

NBER WORKING PAPER SERIES

GROWTH AND STRUCTURAL TRANSFORMATION

Berthold Herrendorf
Richard Rogerson
Ákos Valentinyi

Working Paper 18996
<http://www.nber.org/papers/w18996>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2013

We thank Francesco Caselli, Rachel Ngai, Todd Schoellman, and Gustavo Ventura for commenting on earlier versions of this manuscript and Christopher Herrington for excellent research assistance. For financial support, Herrendorf thanks the Spanish Ministry of Education (Grants ECO2009-11165 and ECO2012-31358), Rogerson thanks both the NSF and the Korea Science Foundation (WCU-R33-10005) and Valentinyi thanks the Hungarian Scientific Research Fund (OTKA) (Project K-105660-ny). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w18996.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2013 by Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Growth and Structural Transformation
Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi
NBER Working Paper No. 18996
April 2013
JEL No. E20,O40

ABSTRACT

Structural transformation refers to the reallocation of economic activity across the broad sectors agriculture, manufacturing and services. This review article synthesizes and evaluates recent advances in the research on structural transformation. We begin by presenting the stylized facts of structural transformation across time and space. We then develop a multi-sector extension of the one-sector growth model that encompasses the main existing theories of structural transformation. We argue that this multi-sector model serves as a natural benchmark to study structural transformation and that it is able to account for many salient features of structural transformation. We also argue that this multi-sector model delivers new and sharper insights for understanding economic development, regional income convergence, aggregate productivity trends, hours worked, business cycles, and wage inequality. We conclude by suggesting several directions for future research on structural transformation.

Berthold Herrendorf
W.P. Carey School of Business
Department of Economics
Arizona State University
Tempe, AZ 85287-3806
berthold.Herrendorf@asu.edu

Richard Rogerson
Woodrow Wilson School of
Public and International Affairs
323 Bendheim Hall
Princeton University
Princeton, NJ 08544
and NBER
rdr@princeton.edu

Ákos Valentinyi
Cardiff Business School
Cardiff University
Aberconway Building, Colum Drive
Cardiff CF10 3EU
United Kingdom
and CEPR
and Institute of Economics CRES HAS
valentinyi.a@gmail.com

1 Introduction

The one–sector growth model has become the workhorse of modern macroeconomics. The popularity of the one–sector growth model is at least partly due to the fact that it captures in a minimalist fashion the essence of modern economic growth, which Kuznets (1973) in his Nobel prize lecture described as the sustained increase in productivity and living standards. By virtue of being a minimalist structure, the one–sector growth model necessarily abstracts from several features of the process of economic growth. One of these is the process of structural transformation, that is, the reallocation of economic activity across the broad sectors agriculture, manufacturing and services. Kuznets listed structural transformation as one of the six main features of modern economic growth. Structural transformation has also received a lot of attention in the policy debate of developed countries where various observers have claimed that the sectoral reallocation of economic activity is inefficient, and calls for government intervention. Understanding whether structural transformation arises as an efficient equilibrium outcome requires enriching the one–sector growth model to incorporate multiple sectors. More generally, this raises the question whether incorporating multiple sectors will sharpen or expand the insights that can be obtained from the one–sector growth model. Several researchers have recently begun to tackle these questions, and the objective of this chapter is to synthesize and evaluate their efforts.¹

The first step in the broad line of research on structural transformation is to develop extensions of the one–sector growth model that are consistent with the “stylized facts” of structural transformation. Accordingly, we begin this chapter by presenting the stylized facts of structural transformation and then develop a multi–sector extension of the growth model that serves as a natural benchmark model to address the issue of structural transformation. Given the prominent role attributed to theories of balanced growth in the literature using the one–sector growth model, we start by asking whether it is possible to simultaneously deliver structural transformation and balanced growth. Recent work has identified several versions of the growth model that achieve this, and we present the results of this work in the context of our benchmark multi–

¹A different aspect of structural transformation that Kuznets also noted is the movement of the population from rural into urban areas, which is typically accompanied by the movement of employment out of agriculture.

sector model.

It turns out that the conditions under which one can simultaneously generate balanced growth and structural transformation are rather strict, and that under these conditions the multi-sector model is not able to account for the broad set of empirical regularities that characterize structural transformation. We therefore argue that the literature on structural transformation has possibly placed too much attention on requiring *exact* balanced growth, and that it would be better served by settling for *approximate* balanced growth instead. Put somewhat differently, we think that progress in building better models of structural transformation will come from focusing on the forces behind structural transformation without insisting on exact balanced growth. While the corresponding efforts to uncover the forces behind structural transformation are relatively recent, we describe some headway that has been made. We argue that the recent work suggests that the benchmark multi-sector model with approximate balanced growth is able to account for many salient features of structural transformation for the US, both qualitatively and quantitatively.

Armed with an extension of the one-sector growth model that incorporates structural transformation in an empirically reasonable fashion, we seek to answer the question of whether modeling structural transformation indeed delivers new or sharper insights into issues of interest. We argue that the answer to this question is yes, and we present several specific examples from the literature to illustrate this. These examples have in common that taking into account changes in the sectoral composition of the economy is crucial for understanding a variety of changes in aggregate outcomes. As we will see, this applies to important issues concerning economic development, regional income convergence, aggregate productivity trends, hours worked, business cycles, and wage inequality.²

²Matsuyama (2008) and Ray (2010) also review the literature on structural transformation (or structural change, as Ray calls it). In contrast to them, we devote a large part of our review to documenting the stylized facts of structural transformation and to assessing whether multi-sector extensions of the standard growth model can account for them. Greenwood and Seshadri (2005) review the literature on economic transformation, which refers to broader issues than structural transformation, for example changes in the patterns of fertility and the movement of women out of the household into the labor market.

2 The Stylized Facts of Structural Transformation

As mentioned in the introduction, structural transformation is defined as the reallocation of economic activity across three broad sectors (agriculture, manufacturing, and services) that accompanies the process of modern economic growth.³ In this section, we present the stylized facts of structural transformation. While a sizeable literature on the topic already exists, including the notable early contributions of Clark (1957), Chenery (1960), Kuznets (1966), and Syrquin (1988),⁴ we think that improvements in the quality of previous data and the appearance of new data sets make it worthwhile for us to summarize the current state of evidence.

Because the process of structural transformation continues throughout development, it is desirable to document its properties using relatively long time series for individual countries. The early studies that we cited above attempted to do this. However, the authors of these studies typically had to piece together data from various sources, necessarily creating issues about the comparability of their results across time and countries. In addition, the time period for which data was available was still relatively short. Recent efforts by various researchers to reconstruct historical data have increased the availability of appropriate long time series data for the purposes of documenting structural transformation. Although one still has to piece together data from different sources to generate long time series for most countries, time coverage has improved and compatibility is much less of a problem than it was in the past. We are going to use the Historical National Accounts Database of the University of Groningen as our primary historical data source, which we complement with several other data sources to increase the length of the periods covered.⁵

While it is conceptually desirable to examine changes for individual countries over long time series, and there is now more opportunity to do so, limiting attention to individual countries narrows the perspective unnecessarily. To begin with, it effectively restricts the set of

³We follow much of the literature and use the term manufacturing in this context to refer to all activity that falls outside of agriculture and services. It might seem to be more appropriate to refer to this category as industry, because manufacturing is just the largest component of it, but we prefer to reserve the term “industry” to refer to a generic production category.

⁴The list of subsequent papers is too large for us to attempt to include it in its entirety.

⁵Appendix A contains a detailed description about the historical data sources that we use. Many of them are also underlying the recent historical studies by Dennis and Iscan (2009) about structural transformation in the United States and by Alvarez–Cuadrado and Poschke (2011) about structural transformation in twelve industrialized countries including the United States.

countries that can be studied to those that are currently rich, and so it leaves open the question of whether currently poor countries show the same regularities that currently rich countries showed when they were poor a century or two ago. Limiting attention to long time series data has the additional disadvantage that despite major improvements in constructing historical time series, they typically do not reach the quality of the best data sets for recent years. Therefore, we document the regularities of structural transformation also for five data sets that cover a relatively large set of developing and developed countries during the last thirty or so years: the Benchmark Studies of the International Comparisons Program as reported by the Penn World Table (PWT), EU KLEMS, the National Accounts from the United Nations Statistics Division, the OECD Consumption Expenditure Data, and the World Development Indicators (WDI).⁶

2.1 Measures of Structural Transformation

Before presenting any data, it is useful to briefly note some aspects of measuring economic development and structural transformation.

The two most common measures of economic development at the aggregate level are GDP per capita and some measure of productivity (typically GDP per worker or GDP per hour, depending upon data availability), each expressed in international dollars. While these two measures often coincide, there are cases in which they differ. For example, several European economies have similar values of GDP per hour as the US, but GDP per capita can be as much as 25 percent lower than in the US because hours per adult are much lower. Without knowing the exact context of the issue being addressed, it is unclear whether one should categorize these European countries as equally or less developed than the United States.

Having raised this issue, in this chapter we choose to always measure the level of development by GDP per capita in 1990 international \$. Three reasons motivate this choice. First, in order to be able to identify threshold effects and the like, we insist on the comparability of the GDP numbers across different data sets, and GDP per capita is the only measure that is available for most countries and most of the time. Second, the standard models of structural

⁶We again refer the reader to Appendix A for the details regarding the data sets and how we use them to construct measures of economic activity at the sector level.

transformation take labor supply to be exogenous, implying that they abstract from differences in hours worked. Third, since some of the models that we will consider emphasize the role of income effects for structural transformation, it seems appropriate to characterize the patterns of sectoral reallocation conditional on income. Irrespective of these three reasons for using GDP per capita, we emphasize that most of our figures would look similar if instead we used one of the productivity measures when they are available.

We now turn to measuring structural transformation. The three most common measures of economic activity at the sectoral level are employment shares, value added shares, and final consumption expenditure shares. Employment shares are calculated either by using workers or hours worked by sector, depending on data availability. Value added shares and final consumption expenditure shares are typically expressed in current prices (“nominal shares”), but they may be expressed in constant prices also (“real shares”). While there is a tendency in the literature to view the different measures as interchangeable when documenting how economic activity is reallocated across sectors over time, one of the issues that we want to emphasize in this chapter is that they are in fact distinct. In particular, as we will discuss in detail later on, it is critical to be aware of the distinctions among the different measures when doing quantitative work because even when they display the same qualitative behavior, the quantitative implications may be quite different. Moreover, there are some striking cases in which they display differences even in the qualitative behavior.

Probably the most important reason for the differences between the measures of structural transformation is that employment shares and value added shares are related to production whereas final consumption expenditure shares are related to consumption. Production and consumption measures may display different behaviors because value added is not the same as final output.

A simple example will help to illustrate the distinction between value added and final goods that is relevant here. Consider the purchase of a cotton shirt from a retail establishment. Because the cotton shirt is a “good” as opposed to a “service”, in terms of final consumption expenditure, the entire expenditure will be measured as final consumption expenditure of the manufacturing sector. However, in terms of value added in production, the same purchase will be broken down

into three pieces: a component from the agricultural sector (i.e., the cotton that was used in making the shirt), a component from the manufacturing sector (i.e., the processing of the cotton and the production of the shirt), and a component from the service sector (i.e., the distribution and retail trade services where the shirt was purchased).

The end result of this is that although the same sectoral labels are used when disaggregating GDP into value added and final expenditure, the resulting measures of sectoral shares are conceptually distinct. It follows that both quantities and prices may differ between value added and final expenditure, implying that there is no reason to expect the implied shares to exhibit similar behavior. This will be of particular relevance when connecting models of structural transformation to the data, which we will discuss in detail below.


The previous discussion emphasized the difference between production and consumption measures. However, even the two measures that focus on production might contain different information. One example comes from Kuznets (1966), who showed for the early part of US development that the employment share of services increased considerably at the same time that the value added share of services remained almost constant.

Having emphasized that each of the three measures of economic activity at the sectoral level is distinct, we also want to note that each of them has its limitations as a singular measure. For the case of sectoral employment shares, a key issue is that employment may not reflect changes in “true” labor input, for example, because there are systematic differences in hours worked or in human capital per worker across sectors that vary with the level of development. For the case of value added and consumption expenditure shares, a key issue arises from the need to distinguish between changes in quantities and prices. This is often difficult empirically because reliable data on relative price comparisons across countries are hard to come by. In addition, consumption and production need not coincide because of the presence of investment and of imports and exports, so that neither measure alone is sufficient.

2.2 Production Measures of Structural Transformation

In this subsection we document the patterns of structural transformation based on examining production measures in several different data sets. We first review the available historical time

Figure 1: Sectoral Shares of Employment and Value Added – Selected Developed Countries 1800–2000



Source: Various historical statistics, see Appendix A.

series evidence for currently rich economies. We then turn to the evidence for currently rich and poor countries.

2.2.1 Evidence from Long Time Series for Currently Rich Countries

We construct individual time series of sectoral employment shares and value added shares over the 19th and 20th century for the following ten countries: Belgium, Finland, France, Japan, Korea, Netherlands, Spain, Sweden, United Kingdom and United States.⁷ Since the early data is sketchy and we want to highlight trends over long periods of time, we report the latest available observation for each decade, if any. We note that for these historical time series we only have measures based on production.

Figure 1 plots the historical time series. The vertical axis is either the share of employment or the share of value added in current prices in the three broad sectors of interest. The horizontal axis is the log of GDP per capita in 1990 international dollars as reported by Maddison. The figures clearly reveal what the literature views as the stylized facts of structural transformation. Over the last two centuries, increases in GDP per capita have been associated with decreases in both the employment share and the nominal value added share in agriculture, and increases in both the employment share and the nominal value added share in services. Manufacturing has behaved differently from the other two sectors: its employment and nominal value added shares follow a hump shape, that is, they are increasing for lower levels of development and decreasing for higher levels of development.

Figure 1 reveals several additional regularities that have been somewhat less appreciated in the context of structural transformation. First, focusing on the agricultural sector, we can see that for low levels of development the value added share is considerably lower than the employment share. This finding is interesting in light of the fact that countries which are currently poor tend to have most of their workers in agriculture although agriculture is the least productive sector.⁸ Second, focusing on the service sector, we see that both the employment share and the nominal value added share for the service sector are bounded away from zero even at very low levels of development; the lowest value added share of services is around 20% and the

⁷For a detailed description of the data sources, see the Appendix A.

⁸See Caselli (2005) and Restuccia et al. (2008) for evidence on this point.

lowest employment share is around 10%.⁹ Third, the figure for the nominal value added share in services suggests that there is an acceleration in the rate of increase when the log of GDP per capita reaches around 9.¹⁰ Inspecting the graphs for the other two nominal value added shares more closely, we also note that the nominal value added share for manufacturing peaks around the same log GDP at which the nominal value added share for the service sector accelerates, and so it appears that the accelerated increase in the value added share of services coincides with the onset of the decrease in the nominal value added share for manufacturing sector.¹¹

2.2.2 Evidence from Recent Panels for Currently Rich and Poor Countries

We now turn to an examination of production measures from several more recent data sets, which tend to be of higher quality than the historical data and which include also countries that are currently poor as well as additional variables (nominal versus real, hours versus employment). The goal of this subsection is to assess the stylized facts of structural transformation that we documented for the historical data, as well as to take advantage of the richer data available so as to examine additional dimensions of structural transformation.

Evidence from EU KLEMS

We start with EU KLEMS, which is compiled at the Groningen Growth and Development Center. The primary strength of EU KLEMS in documenting patterns in employment and value added shares is that it has the most complete information for all variables of interest, including sectoral hours worked, and that its value added data have been constructed from the national accounts of individual countries following a harmonized procedure that aims to ensure cross-country comparability.¹² The primary weakness of EU KLEMS is that its coverage is limited to countries with relatively high income; South Korea during the early 1970s is the country with

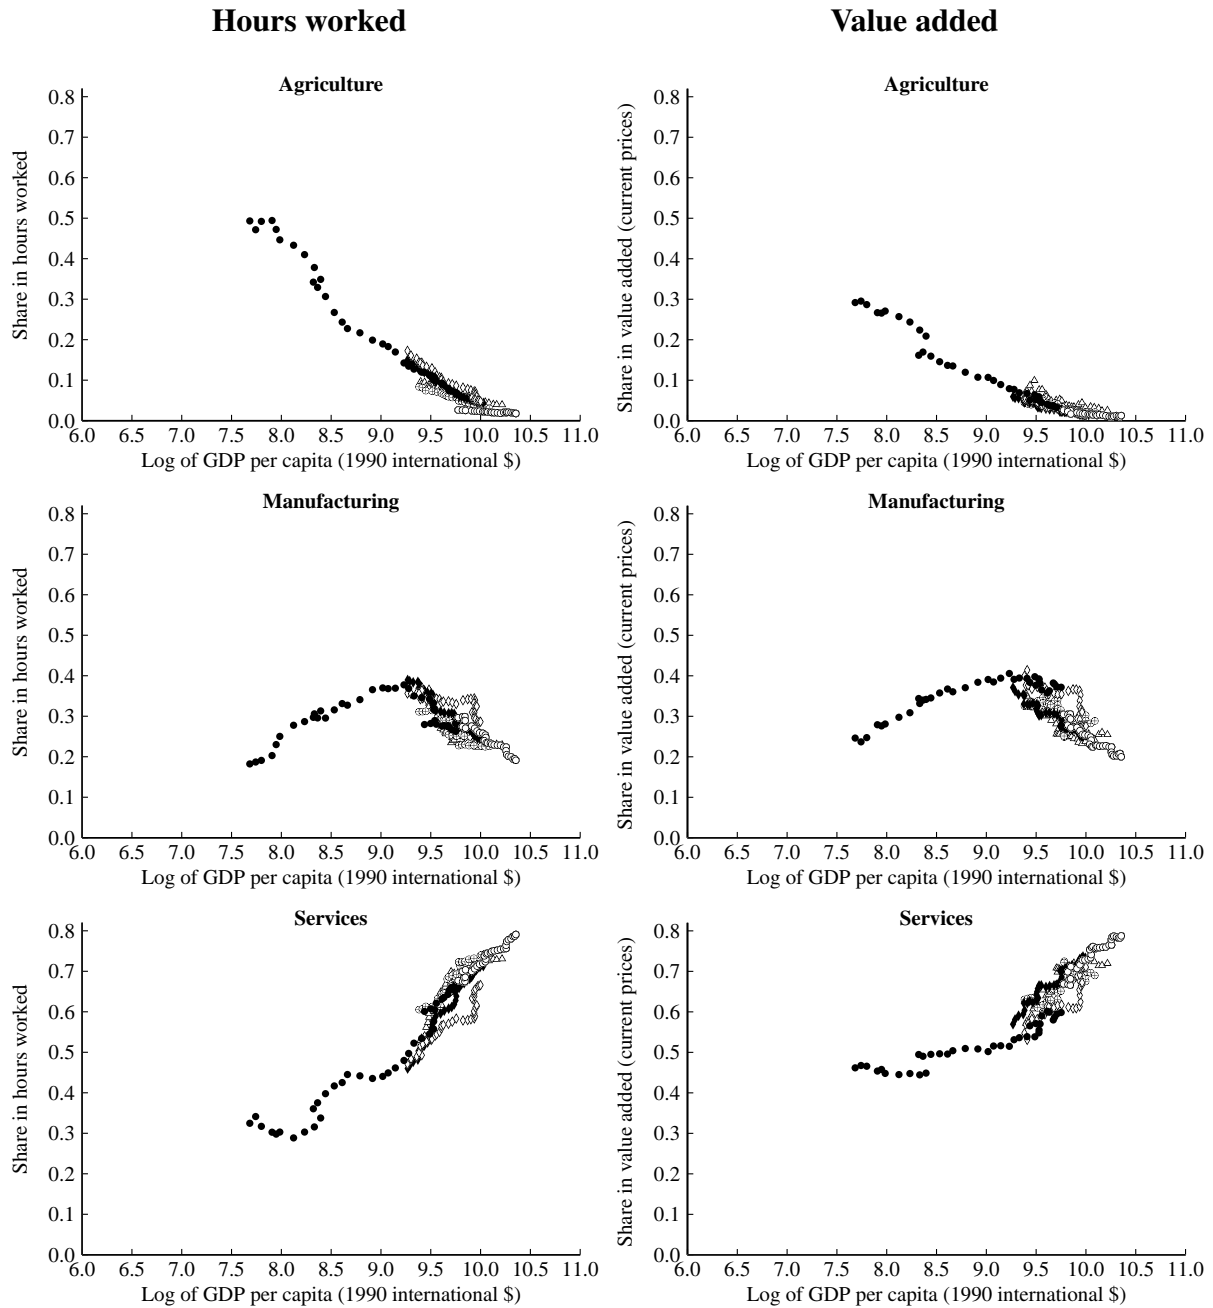
⁹This finding is confirmed by the historical study of Broadberry et al. (2011), who present evidence for England during the 14th century that the employment share of services was around 20%.

¹⁰See Buera and Kaboski (2012a, 2012b) for additional evidence on this point in a larger cross section of countries.

¹¹While we do not develop this issue further here, Buera and Kaboski (2012b) also show that at low levels of GDP per capita the manufacturing sector expands more quickly than does the service sector.

¹²For example, a common industry classification was used and price indices were constructed in a similar way across countries). For more detail see O'Mahony and Timmer (2009) and Timmer et al. (2010).


Figure 2: Sectoral Shares of Hours Worked and Nominal Value Added – 5 Non-EU Countries and Aggregate of 15 EU Countries from EU KLEMS 1970–2007



△△△ Australia ⊕⊕⊕ Canada ◆◆◆ 15 EU Countries ◇◇◇ Japan ●●● Korea ○○○ United States

Source: EU KLEMS, PWT6.3

Figure 3: Sectoral Shares of Hours Worked and Nominal Value Added – 15 EU countries from EU KLEMS 1970–2007



Source: EU KLEMS, PWT6.3

the lowest income in the sample.

We first document the evolution of the shares of sectoral hours worked and nominal value added as functions of the level of development for five non-European countries – i.e., Australia, Canada, Japan, Korea and the United States – as well as for the aggregate of fifteen EU countries.¹³ The data are plotted in Figure 2. The vertical axis is either the share of total hours worked or the share of value added in current prices in the three broad sectors of interest. As before, the horizontal axis is the log of GDP per capita in 1990 international dollars from Maddison.

The plots in Figure 2 confirm several patterns from the historical times series. First, the shares of hours worked and nominal value added for agriculture tend to decrease with the level of development for all countries, whereas the shares for services tend to increase with the level of development for all countries. Second, taken as a whole, the data are consistent with a hump shape for the shares in the manufacturing sector, although all countries except for Korea have decreasing manufacturing shares. Third, the series for both shares as a function of GDP per capita are quite consistent across countries. That is, not only are the qualitative patterns very similar, but so too are the quantitative patterns. This is of particular interest given the considerable attention that has been placed on the role of openness in the growth miracle of Korea (Korea liberalized its manufacturing trade starting in the 1960s and became one of the most open countries in the world). Although to a lesser extent, one could make similar statements for the case of Japan.

Although this last finding might tempt one to conclude that openness is not a quantitatively important determinant of sectoral allocations and structural transformation, we do want to caution the reader against jumping too quickly to this conclusion. Figure 3 shows the same series separately for the 15 EU countries. Although all countries display the same qualitative patterns, there is now substantial heterogeneity in the cross section at any given level of development. This is consistent with the view that these countries form a fairly integrated free-trade zone, thereby allowing for a high degree of specialization, and significant differences in how

¹³These are Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherland, Portugal, Spain, Sweden, and the United Kingdom.

economic activity is allocated across broad sectors.¹⁴

Next, we turn our attention to possible differences between real and nominal shares of sectoral value added, where nominal refers to current prices and real refers to constant prices. Kuznets (1966) concluded that the early available data showed similar qualitative patterns for nominal and real shares. We revisit this comparison because EU KLEMS has more recent and higher quality data than were available to Kuznets. Figure 4 plots the real shares of sectoral value added in the left panel and, for comparison, the nominal shares from Figure 2 in the right panel. The plots show that the qualitative patterns of real and nominal value added shares are fairly similar to each other, confirming what Kuznets found for the earlier data.

One important exception is Korea where the manufacturing share rose to half of real value added, which is considerably higher than in the other countries on the graph. At the same time, the manufacturing share of nominal value added flattened out around the maximum share for the other countries. Moreover, the real service share remained below the service share of the other countries, and actually fell somewhat. At the same time, the nominal service share stayed mostly flat. These observations imply that the price of manufacturing relative to total value added fell by more in Korea than in the other countries. This is consistent with the view that during Korea's massive trade liberalization the relative price of manufactured goods fell considerably at the same time as the real growth rate of manufacturing increased considerably.¹⁵

Evidence from the WDI and the UN Statistics Division


As previously noted, the main shortcoming of both the historical data and of EU KLEMS is that the coverage is limited to countries that have fairly high income today. It is therefore of interest to verify whether the stylized facts of structural transformation extend to data sets that cover countries that are poor today. The two obvious data sets to use in this context are the World Development Indicators (WDI) and the National Accounts that the United Nations Statistics Division collects.

