

IZA DP No. 4177

**A Distribution Dynamics Approach to Regional
GDP Convergence in Unified Germany**

Falko Juessen

May 2009

A Distribution Dynamics Approach to Regional GDP Convergence in Unified Germany

Falko Juessen

*Technische Universität Dortmund
and IZA*

Discussion Paper No. 4177
May 2009

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

A Distribution Dynamics Approach to Regional GDP Convergence in Unified Germany*

This paper uses nonparametric techniques to study GDP convergence across German labor market regions and counties during the period 1992-2004. The main result is that regional convergence in unified Germany has been substantial. In the first years after German unification the distribution of GDP has been characterized by a pronounced bimodality. The dispersion of the GDP distribution has become substantially smaller over time. Although some bimodality remains in most recent years, this bimodality is weak in comparison to previous years. Nevertheless, disparities among regions located in the Eastern and Western part of the country are still apparent.

JEL Classification: O47, R11, C14

Keywords: regional convergence, distribution dynamics, nonparametric econometrics, stochastic kernel

Corresponding author:

Falko Juessen
Technische Universität Dortmund
Department of Economics
D-44221 Dortmund
Germany
E-mail: falko.juessen@tu-dortmund.de

* I thank three anonymous referees for suggesting various improvements to the paper. Moreover, I am grateful to Christian Bayer, Christoph Alsleben, and Michael Roos for their helpful comments. Previous versions of this paper have been presented at the Xth Spring Meeting of Young Economists 2005 in Geneva, the 45th Congress of the European Regional Science Association 2005 in Amsterdam, the Annual Meeting of the Verein für Socialpolitik 2005 in Bonn, the 7th INFER Annual Conference 2005 in London, and the ASSET Annual Meeting 2005 in Rethymnon. I thank participants of these conferences for their helpful comments and suggestions. This paper is part of the project "The International Allocation of Risk" funded by the Deutsche Forschungsgemeinschaft (DFG) in the framework of SFB 475.

1 Introduction

This paper presents an empirical study of GDP convergence across German labor market regions and counties during 1992 to 2004 using nonparametric techniques. The issue of convergence is important for Germany because mitigating regional disparities is regarded as a fundamental objective of German and European regional policy, especially in light of East-West differentials in unified Germany. For instance, the main objective of EU regional policy is to promote the development of regions whose per capita GDP is below 75% of the EU average (see Overman and Puga, 2002).

Convergence means that poor economies catch up with rich ones. Typically, studies on convergence use the neoclassical growth model as their theoretical framework. This model suggests that regional per capita output within a country converge to the same long-run steady-state.¹ By contrast, new theories of industrial location, trade, and integration (see Fujita, Krugman, and Venables, 2000) and most models of the new growth theory cast doubts on the neoclassical optimistic prediction of convergence. In recent years, a number of alternative strategies has been suggested to measure convergence empirically, e.g. traditional cross-sectional regressions of β - and σ -convergence, panel data approaches, and time-series tests (see Magrini, 2004 and Durlauf, Johnson, and Temple, 2005 for survey articles).

Several studies apply these techniques to analyze regional convergence in West Germany. In general, the studies do find evidence for both absolute and conditional convergence. Empirical evidence regarding unified Germany is still scarce, however.

Most studies for unified Germany are limited to an analysis of convergence between the Eastern and Western part of the country.² Although some authors are more pessimistic than others about convergence, the general result is that "East German labor productivity has converged on that in West Germany more slowly than was initially thought but faster than would have been expected on the basis of studies of convergence such as Barro and Sala-i-Martin (1991)" (Barrell and te Velde, 2000, p. 272).

While most studies focus on East-West disparities, other papers, as ours, address regional convergence in unified Germany at a disaggregated geographic level. Using a spatial econometric approach to β -convergence, Kosfeld, Eckey, and Dreger (2002, 2006) find clear evidence for both GDP per capita and labor productivity convergence during the period

¹Regions are by no means small closed economies but instead are highly integrated in terms of mobility of goods, capital, and labor. Barro, Mankiw, and Sala-i-Martin (1995) have extended the neoclassical growth model for partial factor mobility and show that the basic prediction of convergence is not altered if regions are integrated. Razin and Yuen (1997) provide an elaborate analysis of the role of labor mobility in the convergence process.

