

Structural Transformation and the Deterioration of European Labor Market Outcomes

Richard Rogerson*

Arizona State University

October 2, 2003

I have benefitted from the comments of numerous seminar participants. I am particularly grateful to Patrick Kehoe and Ellen McGrattan for comments, to Aspen Gorry for research assistance, and to the NSF for financial support.

1. Introduction

A large literature attempts to determine which factor (or factors) can account for the relative deterioration of labor market outcomes in Europe relative to the US. Even a casual reading of this literature makes two points clear. First, there is a consensus as to what needs to be explained, and second, there is no consensus as to what factor may be most important in providing an explanation. Regarding the facts to be explained, the consensus in this literature is that beginning in the early to mid 1970's most economies of continental Europe experienced a sharp increase in unemployment rates relative to the US that continued throughout the 1980's and lead overall to a sustained increase in relative unemployment of roughly six percent. Regarding explanations for this pattern there are several. Blanchard and Wolfers (2000) argue somewhat generally that a model with common shocks but country specific and time invariant institutions is promising in accounting for the facts. Examples of this general approach include Bertola and Ichino (1996), Ljungqvist and Sargent (1998), Mortensen and Pissarides (1999), den Haan, Hae-fke and Ramey (2002), and den Haan (2003). Others have argued that changing institutions can account for a substantial amount of the deterioration. Examples include Prescott (2002, 2003) and Daveri and Tabellini (1997).

This paper makes several contributions to this literature. First, it challenges the consensus view of the facts to be explained. The consensus view is based on an analysis of relative changes in unemployment rates. This paper argues that labor input is a more informative measure of labor market outcomes and therefore examines the behavior of employment to population rates. Based on this analysis, the deterioration of European labor market outcomes relative to the US begins much earlier, possibly as early as the mid 1950's, and continues at a fairly steady rate until 2000. The overall deterioration measured in terms of employment to population rates is almost 20%. This radically different view of what needs to be explained suggests that existing explanations which stress either institutional changes in the 1970's or differential responses to common shocks in the 1970's are likely to be insufficient.

The second contribution of the paper is to argue that the key to understanding the relative deterioration of European labor market outcomes relative to the US lies in understanding why European economies have failed to develop a market service sector similar to that in the US. In particular, I argue that one must view the development of the US and European labor markets from the perspective of the structural transformation of economic activity that accompanies the process of development. Kuznets argued that this structural transformation was one of

the six main features of the process of development. Basically this reflects the observation that as economies become richer activity moves first from agricultural to manufacturing and then later to services. In the mid 1950's Europe lags the US in terms of development, but closes much of the gap during the subsequent 40 years. Accordingly, it is not surprising that in mid 1950's Europe has a much larger employment rate for agriculture and industry, and much lower employment rate for services, relative to the US. By 2000, however, Europe has largely converged to the US levels for employment rates in agriculture and industry, but has not converged at all in services.

Third, I develop a simple model of structural transformation and within the context of this model provide an account of Europe's labor market performance relative to the US which is based entirely on different evolutions of returns to labor across different activities. The results imply that Europe's labor market deterioration is accounted for by factors which decrease the return to providing services in the market sector relative to the nonmarket sector. As a result, while the US has been moving production of services into the market, Europe continues to provide more services in the home sector. This result is similar to the marketization of production view stressed by Freeman and Shettkat (2002) and supported by time use studies in Germany and the US.

2. Two Views on the Deterioration

This section contrasts two views of the deterioration of European labor market outcomes relative to the US. The distinguishing feature of the two views is the variable which is used to characterize labor market outcomes. The first view, which I shall label the *traditional* view, uses the aggregate unemployment rate as its measure of aggregate labor market outcomes. The second view, which I will label the *modern* view, uses labor input as the measure of aggregate labor market outcomes. Because of issues in obtaining comparable measures of hours of work going far back in time, in what follows the measure of labor input that I use is the ratio of total employment to the size of the population between the ages of 15 and 64. In what follows I shall refer to this as the *employment rate*.

If one accepts the growth model as the dominant conceptual framework used for organizing aggregate observations, then it would follow that observations on labor input are the obvious starting point for characterizing how labor market allocations differ across economies. Conceptual debates aside, however, as a practical matter it is important to ask whether it matters. That is, does looking at the unemployment rate provide us with a different view of labor market outcomes than if we look at the employment rate.

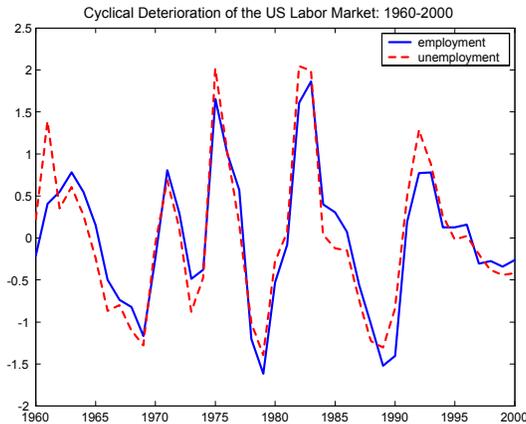


Figure 2.1:

The answer to this question may well depend on the context. First consider the business cycle context. For the case of the post-war US business cycle, it turns out that it does not seem to matter which measure one adopts. Over the business cycle, movements in the unemployment rate and the employment rate are not that far from mirror images of each other, so that both series are providing roughly the same information. Figure 2.1 displays two measures of the cyclical deterioration of the US labor market over the period 1960-2000.

The two curves are basically right on top of each other, implying that our view of what is happening in the labor market over the business cycle is not much affected by which of these two measures we use.

Next we consider the context of the low frequency differences in outcomes

in Europe and the US over the period 1956-2000. For my main set of results I contrast average outcomes in France, Germany and Italy with those in the US, and will refer to the average of France, Germany and Italy as corresponding to Europe. Similar results emerge with a larger set of countries from the continent chosen to represent Europe, but I focus on these three as my benchmark case since they are the three largest economies in continental Europe and hence are of particular interest. Figure 2 presents the findings based on unemployment rates and Figure 3 presents the findings based on employment rates. OECD data is available going back to 1956, so this is the first year in the figures. Each figure contains two lines. One is the raw data, and the other is the Hodrick Prescott trend. In both cases the 1956 trend value is normalized to 0, so that each figure is measuring deterioration relative to 1956. Both figures measure relative deterioration of European outcomes— in the case of unemployment this simply corresponds to the increase in the European unemployment rate relative to the US unemployment rate. In the case of employment rates, this corresponds to the decrease in the European employment rate relative to the US employment rate.