We use the WDI for employment by sector, which it reports since 1980 based on the data

¹⁴Some of the series that we consider later on in this section will reveal differences between Korea and the other countries.

¹⁵Looking at sectoral employment shares, Bah (2008) documents that the process of structural transformation in many developing countries also looks different than the historical experiences of current rich countries.

**Figure 4: Sectoral Shares of Real and Nominal Value Added –
5 Non-EU Countries and Aggregate of 15 EU Countries from EU KLEMS 1970–2007**



△△△ Australia ⊕⊕⊕ Canada ◆◆◆ 15 EU Countries ◇◇◇ Japan ●●● Korea ○○○ United States

Source: EU KLEMS, PWT6.3

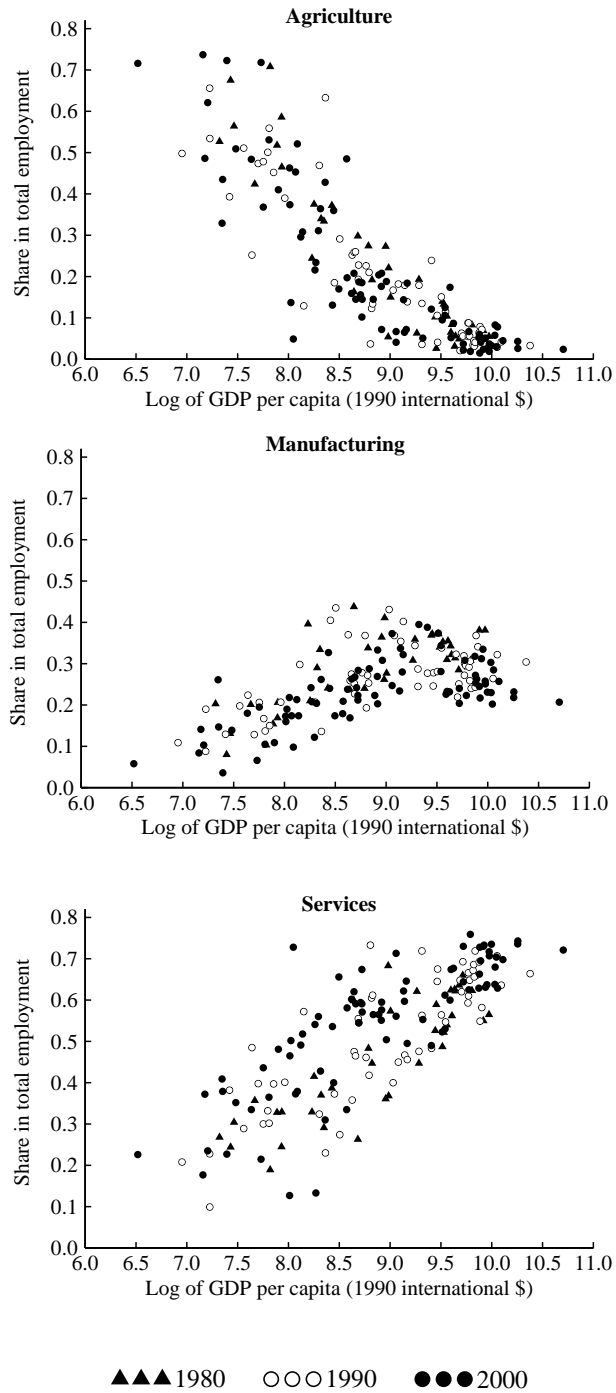
published by the International Labor Organization (ILO). We emphasize that these data are about employed workers instead of hours worked and are of considerably lower quality than those in EU KLEMS because there is much less harmonization underlying the construction of WDI data, which leads to comparability issues. Moreover, WDI employment data are not uniformly available over time for all countries. We use the national accounts of the United Nations Statistics Division for value added by sector. Unlike the WDI, the UN Statistics Division provides continuous coverage for a large number of countries between 1970 and 2007 and makes an explicit effort to harmonize the national accounts data so as to ensure that they are comparable across different countries.

Figure 5 plots the sectoral employment shares from the WDI against GDP per capita from Maddison. The plots confirm that in terms of sectoral employment shares the basic qualitative regularities of structural transformation also hold outside the set of rich countries for which EU KLEMS has data. Specifically, it is the case again that the agricultural employment share decreases in the level of development and that the employment share of services increases in the level of development. Moreover, the employment share in manufacturing is strongly increasing at lower levels of development (log of GDP per worker smaller than 9) before flattening out and then decreasing somewhat for higher levels of development. While this pattern is consistent with a hump shape, we note that the downward sloping part is not very pronounced in the WDI data.

Not surprisingly, the plots also show that employment shares do take on much more extreme values than can be found in EU KLEMS. For example, now the employment share of agriculture can be as high as 70% percent and the employment shares of manufacturing and services can be as low as only 10%. Lastly, for a given level of development the plots show much greater variability in the employment shares relative to what we found in the EU KLEMS data. The extent to which this simply reflects greater measurement error due to lack of comparability and other factors is an open question.


Figure 6 plots nominal value added shares by sector from the UN Statistics Division against GDP per capita from Maddison. Since these data have complete coverage for many rich and poor countries, they come close to a balanced panel. We therefore also plot the fitted nominal

Figure 5: Sectoral Shares of Employment – Cross Sections from the WDI 1980–2000



Source: World Development Indicators 2010

**Figure 6: Sectoral Shares of Nominal Value Added –
Cross Sections from UN National Accounts 1975–2005**



Source: National Accounts United Nations, PWT6.3, own calculations

value added shares from panel regressions. This is intended as a way of summarizing some patterns in the data, instead of as a way of testing any theory. For each sector we regress nominal value added shares on country fixed effects and the level, square, and cube of GDP per worker.¹⁶ We include countries for which no observations are missing, which were not communist, and which had more than a million inhabitants during 1970–2007. Details regarding the construction of the panel of countries and the regression results are contained in Appendix B.

The fitted curves reveal the same qualitative patterns that we have documented previously. It is of particular interest that the hump shape clearly emerges for manufacturing value added. Moreover, it is of interest that the fitted curve for services indicates an acceleration of the service share when the log of GDP per capita reaches a threshold value around 9 and the share of manufacturing value added peaks. Interestingly, this feature occurs at a similar threshold value also for the historical time series which we discussed above.

2.3 Consumption Measures of Structural Transformation

Lastly, we turn to the stylized facts of structural transformation when final consumption expenditure shares are used as a measure of economic activity at the sectoral level. We previously offered two main reasons why final consumption expenditure shares may exhibit different patterns than production value added shares: the presence of investment and of imports and exports and the fact that final consumption expenditure is a fundamentally distinct concept from value added produced. The goal of this subsection is to establish that these differences between consumption and production based measures do not matter much for agriculture and services, but can have important implications for manufacturing.

Comparable cross-country panel data on consumption expenditure by sector are much less available than such data on either employment or value added shares. We begin by presenting relatively long time series evidence for the US and the UK in Figure 7. The main message from the plots is that for these two countries, production and consumption measures display very similar behavior, both qualitatively and quantitatively. Specifically, nominal consumption


¹⁶We report results for a cubic polynomial since adding higher order terms did not have a significant effect on the fitted relationships.

shares for agriculture and services are decreasing and increasing over time, respectively, just as they were in the case for nominal value added shares, and the extent of the changes is quite similar too. Moreover, the consumption share for manufacturing displays a hump shape, just as it did in the case for the nominal value added share for manufacturing. Once again, the quantitative features are also similar, with the peak of the curves occurring at similar values of GDP per capita, and the extent of the decrease after the peak also being similar. One difference between consumption shares and value added shares is that the consumption share for manufacturing tends to be a few percentage points higher than the value added share for manufacturing. This occurs because of the fact that the consumption measure implicitly includes distribution services such as retail trade in its measure of manufacturing consumption.

We next consider two data sets on final consumption expenditure by sector: the OECD Consumption Expenditure Data Base and the Benchmark Studies of the International Comparisons Programme, as reported by the Penn World Table. The OECD data have reasonably long time series for several currently rich countries, namely, Australia, Canada, Japan, Korea, and the United States as well as the seven EU countries Austria, Denmark, Finland, France, Italy, Netherlands, and the United Kingdom. The Benchmark Studies offer relatively large cross sections for the years 1980, 1985, and 1996. We define the sectors for consumption expenditure following the usual conventions; for example, we use food as the category closest to agriculture; for the details see Appendix A. For each data set, we pool the data and plot the nominal consumption expenditure shares of the three sectors against GDP per capita measured in 1990 international dollars.


Figure 8 contains the plots for the OECD data and Figure 9 contains the plots for the Penn World Table data. Two patterns are immediate: the final expenditure share for food tends to decrease with the level of development while the final expenditure share for services tends to increase with development. These two patterns are qualitatively similar to the patterns that we have documented by using the production based measures of economic activity at the sectoral level. However, when we examine the plot for manufacturing consumption we now see some differences. Of particular interest is Korea; whereas it exhibits the same hump shape as the other OECD countries for the nominal production value added share of manufacturing, we see

Figure 7: Sectoral Shares of Nominal Consumption Expenditure – US and UK 1900–2008



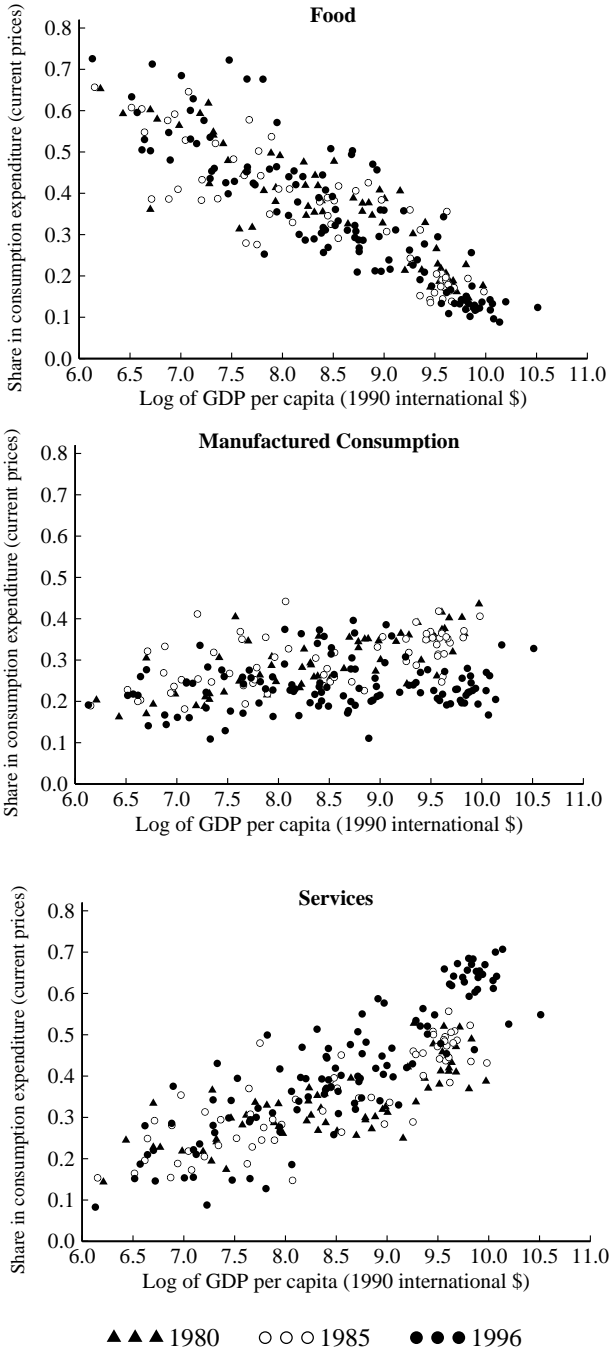
Source: Various historical statistics, see Appendix A.

Figure 8: Sectoral Shares of Nominal Consumption Expenditure – Various Countries, OECD 1970–2007



Source: OECD, EU KLEMS, PWT6.3

Figure 9: Sectoral Shares of Nominal Consumption Expenditure – Cross Sections from the ICP Benchmark Studies 1980, 1985, 1996



Source: International Comparisons Programme (as reported in PWT)

that its consumption share of manufacturing is virtually flat during a period of rapid growth.

The data from the PWT for the manufacturing consumption share effectively show a cloud. While this plot is not necessarily inconsistent with a hump shape for each country coupled with level differences across countries, it suggests that differences between production and consumption measures may be a more common feature of the data in the larger sample of countries. We think this is an important issue that merits further work. If the link between consumption and production measures is different for current developing countries than it was for countries that developed earlier, then this may well have implications for the nature of the development path that these countries follow.¹⁷

3 Modeling Structural Transformation and Growth

In this section we present a natural extension of the one–sector growth model that incorporates structural transformation. We develop our extension in two steps. In the first one, we consider the well known two–sector version of the growth model that has separate consumption and investment sectors. In the second step, we disaggregate consumption into the three components agriculture, manufacturing, and services.

3.1 Background: A Two–Sector Version of the Growth Model

Our presentation of the two–sector growth model closely resembles that in Greenwood et al. (1997), which is a version of Uzawa (1963). We assume an infinitely lived stand–in household with preferences over consumption sequences $\{C_t\}$ given by:

$$\sum_{t=0}^{\infty} \beta^t \log C_t \tag{1}$$

where $0 < \beta < 1$ is the discount factor. Note that, for simplicity, preferences are such that the household does not value leisure. The household is endowed with one unit of productive time and an initial stock of capital, K_0 .

¹⁷We are going to revisit this issue below when we discuss in detail our paper Herrendorf et al. (2013).

There are two constant–returns–to–scale production functions which describe how consumption (C) and investment (X) are produced from capital (k) and labor (n). It is convenient to follow the literature and impose that the production functions are Cobb–Douglas and have the same capital share:

$$C_t = k_{ct}^\theta (A_{ct} n_{ct})^{1-\theta}$$

$$X_t = k_{xt}^\theta (A_{xt} n_{xt})^{1-\theta}$$

where A_{it} represents exogenous labor–augmenting technological progress in sector i . We adopt the notational convention of using upper–case letters to refer to aggregate variables.

Capital accumulates as usual:

$$K_{t+1} = (1 - \delta)K_t + X_t$$

where $0 < \delta < 1$ denotes the depreciation rate.

We assume that capital and labor are freely mobile between the two sectors so that feasibility requires that in each period:

$$K_t = k_{ct} + k_{xt}$$

$$1 = n_{ct} + n_{xt}$$

As is standard, we study the competitive equilibrium for this economy. Although one can obtain the competitive–equilibrium allocations by solving a social planner’s problem, we want to emphasize the role of relative prices and therefore consider a sequence–of–markets competitive equilibrium in which the price of the investment good is normalized to be equal to one in each period. The price of the consumption good relative to the investment good is denoted by P_t , the rental rate for capital is denoted by R_t , and the wage rate is denoted by W_t . We assume that the household accumulates capital and rents it to firms.

We begin our characterization of the equilibrium by establishing that the capital–to–labor ratios are equalized across sectors at each point in time. To see this note that the first–order

conditions for the stand-in firm in sector $i \in \{c, x\}$ are given by:

$$R_t = P_t \theta \left(\frac{k_{ct}}{n_{ct}} \right)^{\theta-1} A_{ct}^{1-\theta} = \theta \left(\frac{k_{xt}}{n_{xt}} \right)^{\theta-1} A_{xt}^{1-\theta}$$

$$W_t = P_t (1 - \theta) \left(\frac{k_{ct}}{n_{ct}} \right)^{\theta} A_{ct}^{1-\theta} = (1 - \theta) \left(\frac{k_{xt}}{n_{xt}} \right)^{\theta} A_{xt}^{1-\theta}$$

Combining these two equations and rearranging gives an expression for the capital-labor ratio in sector $i \in \{c, x\}$:

$$\frac{k_{it}}{n_{it}} = \frac{\theta}{1 - \theta} \frac{W_t}{R_t}$$

It follows that the capital-to-labor ratio in each sector is the same and equals the aggregate capital-to-labor ratio:¹⁸

$$\frac{k_{ct}}{n_{ct}} = \frac{k_{xt}}{n_{xt}} = K_t \quad (2)$$

Next, we establish that the equilibrium value of the relative price P_t is pinned down by technology. To see this, divide the first-order conditions for labor from the two sectors by each other and use the fact that sectoral capital-to-labor ratios are equalized. This gives:

$$P_t = \left(\frac{A_{xt}}{A_{ct}} \right)^{1-\theta} \quad (3)$$

Equations (2) and (3) imply that

$$P_t C_t = \left(\frac{k_{ct}}{n_{ct}} \right)^{\theta} P_t A_{ct}^{1-\theta} n_{ct} = K_t^{\theta} A_{xt}^{1-\theta} n_{ct}$$

It follows that the model aggregates on the production side, that is, we can consider an aggregate production function that produces a single good that can be turned into either consumption or investment via a linear technology with marginal rate of transformation equal to P_t :

$$Y_t = X_t + P_t C_t = K_t^{\theta} (A_{xt})^{1-\theta} (n_{xt} + n_{ct}) = K_t^{\theta} A_{xt}^{1-\theta} \quad (4)$$

¹⁸To see this note that

$$\frac{k_{ct}}{n_{ct}} n_{ct} + \frac{k_{xt}}{n_{xt}} n_{xt} = K_t (n_{ct} + n_{xt}) = K_t.$$

Additionally, equation (2) and the first-order conditions for the firm in the investment sector imply that the marginal products of the aggregate production function determine the rental rate of capital and the wage rate:

$$R_t = \theta K_t^{\theta-1} A_{xt}^{1-\theta} \quad (5)$$

$$W_t = (1 - \theta) K_t^\theta A_{xt}^{1-\theta} \quad (6)$$

To characterize the competitive equilibrium further, we turn to the household side. The household's maximization problem is:¹⁹

$$\max_{\{C_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log C_t \quad \text{s.t.} \quad P_t C_t + K_{t+1} = (1 - \delta + R_t) K_t + W_t$$

Letting μ_t denote the current-value Lagrange multiplier on the period t budget equation, the first-order conditions for C_t and K_t are:

$$\frac{\beta^t}{C_t} = \mu_t P_t$$

$$1 - \delta + R_t = \frac{\mu_{t-1}}{\mu_t}$$

Combining these two equations gives the Euler equation:

$$\frac{1}{\beta} \frac{P_t C_t}{P_{t-1} C_{t-1}} = 1 - \delta + R_t \quad (7)$$

Using equations (4) and (5), equation (7) can be written as a second-order difference equation in the aggregate capital stock K_t . Given a value for the initial capital stock, this second-order difference equation together with a transversality condition determine the equilibrium sequence of capital stocks.

We are now ready to consider the possibility of a balanced growth path in this model. We

¹⁹Note that if total consumption grows at a constant rate γ_c , which will be the case below when we consider generalized balanced growth, then the household's objective function remains finite, and so is well defined. The reason for this is that

$$\sum_{t=0}^{\infty} \beta^t \log C_t = [\log C_0 + \log(1 + \gamma_c)] \sum_{t=0}^{\infty} \beta^t < \infty$$

start by assuming that both technologies improve at constant, though not necessarily equal, rates $\gamma_i > 0$:

$$\frac{A_{it+1}}{A_{it}} = 1 + \gamma_i, \quad i = c, x$$

The standard definition of balanced growth is that endogenous variables grow at constant rates. It turns out that this definition is too strict for models with structural transformation because the very nature of structural transformation is that the sectoral composition changes. We therefore follow the literature and use the weaker concept of *generalized balanced growth path* (GBGP), which only requires that the real interest rate is constant.

The motivation for requiring that the real interest rate be constant is that although it may exhibit short-term fluctuations, it does not show a long-term trend. This, of course, is one of the Kaldor facts. The next result shows that along a GBGP of our two-sector model the other four facts of Kaldor will also hold, that is, K_t and Y_t grow at constant rates and K_t/Y_t and $R_t K_t/Y_t$ are constant.

Proposition 1. *If a GBGP exists, then the Kaldor facts hold along the GBGP.*

Proof. Since R_t is constant along a GBGP, it suffices to show that K_t , Y_t and X_t all grow at rate γ_x .

The fact that R is constant and equation (5) holds in period t and $t + 1$ implies:

$$\frac{A_{xt+1}}{A_{xt}} = \frac{K_{t+1}}{K_t} \tag{8}$$

It follows that K_t also grows at the constant rate of γ_x . Using $Y_t = A_{xt}^{1-\theta} K_t^\theta$ we have:

$$\frac{Y_{t+1}}{Y_t} = \left(\frac{A_{xt+1}}{A_{xt}} \right)^{1-\theta} \left(\frac{K_{t+1}}{K_t} \right)^\theta \tag{9}$$

Using equation (8) this gives:

$$\frac{Y_{t+1}}{Y_t} = (1 + \gamma_x)^\theta (1 + \gamma_x)^{1-\theta} = 1 + \gamma_x \tag{10}$$

Constant growth of K necessarily implies constant growth of X . The fact that the aggregate

technology is Cobb–Douglas implies that factor shares are constant even off a GBGP. \square

If K_t grows at the constant rate γ_x , then the law of motion for capital implies that X_t must grow at the same constant rate. Equation (4) then implies that $P_t C_t$ must also grow at this same rate. Substituting this growth rate into equation (7) pins down the constant value of the rental rate of capital along a GBGP:

$$\frac{1}{\beta}(1 + \gamma_x) = 1 - \delta + R$$

Given a value for A_{x0} , using this version of the Euler equation and the condition on the equilibrium rental rate (5), we obtain the unique value of K_0 along a GBGP:

$$K_0 = \left[\frac{\beta\theta}{(1 + \gamma_x) - \beta(1 - \delta)} \right]^{\frac{1}{1-\theta}} A_{x0} \quad (11)$$

We note several features of this generalized balanced growth path. First, K_t and C_t grow at different rates along the GBGP. In particular, since (3) implies that P_t grows at gross rate $[(1 + \gamma_x)/(1 + \gamma_c)]^{1-\theta}$, and $P_t C_t$ grows at gross rate $(1 + \gamma_x)$, it follows that C_t grows at gross rate $(1 + \gamma_x)^\theta (1 + \gamma_c)^{1-\theta}$, i.e., a weighted average of the two sectoral growth rates in technology. Given that X_t grows at the same rate as both A_{xt} and K_t , it follows that sectoral employment and capital shares are constant along the balanced growth path. In other words, although in this model differential rates of technological progress lead to changes in relative prices of sectoral outputs, these price changes are not associated with any changes in factor allocations over time.

For future reference, it is of interest to note that although we assumed that technological progress in both sectors is constant over time, this is not required for the existence of a GBGP. In fact, because along the GBGP the difference in technological progress only shows up in prices and not in allocations, it follows that the same results would apply even if the growth rate of technological progress in the consumption sector varied over time. This would have no effect on how capital and labor are allocated and would only show up in the behavior of the relative price P_t . While in this case not all variables would grow at constant rates, it would still be true that the rental rate on capital would be constant and that Y_t and K_t would grow at the same constant rate. Thus, there would still be a GBGP.

3.2 A Benchmark Model of Growth and Structural Transformation

We use the model of the previous section as the starting point for our analysis of structural transformation in the context of the growth model.

3.2.1 Set up of the Benchmark Model

As in the previous section, we assume an infinitely lived stand-in household that has preferences characterized by (1) and is endowed with one unit of time and an initial capital stock. Different than in the previous section, we now assume that C_t is a composite of agricultural consumption (c_{at}), manufacturing consumption (c_{mt}) and service consumption (c_{st}):

$$C_t = \left[\omega_a^{\frac{1}{\varepsilon}} (c_{at} - \bar{c}_a)^{\frac{\varepsilon-1}{\varepsilon}} + \omega_m^{\frac{1}{\varepsilon}} (c_{mt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_s^{\frac{1}{\varepsilon}} (c_{st} + \bar{c}_s)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (12)$$

where $\bar{c}_i, \omega_i \geq 0$ and $\varepsilon > 0$. The functional form (12) is a parsimonious choice that allows us to capture two features on the demand side that are potentially important for understanding the reallocation of activity across these three sectors: how the demand of the household reacts to changes in income and in relative prices. In particular, the presence of the two terms \bar{c}_a and \bar{c}_s allows for the period utility function to be non-homothetic and therefore the possibility that changes in income will lead to changes in expenditure shares even if relative prices are constant. The parameter ε influences the elasticity of substitution between the three goods, and hence the response of nominal expenditure shares to changes in relative prices. Note, however, that in the above specification the elasticity of substitution is not equal to ε because it also depends on the non-homotheticity terms.