²See Hallett and Ma (1993), Burda and Funke (1995), Funke and Strulik (2000), and Barrell and te Velde (2000).

1992 to 2000. Kosfeld and Lauridsen (2004) adopt a cross-sectional spatial econometric adjustment model which is based on the concept of spatial error-correction. They find only weak evidence for conditional convergence in the year 2000. The study related closest to ours is the one by Colavecchio, Curran, and Funke (2005), which concludes that the relative income distribution in unified Germany appears to be stratifying into a trimodal/bimodal distribution.

The analysis by Colavecchio, Curran, and Funke (2005) and ours use a similar econometric framework to analyze the convergence hypothesis. The general idea is to examine how the entire cross-sectional distribution of GDP evolves over time.³ This so-called "distribution dynamics" approach to economic convergence has been introduced by Danny Quah in a number of seminal papers, see Quah (1993, 1996a, b, c, 1997, 2001). The main motivation for Quah's approach is that, by focusing on the average behavior of a representative region, traditional approaches to analyze convergence suppress the cross-section dynamics one wishes to investigate (see Durlauf and Quah, 1999). The nonparametric approach introduced by Quah, by contrast, is not limited to an analysis of single moments of the underlying distribution. Consequently, this approach allows one to uncover empirical phenomena such as persistence and the formation of convergence clubs.⁴ Since there are well-known East-West disparities in Germany (see Barel and te Velde, 2000) such patterns can be expected to be relevant in unified Germany.

The main result of our analysis is that regional convergence has been substantial. In the first years after German unification the distribution of GDP has been characterized by a pronounced bimodality, indicating sizeable disparities between regions located in the Eastern and Western part of the country. Over time, the dispersion of the GDP distribution has become substantially smaller. Although some bimodality remains in most recent years, this bimodality is weak in comparison to previous years. Moreover, we find that most convergence in Germany has happened during the period 1992-1995. From 1997 onwards, the GDP distribution across labor market regions may be described as being in equilibrium (from 2000 onwards, respectively, if counties are considered).

The remainder of this paper is organized as follows. Section 2 introduces the data employed. The empirical analysis is divided into two parts. Section 3 examines cross-sectional GDP distributions at different points in time and discusses their implications for convergence. In Section 4, we estimate a Markov-chain model for the dynamic evolution of the GDP distribution over time. This dynamic model is then used to calculate the

³While Colavecchio, Curran, and Funke (2005) examine GDP per capita convergence across German counties, our analysis focuses on GDP per worker convergence across labor market regions and counties (see Section 2).

⁴See Anderson (2004) and Maasoumi, Racine, and Stengos (2007) for related approaches to describe polarization in and distances between distributions in terms of stochastic dominance conditions.

ergodic distribution of GDP. We also offer an alternative approximation of the equilibrium distribution which does not depend on the assumptions of Markovian theory. Section 5 concludes the paper. An appendix provides some background information on the aggregation of counties to labor market regions.

2 Data

For unified Germany, detailed data is available which facilitates a regionally disaggregated analysis at the level of 439 counties. This data is provided by the National Accounts of the Federal States compiled by the Statistical State Office Baden-Wuerttemberg.⁵ The GDP data is measured in current prices, since regional price indices at the county level are not available.

Typically, empirical growth studies use either total population and/or total employment to normalize aggregate measures of GDP. While Colavecchio, Curran, and Funke (2005) examine convergence of GDP per capita across German counties, our main focus lies on GDP per worker, mainly for two reasons.

Firstly, most theoretical growth models are based on production functions. Therefore, their implications relate more closely to GDP per worker than to GDP per capita (see Durlauf, Johnson, and Temple, 2005). A second problem with the ratio of GDP over population is that numerator and denominator are measured on a different basis. While GDP attributes to a region the amount of economic production generated within a region, total population measures the number of people living in a region. Since regions are integrated by commuter flows, this difference induces a distortion of the GDP per capita variable. By contrast, GDP per worker is not subject to this distortion, since both GDP and employment are measured at the place of work. Nevertheless, we will also report some results for GDP per capita.

Besides the appropriate normalization of GDP, another issue of our data is that the borders of German counties are determined by political and historical rather than economic reasons. As examples we can consider counties that are located in the densely populated Ruhr area. While counties such as Bottrop, Gelsenkirchen, Recklinghausen, and Herne are separate administrative units, they form a common labor market in fact. To assess regional convergence, it is preferable to examine coherent and independent economic areas rather than administrative units.