It is clear that Figures 2.2 and 2.3 present two very different pictures of the deterioration of European labor market outcomes relative to the US. Consistent with

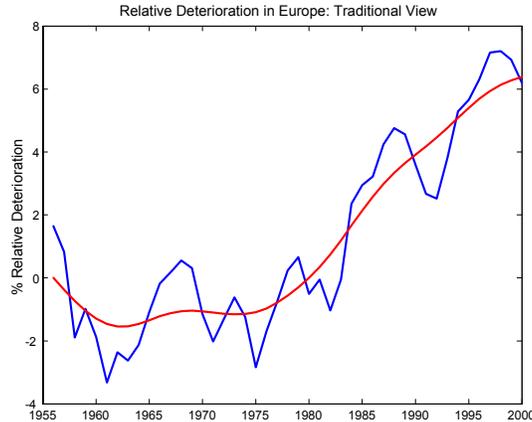


Figure 2.2:

the consensus view described in the introduction, Figure 2.2 shows a concentrated deterioration that begins in the mid to late 1970's, and an overall deterioration of roughly 6%. In contrast, Figure 2.3 shows a deterioration that begins in the mid 1950's, continues throughout most of the period and an overall deterioration of roughly 18%.

The key point here is that when one looks at labor input one is lead to a radically different description of what needs to be accounted for. The timing of changes, the concentration of the changes and the magnitude of the changes are all very different when deterioration is measured in terms of employment rates rather than in terms of unemployment rates. To emphasize this point, Figure 2.4 displays the two trend lines on the same graph to emphasize the difference.

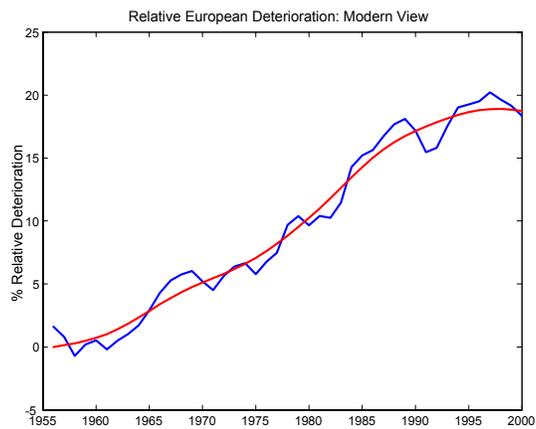


Figure 2.3:

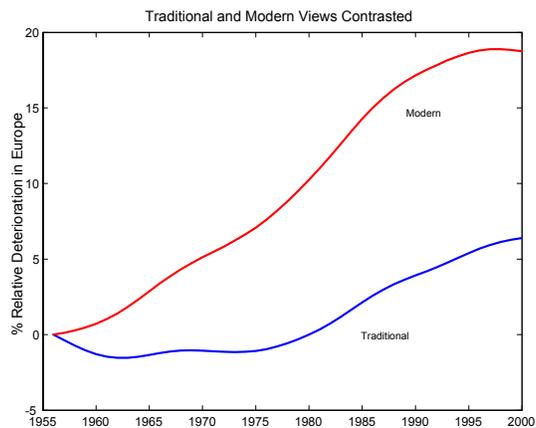


Figure 2.4:

The difference in these views is potentially very significant. Much of the existing literature is based on the picture corresponding to the old view and has led many researchers to look for shocks that occurred during the mid to late 1970's as the potential driving force behind the changes. However, based on the picture corresponding to the modern view, the driving forces must be present going back to the mid 1950's.

3. Structural Transformation and the Role of Services

The previous section documented a steady decline of the employment rate in Europe relative to the US. If less time is being devoted to market production in Europe relative to the US, there are two obvious follow-up questions of interest. The first is which (market) goods are not being produced, and the second is which types of individuals are not working. If the decreases are concentrated among a given set of goods or a particular group of individuals, this information may help to shed light on potential sources of the decrease. In this section we address the first of these two questions, and ask whether the relative decrease in employment is concentrated in any particular activities.

I consider three broad sectors—agriculture, industry and services, and for each sector I compute a sectoral employment rate which is total sectoral employment

divided by total population aged 15-64. The table below shows the deterioration in the sectoral employment rates for Europe relative to the US over the period 1956-2000.¹

Relative Deterioration in European Sectoral

Employment Rates, 1956-2000		
Agriculture	Industry	Services
10.7	2.2	6.1

This table suggests that the bulk of the deterioration occurs in agriculture, with the second most occurring in services and the least in industry. Another perspective on this same issue is to examine the relative employment rates between Europe and the US, both at the aggregate and the sectoral levels. Figures 3.1 through 3.3 illustrate the patterns. Because relative employment rates for agriculture are so much larger it is difficult to fit all series on the same graph without losing detail.

The pattern that appears in these figures is that European relative employment rates for industry and agriculture as well as the aggregate have been trending

¹I note that these three activities do not represent the universe of employment since military is not included, whereas the aggregate numbers presented earlier did include employment in the military.