We generalize the previous model to allow for four Cobb–Douglas production functions, one for each of the three consumption goods and one for the investment good. Formally, the

production functions are given by:²⁰

$$c_{it} = k_{it}^\theta (A_{it} n_{it})^{1-\theta}, \quad i \in \{a, m, s\} \quad (13)$$

$$X_t = k_{xt}^\theta (A_{xt} n_{xt})^{1-\theta} \quad (14)$$

There is a tradition in the literature of working with only three production functions, with the assumption that all investment is produced by the manufacturing sector. Under this assumption, the output of the manufacturing sector can be used as either consumption or investment whereas the output of the other two sectors can only be used as consumption. We have not adopted this specification for two reasons. First, despite the apparent reasonableness of the claim that investment is to first approximation produced exclusively by the manufacturing sector, it turns out that this is not supported by the data. Moreover, such an assumption is becoming increasingly at odds with the data over time, due at least in part to the fact that software is both a sizeable and increasing component of investment, and most software innovation takes place in the service sector. In fact, total investment has exceeded the size of the entire manufacturing sector in the US since 2000. The second reason for considering a separate investment sector derives from evidence that technological progress in the investment sector has been more rapid than in the rest of the economy; see, for example Greenwood et al (1997). Because the possibility of differential rates of technological progress across sectors will play a key role in the subsequent analysis, we want to allow for the possibility that this rate is different in the investment sector.

Capital is accumulated as usual:

$$K_{t+1} = (1 - \delta)K_t + X_t$$

As before we assume that capital and labor are freely mobile.²¹ With four sectors, the feasibility

²⁰We follow much of the literature in abstracting from the differences between physical capital and land and treating land as part of physical capital. We then restrict attention to Cobb–Douglas production functions in capital and labor that have the same capital share in all sectors, which is analytically very convenient, because it implies that we can aggregate the sectoral production functions to an economy–wide Cobb–Douglas production function. In Section 5.1.2 below we will explore to which extent the assumption of equal sectoral capital shares is borne out by the data. For now, we just mention that even if one thinks that sectoral capital shares (where capital includes land) are similar, then there are still important applications for which it is crucial that land is a fixed factor. For such applications one needs to model land and physical capital separately.

²¹We discuss the case of restricted labor mobility in Section 7.2.

conditions now take the form:

$$K_t = k_{at} + k_{mt} + k_{st} + k_{xt}$$

$$1 = n_{at} + n_{mt} + n_{st} + n_{xt}$$

3.2.2 Equilibrium Properties of the Benchmark Model

We again consider a sequence-of-markets competitive equilibrium in which the price of the investment good is normalized to equal one in each period. The prices of the consumption goods relative to the investment good are denoted by p_{it} , $i \in \{a, m, s\}$. We again assume that the household accumulates capital and rents it to firms.

Several key properties of the two-sector model that we established above continue to hold in the four-sector model. Specifically, using the same logic as in the previous section, one can show that the capital-to-labor ratios are equalized across the four sectors at each point in time, and are equal to the aggregate capital-to-labor ratio:

$$\frac{k_{it}}{n_{it}} = K_t, \quad i = a, m, s, x \quad (15)$$

Moreover, as before, relative prices are determined by technology:

$$p_{it} = \left(\frac{A_{xt}}{A_{it}} \right)^{1-\theta}, \quad i = a, m, s \quad (16)$$

Using the above results, one can also show that our multi-sector model aggregates on the production side:

$$Y_t = p_{at}c_{at} + p_{mt}c_{mt} + p_{st}c_{st} + X_t = K_t^\theta A_{xt}^{1-\theta} \quad (17)$$

Lastly, the first-order conditions from the firm problems, (5) and (6), are still valid.

On the household side the model is more involved now. In particular, the household problem

now takes the form:

$$\begin{aligned} \max_{\{c_{at}, c_{mt}, c_{st}, K_{t+1}\}_{t=0}^{\infty}} \quad & \sum_{t=0}^{\infty} \beta^t \log \left[\omega_a^{\frac{1}{\varepsilon}} (c_{at} - \bar{c}_a)^{\frac{\varepsilon-1}{\varepsilon}} + \omega_m^{\frac{1}{\varepsilon}} (c_{mt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_s^{\frac{1}{\varepsilon}} (c_{st} + \bar{c}_s)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \\ \text{s.t.} \quad & p_{at}c_{at} + p_{mt}c_{mt} + p_{st}c_{st} + K_{t+1} = (1 - \delta + R_t)K_t + W_t \end{aligned}$$

In what follows, we show that this problem can be split into two subproblems: (i) how to allocate total income between total consumption and savings; (ii) how to allocate total consumption expenditure between the three consumption goods. We develop a useful representation in which the first subproblem closely resembles the problem of the household in the two–sector model considered previously.

In order to have a well defined household problem, we need to make sure that the consumption of agricultural goods will exceed the subsistence term \bar{c}_a in each period. Even if this is the case, a corner solution may still arise in which the household chooses zero consumption of services. For now, we assume that the household problem is well defined and that its solution is interior in all periods. In Proposition 2 below, we offer a formal condition to ensure that this is the case along the GBGP. Essentially this will boil down to requiring that in each period total consumption is “large enough” relative to the two terms \bar{c}_a and \bar{c}_s .

The first–order conditions for an interior solution for the three consumption categories are:

$$\frac{1}{C_t} \omega_a^{\frac{1}{\varepsilon}} (c_{at} - \bar{c}_a)^{-\frac{1}{\varepsilon}} C_t^{\frac{1}{\varepsilon}} = \lambda_t p_{at} \quad (18)$$

$$\frac{1}{C_t} \omega_m^{\frac{1}{\varepsilon}} (c_{mt})^{-\frac{1}{\varepsilon}} C_t^{\frac{1}{\varepsilon}} = \lambda_t p_{mt} \quad (19)$$

$$\frac{1}{C_t} \omega_s^{\frac{1}{\varepsilon}} (c_{st} + \bar{c}_s)^{-\frac{1}{\varepsilon}} C_t^{\frac{1}{\varepsilon}} = \lambda_t p_{st} \quad (20)$$

where λ_t denotes the current–value Lagrange multiplier on the budget constraint in period t . If one raises each of the equations (18)–(20) to the power $1 - \varepsilon$, adds them, and uses the definition (12) of C_t , then one obtains:

$$\frac{1}{C_t} = \lambda_t \left[\omega_a (p_{at})^{1-\varepsilon} + \omega_m (p_{mt})^{1-\varepsilon} + \omega_s (p_{st})^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}} \quad (21)$$

Given that λ_t is the marginal value of an additional unit of expenditure in period t , it follows that the other term on the right-hand side is naturally interpreted as the price of a unit of composite consumption. In view of this we will define the price index P_t by:

$$P_t \equiv \left[\omega_a (p_{at})^{1-\varepsilon} + \omega_m (p_{mt})^{1-\varepsilon} + \omega_s (p_{st})^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}} \quad (22)$$

If one adds the three first-order conditions (18)–(20) and uses this definition of P_t , one also obtains:

$$p_{at}c_{at} + p_{mt}c_{mt} + p_{st}c_{st} = P_t C_t + p_{at}\bar{c}_a - p_{st}\bar{c}_s \quad (23)$$

It follows that the household's maximization problem can be broken down into two subproblems:

(i) Intertemporal Problem. Allocate total income among the composite consumption good and savings:

$$\max_{\{C_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log C_t \quad \text{s.t.} \quad P_t C_t + K_{t+1} = (1 - \delta + r_t)K_t + w_t - p_{at}\bar{c}_a + p_{st}\bar{c}_s$$

(ii) Static Problem. Allocate the period t consumption expenditure $P_t C_t$ among the three consumption goods:

$$\max_{c_{at}, c_{mt}, c_{st}} \left[\omega_a^{\frac{1}{\varepsilon}} (c_{at} - \bar{c}_a)^{\frac{\varepsilon-1}{\varepsilon}} + \omega_m^{\frac{1}{\varepsilon}} (c_{mt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_s^{\frac{1}{\varepsilon}} (c_{st} + \bar{c}_s)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}$$

$$\text{s.t.} \quad p_{at}c_{at} + p_{mt}c_{mt} + p_{st}c_{st} = P_t C_t + p_{at}\bar{c}_a - p_{st}\bar{c}_s$$

This representation nicely separates out the growth component of the model from the structural transformation component of the model. From the perspective of balanced growth in the aggregates K_t and C_t , the representation looks like the two-sector growth model with the exception of one detail: this economy behaves as if there is a time varying endowment, reflected by the term $-p_{at}\bar{c}_a + p_{st}\bar{c}_s$. If this endowment happens to be zero at all dates, then the equivalence to a standard two-sector model is exact. Be that as it may, the Euler equation is still of the form (7). Moreover, although the expression for the relative price P_t is somewhat more complicated

in the current setting compared to the two–sector model, the equilibrium value of this relative price can still be determined directly from primitives without solving for the full equilibrium.

From the perspective of structural transformation, the above representation implies that we can focus on the solution to the static problem of allocating each period’s consumption expenditure between the three consumption goods. The first–order conditions (18)–(20) characterize the solution to this static problem. For future reference, we note two useful implications of the first–order conditions. First, they impose conditions on the ratios of any two consumption goods:

$$\left(\frac{p_{at}}{p_{mt}}\right)^\varepsilon \frac{c_{at} - \bar{c}_a}{c_{mt}} = \frac{\omega_a}{\omega_m} \quad (24)$$

$$\left(\frac{p_{st}}{p_{mt}}\right)^\varepsilon \frac{c_{st} + \bar{c}_s}{c_{mt}} = \frac{\omega_s}{\omega_m} \quad (25)$$

Second, they impose a condition on the ratio of the expenditure on composite consumption and the expenditure on manufactured consumption:

$$\frac{P_t C_t}{p_{mt} c_{mt}} = \left[\frac{\omega_a}{\omega_m} \left(\frac{A_{mt}}{A_{at}}\right)^{(1-\theta)(1-\varepsilon)} + 1 + \frac{\omega_s}{\omega_m} \left(\frac{A_{mt}}{A_{st}}\right)^{(1-\theta)(1-\varepsilon)} \right] \quad (26)$$

Equations (24)–(26) will play a key role below when we study the details of structural transformation within the framework of our four–sector model.

3.3 Connecting the Benchmark Model to Measures of Structural Transformation

Since we will eventually ask whether versions of this model can help us understand the stylized facts of structural transformation that we documented in Section 2, it is relevant to briefly discuss some issues related to how one connects the model just described to the various measures from the data that we have previously examined. While this might appear obvious, there are a couple of issues that require notice.

In Section 2, we disaggregated total value added into the value added of agriculture, manufacturing, and services and measured the shares of these three sectors in total value added. To

connect our model with these measures of sectoral activity, it is natural to assume that the sectoral production functions that we have specified in the benchmark model represent value added production functions. However, because we have modeled the investment sector as a separate sector, one also needs to allocate the value added from the investment sector among the other three sectors. The literature often assumes that the entire value added of the investment sector belongs to manufacturing. This assumption is inconsistent with the data, for the simple reason that in recent years in the United States the value added of the investment sector has exceeded the value added of the (total) manufacturing sector. An alternative is to allocate investment value added to the other sectors using constant shares. This is also at odds with the data, since as shown in Herrendorf et al. (2013), the increasing importance of software as a component of investment has led to an increase in the share of investment value added occurring in the service sector. Nonetheless, since it serves to facilitate transparency, we will adopt this alternative as a benchmark in the next section when we discuss the qualitative features of balanced growth paths in different special cases of the model. However, it should be kept in mind that movements in the sectoral distribution of investment value added shares could affect the predictions that we highlight. As a practical matter, while this effect can matter, it is probably not so relevant at the quantitative level because total investment is a relatively small share of GDP.

The second issue concerns how to connect the model with production value added data versus consumption expenditure data. Specifically, assuming that the sector production functions are interpreted as value added production functions leads to a difficulty when trying to connect the model with data on consumption expenditure shares. Because equilibrium requires that $c_{it} = k_{it}^{\theta}(A_{it}n_{it})^{1-\theta}$, it would seem natural to identify $p_{it}c_{it}/\sum_j p_{jt}c_{jt}$ as the model's measure of the nominal consumption share of sector i in period t . However, this share is not the appropriate measure for the nominal consumption expenditure share of sector i as measured in the data. To see why, let us return to the example discussed earlier of the purchase of a cotton shirt. To measure the contribution of this shirt to manufactured final consumption expenditure, we need to aggregate all value added that goes into the production of the shirt through the use of intermediate inputs from each of the three sectors. This requires us to take into account the input–output relationships about how value added is aggregated into final consumption expen-

diture. In contrast, the above definition of consumption shares includes only the value added that came from the manufacturing sector itself, and so it does not reflect how final consumption expenditure is measured in a world in which each sector uses intermediate inputs from the other sectors.

To avoid this problem, one could alternatively assume that $p_{it}c_{it}/\sum p_{jt}c_{jt}$ in the model does correspond to the nominal consumption expenditure share of sector i in period t as measured in the data. But since in equilibrium $c_{it} = k_{it}^{\theta}(A_{it}n_{it})^{1-\theta}$, it would then follow that $p_{it}k_{it}^{\theta}(A_{it}n_{it})^{1-\theta}$ is not an appropriate measure of value added from sector i in period t as measured in the data. Returning to the shirt example, this piece of c_{mt} now reflects the value added components from each of the three sectors that went into producing the final product, and so it cannot be the value added from one particular sector. In order to maintain consistency, it must be that the production functions summarize the labor and capital from the various stages of production that are used to produce final consumption expenditure. In order to obtain value added shares one would have to use (inverse) input–output relationships to unbundle the final consumption expenditure into its value added components. Moreover, since n_{it} now reflects all of the labor that went into producing the shirt at each of the various stages of production, it is no longer the case either that n_{it} is an appropriate measure of the employment share of sector i in period t .

The bottom line from this discussion is that if one wants to have a model that can simultaneously address the shares of sectoral employment, value added, and consumption expenditure, then one will need to explicitly include the details of the input–output structure involved in transforming sectoral value added into sectoral consumption expenditure. We have chosen not to do this in order to preserve a greater degree of transparency in the presentation. In view of this, we need to keep in mind that when we discuss the model implications for the measures of structural transformation, we can either connect the production measures (employment shares and value added shares) to the data, implying that the consumption measure (consumption expenditure shares) does not have a close empirical counterpart, or we can connect the consumption measure to the data implying that the two production measures do not have close empirical counterparts. Whichever way we choose, our model will not be able to make statements about all three measures of structural transformation at the same time. Moreover, as we

discuss later on in more detail, one should not assume that preference and technology parameters are invariant to the interpretation that one imposes on the model objects.

4 The Economic Forces Behind Structural Transformation: Theoretical Analysis

The Kaldor facts regarding balanced growth over long periods of time have led the profession to focus on specifications of the one–sector neoclassical growth model that generate balanced growth. The evidence that we presented in Section 2 suggests that the continuing process of reallocation of activity across sectors coexists with the stable behavior of aggregate variables that characterizes balanced growth. It is therefore perhaps not surprising that the theoretical literature on structural transformation has looked for specifications of the previous model that give rise to a generalized balanced growth path along which structural transformation occurs. We begin this section by summarizing the results of this theoretical literature and its predictions for the nature of structural transformation. We close this section with a discussion of whether the focus on specifications that deliver exact balanced growth might be too stringent. Irrespective of whether this is the case, we believe that the search for specifications that deliver balanced growth and structural transformation has proven useful in helping researchers isolate various forces that are potentially important in shaping structural transformation.

4.1 Two Special Cases with Analytical Solutions

Our previous derivations put us in position to easily summarize recent findings in the literature about the joint possibility of generalized balanced growth and structural transformation. In this subsection we focus on two recent papers that emphasize different economic forces behind structural transformation, notably Kongsamut et al. (2001) and Ngai and Pissarides (2007).

4.1.1 Preliminaries

If we are to look for a balanced growth path it is natural to limit ourselves to situations in which technological progress is constant. We therefore assume:

$$\frac{A_{it+1}}{A_{it}} = 1 + \gamma_i, \quad i = a, m, c, x \quad (27)$$

As previously noted, even if all aggregates grow at constant rates, it will typically not be the case that all sector-level variables grow at constant rates. We therefore follow the literature and focus on generalized balanced growth paths (GBGP), which are defined to be equilibrium paths along which the rental rate of capital is constant, i.e., $R_t = R$.

We next turn to the issue of whether there are specifications of the model for which a GBGP exists along which structural transformation occurs. At this stage we will simply pose this question from a qualitative perspective. Specifically, we will say that a GBGP exhibits structural transformation if either sectoral employment shares (n_{it}) or sectoral value added (or consumption expenditure) shares ($p_{it}c_{it}/Y_t$) are not constant for all three consumption sectors. The issue of generating the “right” patterns of structural transformation, both qualitatively and quantitatively, will be taken up later.

As a starting point it is useful to examine two special cases. The first special case makes the extreme assumption that the three consumption goods are perfect substitutes: $\bar{c}_a = \bar{c}_s = 0$, $\omega_a = \omega_m = \omega_s$, $A_{at} = A_{mt} = A_{st}$, and $\varepsilon \rightarrow \infty$. In this case the model is identical at the aggregate level to the two-sector model in the previous section, and so it has a unique balanced growth path in terms of C_t and K_t . However, since the three consumption goods are perfect substitutes and have identical production functions, the allocation of labor and capital between the three sectors is indeterminate, beyond the restriction that capital-to-labor ratios must be the same in all sectors with positive output. Because of this indeterminacy it is obviously the case that one can accommodate whatever patterns one desires in terms of changes in either labor allocations or value added shares across sectors. However, since, as we have seen in Section 2 above, the features of structural transformation appear to be stable over time and across countries, this does not seem a very appealing way to account for structural transformation.

The second special case of interest assumes that $\bar{c}_a = \bar{c}_s = 0$ and $\varepsilon = 1$, so that the preference aggregator is Cobb–Douglas. We do not present the details here, but one can show that the unique balanced growth path has constant sectoral labor and value added shares. This happens despite the fact that we have not restricted the relative rates of productivity growth among the three consumption sectors. Intuitively, with Cobb–Douglas preferences, employment and value added shares are independent of relative productivities. With sectoral employment and capital shares fixed, differences in relative productivities generate differences in relative outputs, but these differences in output are perfectly offset in terms of value added shares by changes in relative prices. While this special case gives rise to balanced growth and avoids the indeterminacy of the previous case, it does not give rise to structural transformation along the balanced growth path.

In what follows we describe two scenarios that can generate structural transformation along a GBGP. Each of them can be understood as a departure from this second special case.

4.1.2 Case 1: Income Effects and Structural Transformation

Case 1 corresponds to the analysis found in Kongsamut et al. (2001) and represents the extreme scenario in which all structural change is driven by *income effects* that are generated by the non–homotheticity terms \bar{c}_a and \bar{c}_s when income changes but relative prices remain the same. For this case we assume that technological progress is uniform across all consumption sectors ($\gamma_i = \gamma_j$ for all $i, j = a, m, s$) and that the parameter governing the elasticity of substitution among consumption goods is unity ($\varepsilon = 1$).²² The consumption aggregator (12) then takes the well known Stone–Geary form:

$$C_t = \omega_a \log(c_{at} - \bar{c}_a) + \omega_m \log(c_{mt}) + \omega_s \log(c_{st} + \bar{c}_s) \quad (28)$$

With \bar{c}_a and \bar{c}_s positive it is easy to see intuitively how one may get structural transformation along a GBGP; as income grows, the non–homotheticity of the demands for the different consumption goods will lead to changes in the value added shares. However, there is a potential issue in obtaining generalized balanced growth when \bar{c}_a and \bar{c}_s are positive. To see this, recall

²²Note that ε equals the elasticity of substitution only if $\bar{c}_a = \bar{c}_s = 0$.

the Euler equation (7) for the household problem. From this equation, if R_t is constant over time, then it must be that $P_t C_t$ grows at a constant rate. From the period–budget equation, (23), and noting that factor payments are equal to output, we have:

$$P_t C_t + p_{at} \bar{c}_a - p_{st} \bar{c}_s = K_t^\theta A_{xt}^{1-\theta} + (1 - \delta)K_t - K_{t+1} \quad (29)$$

Since the right–hand side grows at rate γ_x , $P_t C_t + p_{at} \bar{c}_a - p_{st} \bar{c}_s$ must also grow at rate γ_x . If $p_{a0} \bar{c}_a - p_{s0} \bar{c}_s$ is not zero, then $p_{at} \bar{c}_a - p_{st} \bar{c}_s$ will grow at rate γ_x only if relative prices also grow at rate γ_x . However, this contradicts the fact that p_{at} and p_{st} both grow at gross rate $[(1 + \gamma_{xt})/(1 + \gamma_{ct})]^{1-\theta}$, which is implied by expression (16). Hence, balanced growth requires that $p_{a0} \bar{c}_a - p_{s0} \bar{c}_s = 0$, which is equivalent to:

$$\frac{\bar{c}_a}{\bar{c}_s} = \left(\frac{A_{a0}}{A_{s0}} \right)^{1-\theta} \quad (30)$$

Note that since both relative prices grow at the same rate, this condition implies that $p_{at} \bar{c}_a - p_{st} \bar{c}_s = 0$ at all dates t .²³

Given condition (30), equation (29) simply requires that $P_t C_t$ grows at rate γ_x . From the perspective of balanced growth this economy then looks very much like the two–sector model that we considered in the previous section. In particular, similar to that two–sector model, the share of labor and capital devoted to consumption versus investment is constant along a GBGP.

We make two remarks regarding condition (30). First, note that if either of \bar{c}_a or \bar{c}_s is positive, then they must both be positive. As we discuss in a later section, many papers have implicitly assumed that $\bar{c}_a > 0$ and $\bar{c}_s = 0$, which is inconsistent with condition (30). Second, this condition relates the parameters of preferences and technology to each other, and is therefore somewhat of a “fragile” condition. We shall return to this point later in this section.

Next we consider whether structural transformation occurs along the GBGP. To examine

²³This point illustrates that the assumption of the same rate of technological progress in the agriculture and service sectors is a necessary condition and not merely a simplification.

this note that if $\varepsilon = 1$, then (24)–(25) imply the Stone–Geary demand system:

$$c_{at} = \omega_a \frac{P_t C_t}{p_{at}} + \bar{c}_a \quad (31)$$

$$c_{mt} = \omega_m \frac{P_t C_t}{p_{mt}} \quad (32)$$

$$c_{st} = \omega_s \frac{P_t C_t}{p_{st}} - \bar{c}_s \quad (33)$$

Moreover, the assumption that all consumption sectors grow at the same rate implies that the relative prices of the three consumption goods are constant:

$$\frac{p_{it}}{P_t} = \frac{p_{i0}}{P_0}, \quad i \in \{a, m, s\}$$

Hence c_{at} , c_{mt} , and c_{st} grow at a slower rate, at the same rate, and at a faster rate than C_t , respectively. Given that the relative prices of the three consumption goods are constant, it follows that $p_{it}c_{it}/P_t C_t$ is decreasing for agriculture, constant for manufacturing and increasing for services. Since total consumption expenditures are a constant share of total output, it follows that these properties also carry over to both n_{it} and $p_{it}c_{it}/Y_t$.

In summary, and more formally, we have the following result:

Proposition 2. *Assume that condition (30) holds and that*

$$\bar{c}_s \leq \omega_s \left(\frac{A_{s0}}{A_{x0}} \right)^{1-\theta} \left[K_0^\theta A_{x0}^{1-\theta} - (\gamma_x + \delta)K_0 \right] \quad (34)$$

where K_0 is given by (11).

Then there is a unique GBGP. Along the GBGP, the employment and nominal value added shares of the investment sector are constant. The employment and nominal value added shares are decreasing for agriculture, constant for manufacturing and increasing for services.

Proof. We start by noting that it is straightforward to show that (11) implies that $K_0^\theta A_{x0}^{1-\theta} > (\gamma_x + \delta)K_0$. Hence, $P_0 C_0 = K_0^\theta A_{x0}^{1-\theta} - (\gamma_x + \delta)K_0 > 0$ and condition (34) is well defined. Condition (34) ensures that the right-hand side of (33) is positive at $t = 0$. Since the economy grows while relative prices remain constant, this implies that the right-hand side is positive for

all t . In this case, equations (31)–(33) are well defined and they have a unique interior solution for c_{at}, c_{mt}, c_{st} . The existence of a unique GBGP and the statements about the shares then follow directly from the previous discussion. \square

4.1.3 Case 2: Relative Price Effects and Structural Transformation

The second case that we consider corresponds to the analysis found in Ngai and Pissarides (2007).²⁴ Whereas the previous case generated structural transformation purely via income changes and asked whether this could be consistent with balanced growth, Ngai and Pissarides consider the polar extreme case in which structural transformation is generated purely from changes in relative prices and ask whether this can be consistent with balanced growth. Accordingly, they assume that $\bar{c}_a = \bar{c}_s = 0$. In order to have relative price changes operating it is clearly necessary to have differential rates of technological progress among the three consumption goods sectors, so no restrictions will be placed on the relative values of γ_i . Given our earlier discussion, however, we know that ε will have to take on a value other than unity.

The analysis of this case follows directly from our analysis of the two–sector model. Specifically, if the values of $\gamma_a, \gamma_m,$ and γ_s are different, then the price index P_t will not grow at a constant rate. However, as noted at the end of the section on the two–sector model, this has no bearing on the existence of a unique GBGP; there still is a unique GBGP that features a constant share of labor and capital allocated to total consumption. Along the GBGP the value of $P_t C_t$ will grow at the constant rate γ_x even though neither component grows at a constant rate.

