that under A1, the first term is monotone increasing, and under A2, the second term is monotone increasing.

(iii) To show this, write h_{mt} as:

$$h_{mt} = h_{m1t} \cdot h_{m2t}$$

where

$$h_{m1t} = \frac{(1 - \alpha_a) \alpha_m^\sigma}{\left[\alpha_m^\sigma + (1 - \alpha_m)^\sigma \left(\frac{A_{mt}}{A_{st}} \right)^{1-\sigma} \right]}$$

$$h_{m2t} = \left[1 - \frac{\bar{c}_a}{A_{at}} + \frac{\bar{c}_s}{A_{st}} \right]$$

Note that:

$$\frac{\dot{h}_{mt}}{h_{mt}} = \frac{\dot{h}_{m1t}}{h_{m1t}} + \frac{\dot{h}_{m2t}}{h_{m2t}}$$

Simple algebra gives:

$$\frac{\dot{h}_{m1t}}{h_{m1t}} = -\frac{g(1-\sigma)(1-\alpha_m)^\sigma e^{g(1-\sigma)t}}{\alpha_m^\sigma + (1-\alpha_m)^\sigma e^{g(1-\sigma)t}} \tag{B.1}$$

and

$$\frac{\dot{h}_{m2t}}{h_{m2t}} = \frac{\bar{c}_a g_a e^{-g_a t} - \bar{c}_s g_s e^{-g_s t}}{1 - \bar{c}_a e^{-g_a t} + \bar{c}_s e^{-g_s t}} \tag{B.2}$$

Assumption A1 implies that $\frac{\dot{h}_{m1t}}{h_{m1t}}$ is negative and decreasing, with asymptotic value $-g(1-\sigma)$. Assumption A1 also implies that $\frac{\dot{h}_{m2t}}{h_{m2t}}$ is asymptotically zero, and assumption A2 implies that it is positive and decreasing. Combining these properties gives the result. □

B.2 *Proof of Proposition 2*

(i) To prove this it is sufficient to examine how changes in the three productivity growth rates affect the profile for h_{mt}. Consider first the case of an increase in g_a. At any point in time t, this will imply unchanged values for both A_{mt} and A_{st} but a higher value for A_{at}. Equation (3.3) implies that an increase in A_a holding A_m and A_s constant will increase h_m, so that an increase in g_a implies that the h_{mt} profile will be higher at each point in time. It necessarily follows that the maximum value of h_{mt} must also be larger. The results for g and g_s follow similarly.

(ii) We start by noting that when h_m reaches its peak, it must be that $\dot{h}_m/h_m = 0$.

3

Defining h_{m1} and h_{m2} as previously, this requires:

$$\frac{\dot{h}_{m1}}{h_{m1}} + \frac{\dot{h}_{m2}}{h_{m2}} = 0$$

Substituting from Equations (B.1) and (B.2) this becomes:

$$\frac{g(1-\sigma)(1-\alpha_m)^\sigma e^{g(1-\sigma)t}}{\alpha_m^\sigma + (1-\alpha_m)^\sigma e^{g(1-\sigma)t}} = \frac{\bar{c}_a g_a e^{-g_a t} - \bar{c}_s g_s e^{-g_s t}}{1 - \bar{c}_a e^{-g_a t} + \bar{c}_s e^{-g_s t}} \tag{B.3}$$

This equation implicitly defines a function $t^*(g_a, g, g_s)$ that defines the point in time at which h_m^* is reached. Note that Assumption A1 implies that the left hand side is increasing in t and Assumption A2 implies that the right hand side is decreasing in t. Because the left hand side is increasing in g, it follows directly that t^* is decreasing in g.

Holding g_a and g_s constant, Equation (3.2) implies that the time profile for h_{at} is unaffected. A decrease in t^* thus implies a higher value of h_a^* and therefore a lower value of h_n^*.

(iii) When $\bar{c}_s = 0$, Equation (B.3) becomes

$$\frac{g(1-\sigma)(1-\alpha_m)^\sigma e^{g(1-\sigma)t}}{\alpha_m^\sigma + (1-\alpha_m)^\sigma e^{g(1-\sigma)t}} = \frac{g_a \bar{c}_a e^{-g_a t}}{1 - \bar{c}_a e^{-g_a t}} \tag{B.4}$$

Also, Equation (3.2) that gives the solution for h_{at} as a function of the productivities becomes:

$$h_{at} = \alpha_a + (1-\alpha_a)\bar{c}_a e^{-g_a t} \tag{B.5}$$

Using Equation (B.5) to substitute on the right hand side of Equation (B.4) gives:

$$\frac{g(1-\sigma)(1-\alpha_m)^\sigma e^{g(1-\sigma)t^*}}{\alpha_m^\sigma + (1-\alpha_m)^\sigma e^{g(1-\sigma)t^*}} = \frac{g_a(1-\alpha_a - h_n^*)}{h_n^*} \tag{B.6}$$

We can now use this equation to prove the result. Assume by way of contradiction that

4

h_n^* decreases when g_a increases. Then the right hand side of Equation (B.6) increases, and hence so must the left hand side. The left hand side increasing implies that t^*increases. But g_a increasing and t^* increasing together imply that h_a^* decreases, hence contradicting that h_n^* decreases. We conclude that h_n^* must increase.

The argument for the case when $g_s = 0$ is similar. When $g_s = 0$, Equation (B.3) becomes:

$$\frac{g(1-\sigma)(1-\alpha_m)^\sigma e^{g(1-\sigma)t^*}}{\alpha_m^\sigma + (1-\alpha_m)^\sigma e^{g(1-\sigma)t^*}} = \frac{g_a[1-(h_n^* + (1-\alpha_a)\bar{c}_s)]}{h_n^* + (1-\alpha_a)\bar{c}_s} \tag{B.7}$$

We again argue by way of contradiction. Specifically, assume that h_n^* decreases when g_a increases. The right hand side is increasing in g_a and decreasing in h_n^*, and so will increase. This implies that the left hand side must increase, which implies that t^* must increase, but this contradicts the fact that h_n^* decreases.

The result for g_s follows trivially, since g_s disappears from Equation (B.4) when $\bar{c}_s = 0$.

□