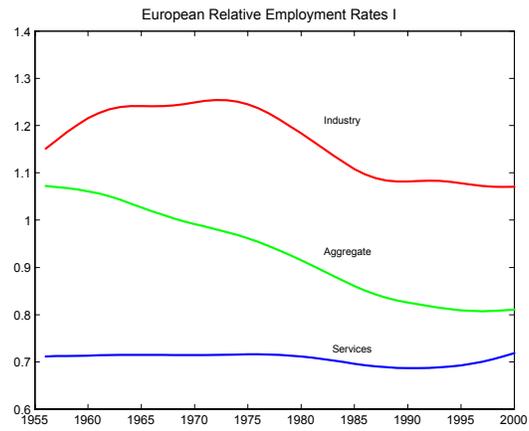


Figure 3.1:

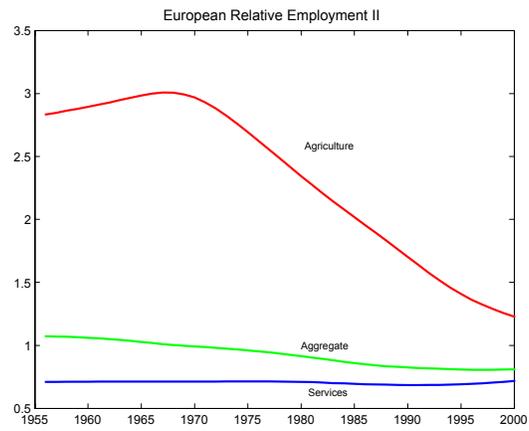


Figure 3.2:

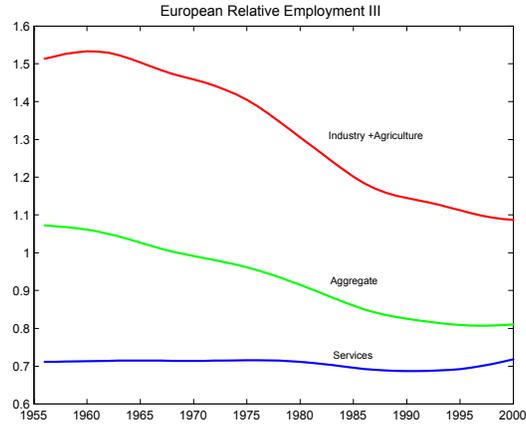


Figure 3.3:

down over time, as has the aggregate relative employment rate, whereas the relative employment rate in services has been relatively flat. Figure 3.3 shows that the combined relative employment rate for agriculture and industry has been decreasing since roughly 1960, while the relative employment rate for services has been roughly constant.

These graphs may lead one to conclude that the key to understanding Europe's relative decline in employment lies in understanding the relative decline in industry and agriculture. However, what I argue next is that such a conclusion is mistaken because it does not properly take into account the presence of the structural transformation in economic activity that accompanies the process of growth. Kuznets claimed that the process of structural transformation is one

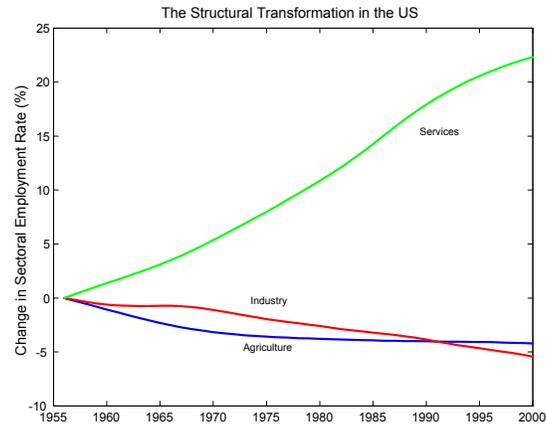


Figure 3.4:

of the six main features of the development process. The basic pattern that is followed by economies is that at low levels of development most resources are devoted to agriculture. As the economy develops resources are transferred to the industrial sector, and at yet higher levels of development resources shift to the service sector.

The next two figures show that both Europe and the US have been experiencing this transformation in the 1956-2000 period.

To understand the significance of the process of structural transformation for the issue at hand, consider the following situation. If one set of economies lags another, we should expect to see higher relative employment rates in industry and agriculture, and lower relative employment rates in services. Moreover, if

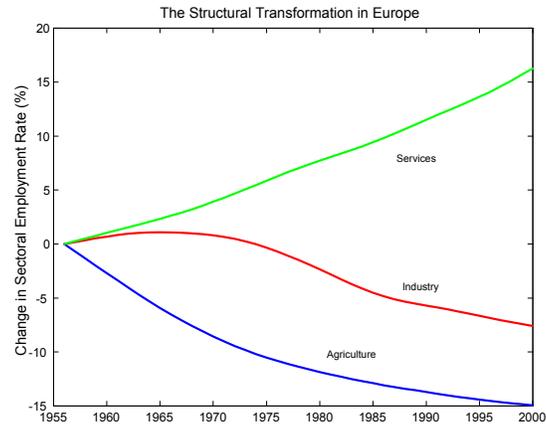


Figure 3.5:

this set of economies catches up, we would expect to see a reduction in relative employment rates in industry and agriculture and a rise in relative employment rates in services. It is well known that Europe experienced a significant degree of catch up to the US over the period 1956-2000, so this is precisely the pattern that should be expected.

With this in mind the next table shows the gap between European employment rates and US employment rates by sector in both 1956 and 2000.

Employment Rate Differential in Europe

Year	Agriculture	Industry	Services
1956	11.2	3.4	-9.4
2000	0.5	1.2	-15.5

This table shows that as of 1956, Europe has higher employment rates in both agriculture and industry, and a lower employment rate in services. Qualitatively this is consistent with the notion that as of 1956 Europe is lagging the US in the development process and hence has a larger amount of economic activity taking place in agriculture and industry. Given that over the period 1956-2000 Europe closes much of the gap between itself and the US in terms of output per hour, at a qualitative level this would lead us to expect that Europe's employment rates for agriculture and industry would decrease and approach those of the US, and in fact this is exactly what we observed. So, although the largest deterioration in employment rates occurs in agriculture, this is consistent qualitatively with the notion that Europe was catching up to the US over this time period. Similarly, much of the gap in industry is also closed over this period. However, from the perspective of structural transformation and catch-up, we would expect that Europe's employment rate in services would have increased relative to that in the US, but in fact it has decreased relative to that in the US.

To summarize, this analysis suggests that the key to understanding the deteri-

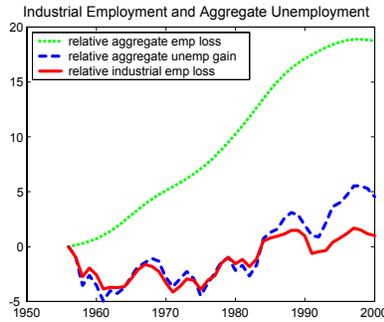


Figure 3.6:

oration of employment rates in Europe relative to the US is the failure of Europe to move workers into the service sector as it closed the gap between itself and the US in terms of productivity.