To assess the implications for structural transformation we again turn to equations (24) and (25). Using equation (16) for relative prices, these two equations can now be written as:

$$\frac{c_{at}}{c_{mt}} = \frac{\omega_a}{\omega_m} \left(\frac{A_{at}}{A_{mt}} \right)^{\varepsilon(1-\theta)} \quad (35)$$

$$\frac{c_{st}}{c_{mt}} = \frac{\omega_s}{\omega_m} \left(\frac{A_{st}}{A_{mt}} \right)^{\varepsilon(1-\theta)} \quad (36)$$

²⁴This work builds on the important earlier contribution of Baumol (1967).

Noting that $c_{it} = K_t^\theta A_{it}^{1-\theta} n_{it}$, we also have:

$$\frac{n_{at}}{n_{mt}} = \frac{\omega_a}{\omega_m} \left(\frac{A_{mt}}{A_{at}} \right)^{(1-\varepsilon)(1-\theta)} \quad (37)$$

$$\frac{n_{st}}{n_{mt}} = \frac{\omega_s}{\omega_m} \left(\frac{A_{mt}}{A_{st}} \right)^{(1-\varepsilon)(1-\theta)} \quad (38)$$

Recalling that labor allocated to the overall consumption sector is constant, it follows that if $\varepsilon = 1$ we have the earlier result that the n_{it} are constant in each of the three consumption sectors. So too are the values of $p_{it}c_{it}/P_tC_t$ and $p_{it}c_{it}/Y_t$. If ε differs from one, then the model can generate structural transformation along a GBGP as long as the rates of technological progress differ among the three consumption sectors. In contrast to Case 1, it is not true in this case that c_{mt} is a constant proportion of C_t , nor is true that C_t grows at a constant rate. Without imposing some additional structure one cannot say more about the nature of structural transformation that occurs.

To simplify exposition, we focus on the special case in which technological progress is strongest in agriculture and weakest in services, that is, $\gamma_a > \gamma_m > \gamma_s$. If in addition we assume that $\varepsilon < 1$, then the above expressions imply that along a GBGP the values of n_{it} , $p_{it}c_{it}/P_tC_t$ and $p_{it}c_{it}/Y_t$ are decreasing for agriculture and increasing for services. The behavior of these values for manufacturing is ambiguous in terms of the direction of change, but the size of the change is bounded by the sizes of the change in the other two sectors. Proposition 5 of Ngai and Pissarides (2007) shows that the evolution of n_m in this case will be either monotonically decreasing or hump shaped.

More formally, we summarize the above discussion with the following proposition.

Proposition 3. *Let $\bar{c}_a = \bar{c}_s = 0$, $\varepsilon < 1$, $\gamma_a > \gamma_m > \gamma_s > 0$, and $\gamma_x > 0$.*

There is a unique GBGP. Along the GBGP, the shares of employment and nominal value added (in current prices) of the investment sector are constant; the shares of employment and nominal value added (in current prices) of the consumption sectors behave as follows: the agricultural shares decline; the services shares rise; the manufacturing shares decrease less than the agricultural shares and increase less than the service shares.

4.1.4 Qualitative Assessment

The previous subsections outlined two different theories of structural transformation in the context of generalized balanced growth. Although we postpone a more rigorous assessment of the economic mechanisms implicit in these two theories until a later section, it is still of interest at this point to assess the extent to which each of the theories taken individually can account for some of the broad patterns that we documented in Section 2. We will see that while each theory can qualitatively account for some of the patterns found earlier, each also has some limitations.

Given the qualifications that we have noted previously in connecting the model with data, we keep in mind that we can either connect the production measures (employment shares and value added shares) to the data, implying that the consumption measure (consumption expenditure shares) does not have a close empirical counterpart, or we can connect the consumption measure to the data, implying that the two production measures do not have close empirical counterparts. Whichever way we choose to proceed, our benchmark model will not be able to make statements about all three measures of structural transformation at the same time.

We begin with the model of Kongsamut et al. (2001). Since the investment sector uses a constant share of labor and accounts for a constant share of (nominal) output, it will not influence the trend behavior of any quantities if it is allocated across the three sectors in constant proportions. Assuming this and starting with the nominal production measures, we conclude that the model can account for the increase in the service sector shares and the decrease in the agricultural sector measures along its GBGP, but it does not generate a hump shape for the manufacturing sector measures. If one allows for the investment share of manufacturing to decrease over time, as is true in the US data, then the model could generate a decline in both production measures for manufacturing. The increasing share of services in investment would only accentuate the rising employment and nominal value added shares for services. Turning to the nominal consumption expenditure measures, the model can account for the increase in the service share, the near constancy of the manufacturing share, and the decrease in the agricultural share.

The model of Kongsamut et al. (2001) has two additional implications that are counterfactual. First, along its generalized balanced growth relative prices need to be constant. It follows

that along a GBGP the real measures of structural transformation must display exactly the same properties as the nominal measures, which means that the model cannot account for the quantitative differences between the nominal and the real measures. Second, the model of Kongsamut et al. (2001) implies that in sufficiently poor economies, the household will consume a zero quantity of services and employment in services will also be zero. In contrast, we saw in Section 2 that even in the poorest countries service employment and value added are bounded away from zero.

Next we turn to the model of Ngai and Pissarides (2007). Once again we note that since along the GBGP the share of labor devoted to investment is constant and the nominal share of investment in output is constant, any constant allocation of investment across the three sectors will not influence of the trend properties. In this case, given the previously assumed ranking for the rates of technological progress, we conclude that structural transformation along the model's GBGP is qualitatively consistent with the evidence for employment and nominal value added shares in both agriculture and services. While the model does not necessarily deliver a hump shape for the manufacturing shares of employment and nominal valued added, it can deliver this for certain parameter values. Turning to the nominal consumption expenditure measures, the model can account for the increase in the service share and the decrease in the agricultural share, and can qualitatively produce hump-shaped dynamics for manufacturing, though this is not guaranteed.

However, the model of Ngai and Pissarides (2007) cannot account for the behavior of all real shares, irrespective of whether we use production or consumption related measures. In particular, given the assumptions about relative TFPs and the CES utility function being inelastic – i.e., $\varepsilon \in [0, 1)$, the model cannot generate the decreases in the real quantities of agriculture and manufacturing relative to services that we documented in Section 2 above. The reason for this is that with a CES utility function, nominal and real shares necessarily move in opposing directions. Given that the model accounts for the relative decline of the nominal shares of agriculture and manufacturing, this implies that it cannot account for the relative decline of the real shares. To see why nominal and real shares move in opposite directions, consider the implications of a decrease in the price of manufacturing relative to services. If $\varepsilon \in [0, 1)$, then

the nominal quantity of manufacturing decreases relative to that of services whereas the real quantity of manufactured goods relative to services remains the same if $\varepsilon = 0$ and increases if $\varepsilon \in (0, 1)$.

In summary, although each of these two specifications can account for some of the qualitative patterns that we documented previously, neither of them is able to match all of the patterns. However, the previous discussion suggests that a model featuring both income and relative price effects might successfully match all of the patterns. For example, adding non-homotheticities to the Ngai–Pissarides model could in principle allow the model to generate an decrease in the quantity of manufacturing relative to services. While such a specification would not permit a balanced growth path, this is a more general issue to which we will return to later in this section.

4.2 Alternative Specifications

In the preceding analysis, we have summarized the results from two papers regarding the possibility of simultaneously having structural transformation and generalized balanced growth. We chose these two papers because they illustrate two different channels through which expenditure shares may change over time: income changes and relative-price changes. In this subsection we describe some alternative formulations of these two channels that have appeared in the literature.

4.2.1 Other Specifications Emphasizing the Effects of Income Changes

Above we chose a specification of preferences where the effects of income changes on expenditure shares were captured by the non-homotheticity terms \bar{c}_a and \bar{c}_s . While we think that this is a tractable and transparent way of introducing income effects, there are several alternative specifications of non-homothetic preferences in the literature. Here we discuss some examples.

In the first quantitative analysis of structural transformation within the framework of the growth model, Echevarria (1997) generated effects from changes in income by using the fol-

lowing alternative specification of the intertemporal utility function:

$$\sum_{t=0}^{\infty} \beta^t \left[\alpha_a \log c_a + \alpha_m \log c_m + \alpha_s \log c_s - \eta \left(\frac{1}{c_a^{\rho_a}} + \frac{1}{c_m^{\rho_m}} + \frac{1}{c_s^{\rho_s}} \right) \right]$$

where $\alpha_i > 0, \eta, \rho_i \geq 0$. If $\eta = 0$ then the preferences reduce to a Cobb–Douglas specification, but if $\eta > 0$ and at least one of the $\rho_i > 0$ then the preferences are not homothetic. To see some of the features of this specification it is useful to examine the properties of the marginal utility of good i , which is given by:

$$MU_i(c_i) = \alpha_i c_i^{-1} + \eta \rho_i c_i^{-1-\rho_i} \quad (39)$$

Note first that the marginal utility of each good will be infinite for zero consumption quantities, implying that the household chooses interior consumption quantities. The second term is positive if $\eta \rho_i > 0$. In this case, it goes to infinity as c_i becomes arbitrarily small and it goes to zero as c_i becomes arbitrarily large.

If, as in Echevarria’s calibration, $\eta > 0$ and $\rho_a > \rho_m > \rho_s = 0$, then at low levels of income (and hence of consumption), there is a force in favor of higher c_a and c_m and of lower c_s , and the force is stronger for c_a than for c_m . In contrast, at high levels of income this force disappears. Intuitively, one can use the parameters η and ρ_i to achieve the same qualitative effects that are generated by the parameters \bar{c}_a and \bar{c}_s in our benchmark model.

The main advantage of Echevarria’s specification of period utility is that an interior solution to the static period problem exists for any positive level of income. This is in contrast to what happens in our benchmark model, since if $\bar{c}_a > 0$ and the present value of income is lower than the present value of $\{p_{at}\bar{c}_a\}$, then the household cannot afford to purchase at least \bar{c}_a units of the agricultural good in all periods and our period utility will not be defined in at least one period. From an analytical perspective, the disadvantage of Echevarria’s specification is that it is not consistent with generalized balanced growth. The reason for this is the presence of the term $\eta c_j^{-\rho_j}$ in the period utility function. If $\eta = 0$, then period utility is of the homothetic log form and a GBGP exists. In contrast, if $\eta > 0$, then it is impossible for the value of total consumption, $\sum_{j \in \{a,m,s\}} p_{jt} c_{jt}$, to grow at the same constant rate at which technological progress grows. As we

saw in Section 3.2 above, this would be required for a GBGP with constant real interest rate to exist.

A recent paper by Boppart (2011) explores more general preferences that are consistent with balanced growth. In particular, Boppart specifies indirect period utility functions that fall into the class of “price-independent-generalized-linearity” preferences defined by Muellbauer (1975, 1976). These preferences are more general than Gorman preferences in that they generate nonlinear Engel curves. Nonetheless they aggregate and allow for a stand-in household. There are two advantages of using “price-independent-generalized-linearity” preferences in the context of structural transformation. First, they avoid the awkward feature of our benchmark specification that can lead to utility not being defined for sufficiently small income. Second, as Boppart establishes, they are consistent with balanced growth if the technology side is as we specified it above.

A different approach to generating effects from changes in income is Foellmi and Zweimüller (2008). Whereas our benchmark model implicitly aggregated individual consumption goods into three broad sectors and defined preferences over the amounts of the three resulting aggregates, these authors specify preferences over an unbounded mass of potential consumption goods. Preferences are such that for each good, marginal utility is finite at zero consumption and decreases to zero at some finite satiation level of consumption. Over time, as income increases, the mass of goods that are consumed increases, so that there is adjustment along both the intensive and the extensive margin. The order in which the goods will be introduced is uniquely determined by the model’s primitives: all of the goods are symmetric from the perspective of production but are given different weights in preferences.²⁵

The fact that new goods are consumed over time implies that labor will necessarily be reallocated across activities over time. In terms of basic economic forces, the key mechanism at work comes from the fact that different goods have different income elasticities. Different than in the specification of our benchmark model, however, any particular good in this model will have an income elasticity of zero asymptotically since at some date satiation will be reached.

In order to connect their model to the standard facts of structural transformation, Foellmi

²⁵This type of preferences is sometimes called “hierarchical preferences”. It was first used by Murphy et al. (1989).

and Zweimüller (2008) need to map individual goods into the three broad sectors. If they assume that agricultural goods are disproportionately the goods with high weights, that services are disproportionately the goods with low weights, and that manufacturing goods lie “in between” these two, then they can match the qualitative patterns presented earlier. As income grows and more of the less weighted goods are consumed, one obtains a declining share for agricultural goods, an increasing share for services, and a hump-shaped pattern for manufacturing. Foellmi and Zweimüller can also generate balanced growth with relatively standard assumptions. Specifically, if they assume that the weighting function on different goods has a power form and there is constant labor augmenting technological progress that is common to the production of all goods, then their model gives rise to a GBGP. As they discuss in their paper, the assumption of a power function for the weighting function is analogous to the assumption of a constant elasticity utility function in the context of the standard one-sector growth model.

Relative to the results that we derived previously about income changes and structural transformation, the specification of Foellmi and Zweimüller (2008) delivers balanced growth and structural transformation in a more robust manner, in the sense that it does not need a condition similar to (30) that imposes a restriction on the parameters of preferences and technology. Moreover, it can also deliver a hump-shaped relationship between GDP per capita and the manufacturing shares. But a limitation of the specification of Foellmi and Zweimüller (2008) is that modeling structural transformation at the level of individual goods does not provide much guidance for how to connect the model with data at the level of broad sectors.²⁶

Hall and Jones (2007) also develop a framework that can give rise to non-homothetic demand functions, though their focus is specifically on the rise of spending on health care, as opposed to the more general process of structural transformation. Nonetheless this is of interest in the current context since increases in health care account for a significant part of the overall increase in the size of the service sector. In the basic model of Hall and Jones, utility in the

²⁶ Buera and Kaboski (2012a, 2012b) adopt a similar preference structure as Foellmi and Zweimüller (2008), except that they stress the introduction of new goods and adjustment along the extensive margin. Other aspects of their analysis are quite different, however. We discuss their model in more detail later in this section and again in Subsection 7.6. For now we simply note that Buera and Kaboski (2012a) derive an explicit mapping from their preferences to a reduced-form representation of preferences over goods and services. The interesting feature of this mapping is that it includes a term that is analogous to our term \bar{c}_s , but rather than being a constant, its value changes over time as technological progress occurs.

current period is derived from a single good that represents all non–health consumption. The period utility function is homothetic and health consumption in period t provides no direct utility flow in period t but does influence the probability of survival to next period. Intuitively, this model has features akin to the model with intensive–extensive margins that we discussed above. Specifically, a household can adjust along the intensive margin by spending more on consumption, or along the extensive margin by spending more on health care and therefore increasing the expected number of periods in which consumption occurs. As the level of consumption increases, the marginal utility from additional consumption at the intensive margin decreases relative to the marginal utility of living an additional period. This can generate an increasing expenditure share for health consumption as incomes rise, and therefore look like a model that features a non–homothetic period utility function over health and non–health consumption.²⁷

4.2.2 Other Specifications Emphasizing Relative Price Effects

In the Ngai–Pissarides model analyzed as Case 2 above, sectoral reallocation of factors of production and nominal value added shares occurred as a result of relative output price changes along the balanced growth path. Relative price changes were in turn generated by having differential rates of technological progress across sectors. The literature has also noted that relative output price changes can result from changes in the relative prices of inputs if sectors vary in the intensity with which they use inputs and there are changes in the relative supply of factors. In this case, one can generate structural transformation via relative price changes even if technological change is neutral.

Two papers in the literature stress this mechanism. Caselli and Coleman (2001) focus on skilled and unskilled workers as the two inputs of interest, noting that non–agriculture is more skill intensive than agriculture. They argue that the effective cost of education decreased in the first half of the 20th century, thereby increasing the relative supply of skilled workers, decreasing the relative price of non-agricultural goods and moving resources out of agriculture.²⁸ Acemoglu and Guerrieri (2008) consider capital and labor as the two inputs of interest, and as-

²⁷In a recent paper, Lawver (2011) uses a version of the model of the model of Hall and Jones (2007) to measure the increase in the quality of health consumption.

²⁸In Subsection 7.2 below, we will revisit this paper and discuss its implications for income convergence between regions.

sume that sectors differ in their capital intensity. Since growth driven by technological change is associated with an increase in the capital–to–labor ratio, changes in relative supplies of capital and labor arise quite naturally.²⁹

Here we sketch the basic idea within our benchmark model. Since the economics of the model of Acemoglu and Guerrieri (2008) is closest to that of Ngai and Pissarides (2007), except that the underlying cause of the relative price movements is different, we illustrate the basic idea by focusing on the implications for structural transformation of differences in the sectoral capital intensities. We assume that technological progress is uniform across the three consumption sectors and define A_t by $A_t \equiv A_{it}^{1-\theta_i}$ for $i \in \{a, m, s\}$. The capital intensities differ across sectors so that the sectoral production functions (13) become:

$$c_{it} = A_t k_{it}^{\theta_i} n_{it}^{1-\theta_i}, \quad i \in \{a, m, s\} \quad (40)$$

All other features of the environment are the same as in the benchmark model described earlier.

The first–order conditions for the stand–in firm in sector $i \in \{a, m, s\}$ are now given by:

$$R_t = p_{it} \theta_i A_t \left(\frac{k_{it}}{n_{it}} \right)^{\theta_i-1} \quad (41)$$

$$W_t = p_{it} (1 - \theta_i) A_t \left(\frac{k_{it}}{n_{it}} \right)^{\theta_i} \quad (42)$$

Dividing these equations by each other gives:

$$\frac{1 - \theta_i}{\theta_i} \frac{k_{it}}{n_{it}} = \frac{1 - \theta_j}{\theta_j} \frac{k_{jt}}{n_{jt}} \quad (43)$$

Two implications follow from this equation. First, sectors with larger capital shares have larger capital–labor ratios; second, the capital–labor ratio grows at the same rate in all sectors.

To derive the implications for relative prices, substitute (43) into (42) and rearrange to

²⁹In a different context, Bar and Leukhina (2010) argue that non–agriculture is more labor intensive than agriculture, and that the increase in population associated with the demographic transition could help explain the initial expansion of the non–agricultural sector in the context of England during the time of the industrial revolution.

obtain:

$$\frac{p_{it}}{p_{jt}} = \Omega_{ij} \left(\frac{k_{it}}{n_{it}} \right)^{\theta_j - \theta_i} \quad i, j \in \{a, m, s\} \quad (44)$$

where Ω_{ij} is a constant that depends on the capital shares. Since the capital–labor ratios of all sectors grow at the same rate, equation (44) implies that for any pair of sectors, the relative price of the sector with the higher capital share decreases as the aggregate capital stock grows. If one assumes:

$$\theta_a > \theta_m > \theta_s \quad (45)$$

it follows that the price of services relative to manufacturing and of manufacturing relative to agriculture will both increase over time. This implication is of course analogous to what we derived in the context of the Ngai–Pissarides model when we assumed that $\gamma_a > \gamma_m > \gamma_s$.

It is important to note that the mechanism of Acemoglu and Guerrieri (2008) relies not only on differences in the sectoral capital intensities, but also on the fact that with Cobb Douglas production functions the elasticity of substitution between capital and labor is equal to one. Indeed, Alvarez-Cuadrado et al. (2012) have recently pointed out that the relative price of sectoral output depends not only on sectoral TFP and capital intensity, but also on the elasticity of substitution between capital and labor. To see how the elasticity of substitution matters in this context, consider first the extreme case in which capital and labor are perfect substitutes. The capital intensity then does not matter at all for relative prices because firms can perfectly substitute labor for capital when capital is relatively expensive. In the other extreme case, capital and labor are perfect complements and the production function is of the Leontief form. The capital intensity then matters crucially for relative prices because one cannot substitute labor for capital when capital is relatively expensive. More generally, the effects of Acemoglu and Guerrieri (2008) are more important if the sectoral elasticity of substitution is smaller.

Although the specification of Acemoglu and Guerrieri (2008) can account for the changes in nominal value added shares, it cannot account for the changes in real value added shares. Moreover, it cannot generate the patterns in sectoral employment shares either.³⁰ To see why,

³⁰A similar issue is also present in Ngai and Pissarides (2007). We will discuss this in more detail later.

note that using (43), it is straightforward to show that:

$$K = \left(\sum_{j=x,a,m,s} \frac{\theta_j}{1-\theta_j} n_j \right) \frac{1-\theta_i}{\theta_i} \frac{k_i}{n_i} \quad (46)$$

Solving this expression for k_i/n_i and substituting the result into equation (40) gives:

$$c_{it} = A_t K_t^{\theta_i} \left(\frac{\frac{\theta_i}{1-\theta_i}}{\sum \frac{\theta_j}{1-\theta_j} n_j} \right) n_{it}, \quad i \in \{a, m, s\}$$

In the polar case of Leontief utility, c_{it}/c_{jt} is constant, so the previous equation implies that n_{it}/n_{jt} is constant too. For positive elasticities of substitution, changes in relative quantities are in the opposite direction of changes in relative prices. In other words, in the model of Acemoglu and Guerrieri there cannot be structural transformation in terms of employment that is consistent with the fact that service employment increased at the same time as which its relative price increased too.

One important additional difference relative to the specification of Ngai and Pissarides (2007) is that the model of Acemoglu and Guerrieri (2008) has exact GBGP only asymptotically, and so the best we can hope for in this model is *approximate* generalized balanced growth. Below we discuss the difference between approximate and exact generalized balanced growth in more detail.

4.2.3 An Alternative View of Structural Transformation

In two recent papers, Buera and Kaboski (2012a, 2012b) have offered a novel representation of structural transformation that implicitly involves elements of both of the special cases discussed previously. Here we offer a simple version of their framework to illustrate the forces at work. In Subsection 7.6 we discuss their specific implications in more detail.

They consider an economy in which there are a continuum of services and a continuum of goods. For simplicity, in their economy goods are only useful as an input into the production of services, and each good is uniquely associated with the production of a specific service. Specifically, each good is produced using labor, and each service is produced using labor and

its corresponding specialized good. They adopt a similar preference structure as Foellmi and Zweimüller (2008), but they assume that each service can only be consumed in the amounts of zero or one, so that increasing consumption will necessarily manifest itself along the extensive margin. From the consumer's perspective all services are symmetric. Consider the following special case of this structure as a benchmark. Assume a single household with one unit of time. Index the continuum of goods and services by z . The technology for producing each good z at time t is $g(z) = A(t)h_g(z)$, where $A(t)$ captures labor augmenting technological change. The technology for producing each service z at time t is Leontief: $s(z) = \min\{A(t)^{\frac{1}{a}}h_s(z), g(z)\}$, where $A(t)$ is the same in both production functions. Because each service is consumed in amount 1, it takes $(1 + a)/A(t)$ units of labor to produce one unit of service, so that total consumption (i.e., the total number of services that are consumed) will be given by $A(t)/(1 + a)$, and a fraction $1/(1 + a)$ of labor will be devoted to the goods sector. So, in this benchmark economy there is no structural transformation in terms of labor allocations between goods and services.

Buera and Kaboski generate interesting implications in this setting by extending it along two dimensions. First, they introduce the possibility of home produced services which also require labor and the specialized good. To create an interesting tradeoff between the choice of whether to produce a given service in the home or in the market, they assume that market production of services is more efficient. This could be modeled in different ways and differs in their two papers. To illustrate some basic workings of the model we assume that market production takes less of the good per unit of output, but that home produced services supply a proportionately higher utility flow. An illustrative example would be the choice between home produced transportation services (buying a car and driving yourself) versus market provided transportation services (buses, or taxis). While having a car increases convenience, the car will also be idle for considerable periods. Second, they introduce heterogeneity into the production side of the economy by assuming that higher z goods require more labor to be produced. This heterogeneity interacts with the choice of whether to produce a given service in the market or the home, since the more expensive it is to produce the durable, the greater is the penalty for home production which requires more of the durable per unit of output. Whether a good is produced in the home or the market in turn has implications for observed allocations of labor

and market value added across market sectors, since having home produced services requires labor from the goods sector, but will not use any labor in the market service sector. In their model, as an economy develops the marginal services that are added represent services with higher benefits to market versus home production. The combination of technological change plus the changing nature of the marginal services being brought into the economy can introduce interesting dynamics for how activity shifts between the market and home sectors. If production shifts toward the market and away from the home, this will be recorded as an increase in the size of the market service sector relative to the goods sector.

As noted earlier, models with these types of preferences necessarily embody a non-homotheticity. But the production heterogeneity in this model implicitly acts like differential technological growth across sectors since the marginal services that are added as an economy grows have differing relative productivity for home versus market production. A general message from this framework is that when thinking about growth and structural transformation it is important to think about the new goods and services that are associated with growth, and the movement of delivery of certain services between the home and market sectors, since the changing nature of activities in the market sector can have important implications for the measured sectoral allocation of market activity.

4.3 Approximate versus Exact Generalized Balanced Growth

Up to this point, our discussion has focused on analytic results concerning the possibility of jointly having generalized balanced growth and structural transformation. This is a natural starting point given the emphasis that the literature using the one-sector growth model places on balanced growth and that conditions under which balanced growth results in the one-sector model are relatively weak – constant returns to scale production with labor augmenting technical change and a period utility function with a constant intertemporal elasticity of substitution. The results that we have presented above for multi-sector models, however, have made it apparent that the conditions for jointly having generalized balanced growth and structural transformation become considerably more stringent – we now need that all production functions are Cobb–Douglas with the same capital share, that the period utility function exhibits a uni-

tary elasticity of substitution, and in some cases that there is a particular relationship between preference and technology parameters. To the extent that there is good reason to believe that many of these conditions are not satisfied, models that impose them may be missing some key features of reality. In fact, some authors have dismissed income changes as an important source of structural transformation on the grounds that they are consistent with generalized balanced growth only under very fragile cross–restrictions on technology and preferences such as the one imposed in (30).

The previous discussion suggests that it may be ill advised to insist on generalized balanced growth in the context of structural transformation. To the extent that (generalized) balanced growth is merely a good approximation to what we see in the data in various countries over long periods of time, the more relevant question is whether there are specifications that can deliver structural transformation and *approximate* generalized balanced growth, which may occur under much less stringent conditions than exact generalized balanced growth.

To date there has not been much systematic analysis of the extent to which approximate generalized balanced growth is a robust feature of multi–sector versions of the growth model along the lines of those that we have considered. But several cases in the literature suggest that approximate generalized balanced growth may in fact be quite robust. To begin with, Kongsamut et al. (2001) consider numerical examples that depart from the exact conditions needed for generalized balanced growth in their setting and find that the equilibrium path does not deviate much from generalized balanced growth. In a similar context, Gollin et al. (2002) study a two–sector model with subsistence consumption in the agricultural sector but not in the other sector – a clear violation of the conditions needed to generate GBGP, but find relatively small variations of the interest rate when their model is calibrated to match the US data over the post 1950 period. Moreover, although the model in Acemoglu and Guerrieri (2008) only has an asymptotic GBGP, the results that they report for numerical simulations suggest that the model’s behavior along a transition path is not that different from balanced growth.

The models just discussed have the feature that asymptotically structural transformation ceases to occur. For example, if structural transformation occurs as the result of the non–homothetic terms \bar{c}_a and \bar{c}_s , then productivity increases will imply that in the limit the size of

the two non-homothetic terms becomes arbitrarily small relative to consumption. Since we observe (approximate) balanced growth and structural transformation over very long periods in the data, it follows that any model that generates structural transformation purely while it is converging to an exact balanced growth path must have very long lived dynamics in order to capture reality.³¹

5 The Economic Forces Behind Structural Transformation: Empirical Analysis

The previous section has focused on models that could generate (approximate) generalized balanced growth and structural transformation as simultaneous outcomes. The various models that we reviewed emphasize different theories for the reallocation of activity across sectors that accompanies growth. In one class of theories, the key driving force is uniform technological progress, and the key propagation mechanism comes from income effects. In another class of theories, the key driving force is technological progress that differs across sectors and the key propagation mechanism comes from relative price effects in consumption. In a third class of models, the driving force is again uniform technological progress, but the propagation mechanism is a combination of different capital intensities or elasticities of substitution in production and relative price effects in consumption.

Rather than focusing narrowly on the conditions required to generate exact balanced growth, we believe that the key to developing quantitative theories of structural transformation is to develop quantitative assessments of the various driving forces and propagation mechanisms that the literature has identified as potentially important. In this section we summarize the recent progress in this effort. We break this section into two subsections. The first subsection considers the direct evidence regarding differences in rates of technological progress, capital intensities, and elasticities of substitution. The second subsection considers the more general issue of the relative importance of the effects coming from changes in income and changes in

³¹Note that this statement does not apply to the model of Ngai and Pissarides (2007) which exhibits structural transformation both along the exact balanced growth path and in the limit.

relative prices.

5.1 Technological Differences Across Sectors

In this subsection we consider the evidence regarding technological differences across sectors along the two dimensions highlighted by the previous theories: differences in technological progress and differences in capital shares and in elasticities of substitution. We also assess the extent to which these differences are appropriate to generate the qualitative features found in the data regarding structural transformation.

5.1.1 Sectoral TFP Growth

Assumptions about TFP growth at the sectoral level played an important role in both of the theories of structural transformation that we highlighted. It is therefore of interest to ask what the empirical evidence is regarding relative growth rates in sectoral TFP. Although this would seem to be a relatively straightforward exercise, it is actually challenging to verify the properties of TFP growth in sectoral value added production functions in a cross-country setting. The main reason is that calculating sectoral TFPs requires data on real value added, capital and labor inputs, and the factor shares at the sector level. Unfortunately, these data are unavailable for most countries. One of the many issues is that in order to compute real value added one must have data on the real quantity of intermediate inputs, not just the value of intermediate inputs.

One data set that has the necessary information for a set of countries is EU KLEMS.³² We begin, therefore, by using the EU KLEMS data starting in 1970 to compute TFP in the production of value added in agriculture, manufacturing, and services for the same set of countries as in Section 2: Australia, Canada, Japan, Korea and the US as well as the aggregate of 10 EU countries.³³ Figure 10 plots the sectoral TFPs for these countries. Given that we are interested in growth rates of TFP, we normalize TFP in 1990 for all sectors in all countries to be one.

³²See Timmer et al. (2010), particularly the chapter on structural change, for further discussion of the details of the EU KLEMS data on multifactor productivity. See also Duarte and Restuccia (2010) who document similar facts about TFP as we do here.

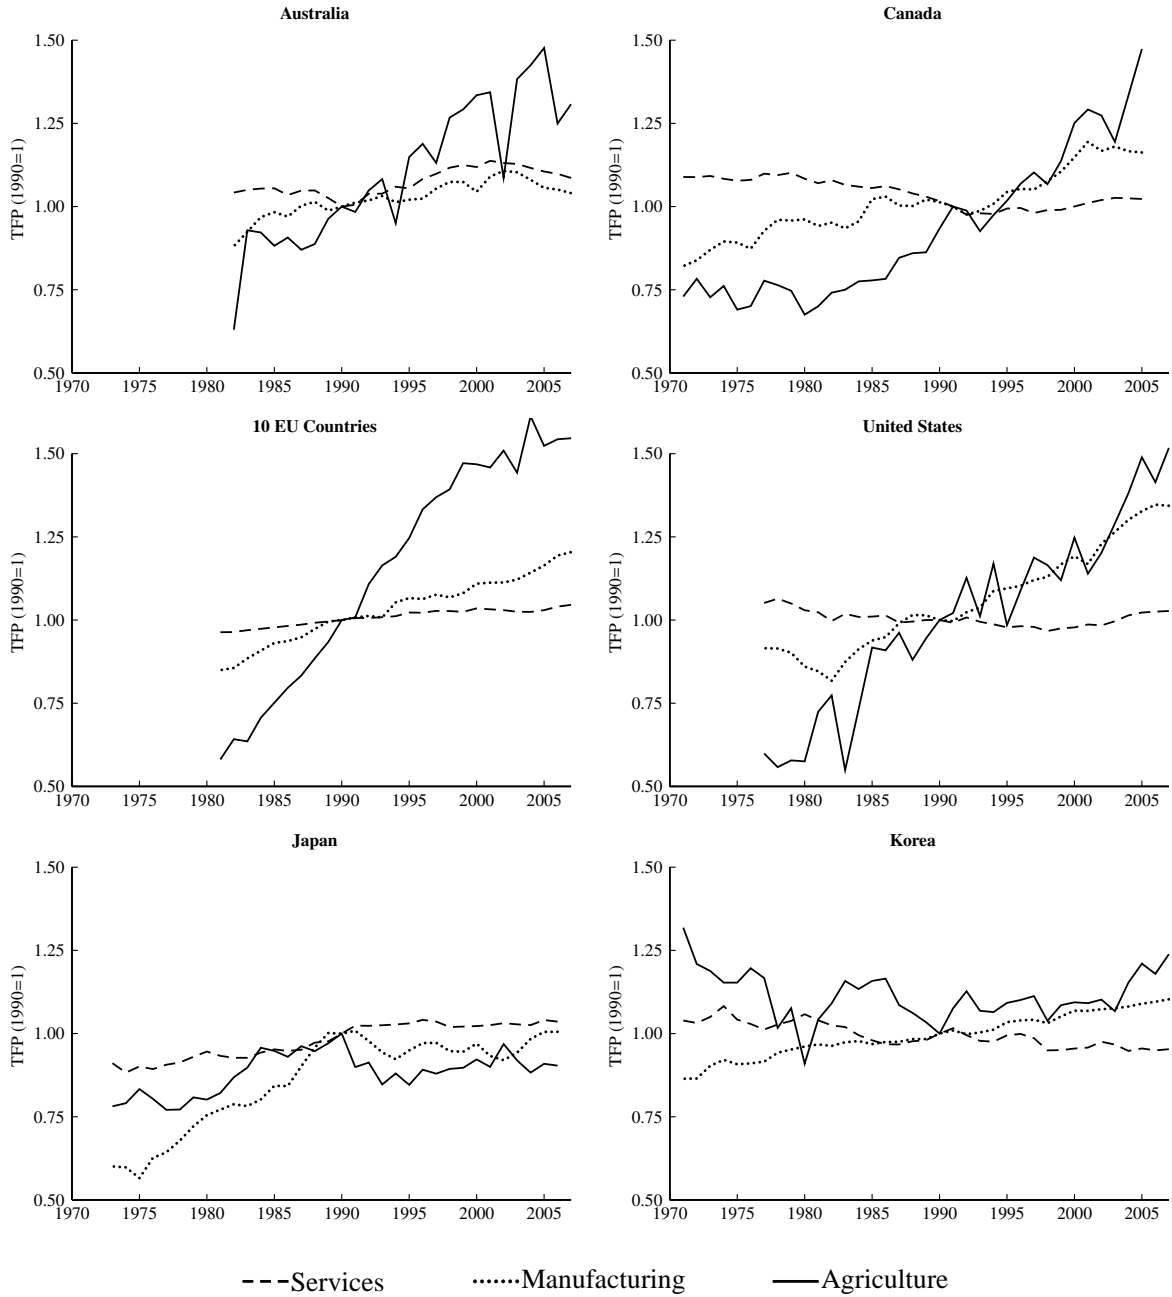
³³The 10 EU countries are the EU member states for which EU KLEMS performs growth accounting: Austria, Belgium, Denmark, Spain, Finland, France, Germany, Italy, the Netherlands, and the United Kingdom.

One message that emerges from Figure 10 is that there are indeed substantial differences in the growth rates of TFP across sectors. Moreover, we can see that the conditions of Ngai and Pissarides (2007) broadly hold for Australia, Canada, the EU 10, and the United States: averaging over the time period 1970–2007, TFP in agriculture shows the strongest growth while TFP in services shows the weakest growth. This is exactly what is needed for the observed reallocation of employment out of agriculture and manufacturing into the service sector in the model of Ngai and Pissarides (2007).

While data limitations make it difficult to obtain long time series evidence on sectoral TFP for a large sample of countries, our theory suggests an alternative method which requires fewer data. Specifically, in the analysis of our benchmark model we highlighted the fact that if sectoral production functions are Cobb–Douglas with equal capital shares then there is a direct inverse relationship in equilibrium between changes in relative prices and changes in relative productivities. Given appropriate data on prices, one could use this relationship to infer changes in relative productivity. Since long time series of price data is much more readily available than the data needed to measure TFP directly, this is an appealing alternative. However, in addition to requiring the assumption of Cobb–Douglas production functions with equal capital shares, there are two limitations to be noted. First, in our model we assumed that technological change was the only factor that varied over time. One can easily imagine policies or regulations that may also affect relative prices across sectors. If these factors are important for some countries during some periods, it may be misleading to assume that all relative price changes are driven by changes in relative productivities. Second, although price data do exist going quite far back in time, the price data that is required to infer relative productivity growth in *value added* production functions is the price per unit of *value added*. In contrast, in practice most available price indices correspond to final goods or to gross output.

Having noted these qualifications, we turn to the evidence documented by Alvarez–Cuadrado and Poschke (2011) about time series changes in the relative price of agriculture to non–agriculture for eleven advanced countries over the last two centuries. A key feature of these data is that the price of agriculture relative to non–agriculture changed its behavior during the last two centuries: while before World War II it showed an increasing trend, after World War

**Figure 10: Sectoral TFP for Selected Countries
– Time Series from EU KLEMS 1970–2007**



Source: EU KLEMS, WORLD KLEMS for Korea

II it started to follow a decreasing trend. Interpreting these changes in relative prices as indicative of changes in relative TFPs, the implication is that prior to World War II, TFP growth in agriculture was actually lower than in non-agriculture.³⁴ The period before World War II also corresponds to the period that saw the largest movement out of agriculture. In contrast to the findings for data since 1970, the longer time series does not seem to be consistent with relative TFPs driving the labor reallocation from agriculture to non-agriculture.

By way of summary, we think there are two main conclusions that can be drawn from this evidence. First, there are systematic differences in TFP growth rates across sectors. After World War II these differences appear to be consistent with what is needed to obtain the observed reallocation of employment out of agriculture and manufacturing into the service sector in the model of Ngai and Pissarides (2007). Second, the differences in TFP growth rates across sectors do not appear to be stable over very long periods of time, at least in the case of agriculture versus non-agriculture, which does not bode too well for the models of structural transformation and exact balanced growth that we highlighted previously.

5.1.2 Sectoral Differences in Capital Shares and Elasticities of Substitution

Next we consider the existing evidence regarding the potential role of differences in sectoral capital shares, as emphasized by Acemoglu and Guerrieri (2008), and of differences in sectoral elasticity of substitution, as emphasized by Alvarez-Cuadrado et al. (2012). Herrendorf et al. (2012a) speak to these questions by assessing how structural transformation is affected by sectoral differences in labor-augmenting technological progress, substitutability between capital and labor, and capital intensity. Using postwar US data on sectoral value added, capital, and labor, they estimate CES production functions and compare them with Cobb-Douglas production functions with different and with equal capital shares. They find that: labor-augmenting technological progress is faster in agriculture than in manufacturing and faster in manufacturing than in services; capital and labor are more easily substitutable in agriculture than in manufacturing and more easily substitutable in manufacturing than in services; agriculture is more

³⁴It should be noted that the evolution of agricultural TFP in Korea between 1970 and 2007 shows a similar u-shaped pattern (see Figure 10).

capital-intensive than services and services are more capital intensive than manufacturing;³⁵

The findings of Herrendorf et al. (2012a) have two implications for the importance of sectoral differences in capital shares and elasticity of substitution as driving forces behind structural transformation. First, in the face of an increasing capital-to-labor ratio, differences in capital shares cause reallocation from agriculture to manufacturing and from services to manufacturing. Second, differences in the elasticity of substitution partly neutralize the differences in the capital shares. In particular, while agriculture has by far the largest capital share it also has the highest substitutability between capital and labor, and in fact agriculture is the only sector for which capital and labor are more substitutable than the Cobb–Douglas case. Herrendorf et al. (2012a) show that, as a result, sectoral differences in labor-augmenting technological progress turn out to be the main quantitative force on the technology side behind the postwar US structural transformation, and that this force is well captured by Cobb–Douglas production functions with equal capital shares but different TFP processes.

5.2 The Importance of Changes in Income and Relative Prices

Since the theoretical literature has emphasized the effects that result from changes in income and relative prices, it is natural to ask what the data say about these two effects. There are two natural and complementary approaches to this question. In the spirit of our earlier analysis, one approach starts with a stand-in household and uses aggregate data to infer the relative importance of the two different mechanisms. The second approach uses data on individual households to estimate properties of preferences and then assesses the implications for aggregate behavior. In the interest of space, we will focus on the first approach, though we will briefly mention some results from the analysis of micro data. We discuss two recent contributions: Dennis and Iscan (2009) and Herrendorf et al. (2013). The former studies the forces leading to the movement of activity out of agriculture in the United States over the last two centuries, whereas the latter focuses specifically on the reallocation of activity across all three sectors in the United States since 1947. We describe each in turn.

³⁵In order to avoid confusion, we stress that these capital shares refer to value added, and not to final expenditure. The capital shares for final expenditure at the sector level can be found in a related paper, Valentinyi and Herrendorf (2008).

5.2.1 The Movement out of Agriculture in the US since 1800

Dennis and Iscan (2009) seek to assess the relative importance of income effects, relative TFP growth and capital deepening on the movement of labor out of agriculture in the US over the last two centuries. Their framework is very similar to our benchmark model with the exception of three details. First, they have only two sectors, agriculture and non-agriculture. Second, they assume that all investment comes from the non-agricultural sector. Third, they do not impose that the capital share is the same in both sectors. Initially, Dennis and Iscan write the utility function as the two-sector analogue of our utility function, but in their empirical analysis they also allow for the possibility that the subsistence term \bar{c}_a changes over time. Given our earlier discussion, we note that while this general specification is not consistent with generalized balanced growth, it captures the basic forces that the theoretical literature has emphasized.

Dennis and Iscan (2009) derive an equilibrium relationship that expresses the share of labor devoted to agriculture as a function of three factors, which in turn reflect income effects through the subsistence term, relative productivity effects via differential growth rates of TFP, and capital deepening effects. Expressed in terms of our notation, this equilibrium relationship is:³⁶

$$1 - n_{at} = \frac{1 - s_a(c_{at})}{1 + p_R(A_{at}, A_{nt})s_k(k_{at}, k_{nt})s_X(c_{nt}, X_t)} \quad (47)$$

where

$$s_a(c_{at}) = \frac{\bar{c}_a}{c_{at}}, \quad p_R(A_{at}, A_{nt}) = \frac{\omega_a}{\omega_n} \left(\frac{A_{nt}}{A_{at}} \right)^{1-\varepsilon},$$

$$s_k(k_{at}, k_{nt}) = \left(\frac{1 - \theta_a}{1 - \theta_n} \right)^\varepsilon \left(\frac{k_{nt}^{\theta_n}}{k_{at}^{\theta_a}} \right)^{1-\varepsilon}, \quad s_X(c_{nt}, X_t) = \frac{X_t}{c_{nt} + X_t}.$$

The term $1 - s_a(c_{at})$ captures the income effect that operates through the subsistence term \bar{c}_a . The terms $p_R(A_{at}, A_{nt})$ and $s_k(k_{at}, k_{nt})$ capture the relative price effects that arises from differential technological progress and capital deepening, respectively, while the term $s_X(c_{nt}, X_t)$ captures the effects associated with changes in the investment rate.

³⁶We use the index n for the non-agricultural sector.

Dennis and Iscan (2009) calibrate the key parameters of the model (elasticity of substitution, subsistence terms, preference weights, and capital shares) and then assess the extent to which equation (47) holds in the data. In particular, they substitute actual values into the right-hand side of equation (47), solve for the implied share of labor allocated to agriculture and compare this to the actual series from the data. To assess the importance of the different factors they carry out the same exercise but only allow one of the factors to change over time.

The main findings of Dennis and Iscan (2009) are as follows. First, the model does a reasonable job of capturing the time series changes in the employment share of agriculture since 1800. If the value of \bar{c}_a is held fixed throughout, the model somewhat under-predicts the employment share for agriculture in the 1800s, but does fine in the post 1950 period. A small time trend in \bar{c}_a over the period 1800–1950 yields a better fit over the entire period. Second, prior to 1950 the income effect is the dominant factor in accounting for the movement of employment out of agriculture, whereas the relative productivity effect is working in the opposite direction. Only in the post 1950 period do the effects of relative productivity and capital deepening play even a modest role in accounting for the change in the employment share of agriculture. They also consider various extensions to their analysis, such as incorporating trade, and they show that the results are robust to these extensions.

We want to stress three key implications of the results of Dennis and Iscan (2009). First, the fact that their model does a reasonable job of capturing the movement of labor out of agriculture over a long time period suggests that our benchmark model is sufficiently rich to capture some key features in the data. Second, the fact that a time varying subsistence term, \bar{c}_{at} , improves the model's ability to account for the movement out of agriculture is notable, and suggests that a deeper theory of how income effects arise may be warranted. Third, at least for the movement of labor out of agriculture in the United States, income effects are effectively the sole driving force behind this decline; even though the other factors play a role after 1950, this occurs when almost all of the decline in the employment share for agriculture has already happened.

It is also relevant to note some limitations of the analysis in Dennis and Iscan (2009). First, it only focuses on the movement of labor out of agriculture and does not address the issue of what forces shape the allocation of employment between manufacturing and services. Second,

all of their results come from a calibration exercise, but there is little direct evidence on some of the key parameters they use for this exercise. Additionally, they connect their model to the data in a somewhat inconsistent fashion, in that they interpret their production functions as value added production functions, but when they look at consumption of agriculture they interpret it as consumption of final goods. In the next subsection, we discuss in detail why this is inconsistent. Third, they focus only on the changes in employment shares, and so do not address the issue of the discrepancy between value added shares and employment shares that we documented earlier. Nonetheless, we think that this paper makes an important contribution to the effort to identify the key economic forces behind structural transformation.

A related exercise was carried out by Buera and Kaboski (2009). Specifically, they assessed the ability of a calibrated version of our three–sector benchmark model to account for the broad patterns of structural transformation in the US from the 1800s to the present. One difficulty that they noted was the ability of the model to account for the acceleration in the nominal value added share of the service sector in the post World War II period.

5.2.2 Structural Transformation in the US since 1947

Herrendorf et al. (2013) offer a related but distinct approach to uncovering the importance of income and relative price effects in accounting for structural transformation. In contrast to Dennis and Iscan (2009), who considered the allocation of employment between agriculture and non–agriculture in the US since 1800, Herrendorf et al. (2013) consider the reallocation among consumption expenditure shares for all three sectors in the US since 1947. Specifically, starting with a stand–in household, they asked whether the utility function in (1) provides a good fit to the US data on expenditure shares in the post World War II period, and if so, what this implies for the values of the key parameters \bar{c}_a , \bar{c}_s and ε , and the implied importance of income and relative price effects.

Although this seems to be a simple question, Herrendorf et al. (2013) argued that the question is not even properly specified. The reason for this is related to the difference between value added and final expenditure, which we have previously discussed. In particular, if one interprets the sectoral production functions as value added production functions then the arguments

of the utility function necessarily represent the corresponding consumption of sectoral value added. In terms of our previous example of the purchase of a cotton shirt, this implies that the shirt is broken into three value added pieces, each of which the household values as they contribute to the three different categories of value added. Herrendorf et al call this the *value added approach*. Alternatively, one may interpret the commodities in the utility function as final expenditure categories, as is typically done in household expenditure studies. The outputs of the production functions must then be viewed as final expenditure rather than value added. In terms of the purchase of a cotton shirt, the consumer simply derives utility from the shirt as a whole as it contributes to the single category of manufacturing consumption. Herrendorf et al call this the *final expenditure approach*. It is important to note that there is no right or wrong in terms of these two approaches. From the perspective of preferences, these are simply two different ways of aggregating across the many characteristics that consumers value. As is true with any attempt to aggregate individual characteristics into broader groups, one can imagine examples where one approach seems preferable.

The choice of interpretation matters if the relative prices and quantities are not the same for the two different interpretations. In particular, even if the two different approaches display similar qualitative properties in terms of changes over time, differences in quantitative properties may have important implications for parameters of the utility function and the importance of income and relative price effects. Herrendorf et al. (2013) carry out the manipulations necessary to have consistent sets of data for the two approaches and they provide the following answers.

One possible outcome from this exercise is that one of the approaches provides a better fit to the data, in which case one might use this as evidence in support of one approach over the other. However, Herrendorf et al. (2013) found that for both approaches the preferences represented by (1) yield very good fits to the postwar US data on relative prices and expenditure shares. However, the two approaches yield very different parameter estimates for the utility functions and very different assessments of the relative importance of the effects of relative prices and income.

For the final expenditure approach, income effects are the dominant source of changes in expenditure shares, and the Stone–Geary utility function (28) of Kongsamut et al. (2001) pro-

vides a good fit to the data.³⁷ For the value added approach, it turns out that relative price effects are a much more important source of changes in expenditure shares. Moreover, the homothetic Leontief utility function $\min_{c_{at}, c_{mt}, c_{st}} \{\omega_a c_{at}, \omega_m c_{mt}, \omega_s c_{st}\}$, which results in $\varepsilon = \bar{c}_a = \bar{c}_s = 0$, provides a reasonable fit to the data. Interestingly, this utility function is a special case of the class of inelastic CES utility functions that Ngai and Pissarides (2007) considered.³⁸

It is important to emphasize what these results mean. They are not an example of researchers obtaining different estimates for a given parameter from different data sets, suggesting that further work is needed to narrow down the set of possible values. Instead, the implication is that there are two different ways to interpret commodities in the utility function in multi-sector models. It turns out that being explicit about which interpretation is adopted is of critical importance, in that it has implications for what data is required to connect the model with the data, and as just shown, this has very important implications for implied preference parameters. Furthermore, note that the two approaches are just two different aggregate representations of the same underlying economic data. The key message is that one cannot talk about the importance of income or relative price effects as drivers of structural transformation without specifying what representation of the data one is adopting. What shows up as income effects in one representation may manifest itself as relative price effects in the other representation. Different representations are connected via the complex input–output relationships in the economy. Herrendorf et al. (2013) show how one can construct the mapping between the two representations for a given input–output structure.