In the previous section we noted that the evolution of the deterioration is quite different if one examines unemployment rates rather than if one examines employment rates. We close this section by noting that the deterioration in unemployment rates follows very closely the deterioration in employment rates in industry. The next figure illustrates this fact.

What is of particular interest about this figure is that it indicates that from an empirical perspective, those researchers who have been trying to explain the deterioration of European labor market outcomes as measured by unemployment rates were effectively seeking to explain the deterioration of employment in indus-

try. Understanding the factors that account for the specific pattern found for the relative employment decline in industry is an interesting topic, but as the above figure makes clear, this relative employment decline is only a small piece of the overall relative employment decline.

4. Theoretical Analysis

This section develops a simple model of structural transformation for the period 1956-2000 and then uses it to rationalize the evolutions documented in the previous section. The model developed here purposefully abstracts from many details in order to focus attention on one particular set of economic forces: how changes in the returns to various activities influence the equilibrium time allocation across these activities. In focusing attention exclusively on this set of forces, the model is best viewed as the minimal structure required in order to articulate these forces. It is important to note many of these abstractions and to understand how they may influence the interpretation of the model, but I postpone this discussion until the model has been completely described.

4.1. Model

The model is closely related to the those studied by Echevarria (1997) and Kongsamut, Rebelo and Xie (2001). Similar to those models, this one will be designed so that technological change which is neutral across sectors can produce an ongoing reallocation of activity across sectors. Central to generating this reallocation will be a particular form of non-homotheticity in preferences.

There is a representative household with preferences given by:

$$\sum_{t=0}^{\infty} \beta^t (U(C_t, L_t) + G(A_t))$$

where C_t is a composite good representing consumption of non-agricultural goods and services in period t , L_t is leisure in period t , and A_t represents consumption of agricultural goods in period t . The utility function U will be assumed to be log-linear in C_t and L_t :

$$U(C_t, L_t) = a \log(C_t) + (1 - a) \log(L_t)$$

The composite consumption good C_t is defined by:

$$C_t = C(M_t, S_t) = [a_c(M_t - \bar{m})^{\varepsilon_c} + (1 - a_c)(S_t + \bar{s})^{\varepsilon_c}]^{1/\varepsilon_c}$$

where M_t is consumption of nonagricultural goods and S_t is consumption of services. In what follows we will refer to nonagricultural goods as the manufacturing good, though the empirical counterpart of this good is all output from the industrial sector, and not just the manufacturing component. The aggregator for M_t and S_t is a standard CES except for the presence of the terms \bar{m} and \bar{s} . These two (positive) constants play a key role in the model. In particular, they make preferences non-homothetic and thereby allow the model to generate a structural transformation of economic activity in the presence of proportional improvements in productivity in manufacturing and services.² It is intuitive that these preferences can generate this result. Abstracting from the consumption of agricultural goods for the time being, if \bar{m} and \bar{s} were both zero, a consumer with these preferences facing constant prices and changing income would simply scale consumption of both goods and services proportionately to changes in income. In particular, at low levels of income the individual would choose to consume low amounts of both goods and services. However, a positive value of \bar{m} effectively implies that a minimum level of the manufacturing good needs to be consumed in order to keep marginal utility of these goods bounded, whereas a positive value of \bar{s} im-

²This specification differs from the one used by Rebelo et al. This will be discussed in more detail in the following subsection.

plies that even as consumption of services goes to zero, marginal utility of services will remain bounded. Hence, from the perspective of the consumer, at low levels of income this will require a disproportionate amount of income is devoted to purchasing manufacturing goods rather than services. From the economy's perspective, this implies that a disproportionate number of resources will be devoted to producing manufacturing goods rather than services. As productivity increases, the impact of \bar{m} and \bar{s} diminishes—in the region of commodity space where the levels of M_t and S_t are large relative to \bar{m} and \bar{s} , the extent of non-homotheticity diminishes.

A key feature of the model is that we allow for the possibility that services can be produced both in the market and at home.³ In particular, we assume that aggregate consumption of services S_t is given by a CES aggregator of market and home produced services:

$$S_t = [a_s S_{mt}^{\varepsilon_s} + (1 - a_s) S_{nt}^{\varepsilon_s}]^{1/\varepsilon_s}$$

where S_{mt} is consumption of market produced services in period t and S_{nt} is consumption of home produced services in period t .

³We could also allow for the possibility that manufactured goods can also be produced at home. As an empirical matter, this does not seem to be a particularly important component of production in rich economies, and hence we abstract from it.

Utility derived from the consumption of agricultural products, $G(A_t)$, is assumed to take a very simple form. Specifically, it is assumed that individuals do not desire more than \bar{a} units of the agricultural good, but that they also require consumption of at least this amount. This is accomplished by assuming that G satisfies $G(A_t) = \min\{A_t, \bar{a}\}$ if $A_t \geq \bar{a}$ and that $G(A_t) = -\infty$ if $A_t < \bar{a}$. In equilibrium this will imply that the movement of labor out of agriculture is a very simple function of improvements in agricultural productivity.

Next we turn to a description of technology. For simplicity we abstract from capital and assume that all technologies are linear in labor. We assume that average productivities change over time, and in particular, we write:

$$M_t = \theta_{mt} h_{mt}$$

for the manufacturing sector, and

$$A_t = \theta_{at} h_{at}$$

for the agricultural sector. As noted previously there are two technologies in the service sector, a market technology and a home technology. These are specified

as:

$$S_{it} = \theta_{sit} h_{sit}, \quad i = m, n$$

Leisure for the representative household in period t is given by:

$$L_t = 1 - h_{at} - h_{mt} - h_{smt} - h_{snt}$$

Note that the model is effectively static, in that equilibrium outcomes in period t will depend only upon period t variables. Despite this, the interaction of productivity increases with non-homothetic preferences will potentially induce an interesting dynamic evolution in labor allocations in this economy, which are the focus of this study.

4.2. Remarks About the Model

Several remarks should be noted concerning the model and some of the simplifying assumptions. First, the production structure assumes that there are four technologies that all produce final goods and services using labor as the only input. In reality some business services, for example, are inputs into the production of manufactured goods, and some manufactured goods are inputs into the production of some services. Many agricultural goods serve as inputs into the manufacturing

sector. The final goods and services offered by supermarkets entail inputs from agriculture and manufacturing, but workers in supermarkets are recorded as employed in the service sector. While a more complex production structure would allow for a richer set of interactions, the simpler structure here allows us to better isolate the role of some basic forces. Though it is possible that some of these interactions are significant quantitatively, the production structure adopted here seems a useful starting point.