We stress two key results. First, the fact that the model is able to account for changes in expenditure shares for the US since 1947 is again support for the parsimonious model that we have adopted as our benchmark. Second, it highlights that empirical researchers working with multi-sector models must take care to be explicit about how commodities in utility functions are to be interpreted. Different interpretations have dramatically different implications for how

³⁷Many other papers have estimated linear expenditure systems implied by the Stone–Geary utility specification. A review of this literature is Blundell (1988).

³⁸While Buera and Kaboski (2009) independently reached the conclusion that a low σ is required to match value added data, they also found that the benchmark model cannot account for the increase of the share of services in the last thirty years. Herrendorf et al. (2013) show that the reason for the different conclusions is that Buera and Kaboski (2009) assume that all investment is produced in manufacturing. This implies that they do not take into account that the investments produced in services have risen sharply since World War III.

the models are to be connected with the data and what the implied parameters of the utility function.

One of the limitations of this study is that it only focuses on the post 1947 period for the US, and this is a period in which the US has already experienced much of the reallocation out of agriculture. While it is of interest to extend this type of analysis to longer time periods and different countries, a key issue is data availability.³⁹

6 Extensions of the Benchmark Model

In this section we discuss relaxing three features present in the analysis of the benchmark model. The first is the assumption that there is no international trade (“closed economy”). The second is the assumption of that there is no cost of moving labor across sectors (“perfect labor mobility”). The third is the assumption that there are no costs of moving goods across sectors (“zero transportation costs”).

6.1 International Trade

Thus far our theoretical analysis has taken place under the assumption of a closed economy. A key implication of being a closed economy is that the production of each of the four sectors must equal the corresponding household choices (either of investment or of one of the three consumption goods). The equality between sectoral productions and consumption/investment played a key role in generating the results concerning structural transformation that we obtained in the benchmark model. For example, in the model of Ngai and Pissarides (2007), we saw that labor moved out of the consumption sector that had the highest productivity growth because of the household’s desire to maintain the composition of its consumption allocation (inelastic demand). In the model of Kongsamut et al. (2001), technological progress was uniform across sectors, but labor moved out of agriculture because of the household’s desire to change the com-

³⁹This is relevant for the analysis of Buera and Kaboski (2009). They carry out a calibration exercise for the US over a longer time period, but need to use different sources for relative prices in the pre 1947 period. Given that prices for value added consumption and final consumption are quite different in the post 1947 period and have very different implications for preference parameters, an issue arises with how to interpret results that use a mixture of prices.

position of its consumption allocation towards manufactured goods and services (differences in income elasticities).

In this subsection we discuss the extent to which openness changes the results about structural transformation. We begin with the simple observation that the competitive equilibrium of a model in which all commodities are tradeable without costs will have a complete separation between the decisions of firms and households. This observation implies that in an open–economy version of our benchmark model without trade costs the production measures of structural transformation (i.e. employment and value added shares) would generically follow a different pattern than the consumption expenditure share. This is relevant because, as we have documented in Section 2, there is a discrepancy between production and consumption shares in some instances, most notably for the share of manufacturing in Korea.

Matsuyama (2009) was the first to analytically work out the idea of the previous paragraph for a simple two–country model. He abstracts from capital and considers a Stone–Geary utility function over the three consumption goods food, manufactured goods, and services. He assumes that agricultural goods are an endowment whereas manufactured goods and services are produced with technologies that are linear in labor, and that agricultural and manufactured goods can be traded with the rest of the world at zero trade costs whereas services cannot be traded. Matsuyama shows two results for this simple model. First, if there is technological progress in manufacturing then the total manufacturing labor of both countries declines. Second, if one of the two countries experiences stronger technological progress in manufacturing than the other, then manufacturing labor in the first country may initially increase while manufacturing labor in the second country decreases unambiguously. Eventually, when technological progress in the manufacturing sector has been sufficiently strong, the share of manufacturing labor in the first country will decrease also. These results suggest that a hump–shaped relationship may occur in the country which experiences the stronger technological progress in manufacturing.

Yi and Zhang (2010) generalize the idea of Matsuyama to a two–country version of our benchmark model of structural transformation, in which all goods are produced with labor only. The assumption that agricultural and manufactured goods are tradeable without costs would

then lead to the counterfactual implication that each country specializes in either agriculture or manufacturing. They therefore assume that each of the three sectors is the aggregate of a continuum of goods as in Eaton and Kortum (2002). Yi and Zhang (2010) simulate their model under the assumption that one country has higher productivity growth in manufacturing than the other country. They provides examples for which the country with the higher productivity growth in manufacturing experiences a hump shape in the shares of manufacturing employment and value added while the other country experiences a downward sloping shape in the shares of manufacturing labor and value added.

From the empirical perspective it is of interest to ask whether there is evidence for the effects of openness on structural transformation, besides the hump shape of manufacturing employment and value added. One clear prediction of the models of Matsuyama (2009) and Yi and Zhang (2010) is that the labor shares of sectors that produce tradeable goods should differ across countries that have different sectoral productivities. In Section 2 we noted that there was some evidence of dispersion in sectoral labor shares across countries in the European Union and Japan, with Germany and Japan having unusually large share of manufacturing hours worked and Korea having an unusually large share of real manufactured value added. Betts et al. (2011), Spasi (2011), and Teignier (2012) study the role of international trade in Korea's industrialization. They find that international trade played a crucial role for the rapid rise in the manufacturing value added and employment shares. Teignier (2012) finds in addition that international trade could have played a much larger role if South Korea had not introduced agricultural protection policies. While such a story may be consistent with various accounts regarding the importance of trade in the development of South Korea, it is hard to reconcile with the patterns we found in Section 2. Specifically, we found there that South Korea did not display any distinctive behavior for the labor allocations.

We conclude that the effects of openness on structural transformation show up in a discrepancy between production and consumption in sectors that trade with the rest of the world. In the past, this applied to manufacturing, and to a lesser extent to agriculture. In recent years, however, there has been an increasing trend toward trade in services. An open question moving forward concerns the extent to which increased trade in services will influence the nature of

structural transformation. For example, will increased trade in services hasten the movement of resources out of manufacturing in a country like the US which has relatively high productivity in many service industries, and is therefore thought to have a comparative advantage in services?

6.2 Labor Mobility

Our benchmark model assumed that labor was homogeneous and could be allocated across sectors without any labor mobility costs. There are several interesting issues that arise when there are labor mobility costs. In this subsection we discuss the most relevant ones.

We begin with the paper by Lee and Wolpin (2006) about the large reallocation of labor from manufacturing to services in the United States over the period from 1968 to 2000. The goals of this paper are to measure the costs associated with sectoral labor reallocation and to assess the relative importance of labor demand and supply factors for sectoral labor reallocation, where labor demand factors are defined as changes in sectoral productivity and relative prices and labor supply factors are defined as changes in demographics, fertility, and educational attainment. To reach these goals, they develop a framework with a detailed labor market. To begin with, there are three occupational choices in each sector: blue collar, white collar, and pink collar (i.e., secretarial, clerical etc.). Moreover, workers differ in their educational attainment and they can accumulate sector-specific and occupation-specific human capital while working. Lastly, there are various types of technological change and the production functions have a constant elasticity of substitution between capital and labor.

Lee and Wolpin (2006) estimate their model using micro data. Their main findings are as follows. First, labor demand factors are the key driving forces behind the reallocation of labor across sectors. In contrast, labor supply factors do not play much of a role. This finding is consistent with the emphasis that our benchmark model puts on technological factors. Second, and in contrast to our benchmark model, the mobility costs associated with moving across sectors are large; for example, the monetary cost of changing sectors can be as large as 75 percent of annual earnings. Moreover, changing occupations within a sector is significantly less costly than changing sectors while maintaining the same occupation.

Lee and Wolpin (2006) carry out several counterfactuals regarding how changes in mobility costs would have affected the evolution of labor market outcomes. Interestingly, they find that if mobility costs had been zero, aggregate productivity would have been higher and the labor market histories of individual workers would have been different, but the evolution of sectoral employment shares and value added shares would not have changed much. The economics behind this result is that with lower mobility costs workers can better allocate their time to the sector in which their idiosyncratic productivity is highest. This raises aggregate productivity and changes the labor market histories of individual workers. However, since it leads to flows of workers in *both* directions, the effect on relative sectoral employment is relatively small. This result suggests that abstracting from mobility costs in our benchmark model does not have large quantitative effects on the sectoral employment allocation.

Lee and Wolpin (2006) also ask what would have happened if sectoral labor mobility had been more costly. They find that while there would have been little effect on trend changes in employment shares, the level of the employment share of services would have shifted upward. This result runs counter to the intuition that increased mobility costs will decrease the flow of workers into the expanding service sector. To understand this, it is important to realize that this intuition is based on how mobility costs affect the response to an unanticipated shock. In contrast, what matters for Lee and Wolpin's exercise are the choices that forward looking new entrants make in the face of the trend that the service sector is becoming more attractive in comparison to the goods sector. If we increase the size of mobility costs, then more entrants move directly into the service sector, instead of first going to the manufacturing sector and later switching to the service sector.

There is more evidence that the role of new entrants is crucial for the labor reallocation across sectors in the context of structural transformation. For example, Kim and Topel (1995) show that during Korea's rapid industrialization almost all of the changes in the sectoral employment shares of agriculture and manufacturing resulted from changes in the behavior of new entrants. As a result, the large decrease in the agricultural employment share and the large increase in the manufacturing employment share were accomplished with little reallocation of

existing workers.⁴⁰ To the extent that new entrants are an important source of labor market flexibility one might conjecture that economies with different rates of growth in the labor force might experience different patterns of structural transformation. However, we are not aware of existing evidence that supports this conjecture.

While some mobility costs might reflect technological factors, there is also the possibility that policies, regulations and institutional factors lead to the barriers to labor mobility. Examples include implicit or explicit firing costs levied on employers, subsidies to establishments in declining industries, entry barriers that make it costly for firms to start up new establishments, generous unemployment benefits or early retirement schemes that are offered to displaced workers, and direct restrictions on the mobility of workers.⁴¹ There are many studies of these types of factors, but most of them make no reference to the process of structural transformation. The reason for this is that most job creation and destruction occurs within rather than across narrow industrial classifications, and so the main effects come from the reallocation of resources across establishments when jobs are created and destroyed.

Three exceptions that study the effects of labor mobility costs in the context of structural transformation are Nickell et al. (2002), Messina (2006) and Hayashi and Prescott (2008). Nickell et al. (2002) examine the correlations between sectoral composition and various policy and institutional factors in a panel data set panel of 14 OECD countries and 5 one-digit industries during the period 1975–94. One of their findings is that countries with more stringent employment protection policies have larger industrial sectors, suggesting that employment protection policies might impede the reallocation of employment from manufacturing into services. Messina (2006) considers the role of entry barriers. One distinguishing feature of structural transformation in Europe is that conditional on aggregate productivity (i.e., output/hour), Europe has a much lower employment share for services than do other rich countries.⁴² Messina argues that this is the result of higher entry barriers in Europe, including such factors as direct costs associated with licensing and indirect costs associated with zoning restrictions or regula-

⁴⁰Matsuyama (1992) and Rogerson (2006) both present models of sectoral reallocation that have this property.

⁴¹China is a clear example of an economy that has direct restrictions on the mobility of workers. Dekle and Vandenbroucke (2012) that these restrictions slowed the Chinese movement out of agriculture.

⁴²This was not apparent in Section 2 since we plotted the service share of hours worked versus per capita income rather than output per hour.

tions that restrict shopping hours, etc. Because the reallocation of workers into services requires additional entry of establishments into the service sector, these barriers retard the movement of economic activity into the service sector. Hayashi and Prescott (2008) study the movement of labor out of agriculture in Japan before World War II. They argue that the prewar patriarchy that forced the son designated as heir to stay in agriculture effectively amounted to a barrier to the movement of labor out of the agriculture sector. Using a standard neoclassical two-sector growth model, they show that the barrier-induced sectoral distortion and the implied lack of capital accumulation account well for the depressed output level of Japan's prewar economy.

Although Lee and Wolpin (2006) incorporated a range of factors that make mobility costly for individual workers, their model still shares the feature of our benchmark model that all labor reallocation was voluntary from the perspective of the worker. A large literature has documented the large earnings losses that older workers face when they are displaced; see, for example, Jacobson et al. (1993). To many policymakers and commentators, the reallocation of labor from manufacturing to services that is part of the process of structural transformation is synonymous with the displacement of older, high-tenure workers in the manufacturing sector and either unemployment or large losses in earnings. While the connection may seem clear cut, direct evidence on this point is much less clear cut. As noted by many authors, most job creation and destruction occurs within narrow industry classifications, and so is not directly related to the reallocation of activity across broad sectors.⁴³

6.3 Goods Mobility

If openness matters for the process of structural transformation in some settings then it follows that the cost of moving goods may influence structural transformation as well through their effect on trade. More interesting is the possibility that transport costs might influence structural transformation in a closed economy setting. One simple idea in this literature stems from noting that while agriculture is predominantly rural, much of the activity outside of agriculture takes place in cities. It follows that food consumed by non-agricultural workers needs to be transported from rural to urban areas. If this is the case then high costs of moving food from

⁴³See for example Davis and Haltiwanger (1992).

rural areas could exert a negative influence on the movement of labor out of agriculture.

Herrendorf et al. (2012b) study this idea in the context of the transport revolution in the US before the civil war, during which the construction of the railroads reduced dramatically the transportation costs to the most fertile farm land in the Midwest. They build a model with two regions (Midwest and Northeast) and three sectors (agriculture, manufacturing, services). Consistent with our benchmark model, their model also allows for both income effects via a subsistence term in the utility from agriculture, and productivity effects in terms of the factors that determine the allocation of labor to agriculture. They show that the reduction in transportation costs between the two regions leads to the settlement of the most fertile farm land in the Midwest, which is followed by a reduction in the agricultural labor force.

Adamopoulos (2011) and Gollin and Rogerson (2010) study this idea further in the context of a static model with agriculture and non-agriculture and different locations. Adamopoulos shows that transportation costs between locations can exert an important influence on the allocation of resources across locations and between agriculture and non-agriculture. Gollin and Rogerson carry out some numerical exercises to suggest that there is a strong interaction between increases in productivity and reductions in transportation costs in terms of their impact on labor moving out of agriculture.

7 Applications of Structural Transformation

In this section we return to the question we posed in the introduction to this chapter: Does incorporating structural transformation into the standard growth model deliver new insights? In other words, is there a substantive payoff to working with versions of the growth model that account for structural transformation? We discuss several issues where changes in the sectoral composition of the economy matter has been shown to matter. We conclude that explicit modeling of structural transformation offers important additional insights in these cases.

7.1 Structural Transformation and Economic Development

Caselli (2005) and Restuccia et al. (2008) argue that the proximate cause of much of the large differences in living standards across countries is attributable to two simple facts: (1) developing countries are much less productive in agriculture relative to developed countries, and (2) developing countries devote much more of their labor to agriculture than do developed countries. These two facts suggest that in order to understand why developing countries are so poor it is of first-order importance to understand the forces that shape the allocation of resources between agriculture and the other sectors. A version of the growth model extended to incorporate structural transformation is the natural framework to be used in this context.

Work by Gollin, Parente, and Rogerson (2002, 2007) illustrates how low agricultural productivity can be the source of large cross-country differences in aggregate productivity. For ease of exposition we focus on the simpler presentation in the 2002 paper, which uses a two-sector version of our benchmark model, with the two sectors being agriculture and non-agriculture. They assume that the population is constant and normalize it to one. Preferences are such that there is a subsistence level \bar{c}_a of agricultural consumption at which individuals are also satiated. The non-agricultural production function is essentially a Cobb–Douglas production function in capital and labor. In contrast, there are two agricultural production functions: a traditional and a modern one.⁴⁴ Both agricultural production functions are linear in labor, though the analysis would be unaffected by assuming a fixed quantity of land and decreasing returns to scale in labor. The traditional production is assumed to be the same across countries and to be sufficiently productive to exactly meet subsistence agricultural needs when all labor is allocated to it. The modern production function has a country-specific TFP parameter and it is the only production function that is subject to technological progress.

In this model, only the agricultural technology with the larger productivity will be used in equilibrium. Initially this is the traditional technology. Since the modern technology is subject to technological progress, at some point the modern technology will replace the traditional technology as the only technology that will be used. The somewhat extreme structure of the model then yields a very simple solution method for determining the equilibrium. Total food

⁴⁴Hansen and Prescott (2002) use a similar assumption but at the aggregate level.

production must be \bar{c}_a . As long as the traditional technology is used, this means that all labor will be in agriculture. When the modern technology starts to dominate the traditional technology, labor will start to flow from agriculture to non-agriculture. With the time series for labor allocations determined, the remainder of the model becomes a standard growth model with an exogenously given process for labor. The growth rate of labor in the non-agricultural sector is completely determined by the exogenous growth rate of labor productivity in the modern agricultural sector. Since all countries have the same output of agriculture, cross-country differences in aggregate output are entirely driven by differences in non-agricultural output.

Several implications follow. First, countries that use the modern technology in agriculture but have low productivity in it will have to devote more labor to agriculture. This leads to less labor, and capital, in non-agriculture, and hence to less aggregate output. Given the observed differences in the share of labor that is allocated to agriculture, Gollin et al. (2002) show that this mechanism can account for a large part of the cross-country differences in aggregate output. This is interesting because in their model the only difference across countries is the level of productivity of agriculture.

Second, assuming that productivity growth rates are constant over time, the model necessarily implies that transition dynamics will be long-lived, thereby addressing a point emphasized by King and Rebelo (1993) that in a standard one-sector growth model transition to the steady state capital level is rapid.⁴⁵ This point does not carry over to the two-sector model because labor allocated to the non-agricultural sector only slowly converges to its asymptotic level. Third, the model implies that in a closed economy setting advances in agricultural productivity are a precondition for growth. This view was a central argument of Schultz (1953), and figured prominently in later contributions by Johnston and Mellor (1961), Johnston and Kilby (1975), and Timmer (1988), among others. More recently, it has taken a central state in the writing of non-economists such as Diamond (1997).⁴⁶

Laitner (2000) considers a similar framework as Gollin et al. (2002) but focuses on a differ-

⁴⁵Chang and Hornstein (2011) make a related point about Korea. They show that two modifications of the one-sector growth model are essential to account for the long-lived transition dynamics since 1960 during which Korea continued to accumulate capital. The first one is to distinguish between agriculture and non-agriculture and to take into account that Korean agriculture used relatively little physical capital. The second essential modification is to model that the relative price of capital remained high during most of the transition dynamics.

⁴⁶See Tiffin and Irz (2006) for a recent empirical assessment.

ent issue. He notes that in the time series data there is evidence of an increase in savings rates early in the industrialization process. Whereas some have argued that the increase in savings rate is the driving force behind the industrialization process, Laitner shows that, in a model of structural transformation, this apparent increase in savings rate is simply an artifact of how NIPA measures saving. Early in the development process most labor is employed in agriculture, and so most savings take the form of realized capital gains in the value of land, which is not recorded as savings by the NIPA. As labor moves out of agriculture and agriculture becomes a smaller part of aggregate output, this issue becomes less important quantitatively. Laitner argues that viewed from the perspective of his model of structural transformation, one should not attach any significance to the apparent increase in savings rates that occur in the early stages of development.

7.2 Structural Transformation and Regional Income Convergence

One of the dramatic secular changes in the US economy over the post World War II period is the convergence of incomes across regions; see, for example, Barro and Sala-i-Martin (1992). In the context of the standard one-sector neoclassical growth model, this convergence in incomes would be attributed to changes in either regional TFP or regional factor accumulation. Caselli and Coleman (2001) show that a model of structural transformation provides a richer understanding of the economic forces at work. The motivation for their analysis is provided by the fact that the convergence in regional incomes between the North and the South of the United States coincided with a dramatic narrowing of regional differences in the employment share in agriculture. They use a model that differs from our benchmark model along several dimensions. First, they consider a two-sector version of the model, with the two sectors being agriculture and non-agriculture. Second, they consider a two-region version of the model, where each region has the same structure as our model and there is free mobility of goods across regions. They assume that the technologies are such that the North has a comparative advantage in manufacturing and the South has a comparative advantage in agriculture. They focus on the special case in which the technologies in manufacturing are the same in both regions and the South has higher in TFP agriculture (for simplicity, they assume that the North has zero TFP in

agriculture). Third, they assume that there are mobility costs in terms of sectoral reallocation of labor. Specifically, all workers begin in the agricultural sector, and they must pay a cost if they are to move to the non-agricultural sector. They interpret this mobility cost as the cost of acquiring skills that are needed in the non-agricultural sector and argue that it is necessary if one is to account for the secular changes in labor allocations and relative wages.

The basic economics of their analysis is the following. When the United States was relatively poor, more of its workers were engaged in agriculture, due to non-homothetic preferences which imply a large share for agricultural expenditures at low levels of income. Because the South had a comparative advantage in agriculture, the South was doing relatively more agriculture. Because of mobility costs, wages were higher in non-agriculture. Putting these features together, incomes were lower in the South. Over time, production technology in non-agriculture advanced, leading to a decline in the share of workers in agriculture. They also posit that in addition mobility costs decreased, therefore leading to convergence between agricultural and non-agricultural wages.⁴⁷

7.3 Structural Transformation and Aggregate Productivity Trends

Our model of structural transformation allows for the possibility that different sectors have different levels as well as growth rates of labor productivity. Herrendorf and Valentinyi (2012) provide evidence from the 1996 Benchmark Study of the Penn World Tables on sectoral TFP differences across countries. They find that there are large sectoral TFP differences relative to the United States not only in agriculture, but also in manufacturing, and that the sectoral TFP differences in these two sectors are much larger than in the service sector. Aggregate labor productivity may then be affected by the sectoral composition of the economy. In particular, to the extent that different countries are at different stages of the process of structural transformation, sectoral reallocation associated with structural transformation could generate significant changes in aggregate productivity growth [Echevarria (1997)]. In principle, episodes of acceleration or slowdown in aggregate productivity growth may occur even if in each country sectoral

⁴⁷In related work, Hnatkovska and Lahiri (2012) show that structural transformation importantly contributed to the narrowing of the urban-rural wage gap in India during 1983–2010.

productivities are growing at constant rates.

In a recent paper, Duarte and Restuccia (2010) have investigated the importance of these effects in a sample of 29 countries for the period of 1956–2004. They employed a somewhat simplified version of our benchmark model in which labor is the only factor of production (and production functions are linear in labor). They assumed that each sector’s labor productivity grows at a constant rate, but that level and growth rates differ across economies as dictated by the data.

The preference structure of Duarte and Restuccia (2010) assumes a period utility function which is a two–good version of (28):

$$C_t = \omega \log(c_{at} - \bar{c}_a) + \omega_n \log(c_{nt})$$

c_{nt} stands for non–agricultural consumption and it is a CES aggregator of manufactured goods and services. Preference parameters are calibrated so as to match the behavior of the US economy and are assumed to be the same across countries. The initial productivity levels of all countries relative to the US are inferred from the model by requiring that it match the observed employment shares in the initial period. Inputting the sectoral productivity growth rates from the data, Duarte and Restuccia (2010) then simulate the model and compute the implied series for aggregate labor productivity.

Even though their model assumes constant productivity growth rates at the sectoral level of each country, it generates large movements in relative aggregate productivity across countries over time. Key to this finding is that differences in the levels and growth rates of labor productivity between rich and poor countries are larger in agriculture and services than in manufacturing. This implies that during the process of structural transformation, the reallocation of labor from agriculture to manufacturing led to a catch up of aggregate productivity relative to the US, and the reallocation from manufacturing to services leads to a falling behind of aggregate productivity relative to the US.

In related research, Bah and Brada (2009) study the countries from Central Europe which have recently entered the European Union. The point of departure of their analysis is the

stylized fact that central planning during communist times resulted in “over–agrarianism” and “over–industrialization”, and the neglect of the service sector in these countries. Bah and Brada document that even today employment in the service sector is considerably smaller in Central Europe than in the core countries of the European Union. Moreover, they find that in all of these countries the service sector has lower TFP than the manufacturing sector. This implies that structural transformation into the service sector will lead to losses in GDP per capita, unless reforms are implemented that make the service sectors more productive.

7.4 Structural Transformation and Hours Worked

Following Prescott (2004), there is a sizeable literature that seeks to understand the large differences in hours worked that have emerged over time between the US and countries in continental Europe. In order to be able to compare hours worked across countries of different size, Prescott divided total hours worked by the working–age population. Prescott used the standard one–sector growth model to demonstrate that changes in labor taxes could account for much of the emerging difference.

Rogerson (2008) argued that a model of structural transformation provides additional insights into the evolution of hours. In particular, he compared the evolution of hours worked per working–age person in the US to those in an aggregate of five continental European economies (Belgium, France, Germany, Italy, Netherlands) since 1956. Whereas hours worked were about 5% higher in Europe in 1956, by 2003 they were more than 30% lower. Looking at the sectoral evolution of hours worked reveals an interesting pattern. During the period in which hours worked in these European economies fell by more than 35% relative to the US, one observes that the relative level of hours worked in the goods sector in Europe fell dramatically, whereas the relative level of hours worked in services remained relatively flat.⁴⁸ One might be tempted to conclude that the key to understanding the relative decline in hours worked in Europe lies in understanding the relative decline in hours worked in the goods sector. However, when one views the sectoral evolution of hours worked in the context of structural transformation one is lead to exactly the opposite conclusion. Specifically, in 1956 Europe was considerably behind

⁴⁸Hours worked in a sector again is defined as total hours worked divided by the working–age population.

the US in terms of development, and consistent with our earlier empirical analysis, had a larger share of hours in the goods sector and a smaller share in the service sector than the United States. By 2003 Europe has basically caught up to the United States in terms of productivity. Holding all else constant, one would expect that the sectoral hours worked distribution in Europe in 2003 would look similar to that in the United States. That is, the process of structural transformation leads us to expect that while hours in the goods sector in Europe should have decreased relative to the US, hours in the service sector in Europe should in fact have increased. Put somewhat differently, the issue of understanding why hours worked are so much lower in Europe reduces to the issue of understanding why the European service sector has failed to grow like its counterpart in the US. In fact, this dynamic was apparent in the hours plots in Figure 2.

In addition to simplifying the analysis by aggregating agriculture and manufacturing to one category and by abstracting from capital, Rogerson's model differs from our benchmark model along two key dimensions: he adds a labor supply decision and he allows for home production, which he assumes to be substitutable with the output of the service sector. His model combines both income and price effects to generate structural transformation. Taking changes in productivity and labor taxes as given, he calibrates the preference parameters so as to match the changes in the US economy between 1956 and 2003, including the change in time devoted to home production.⁴⁹ He then feeds in European values for productivity and taxes in both 1956 and 2003 and examines the ability of the model to account for aggregate and sectoral observations in Europe in 1956 and 2003. Overall, Rogerson finds that the model accounts well for the sectoral European labor allocations.

Rogerson assumes that the utility function is non-homothetic in that it has a subsistence level of goods consumption. This turns out to be important for understanding relative hours worked in Europe in the initial year of his study, 1956. At that time, Europe already had higher tax rates than the US, yet they had higher hours of work. The non-homotheticity acts like a negative income effect, and this effect is larger the lower is aggregate productivity. Given that Europe lagged the US in aggregate productivity in 1956, this effect serves to increase hours in

⁴⁹See Aguiar and Hurst (2007) and Ramey and Francis (2009) for evidence on the decline of home production time in the US.

Europe relative to the US. Additionally, because the model generates structural transformation, Europe devoted more labor to goods production than the US in 1956. Because there are fewer non-market substitutes for goods, this effect also serves to increase the amount of time devoted to market work.

In related work, Ngai and Pissarides (2008) add a home production sector to their earlier model of structural transformation that we have discussed above, Ngai and Pissarides (2007). They showed that over time the model with home production generates a shallow u-shaped curve for hours devoted to market work, and that it leads to the marketization of home production, i.e., the movement of time out of home production and into market production of services. Both of these patterns are found in the US data. The initial decrease in market work is associated with the movement of activity into services, which have better home produced substitutes. But as time advances, a higher rate of growth in the productivity of market produced services relative to home produced services leads to the movement of activity out of the home sector and into the market sector, which results in the increase in market hours.

Another dramatic trend in labor market outcomes has been the rise of female labor force participation. Several authors have argued that the process of structural transformation is an important factor in accounting for this change. The basic idea is that jobs in the goods sector (i.e., agriculture and manufacturing) and the service sector tend to have different weights on various dimensions of labor input. In particular, the goods sector places more emphasis on “brawn” while the service sector places more emphasis on “brains”. If men and women have different relative endowments of these two factors, then the movement of activity from one sector to the other could plausibly affect the desire of women to seek employment in the market sector. Fuchs (1968) noted this explanation for the rise of female labor force participation.⁵⁰

Rendall (2010) builds a two-sector model in which she can quantitatively evaluate the difference between men and women and argues that structural transformation is an important quantitative factor in accounting for the rise of female labor force participation. In related work, Akbulut (2011) also argues that the rise of the service sector has been an important factor

⁵⁰Galor and Weil (1996) also note the changing demands for brain and brawn, though not in the specific context of structural transformation. See also the papers by Goldin (1995, 2006) for additional analysis of the evolution of female labor force participation patterns.

in accounting for the rise of female labor force participation in the US, but the key reallocation in her model is the movement of labor out of home produced services and into market produced services in response to a more rapid rate of technological progress in market services relative to home produced services.

Olivetti (2012) provides evidence from a large sample of developed and developing countries that connects the u-shaped profile for female labor force participation to structural transformation, extending the earlier work by Goldin (1995). Specifically, she finds that as countries develop, the share of women who work in agriculture relative to all working women decreases faster than the share of men who work in agriculture relative to all working men, the share of women who work in services increases faster than the share of men, and the share of women who work in manufacturing remains flatter than the share of men.

7.5 Structural Transformation and Business Cycles

There are many different ways in which theories of structural transformation and business cycles might overlap. One idea which frequently recurs is that some business cycles are the result of periods of greater reallocation of economic activity across sectors. To the extent that this reallocation of activity occurs at the broad sectoral level emphasized by models of structural transformation, structural transformation and business cycles could be intimately related.

Using the search model of Lucas and Prescott (1974) as a reference point, Lilien (1982) argued that if it takes time for labor to move from one sector to another, then periods of above average reallocation will also be periods of above average unemployment. He then argued that business cycles in the post World War II US were characterized as periods of above average reallocation of labor among two-digit sectors, as measured by the variance in employment growth rates at the two sector level. However, subsequent work by Abraham and Katz (1986) argued that Lilien's statistical finding about changes in the variance of sectoral growth rates could simply be due to the fact that sectors vary in their response to aggregate shocks, and that data on vacancies supported this latter explanation over the sectoral shifts explanation.

The idea of Lilien (1982) has experienced a recent resurgence in popularity in the face of the current recession, with various economists suggesting that "mismatch" is an important

element of the current high level of unemployment, and that the decline of broad sectors such as manufacturing and construction is an important element of this mismatch. However, despite its popularity, recent empirical research by Sahin et al. (2011) and Herz and van Rens (2011) finds little evidence for this explanation.

We note that even if reallocation were concentrated during recessions, it would not follow that recessions are caused by the reallocation. Rather, it may be that recessions are caused by a second factor, and that the decisions that lead to reallocation are made in such a way that reallocation coincides with the recession. That is, for example, it may be that steel mills go out of business permanently during recessions, but this may simply reflect that the optimal timing of exit for a steel mill is during a downturn in economic activity. Rogerson (1991) argued that movement out of agriculture in the US has been concentrated during upturns in economic activity, whereas the movement of workers out of manufacturing has been concentrated during downturns.

Even if structural transformation is not the cause of business cycles, it may still exert an influence on business cycles. For example, to the extent that value added varies in volatility across sectors, the sectoral composition of aggregate output is a potentially important determinant of business cycle fluctuations. In what follows, we mention two examples of this idea.

The first example is Da Rocha and Restuccia (2006), who disaggregate the economy into agriculture and non-agriculture and document that indeed there are important differences between the two sectors. In particular, they find that the agricultural sector is more volatile than the rest of the economy, is not correlated with the rest of the economy, and has counter-cyclical employment. They show that this implies that countries with a larger agricultural sector have more volatile aggregate output and less volatile employment. Moreover, it implies that as structural transformations out of agriculture occur, business cycle properties across countries converge.

The second example of how the sectoral composition matters is due to Carvalho and Gabaix (2013) and Moro (2012). They disaggregate the economy into services and manufacturing, largely ignoring agriculture. They document that the volatility of services is lower than in manufacturing. Moro (2012) argues that the reason for this is that the share of intermediate inputs is larger in manufacturing than in services. Irrespective of why the volatilities differ

between the two sectors, the implication is that the volatility of aggregate output declines as the share of services increases along the path of structural transformation. Carvalho and Gabaix (2013) find that this accounts for most of the “great moderation” and its recent undoing. In particular, the great moderation is due to a decreasing share of manufacturing between 1975 and 1985, and its recent undoing in the form of rising aggregate volatility is due to the increase of the size of the financial sector.

7.6 Structural Transformation and Wage Inequality

One of the dramatic secular changes in the US economy over the last fifty years has been the marked increase in wage inequality that is associated with the return to skill. In a recent paper, Buera and Kaboski (2012a) argue that this rising return to skill is intimately connected to the structural transformation of economic activity towards services. They document in time series data the same threshold behavior of value added in services that we have found above, that is, there is a threshold for per capita income at which one observes an acceleration in the increase in the value added share for services. Interestingly, at that threshold there is also an increase in the fraction of the workforce that becomes skilled and of the skill premium. In the context of the US they also document that the entire rise in the service sector’s share of value added in the last fifty years is accounted for by growth in sub-sectors that have higher than average shares of skilled labor. They go on to build a model that links these patterns as the outcome of structural transformation that is driven by neutral productivity growth.

We previously described some general features of the framework that they use. Relative to our earlier discussion, the key modification in this paper is that there are two types of labor: skilled and unskilled. Skilled labor is specialized to a particular service, is costly to acquire and is subject to an increasing cost curve. To capture the fact that home production is necessarily less specialized, they assume that skilled labor is equivalent to unskilled labor in home production. Services differ in “complexity”, where complexity captures both the amount of labor that is required to produce them and the relative productivity advantage of skilled labor in producing the service.

As the economy develops it produces services that are increasingly complex, thereby creat-

ing additional incentives for both market production of wants and skill accumulation. Because there is an upward sloping supply curve for skilled workers, the skill premium is also increasing. The structure of their model is such that the relative advantage of skilled labor in producing more complex services only emerges beyond a critical threshold level of complexity, so that these patterns also emerge beyond a threshold. A key fact that this model is able to account for that our benchmark model cannot is that this model predicts that the share of services in nominal value added is flat below some threshold.

An important implication of this work is that adding the different roles of human capital in various activities is an important ingredient in understanding some key features of structural transformation.

8 Conclusion

Our goal in this chapter has been to summarize the basic facts about structural transformation, and to present simple versions of the growth model that serve as the benchmark models being used to organize our thinking about these facts. Much of the early literature has focused on trying to identify multi-sector versions of the growth model that can generate structural transformation while simultaneously generating balanced growth. While the search for specifications that can simultaneously yield structural transformation at the sectoral level and balanced growth have proven to be useful in organizing research, we believe that focusing on frameworks that yield exact balanced growth is probably overly restrictive. The literature should instead focus on building models that can quantitatively account for the properties of structural transformation and in the process assess the importance of various economic mechanisms. We use this concluding section to highlight what we view as important priorities for future research in this area.

While we have a substantial amount of data regarding the process of structural transformation in today's advanced economies, it would be good to know more about the nature of structural transformation in today's less developed economies. To what extent are they following different paths from today's developed economies? And if so, what are the factors that give

rise to these differences?

Two economies of particular current interest in this regard are China and India, both because of their size and because they have been experiencing very rapid growth. What role does structural transformation play in these countries' growth? Dekle and Vandenbroucke (2012) have studied structural transformation in China during 1978–2003. They have found that differential sectoral productivity growth and the reduction of the relative size of the Chinese government caused most of the structural transformation, but that mobility frictions (like the hukou system) slowed the movement out of agriculture. Rubina (2012) has studied structural transformation in India during 1980–2005. Contrary to the patterns that we have documented above, she has found that TFP growth was fastest in services. Moreover, she has found that a three-sector model can account for changes in sectoral value added but not in employment shares.

The growth miracle episode in South Korea has also attracted recent attention, specifically as it relates to the issue of openness, structural transformation and growth. Betts et al. (2011), Sposi (2011), and Teignier (2012) have studied structural transformation in South Korea during its growth miracle. They argue that international trade accelerated the transition out of agriculture into industry and services. Teignier (2012) argues in addition that international trade could have played an even larger role if South Korea had not simultaneously introduced agricultural protection policies.

Üngör (2011) has compared Latin America with East Asia. He has found that differences in sectoral productivity growth rates account well for the different sectoral reallocations in the two regions, and in particular for the fact that Latin America has moved much more slowly out of agriculture.

We think that more quantitative case studies of structural transformation in currently poor countries will help to sharpen our understanding of the forces behind structural transformation in such countries. Additionally, we think it will be useful to think about the factors that influence productivity growth. Virtually all of the literature on structural transformation takes productivity changes as given, and effectively considers the implications of the exogenously given paths for productivity on the process of structural transformation. But if the paths of productivity differ significantly across countries, then it is important to ask what factors are

responsible for these differences? If the differences are more pronounced in particular sectors in particular countries, what are the factors that account for this? Is it policies that influence the diffusion of technology, or perhaps policies that generate misallocation of inputs across producers?

Moving forward, we also think it will be useful to refine the standard three–sector focus of the literature. As today’s advanced economies are increasingly dominated by services, it will be important to distinguish between different activities within services. For example, education and health care are very different activities than retail trade, in that they both represent an investment and tend to use very different skill intensities for the labor that they employ. The work of Jorgenson and Timmer (2011) and Duarte and Restuccia (2012) is a first step in this direction. Using data from EU KLEMS, Jorgenson and Timmer (2011) document for the European Union, Japan, and the US that there is substantial heterogeneity among services. Personal, finance, and business services have low productivity growth and increasing shares in employment and GDP whereas distribution services have rapid productivity growth and constant shares. Using data from the International Comparison Program 2005, Duarte and Restuccia (2012) study for a large cross section of countries the difference between traditional and non–traditional services where traditional services comprise mostly non–market services such as domestic and household services, education, health, and housing and non–traditional services comprise communication and transport services, insurance and financial services, and recreational and cultural services. For traditional services, they find that the relative price increases and the real expenditure share decreases with income whereas for non–traditional services they find the opposite. An important task for future work is to build models that are consistent with these facts and to explore to implications that these models have for structural transformation and for aggregate outcomes.

There are many issues that we have not addressed or only touched upon in passing. One such issue is the role that human capital plays in the process of structural transformation. Buera and Kaboski (2012a) emphasize the fact that effectively all of the growth in the service sector in the US in the post WW II period occurs in high skill services. While they emphasized the role of human capital in the movement of resources from the goods producing sector to the service sector, it is also plausible that human capital may be important in understanding the movement

of workers between the agricultural sector and the non-agricultural sector. In fact, the work by Caselli and Coleman (2001) that we described earlier is one paper that emphasized the role that human capital plays in this part of the structural transformation process. Recent work by Herrendorf and Schoellman (2012) provides additional evidence on this point. Using the CPS, they document for the US that wages per hour are considerably higher in non-agriculture than in agriculture. They show that this is accounted for by two main facts: non-agricultural workers are positively selected in that they have more years of schooling; the returns to schooling and experience are higher in non-agriculture. An open question is to what extent similar findings hold in poorer countries than the US. In a recent paper, Lagakos and Waugh (2013) argue that accounting for the heterogeneous quality of labor across sectors is important in understanding the fact that poor countries seem to have particularly low labor productivity in agriculture.

Another issue that we have not addressed is the role of industrial policy, broadly conceived. Specifically, we have chosen to discipline the analysis by assuming sectoral production functions with constant returns to scale and by abstracting from spillovers or externalities. As a result, we have interpreted structural transformation as a feature of the efficient equilibrium path, implying that there is no meaningful role for government policy. While our model framework can be used to understand how particular policies might distort the allocation of resources across sectors, there is no positive prescription for policy.

There is a sizeable literature that discusses structural transformation when there are increasing returns to scale and the equilibrium path is inefficient; see Matsuyama (2008) for specific references. The typical assumption in this literature is that non-agricultural production is subject to increasing returns, which accrue at the sectoral level, perhaps as the result of learning by doing, and which are not taken into account by households and firms. Multiple steady states then arise naturally and initial conditions determine the equilibrium path, and in particular whether the economy ends up in a “poverty trap”, that is, a steady state with low GDP per capita and the majority of the labor force in agriculture. These types of models suggest that policy may provide the “big push” that lets the economy escape from its poverty trap and leads to industrialization and self-sustaining economic growth. We have not discussed this theoretical possibility in more detail above because the empirical evidence on the success of

“big–push” policies in particular, and industrial policies more generally, is mixed at best. But more generally, the extent to which externalities, public goods, market power, or other factors associated with inefficient equilibrium outcomes shape the process of structural transformation remains largely unresolved.

Appendix A: Data Sources and Sector Assignments

Historical data 1800–2008

- Data source: GDP per capita at international dollars
 - Data on GDP per capita at 1990 international dollars are from Maddison (2010) for all countries and most years. There are some years in the early 19th century for Belgium, Netherlands, Sweden, United Kingdom, and the United States when there are data on value added and employment shares, but Maddison does not report data on GDP per capita. We calculated GDP per capita at international dollars for these years in the following way. From alternative sources, we first calculated real GDP per capita for the missing years and for the first year for which Maddison’s data is available. We then calculated the growth rates between the missing years and the first year for which the Maddison data is available. Lastly, we combined the growth rates with the Maddison’s data to calculate the per capita GDP at international dollars for the missing years. Next we list the data sources for these calculations.
 1. *Belgium*. 1835–1845: real GDP from Groningen Growth and Development Centre, Historical National Accounts Database 2009, and population from Maddison (2010).
 2. *Netherlands*. 1807–1830: real GDP per capita from Smits et al. (2007).
 3. *Sweden*. 1800–1820: real GDP per capita from Krantz and Schn (2007).
 4. *United Kingdom*. 1800–1830: real GDP per capita from Clark (2009).
 5. *United States*. Louis Johnston and Samuel H. Williamson, “What Was the U.S. GDP Then?” MeasuringWorth, 2011.
- Data source: Value added at current prices
 - *Belgium*. 1835–1990: Groningen Growth and Development Centre, Historical National Accounts Database 2009. 1991–2007: EU KLEMS 2009.

- *Spain*. 1885–1940: Groningen Growth and Development Centre, Historical National Accounts Database 2009, 1953–2004: Groningen Growth and Development Centre 10–sector Database 2007.
 - *Finland*. 1860–2001: Groningen Growth and Development Centre, Historical National Accounts Database 2009.
 - *France*. 1815–1938: Groningen Growth and Development Centre, Historical National Accounts Database 2009, 1950–1960: Mitchell (2007) Table J2, 1970–2005: Groningen Growth and Development Centre 10–sector Database 2007.
 - *Japan*. 1885–1940: Groningen Growth and Development Centre, Historical National Accounts Database 2009, 1953–2004: Groningen Growth and Development Centre 10-sector Database 2007.
 - *Korea*. 1911–1940: Groningen Growth and Development Centre, Historical National Accounts Database 2009, 1953–2005: Groningen Growth and Development Centre 10–sector Database 2007.
 - *Netherlands*. 1807–1913: Smits et al. (2007), 1970–2005: Groningen Growth and Development Centre 10–sector database, August 2008.
 - *Sweden*. 1800–2000: Krantz and Schn (2007), 2000–2005: Groningen Growth and Development Centre 10-sector Database, August 2008.
 - *United Kingdom*. 1801, 1941–1851: Broadberry et al. (2011) Table 8–9, 1811–1831, 1860–1910, 1950: Mitchell (2007) Table J2, 1920–1938: Feinstein (1972) Table 9, 1960–2005: Groningen Growth and Development Centre 10–sector Database 2007.
 - *United States*. 1800–1900: Agriculture and Manufacturing, Gallman (1960), Services, Gallman and Weiss (1969), 1909–1918: King (1930), 1919–1928: Kuznets et al. (1941), 1929–1946: Carter et al., eds (2006) Table Ca35–53, 1947–2008: Value Added by Industry, Gross Domestic Product by Industry Accounts, Bureau of Economic Analysis.
- Data source: Employment

- *Belgium*. 1846–1961: Mitchell (2007) Table B1, 1970–2007: EU KLEMS 2009.
 - *Spain*. 1860–1964: Mitchell (2007) Table B1, 1970–2007: EU KLEMS 2009.
 - *Finland*. 1805–1960: Mitchell (2007) Table B1, 1970–2007: EU KLEMS 2009.
 - *France*. 1856–1968: Mitchell (2007) Table B1, 1970–2007: EU KLEMS 2009.
 - *Korea*. 1953–2005: Groningen Growth and Development Centre 10–sector Database 2007.
 - *Netherlands*. 1807–1913: Smits et al. (2007), 1920–1947: Mitchell (2007) Table B1, 1970–2005: Groningen Growth and Development Centre 10–sector Database 2008.
 - *Sweden*. 1850–2000: Krantz and Schn (2007), 2000–2005: Groningen Growth and Development Centre 10–sector Database 2008.
 - *United Kingdom*. 1801, 1813–1820 average assigned to 1817, 1851: Broadberry et al. (2011) Table 1 and Table 12, 1841: Mitchell (2007) Table B1, 1861–1938: Feinstein (1972) Table 59–60, 1948–2005: Groningen Growth and Development Centre 10–sector Database 2007.
 - *United States*. 1840–1920: Carter et al., eds (2006) Table Ba814–830, 1929–2008: NIPA Table 6.8 Persons Engaged in Production, Bureau of Economic Analysis.
- Sector assignments
 1. Agriculture corresponds to the sum of International Standard Industrial Classification (ISIC) sections A–B. If ISIC classification was not available, we assigned industries to agriculture if the source table heading said “Agriculture” or “Agriculture, forestry and fishing”
 2. Manufacturing corresponds to the sum of ISIC sections C, D, F and includes mining, manufacturing, and construction. If ISIC classification was not available, we assigned industries to manufacturing if the source table heading said “Mining” or “Extractive industries” or “Manufacturing” or “Construction” or “Electricity, Gas and Water Supply” or “Utilities”.

3. Services correspond to the sum of ISIC sections E, G–P and include utilities, wholesale, retail trade, hotels and restaurants, transport, storage and communication, finance, insurance, real estate, business services, and community social and personal services. If ISIC classification was not available, we assigned industries to services if the source table heading said “Commerce” or “Finance” or “Trade” or “Transport” or “Communication” or “Services”

EU KLEMS 2009

- Data sources (EU KLEMS series code in brackets)
 1. Employment
 - Total hours worked by persons engaged in millions (H.EMP)
 2. Value added
 - Gross value added at current basic prices (VA)
- Sector assignment
 1. Agriculture corresponds to the sum of International Standard Industrial Classification (ISIC) sections A–B.
 2. Manufacturing corresponds to the sum of ISIC sections C, D, F and includes mining, manufacturing, and construction.
 3. Services correspond to the sum of ISIC sections E, G–P and includes utilities, wholesale, retail trade, hotels and restaurants, transport, storage and communication, finance, insurance, real estate, business services, and community social and personal services.

World Development Indicators 2010

- Data sources (WDI series code in brackets)
 1. Employment

- Employment in agriculture (% of total employment) (SL.AGR.EMPL.ZS)
 - Employment in industry (% of total employment) (SL.IND.EMPL.ZS)
 - Employment in services (% of total employment) (SL.SRV.EMPL.ZS)
2. Value added
- Agriculture, value added as % of GDP (NV.AGR.TOTL.ZS)
 - Industry, value added as % of GDP (NV.IND.TOTL.ZS)
 - Services, etc., value added as % of GDP (NV.SRV.TETC.ZS)
- Oil production
 1. Oil rents as % of GDP, (NY.GDP.PETR.RT.ZS)
 - Sector assignment
 1. Agriculture corresponds to the sum of ISIC divisions 1–5 and includes forestry, hunting, and fishing, as well as the cultivation of crops and livestock production.
 2. Manufacturing corresponds to the category “Industry” in the WDI, which is the sum of ISIC divisions 10–45 and includes mining, manufacturing, construction, electricity, water, and gas.
 3. Services correspond to the sum of ISIC divisions 50–99 and include value added in wholesale and retail trade (including hotels and restaurants), transport, and government, financial, professional, and personal services (such as education), health care, and real estate services. They also include imputed bank service charges, import duties, and statistical discrepancies as well as discrepancies arising from rescaling.

National Accounts of the United Nations Statistics Division

- Data sources
 1. Gross value added by economic activity at current prices in national currency
- Sector assignment

1. Agriculture corresponds ISIC sections A–B.
2. Manufacturing corresponds to the sum of ISIC sections C–F and includes mining, manufacturing, utilities, and construction.
3. Services correspond to the sum of ISIC sections G–P and includes wholesale, retail trade, hotels and restaurants, transport, storage and communication, finance, insurance, real estate, business services, and community social and personal services.