The model also assumes that there is a single good produced by each sector. In reality there is tremendous heterogeneity of goods even within each sector, and this is probably especially true for the service sector. The service sector includes such diverse items as health services, education, legal services, restaurant meals and cleaning services. Undoubtedly the factors that influence how many advanced medical procedures are provided in the market and how many home cleaning services are provided in the market are quite different. And a model that was designed to focus on the former would probably include some features different than a model designed to shed light on the latter. While one strategy to deal with this would be to allow for heterogeneity of goods within a given sector, we have chosen not to do this to minimize the dimensionality of the model. As will be seen in the next section however, the model will be calibrated so as to

emphasize that part of the service sector for which there are good nonmarket substitutes.

Having assumed a single agent representative household, the model also abstracts from heterogeneity in the distribution of human capital as a factor in influencing the allocation of time across activities. In particular, one might suspect that greater heterogeneity in the ratio of market productivity to home productivity across individuals would influence how an economy allocates time across market and nonmarket activities. This factor is stressed in a recent paper by Davis and Henreksson (2003). In such a framework, differences in market wage structures holding the distribution of skills constant would also presumably play a role in shaping how time is allocated between market and nonmarket activities. Once again, although it is possible that these interactions are quantitatively significant, the view here is that the simpler model remains a useful starting point.

An additional factor that the model abstracts from is trade, i.e., the model just described is for a closed economy. Over the last several decades there has possibly been a significant interaction between trade among countries and the process of structural transformation, as production of certain goods migrates to less-developed countries. These interactions are definitely of interest, but the underlying premise of the analysis done here is that trade *per se* is not the driving

force behind the process of structural transformation, and that we can therefore learn something about the underlying forces that shape the structural transformation without allowing for trade. Again, while it is of interest to consider how factors that influence trade have influenced the observed pattern of structural transformation, a model without trade seems a useful starting point.

There is also a measurement issue that should be noted at this point. The model takes sectoral average labor productivity as a primitive. However, from the model's perspective, this same value can be interpreted as the private return to working in a particular activity, which is not the same as measured labor productivity since taxation can create a wedge between these two. Additionally, there is an important issue of aggregation across different types of good within a sector. For example, consider the following situation. Assume that those services that are offered in both the market and nonmarket sectors tend to have lower productivity. If an economy in which the market productivity for producing these services is relatively higher tends to produce them in the market whereas another economy with relatively lower market productivity for producing these goods tends to produce them at home, then the economy with higher productivity at producing individual services may appear to have lower market productivity because of the aggregation bias induced by the selection of which services are provided in the

market. If a significant part of the cross-country differential in market provided services corresponds to variation in the extensive margin rather than the intensive margin, this may be a significant factor in terms of comparing aggregate sectoral productivities.

There are of course a variety of factors that can influence measured labor productivity, including product and labor market regulation, tax policies which influence the incentive to invest in capital or adopt new technologies, and labor unions just to mention a few. Ultimately, it is important to understand which particular factors are responsible for differences in average labor productivity. The model as currently set up is only intended to assess the consequences for time allocations of a given exogenous pattern of productivity differences. Again, this seems a useful first step.

Lastly, I note that the model described will be used to study the process of structural transformation in a group of advanced economies in the period 1956-2000. As currently specified, this model would not be able to address the structural transformation over a much longer period. The reason for this is that if a country has sufficiently low productivity in agriculture and manufacturing, they may not be able to simultaneously produce enough to meet both the \bar{a} and \bar{m} requirements, or possibly they could do so only by devoting all of their time endowment to

working. This would produce a sharp decline in time spent working associated with sustained productivity increases at low levels of productivity. It is possible that this can be remedied with a variation of the utility function in which utility derived from manufacturing goods is increasing in M at all levels but that there is no utility from consuming services until M reaches the level of \bar{m} . For the period studied, these issues do not arise. I note that the work of Kongsamut et al mentioned earlier has a significantly different assumption on preferences. They assume that agriculture has a term \bar{a} that is equivalent to the \bar{m} in manufacturing, and that there is no equivalent term that enters into preferences over consumption of the manufacturing good. A feature of their structural transformation is that employment in manufacturing stays constant while labor moves from agriculture into services. This prediction is clearly at odds with the data and hence some modification is necessary in order to match the evolution of sectoral employments for industrialized countries in the period 1956-2000.

4.3. Calibrating to the US Structural Transformation

In this subsection we describe a procedure aimed at providing a parameterization of the model just described so that it matches the key features of the US structural transformation over the period 1956 to 2000.

All productivities are normalized to one for the initial period, i.e., $\theta_{m1} = \theta_{sm1} = \theta_{sn1} = \theta_{a1} = 1$. This corresponds to a choice of units. We also impose the restriction that $\bar{m} = \bar{s}$. While this may seem arbitrary, at low levels of time spent in agriculture, based on computations, this restriction implies that a neutral increase in productivity leads to roughly no change in total time spent working, which is the restriction commonly imposed in macroeconomic analyses.

The two substitution elasticities in the utility function, ε_s and ε_c are set up front. The elasticity of substitution between market produced services and home produced services, which is determined by the parameter ε_s , will be a key parameter in the analysis that follows. In the benchmark specification this value will be set to .8, implying a fairly high elasticity of substitution between the two types of goods. The justification for this is the observation that many of the services produced outside of the market do have good substitutes at home, including for example, cooking, cleaning, laundry, childcare, elderly care, home repair, home improvements, and yardwork. The reverse is not true—there are many market services, for example, advanced medical procedures, that do not have good non-market produced substitutes. As noted earlier, a richer model would incorporate heterogeneity within services to reflect this fact, but the above value is chosen to reflect that at the margin there is substantial opportunity for substitution between

home and market produced services, and given the above choice of ε_s , the model will be interpreted as reflecting that particular margin.

The parameter ε_c determines the amount of substitution between manufactured goods and services. In the limit as ε_c goes to zero, the aggregator between these two commodities becomes Cobb-Douglas. This seems a natural benchmark. It implies that any changes in expenditure shares across commodities are due to the presence of the terms \bar{m} and \bar{s} and not due to changes in relative prices *per se*.

The remaining parameters are chosen to match certain aspects of the time allocation across activities in the US. In the previous sections we have looked at the evolution of sectoral employment rates. It was noted earlier that a preferred measure would be time devoted to work rather than simply the number of people employed, but that there were issues in obtaining consistent cross-country measures of hours worked over long time periods. This issue is also relevant in relating the model to the data. In particular, the model makes predictions about time devoted to various activities. If hours per worker did not vary systematically across sectors then this would not be a problem, since relative employment levels would be the same as relative hours. Unfortunately, this is not the case. Hours per worker in the service sector are systematically less than hours per worker in the

goods producing sector, and moreover this ratio has a significant negative trend. In particular, while there has been relatively little trend in hours per worker in the goods producing sector, hours per worker in services has declined by roughly 20%. To take this into account in the calibration exercise, I compute the ratio of hours per worker in the service sector relative to hours per worker in the goods producing sector and use this to adjust the employment level in services. This adjustment factor declines from about .95 in 1956 to about .75 in 2000. I assume that hours per worker in agriculture are the same as hours per worker in the goods producing sector. These adjusted numbers are used in conjunction with the sectoral employment data used earlier to compute relative hours worked across the sectors.

Consistent with the above computed ratios of hours worked across sectors, initial values are chosen so that total time devoted to market work in the initial period is .32, and total time devoted to non-market work is .08.⁴ This determines targets for initial hours of $h_{m1} = .117$, $h_{sm1} = .162$, $h_{sn} = .08$ and $h_{a1} = .0321$.

The value of \bar{a} is chosen so that $h_{a1}\theta_{a1} = \bar{a}$, and since $\theta_{a1} = 1$ this implies

⁴This value is somewhat low relative to estimates of total time spent in nonmarket work that come from time use studies. The justification for this is to allow for the possibility that some fraction of nonmarket work may not have good market substitutes, and the way in which the model is formulated it is this segment of nonmarket work that is most relevant. However, I have also redone the analysis with larger values of nonmarket work and the key results are similar.

$\bar{a} = h_{a1} =$. The remainder of the profile for θ_{at} is set so that we perfectly match the time series for time devoted to agriculture in the US. This implies an average growth rate of θ_{at} equal to 2.6%, but it turns out that the growth is faster in the early part of the period than in the later part.

It remains to specify values for the four preference parameters a , a_c , a_s , \bar{m} and \bar{s} (recall that we have imposed $\bar{m} = \bar{s}$), and the productivity profiles for manufacturing, market services and home services. For simplicity, productivity profiles are restricted to have constant growth rates. Output per hour in the US has grown by roughly 2% per year between 1956 and 2000. We assume that the manufacturing and market service sector productivities in the model grow at constant annual rates of 2%. We note that there is no problem in assuming that productivity in the market service sector grows more slowly than this value, and that the key conclusions reached below are not affected by this. Basically, whatever changes are generated by a different growth rate for market service productivity is undone by a compensating change in productivity growth for home services. At this point we leave the growth rate of productivity in nonmarket services as a free parameter to be determined later.

The strategy to determine the remaining four preference parameters is loosely as follows. Given a value for \bar{m} and \bar{s} , the values of a , a_c and a_s are pinned down

by requiring that we match the 1956 values for relative hours across activities. Intuitively, this works as follows. Given a value for ε_s and relative productivities in the home and market service sectors, the ratio of home to market hours in services is completely determined by a_s . Given values for a_s , \bar{m} , \bar{s} and all of the productivities, a_c determines the ratio of hours in manufacturing to hours in services. And the parameter a then scales total hours of work up or down.

Given this procedure, the value of \bar{m} and \bar{s} is pinned down by requiring that the reduction in time devoted to manufacturing in the model matches that found in the data when the model is solved for 45 periods. In principle this value is influenced by the growth rate of productivity for nonmarket services, but as a practical matter this effect is very small. The implied value of \bar{m} is .05.

The growth rate of productivity of nonmarket services is pinned down by requiring that the increase in time devoted to market services in the model matches the increase found in the data. This requires a growth rate in home service productivity of 1.75% per year. The resulting parameterization is intuitive: in the data there has been a trend increase in total hours of market work and it is driven by the increase in market services. The model rationalizes this by having market service productivity increase at a higher rate than home service productivity, thereby leading individuals to consume more market services and less home services.

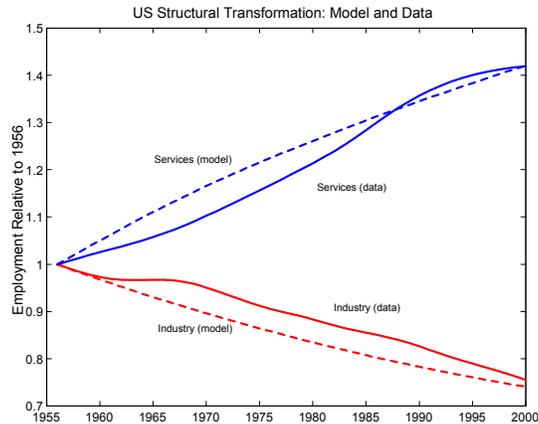


Figure 4.1:

The next figure shows the evolution of hours worked in each sector and those predicted by the calibrated model for manufacturing and services. (Recall that hours in the agricultural sector in the model and the data match perfectly by construction.)

The calibrated model matches the overall patterns found in the data. Restricting productivity processes to have constant growth rates does prevent the model from obtaining a closer fit to the actual patterns, the model does replicate the overall changes in labor inputs.

4.4. Application to Europe

In this section we take the preference structure from the calibrated model as given and ask what productivity processes would be required in order for the model to predict a structural transformation that matches that observed in Europe. There are two key components to this calculation. One concerns the initial level of productivity in Europe relative to the US, and the second component is the evolution of productivity in Europe relative to the US.

In carrying out this calculation I assume that the relative sectoral values of hours per worker are the same as those used in the US calibration. In particular, I assume that hours per worker in industry and agriculture are constant over time and that hours per worker in services decreases over time. Ideally this calculation would include the actual values for Europe, but I have not been able to find a long time series with the required values. Also, this choice implies that the analysis is implicitly abstracting from the issue of longer paid vacations in Europe relative to the US.