Historical Consumption Shares UK and US

- Data source: GDP per capita at international dollars at 1990 international dollars are from Maddison (2010)
- Data source: US Consumption share in current prices
 - 1900–1928: Carter et al., eds (2006)
 - 1929–2008: BEA
- Data source: UK Consumption share in current prices
 - 1900–1964: Feinstein (1972)
 - 1965–2008: Office of National Statistics (ONS)

Penn World Tables

- Data source: PWT6.3 (PWT series code in brackets)
 1. Real Gross Domestic Product per Capita Relative to the United States (G-K method, current price) (y)
 2. Real GDP per capita in constant prices: chain series (rgdpch)
 3. Real GDP per worker in constant prices: chain series (rgdpwok)
 4. Population (pop)

- Data source: PWT benchmark 1980
 - Sector assignment
 1. Agriculture corresponds to the sum of PWT80 items 1–50
 2. Manufacturing corresponds to the sum of PWT80 items 51–54, 56–58, 63–66, 68–78, 81–83, 91–93, 95–97, 103–108, 112–113, 118–122
 3. Services correspond to the sum of PWT items 55, 59–62, 67, 79–80, 84–90, 94, 98–102, 109–111, 114–118, 123–125

- Data source: PWT benchmark 1985
 - Sector assignment
 1. Agriculture corresponds to the sum of PWT80 items 1–41
 2. Manufacturing corresponds to the sum of PWT80 items 42–47, 49–51, 56–61, 63–68, 70–72, 75–77, 82–84, 86–87, 94–97, 101, 107–109
 3. Services correspond to the sum of PWT items 48, 52–55, 62, 69, 73–74, 78–81, 85, 88–93, 98–100, 102–106

- Data source: PWT benchmark 1996
 - Sector assignment
 1. Agriculture corresponds to bread and cereals, meat, fish, milk, cheese and eggs, oils and fats, fruit, vegetables and potatoes, other food, non-alcoholic beverages, alcoholic beverages.
 2. Manufacturing corresponds to tobacco, clothing including repairs, footwear including repairs, fuel and power, furniture, floor coverings and repairs, other household goods incl. household textiles, household appliances and repairs, personal transportation equipment.
 3. Services correspond to gross rent and water charges, medical and health services, operation of transportation equipment, purchased transport services, com-

munication, recreation and culture, education, restaurants, cafes and hotels, other goods and services.

OECD Consumption Expenditure Data

- Data source:
 - Final consumption expenditure of households, national currency, current prices, OECD National Accounts Statistics. This data set includes the final consumption expenditure of households broken down by the COICOP (Classification of Individual Consumption According to Purpose) classification and by durability.
- Sector assignment (COICOP codes in brackets)
 1. Food: “Food and non–alcoholic beverages” (P31CP010)
 2. Manufactured goods: “Durable goods” plus “Semi–durable goods” plus “Non–durable goods” minus “Food and non–alcoholic beverages” (P311B+P312B+P313B–P31CP010)
 3. Services: Services (P314B)
- Construction of the data for E7 countries (Austria, Denmark, Finland, France, Italy, Netherlands, United Kingdom) for the period 1980–2009. Consumption expenditure data are from the National Accounts of Eurostat both in local currency and euro. Then, for each year and each country, a conversion rate between local currency and euro was calculated by dividing total consumption expenditures in local currency with total consumption expenditures in euros. The three expenditure items expressed in local currency were converted into euros using this conversion rate, and then they were aggregated.

Real GDP per capita at 1990 international \$

- Prior to 1970 the data on GDP per capita at 1990 international dollars are from Maddison (2010) for all years and countries if it was available,

- After 1970 we constructed real GDP per capita at 1990 international \$ in the following ways. The data on GDP per capita at 1990 international dollars for the United States were taken from Maddison (2010). The real GDP per capita of the United States was multiplied by the data on real GDP per capita relative to the United States to calculate the real GDP per capita at 1990 international \$ for each country and each year.

Appendix B: Panel Regressions

To get a balanced panel, we only include countries with data over the entire period 1970–2007. In addition, we restrict the sample in three ways: we exclude countries in which the average ratio of oil rent to GDP exceeds 20% during 1970–2007;⁵¹ we exclude countries with average populations of fewer than a million during 1970–2007; we exclude the former communist countries. The reason for these exclusion criteria is that the sector composition in these countries may be distorted. This leaves 103 countries.

Table 1: Panel Data Analysis Agriculture, 1970–2007

	Dependent variable: Agricultural share in value added					
	(1)	(2)	(3)	(4)	(5)	(6)
log GDP per capita	-0.121** (0.001)	-0.489** (0.021)	0.450* (0.184)	-0.126** (0.015)	-0.396** (0.067)	0.169 (0.274)
(log GDP per capita) ²		0.022** (0.001)	-0.096** (0.022)		0.017** (0.004)	-0.056 (0.035)
(log GDP per capita) ³			0.005** (0.001)			0.003* (0.001)
Country fixed effects	No	No	No	Yes	Yes	Yes
R^2	0.751	0.783	0.786	0.751	0.781	0.784
N	3914	3914	3914	3914	3914	3914

Notes: Heteroscedasticity robust standard errors in parentheses. Significance levels are indicated by † $p < 0.10$, * $p < 0.05$, ** $p < 0.01$.

⁵¹The oil-rent-to-GDP ratio is taken from the WDI.

Table 2: Panel Data Analysis Manufacturing, 1970–2007

	Dependent variable: Manufacturing share in value added					
	(1)	(2)	(3)	(4)	(5)	(6)
log GDP per capita	0.043** (0.001)	0.447** (0.021)	-1.196** (0.144)	0.054** (0.017)	0.497** (0.078)	-1.252** (0.446)
(log GDP per capita) ²		-0.025** (0.001)	0.182** (0.018)		-0.028** (0.005)	0.198** (0.058)
(log GDP per capita) ³			-0.009** (0.001)			-0.009** (0.002)
R^2	0.234	0.331	0.352	0.234	0.331	0.348
N	3914	3914	3914	3914	3914	3914

Notes: Heteroscedasticity robust standard errors in parentheses. Significance levels are indicated by [†] $p < 0.10$, * $p < 0.05$, ** $p < 0.01$.

Table 3: Panel Data Analysis Services, 1970–2007

	Dependent variable: Service share in value added					
	(1)	(2)	(3)	(4)	(5)	(6)
log GDP per capita	0.078** (0.001)	0.041* (0.019)	0.745** (0.170)	0.072** (0.012)	-0.101 (0.089)	1.084* (0.417)
(log GDP per capita) ²		0.002* (0.001)	-0.086** (0.021)		0.011 [†] (0.006)	-0.142* (0.055)
(log GDP per capita) ³			0.004** (0.001)			0.006** (0.002)
R^2	0.493	0.493	0.496	0.493	0.485	0.476
N	3914	3914	3914	3914	3914	3914

Notes: Heteroscedasticity robust standard errors in parentheses. Significance levels are indicated by [†] $p < 0.10$, * $p < 0.05$, ** $p < 0.01$.

References

- Abraham, Katharine and Larry Katz**, “Cyclical Unemployment: Sectoral Shifts or Cyclical Unemployment?,” *Journal of Political Economy*, 1986, 94, 507–522.
- Acemoglu, Daron and Veronica Guerrieri**, “Capital Deepening and Non–Balanced Economic Growth,” *Journal of Political Economy*, 2008, 116, 467–498.
- Adamopoulos, Tasso**, “Transportation Costs, Agricultural Productivity, and Cross–Country Income Differences,” *International Economic Review*, 2011, 52, 489–521.
- Aguiar, Mark and Eric Hurst**, “Measuring Leisure: The Allocation of Time Over Five Decades,” *Quarterly Journal of Economics*, 2007, 122, 969–1006.
- Akbulut, Rahsan**, “Sectoral Changes and the Increase in Women’s Labor Force Participation,” *Macroeconomic Dynamics*, 2011, 15, 240–264.
- Alvarez–Cuadrado, Francisco and Markus Poschke**, “Structural Change out of Agriculture: Labor Push versus Labor Pull,” *American Economic Journal: Macroeconomics*, 2011, 3, 127–158.
- Alvarez-Cuadrado, Francisco, Ngo Van Long, and Markus Poschke**, “Capital–Labor Substitution, Structural Change, and Growth,” Manuscript, McGill University, Montreal 2012.
- Bah, El–Hadj**, “Structural Transformation in Developed and Developing Countries,” Manuscript, Arizona State University., Tempe, AZ 2008.
- Bah, El–hadj and Josef C. Brada**, “Total Factor Productivity Growth, Structural Change and Convergence in Transition Economies,” *Comparative Economic Studies*, 2009, 51, 421–446.
- Bar, Michael and Oksana Leukhina**, “Demographic Transition and Industrial Revolution: A Macroeconomic Investigation,” *Review of Economic Dynamics*, 2010, 13, 424–451.
- Barro, Robert J. and Xavier Sala-i-Martin**, “Convergence,” *Journal of Political Economy*, 1992, 100, 223–251.

- Baumol, William J.**, “Macroeconomics of Unbalanced Growth: The Anatomy of the Urban Crisis,” *American Economic Review*, 1967, 57, 415–426.
- Betts, Caroline M., Rahul Giri, and Rubina Verma**, “Trade, Reform, and Structural Transformation in South Korea,” Manuscript, University of Southern California 2011.
- Blundell, Richard**, “Consumer Behaviour: Theory and Evidence – A Survey,” *Economic Journal*, 1988, 98, 16–65.
- Boppart, Timo**, “Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences,” Working Paper 2, University of Zürich 2011.
- Broadberry, Stephen, Bruce M. S. Campbell, and Bas van Leeuwen**, “The Sectoral Distribution of the Labour Force and Labour Productivity in Britain,” Manuscript, London School of Economics 2011.
- Buera, Francisco J. and Joseph P. Kaboski**, “Can Traditional Theories of Structural Change Fit The Data?,” *Journal of the European Economic Association*, 2009, 7, 469–477.
- **and** ———, “The Rise of the Service Economy,” *American Economic Review*, 2012, 102, 2540–2569.
- **and** ———, “Scale and Origins of Structural Change,” *Journal of Economic Theory*, 2012, 147, 684–712.
- Carter, Susan B., Scott Sigmund Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, and Gavin Wright, eds**, *Historical Statistics of the United States, Earliest Times to the Present: Millennial Edition*, New York: Cambridge University Press, 2006.
- Carvalho, Vasco M. and Xavier Gabaix**, “The Great Diversification and its Undoing,” *forthcoming: American Economic Review*, 2013.
- Caselli, Francesco**, “Accounting for Cross-Country Income Differences,” in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1A, Amsterdam and New York: North Holland, 2005, chapter 9, pp. 679–742.

—— and **Wilbur John Coleman**, “The U.S. Structural Transformation and Regional Convergence: A Reinterpretation,” *Journal of Political Economy*, 2001, 109, 584–616.

Chang, Yongsung and Andreas Hornstein, “Transition Dynamics in the Neoclassical Growth Model: The Case of South Korea,” Working Paper 11–04, Federal Reserve Bank of Richmond, Richmond 2011.

Chenery, Hollis B., “Patterns of Industrial Growth,” *American Economic Review*, 1960, 50, 624–653.

Clark, Colin, *The Conditions of Economic Progress*, 3 ed., London: Macmillan, 1957.

Clark, Gregory, “The Macroeconomic Aggregates for England, 1209-2008,” manuscript, University of California, Davis 2009.

Da Rocha, Jos Maria and Diego Restuccia, “The Role of Agriculture in Aggregate Business Cycles,” *Review of Economic Dynamics*, 2006, 9 (3), 455–482.

Davis, Steven J. and John C. Haltiwanger, “Gross Job Creation, Gross Job Destruction, and Employment Reallocation,” *Quarterly Journal of Economics*, 1992, 107, 819–63.

Dekle, Robert and Guillaume Vandenbroucke, “A Quantitative Analysis of China’s Structural Transformation,” *Journal of Economic Dynamics and Control*, 2012, 36, 119–135.

Dennis, Benjamin N. and Talan B. Iscan, “Engle versus Baumol: Accounting for Structural Change using two Centuries of U.S. Data,” *Explorations in Economics History*, 2009, 46, 186–202.

Diamond, Jared M., *Guns, Germs, and Steel: The Fates of Human Societies*, New York: W.W. Norton, 1997.

Duarte, Margarida and Diego Restuccia, “The Role of the Structural Transformation in Aggregate Productivity,” *Quarterly Journal of Economics*, 2010, 125, 129–173.

—— and ——, “Relative Prices and Sectoral Productivity,” Manuscript, University of Toronto 2012.

Eaton, Jonathan and Samuel Kortum, “Thechology, Geography, and Trade,” *Econometrica*, 2002, 70, 1741–1779.

Echevarria, Cristina, “Changes in Sectoral Composition Associated with Economic Growth,” *International Economic Review*, 1997, 38, 431–452.

Feinstein, Charles H., *National Income, Expenditure and Output of the United Kingdom 1855–1965*, Cambridge: Cambridge University Press, 1972.

Foellmi, Reto and Josef Zweimüller, “Structural Change, Engel’s Consumption Cycles, and Kaldor’s Facts of Economic Growth,” *Journal of Monetary Economics*, 2008, 55, 1317–1328.

Fuchs, Victor, *The Service Economy*, New York: Columbia University Press, 1968.

Gallman, Robert E., “The United States Commodity Output, 1839-1899,” in William N Parker, ed., *Trends in the American Economy in the Nineteenth Century*, NBER Studies in Income and Wealth, Princeton: Princeton University Press, 1960.

——— **and Thomas J. Weiss**, “Production and Productivity in the Service Industries,” in Victor R. Fuchs, ed., *The Service Industries in the Nineteenth Century*, NBER Studies in Income and Wealth, New York: Columbia University Press, 1969.

Galor, Oded and David N. Weil, “The Gender Gap, Fertility, and Growth,” *American Economic Review*, 1996, 86, 374–387.

Goldin, Claudia, “The U-Shaped Female Labor Force Function in Economic Development and Economic History,” in T. Paul Schultz, ed., *Investment in Women’s Human Capital and Economic Development*, Chicago, IL: The University of Chicago Press, 1995.

———, “Ely Lecture: The Quiet Revolution that Transformed Women’s Employment, Education, and Family,” *American Economic Review Papers and Proceedings*, 2006, 96, 1–21.

Gollin, Douglas and Richard Rogerson, “Agriculture, Roads and Economic Development in Uganda,” 2010. Working Paper.

———, **Stephen L. Parente, and Richard Rogerson**, “The Role of Agriculture in Development,” *American Economic Review, Papers and Proceedings*, 2002, 92, 160–164.

———, ———, and ———, “The Food Problem and the Evolution of International Income Levels,” *Journal of Monetary Economics*, 2006, 54, 1230–1255.

Greenwood, Jeremy and Ananth Seshadri, “Technological Progress and Economic Transformation,” in Philippe Aghion and Steven Durlauf, eds., *Handbook of Economic Growth*, Vol. 1B, Amsterdam and New York: North Holland, 2005, chapter 19, pp. 1225–1273.

———, **Zvi Hercowitz, and Per Krusell**, “Long-run Implication of Investment-Specific Technological Change,” *American Economic Review*, 1997, 87, 342–362.

Hall, Robert E. and Charles I. Jones, “The Value of Life and the Rise in Health Spending,” *The Quarterly Journal of Economics*, 2007, 122, 39–72.

Hansen, Gary .D. and Edward C. Prescott, “Malthus to Solow,” *The American Economic Review*, 2002, 92 (4), 1205–1217.

Hayashi, Fumio and Edward C. Prescott, “The Depressing Effect of Agricultural Institutions on the Prewar Japanese Economy,” *Journal of Political Economy*, 2008, 116, 573–632.

Herrendorf, Berthold and Ákos Valentinyi, “Which Sectors Make Poor Countries so Unproductive?,” *Journal of the European Economic Association*, 2012, 10, 323–341.

——— and **Todd Schoellman**, “Why is Measured Labor Productivity so Low in Agriculture?,” Manuscript, Arizona State University 2012.

———, **Christopher Herrington, and Ákos Valentinyi**, “Sectoral Technology and Structural Transformation,” Manuscript, Arizona State University 2012.

———, **James Schmitz Jr, and Arilton Teixeira**, “The Role of Transportation in U.S. Economic Development: 1840-1860,” *International Economic Review*, 2012, 53, 693–715.

———, **Richard Rogerson, and Ákos Valentinyi**, “Two Perspectives on Preferences and Structural Transformation,” *forthcoming: American Economic Review*, 2013.

- Herz, Benedikt and Thijs van Rens**, “Structural Unemployment,” 2011. Manuscript.
- Hnatkovska, Viktoria and Amartya Lahiri**, “Structural Transformation and the Rural–Urban Divide,” Manuscript, University of British Columbia 2012.
- Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan**, “Earnings Losses of Displaced Workers,” *American Economic Review*, 1993, 83, 685–709.
- Johnston, Bruce F. and John W. Mellor**, “The Role of Agriculture in Economic Development,” *American Economic Review*, 1961, 51, 566–593.
- Johnston, Bruce F and Peter Kilby**, *Agriculture and Structural Transformation: Economic Strategies in Late–Developing Countries*, Oxford University Press, 1975.
- Jorgenson, Dale W. and Marcel P. Timmer**, “Structural Change in Advanced Nations: A New Set of Stylised Facts,” *Scandinavian Journal of Economics*, 2011, 113, 1–29.
- Kim, Dae-Il and Robert H. Topel**, “Labor Markets and Economic Growth: Lessons from Korea’s Industrialization, 1970–1990,” in Richard B. Freeman and Lawrence F. Katz, eds., *Differences and Changes in Wage Structure*, Chicago: University of Chicago Press for NBER, 1995, pp. 227–264.
- King, Robert G. and Sergio Rebelo**, “Transitional Dynamics and Economic Growth in the Neoclassical Model,” *American Economic Review*, 1993, 83, 908–931.
- King, Wilford Isbell**, “Industrial Origin of Total Realized Income,” in Wilford Isbell King, ed., *The National Income and Its Purchasing Power*, National Bureau of Economic Research, 1930.
- Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie**, “Beyond Balanced Growth,” *Review of Economic Studies*, 2001, 68, 869–882.
- Krantz, Olle and Lennart Schn**, *Swedish Historical National Accounts 1800-2000* Lund Studies in Economic History, Lund: Lund University, 2007. [http://www.ekh.lu.se/database/lumadd/National Accounts/default.htm](http://www.ekh.lu.se/database/lumadd/National%20Accounts/default.htm).

Kuznets, Simon, *Modern Economic Growth*, New Haven: Yale University Press, 1966.

———, “Modern Economic Growth: Findings and Reflections,” *American Economic Review*, 1973, 63, 247–258.

———, **Lillian Epstein, and Elizabeth Jenks**, “Distribution by Industrial Source,” in Simon Kuznets, ed., *National Income and Its Composition, 1919-1938*, Vol. 1, National Bureau of Economic Research, 1941.

Lagakos, David and Mike Waugh, “Selection, Agriculture, and Cross-Country Productivity Differences,” *forthcoming: American Economic Review*, 2013.

Laitner, John, “Structural Change and Economic Growth,” *Review of Economic Studies*, 2000, 67, 545–561.

Lawver, Daniel, “Measuring Quality Increases in the Medical Sector,” Manuscript, Arizona State University 2011.

Lee, Donghoon and Kenneth I. Wolpin, “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 2006, 74, 1–46.

Lilien, David, “Sectoral Shifts and Cyclical Unemployment,” *Journal of Political Economy*, 1982, 90, 777–793.

Lucas, Robert and Edward Prescott, “Equilibrium Search and Unemployment,” *Journal of Economic Theory*, 1974, 7, 188–209.

Maddison, Angus, *Statistics on World Population, GDP and Per Capita GDP, 1-2008 AD*, Groningen: University of Groningen, 2010.

Matsuyama, Kiminori, “A Simple Model of Sectoral Adjustment,” *Review of Economic Studies*, 1992, 59, 375–388.

———, “Structural Change,” in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2 ed., Palgrave Macmillan, 2008.

- , “Structural Change in an Interdependent World: A Global View of Manufacturing Decline,” *Journal of the European Economic Association*, 2009, 7 (2-3), 478–486.
- Messina, Julian**, “The Role of Product Market Regulations in the Process of Structural Change,” *European Economic Review*, 2006, 50, 1863–1890.
- Mitchell, Brian**, *International Historical Statistics*, London: Palgrave Macmillan, 2007.
- Moro, Alessio**, “The Structural Transformation between Manufacturing and Services and the Decline in the U.S. GDP Volatility,” *Review of Economic Dynamics*, 2012, 15, 402–415.
- Muellbauer, John**, “Aggregation, Income Distribution and Consumer Demand,” *Review of Economic Studies*, 1975, 62, 526–543.
- , “Community Preferences and the Representative Consumer,” *Econometrica*, 1976, 44, 979–99.
- Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny**, “Income Distribution, Market Size, and Industrialization,” *Quarterly Journal of Economics*, 1989, 104, 537–564.
- Ngai, L. Rachel and Christopher A. Pissarides**, “Structural Change in a Multisector Model of Growth,” *American Economic Review*, 2007, 97, 429–443.
- **and** ———, “Trends in Hours and Economic Growth,” *Review of Economic Dynamics*, 2008, 11, 239–256.
- Nickell, Stephen, Stephen Redding, and Joanna Swaffield**, “Educational Attainment, Labour Market Institutions, and the Structure of Production,” Working Paper, Centre for Economic Performance, London School of Economics, London, UK 2002.
- Olivetti, Claudia**, “The Female Labor Force and Long–run Development: The American Experience in Comparative Perspective,” Manuscript, Boston University, Boston 2012.
- O’Mahony, Mary and Marcel P. Timmer**, “Output, Input, Productivity Measures at the Industry Level: the EU KLEMS Database,” *Economic Journal*, 2009, 119OMahony-Timmer-2009, F374–F403.

- Prescott, Edward C.**, “Why Do Americans Work So Much More Than Europeans?,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 2004, 2-13, 28.
- Ramey, Valery A. and Neville Francis**, “A Century of Work and Leisure,” *American Economic Journal: Macroeconomics*, 2009, 1, 189–224.
- Ray, Debraj**, “Uneven Growth: A Framework for Research in Development Economics,” *Journal of Economic Perspectives*, 2010, 24, 45–60.
- Rendall, Michelle**, “The Rise of the Service Sector and Female Market Work: Europe vs. US,” Technical Report, University of Zurich 2010. Manuscript.
- Restuccia, Diego, Dennis Tao Yang, and Xiaodong Zhu**, “Agriculture and Aggregate Productivity: A Quantitative Cross–Country Analysis,” *Journal of Monetary Economics*, 2008, 55, 234–50.
- Rogerson, Richard**, “Sectoral Shifts and Cyclical Fluctuations,” *Revista de Analisis Economico*, 1991, 6, 37–46.
- , “Understanding Differences in Hours Worked,” *Review of Economic Dynamics*, 2006, 9, 365–409.
- , “Structural Transformation and the Deterioration of European Labor Market Outcomes,” *Journal of Political Economy*, 2008, 116, 235–259.
- Rubina, Verma**, “Can Total Factor Productivity Explain Value Added Growth in Services,” *Journal of Development Economics*, 2012, 99, 163–177.
- Sahin, Aysegul, Joseph Song, Giorgio Topa, and Gianluca Violante**, “Measuring Mismatch in the US Labor Market,” 2011. Manuscript.
- Schultz, Theodore W.**, *The Economic Organization of Agriculture*, New York: McGraw–Hill, 1953.

- Smits, Jan-Pieter, Edwin Horlings, and Jan Luiten van Zanden**, *Dutch GNP and its Components 1800-1913* Lund Studies in Economic History, Groningen: University of Groningen, 2007. <http://nationalaccounts.niwi.knaw.nl>.
- Sposi, Michael**, “Evolving Comparative Advantage, Structural Change, and the Composition of Trade,” Manuscript, University of Iowa 2011.
- Syrquin, Moshe**, “Patterns of Structural Change,” in Hollis Chenery and T.N. Srinivasan, eds., *Handbook of Development Economics*, Vol. 1, Amsterdam and New York: North Holland, 1988, chapter 7, pp. 203–273.
- Teignier, Marc**, “The Role of Trade in Structural Transformation,” Manuscript, Universidad de Alicante 2012.
- Tiffin, Richard and Xavier Irz**, “Is Agriculture the Engine of Growth?,” *Agricultural Economics*, 2006, 35, 79–89.
- Timmer, Marcel P., Robert Inklaar, Mary O’Mahony, and Bart van Ark**, *Economic Growth in Europe: A Comparative Industry Perspective*, Cambridge University Press, 2010.
- Timmer, Peter C.**, “The Agricultural Transformation,” in “Handbook of Development Economics,” Vol. 1, Amsterdam and New York: North Holland, 1988, chapter 8, pp. 275–331.
- Üngör**
- Üngör, Murat**, “Productivity Growth and Labor Reallocation: Latin America versus East Asia,” Manuscript, Central Bank of Turkey, Istanbul 2011.
- Uzawa, Hirofumi**, “On a Two–sector Model of Economic Growth II,” *Review of Economic Studies*, 1963, 30, 105–118.
- Valentinyi, Ákos and Berthold Herrendorf**, “Measuring Factor Income Shares at the Sectoral Level,” *Review of Economic Dynamics*, 2008, 11, 820–835.
- Yi, Kei-Mu and Jing Zhang**, “Structural Transformation in an Open Economy,” Manuscript, University of Michigan 2010.