As noted previously, the formulation of the model implies a trivial mapping between time devoted to agricultural production and productivity, and hence it is straightforward to derive the agricultural productivity series needed to generate the European series for time devoted to agricultural productivity. Initial agricul-

tural productivity is set so as to match the level of agricultural employment in Europe relative to that in the US in 1956. And as was done in the case of the US, the profile of productivity change in agricultural is then chosen so as to perfectly match the evolution of employment in agriculture. It turns out that assuming a constant rate of growth of 4.39% per year accomplishes this.

Next we consider the productivity profile in manufacturing. For simplicity I again restrict attention to a profile with a constant growth rate. The goal is to choose this growth rate so that the relative employment level in Europe matches that found in the data. Figure 3.1 shows the actual series. A feature of this evolution is that the relative level in Europe increases until the early 1960's. With productivity increasing at a constant rate over time this particular feature is difficult to match, so I choose an initial productivity level to match the peak of this ratio and then choose the growth rate of productivity in manufacturing so that the relative level in 2000 matches that found in the data.⁵ (Strictly speaking these values cannot be determined without also assuming something about the productivity processes in services, but as a practical matter this is not particularly important.) The values that come from this procedure imply that

⁵The workweek in manufacturing in Europe did decrease in the earlier part of this time period, in which case it may be best to consider a larger ratio as the appropriate target for relative hours of work. I will try and incorporate some evidence on this in later versions of the paper.

in 1956 manufacturing productivity in Europe is 62% of the level in the US, and that over the period it grows at the rate of 2.6%, so that in the year 2000 the productivity gap is closed to roughly 20%.

Next we consider the productivity processes for home and market services in Europe relative to the US. In the data, relative employment in (market) services in Europe relative to the US is roughly equal to about .71 of the US value throughout the entire 1956-2000 period. Although there is a dip and a recovery in the 1980's and 1990's, the pattern I will focus on here is a constant relative rate of service employment, equal to .71. As commented in the calibration of the model to US data, there are many choices of the profiles for market and service productivities that are consistent with a given time series for time devoted to market services. That is also true here. However, although these different choices imply different levels of the two productivity series, their implications for growth rates in these productivities relative to the US are roughly the same.

For the benchmark calculation we adopt the following procedure to deal with the issue of multiplicity. The initial level of market service productivity is chosen as free parameter. Given a choice for this parameter, there is a unique initial value for the productivity of home services that will be consistent with the actual labor input in market services in Europe relative to the US. One reason for treating the

initial level of market service productivity as a free parameter is that although there is no direct measure of this productivity, some researchers (see, e.g., Bernard and Jones) have claimed that there has been greater convergence in productivities in services than in other sectors. Although measurement issues make such conclusions somewhat tentative, it turns out that varying this initial productivity level provides a way to vary the extent of relative convergence in service productivity relative to manufacturing productivity. In the benchmark setting I will assume that market services in Europe in 1956 are an additional 10% lower relative to the US than is European productivity in manufacturing. Other values considered for this additional gap in market service productivity were 0 and 20%.

With no additional restrictions, the model can effectively produce a continuum of series for home and market productivities that would produce a given time series for time devoted to market services. To a first approximation these series imply roughly the same ratio of the two productivities but different levels. To select one particular element from this set we do the following. We choose a terminal value for productivity of home services in Europe relative to the US and then assume a constant growth rate of this productivity consistent with the initial and terminal values. The motivation for this choice is that to the extent that home productivity is less affected by regulations, taxes, and other institutional factors,

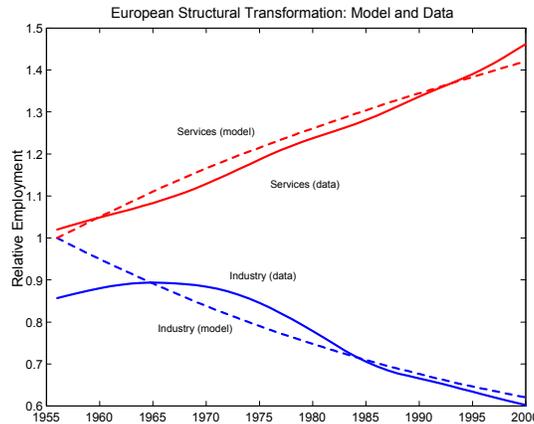


Figure 4.2:

we might expect greater convergence in these levels over time even in the face of substantial differences in regulations, tax rates and institutions that impact on market productivity. In the benchmark calculation I assume that in 2000, European productivity in home produced services is equal to 90% of its level in the US, though we also consider levels of this productivity relative to the US of 100% and 80%.

The next figure shows the structural transformation for Europe in the calibrated model and in the data.

Assuming a constant growth rate for manufacturing productivity prevents the model from matching the exact shape of the relative employment series for industry. Note that in the calibration procedure the series that were targeted were the

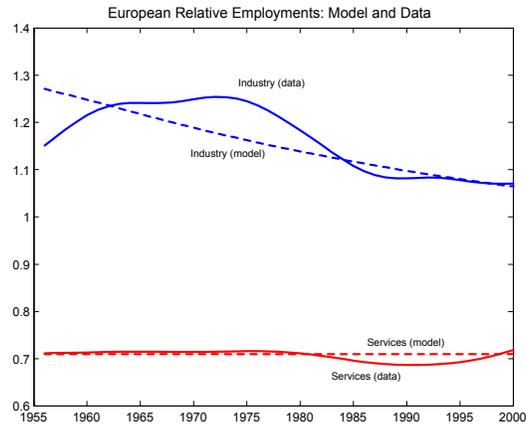


Figure 4.3:

sectoral employment levels in Europe relative to the US. Hence, any discrepancy between the calibrated model and the data for the US will necessarily translate into a discrepancy in this picture as well.

Given this comment another picture of interest is the one that shows relative sectoral hours across Europe and the US. This is shown in the next figure.

The approach taken in this section was to ask what productivity processes could generate patterns similar to those observed in the data. It follows that a key output of the procedure are the productivity series for Europe relative to the US. These are shown in the next figure.

The results in the graph show that Europe experiences catch-up relative to the US in all activities. However, the catch-up is greatest in terms of home

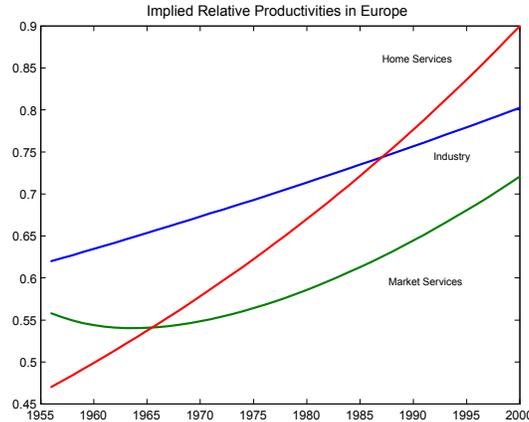


Figure 4.4:

productivity. Overall, market service and industry productivity close roughly half of the initial gap relative to the US. Although catch-up in market services is roughly the same as in industry over the entire period, the calibration implies that the catch-up in services has been faster since 1970. Perhaps most significantly, assuming a current gap of only 10% in home service productivity, one can explain the lack of catch-up in market service employment with only a 10% gap in market service productivity relative to industry productivity.

As noted earlier, several variations of this benchmark case have also been computed. The basic message is very similar to the one displayed above. Slower growth in home productivity simply translates into slower growth in market service productivity, but with little impact on relative levels.

4.5. Discussion

The calculation in the previous subsection produced series for productivities in Europe that could reconcile the employment evolutions in Europe and the US. As noted earlier, the approach followed here is to simply treat these productivities as primitives and not to model how various underlying factors may induce these differences. Nonetheless, given the results it is of interest to try and assess whether the required differences in productivity might be in line with actual differences. If so, the productivity story formalized here should be taken more seriously as a potential explanation.

Let's begin with the productivity differences in industry. Specifically, the calibration required that as of 2000 Europe lags the US by 20%. Is such a difference plausible? I think the answer is a clear yes. Recall our earlier discussion which pointed out that taxes are one factor that can drive these differences. Differences in effective tax rates on labor income between the US and these European countries are in the neighborhood of 20%. It follows that one may be able to rationalize differences of this magnitude without assuming any true differences in productivity.

Next consider the market service sector. According to the results in the previous subsection, as of 2000 we need to rationalize a level of productivity relative

to the US of roughly 28%. The argument just given above regarding taxes can account for roughly 20% of this. However, there is an additional tax factor that may help to account for the remaining 8%. In the US many services are tax exempt, whereas this tends not to be the case in Europe. Given that average consumption taxes are about 5%, this may account for the bulk of the remainder. However, there are a host of regulations and policies in Europe that one could expect to plausibly reduce productivity by a few percent relative to the US. An example that has received considerable attention is the restriction on hours of operation for shops in Germany. Employment protection policies may also account for some productivity differences. Moreover, although these policies apply to services and industry, it is possible that they have differential impact across sectors since underlying turnover rates differ across sectors.

A key element in the above calculations was the assumption of greater catchup in home productivity. I do not have any direct evidence on this point, but I note that such a pattern could arise if adoption rates for many household appliances was slower in Europe than in the US.

The analysis carried out has focused only on one aspect of taxation—the fact that taxes drive a wedge between private and actual returns to some activity. Another set of effects arises from potential income effects associated with redis-

tribution of tax revenues, or additional substitution effects if the tax revenues are used to subsidize various activities that do entail market work. These additional effects can also be large, as shown in the recent work of Prescott (2002, 2003).

I conclude from this discussion that the productivity differences required by the previous calculations do not seem particularly large, and that a pure tax story may even be able to account for the bulk of the required differences. But given the multitude of policies and regulations that may also plausibly affect productivity, the differences do not seem large.

5. Conclusion

This paper makes three key points. First, it argues that much of the literature on the European labor market problem has misdiagnosed the problem by focusing on relative unemployment rather than relative employment levels. Specifically, the European labor market problem seems to date back to the mid 1950's. Second, the key to the understanding the source of the European labor market problem is to understand why Europe has not developed a market service sector more similar to that of the US as it has closed the gap with the US in terms of output per hour. Third, it shows that a story in which productivity differences and/or taxes are central can potentially go a long way to accounting for the relative deterioration

of European labor market outcomes. To be sure, the model analyzed here is very simple and it will be important to see the extent to which the quantitative conclusions are affected by adding various features.

References

- [1] Bertola, G., and A. Ichino, “Wage Inequality and Unemployment: United States vs. Europe” NBER Macro Annual, 1995, 13-53.
- [2] Blanchard, O., and J. Wolfers, “The Role of Shocks and Institutions in the Rise of European Unemployment: The Aggregate Evidence”, *Economic Journal* 110 (2000), 1-33.
- [3] Daveri, F., and G. Tabellini, “Unemployment, Growth and Taxation in Industrial Countries”, *Economic Policy* 15 (2000), 47-104.
- [4] Davis, S., and M. Henrekson, “Taxes and Work Activity”, mimeo, 2003.
- [5] Den Haan, W., “Temporary Shocks and Unavoidable Transitions to a High-Unemployment Regime”, mimeo, 2003.
- [6] Den Haan, W., C. Haefke, and G. Ramey, “Shocks and Institutions in a Job-Matching Model”, NBER Working Paper, 2002.
- [7] Echevarria, M., “Changes in Sectoral Composition Associated With Growth”, *International Economic Review* 38 (1997), 431-452.

- [8] Freeman, R., and R. Schettkat, “Marketization of Production and the US-Europe Employment Gap”, NBER Working Paper, 2002.
- [9] Kongsamut, P., S. Rebelo, and D. Xie, “Beyond Balanced Growth”, Review of Economic Studies 2001.
- [10] Ljungqvist L., and T. Sargent, “The European Unemployment Problem,” Journal of Political Economy
- [11] Mortensen, D. and C. Pissarides, “Unemployment Responses to Skill-Biased Technology Shocks,” Economic Journal 109 (1999) 242-265.
- [12] Prescott, E., “Prosperity and Recessions”, Ely Lecture, American Economic Review, 2002.
- [13] Prescott, E., “Why Do Americans Work So Much More than Europeans?”, 2003 Nemmers Lecture.