In the Appendix, we provide some background information on how counties are aggregated to functionally defined labor market regions. We employ the regional delineation used by the most important regional policy programme in Germany, the so-called "Joint

⁵All data can be downloaded from: <http://www.statistik-portal.de/Statistik-Portal/publ.asp>.




Figure 1: GDP per worker across German labor market regions, relative to the German-wide average GDP per worker.

Task: Improvement of the Regional Economic Structure” (GRW). Under this definition, the 439 German counties are aggregated to 271 local labor market regions. In order to facilitate a comparison between types of regions, we will also present results for the raw county data, i.e. for the largest available sample size of 439 regions.

To account for the growth of the German economy that occurred during the period that we study, we have formulated the data for each labor market region (and county) relative to the German-wide aggregate. This normalization also accounts for common changes in inflation.⁶ The empirical analysis employs annual data in the period from 1992 to 2004.⁷ Thus, our dataset consists of a balanced panel of 271 (439) regions observed over 13 years (3523 and 5707 observations, respectively).

As a starting point, we document regional disparities in relative GDP per worker across labor market regions, for 1992 and 2004. Figure 1 illustrates apparent East-West disparities in unified Germany. It is instructive to compare the ten most and least productive regions at the beginning and end of the sample period (which need not be the same in each case). The average relative GDP in the ten most productive regions is 1.30 in 1992 and 1.25 a decade later. The ten least productive regions have an average of only 0.45 in 1992 and

⁶The average GDP per worker for Germany as a whole is calculated as an arithmetic weighted average, not as a simple average. The normalization using average GDP per worker does not affect the shape of the GDP distributions.

⁷Since there is no data for 1993, the data for this year has been interpolated using spline interpolation techniques.

0.72 in 2004. If the same group of regions is compared in 1992 and 2004, the values in 2004 are 1.21 for the ten most productive and 0.78 for the ten least productive regions. These numbers suggest that regional disparities are shrinking over time and that considerable catching up has taken place. However, they also illustrate that regional disparities are still very pronounced in unified Germany.

3 Cross-sectional analysis

In a first step of the empirical analysis we examine kernel density estimates of our key variable relative GDP per worker for different years. We also consider alternative transformations of GDP, such as GDP per capita, or employment-weighted GDP per worker.

One particular aspect of the densities we are interested in is the number of distinct modes. By evaluating whether the densities are unimodal or not we can test the convergence hypothesis. For example, if one starts with a bi- or multimodal density at a given point in time (e.g. a group of very productive and a group of less productive regions), convergence would imply a tendency of the distribution to move towards unimodality (see Bianchi, 1997). If the density is unimodal, its dispersion should decline over time if there is convergence. Another aspect of the densities we are interested in is whether the densities observed at different points in time are statistically different from each other.

In order to address these aspects we start with a graphical inspection of the estimated kernel densities. Thereafter, we employ statistical tests to formally test for the number of modes and time-invariance of the densities.

3.1 Kernel density estimation

3.1.1 Pooled estimations for unified Germany


Given a sample realization, $X = \{x_1, x_2, \dots, x_n\}$, from an unknown population with density f , a nonparametric estimate of this density is given by

$$\hat{f}_h(x) = (nh^{-1}) \sum_{i=1}^n K\left(\frac{x - x_i}{h}\right),$$

where $K(\bullet)$ is the Gaussian kernel function and h is the bandwidth (see Silverman, 1986, and Pagan and Ullah, 1999 for introductions to this field of data analysis).

Figure 2 summarizes kernel densities of various transformations of GDP. Each row in the figure refers to a different transformation of GDP, while the two columns display the results for 271 local labor market regions (left) and 439 counties (right). For each variable and region type we display the cross-sectional densities for three selected years of our

Figure 2: Kernel densities of various transformations of relative GDP across 271 labor market regions (left) and 439 counties (right)



sample period: the first year 1992, 1997, and the final year, 2004.⁸

In our benchmark estimations displayed in Figure 2, we have chosen the bandwidth h optimally according to Silverman's (1986, formula 3.31) rule of thumb. Alternatively, we tried a cross-validation procedure, see Silverman (1986, Section 3.4.3) (results are not reported but available on request from the author). The estimations carried out using cross-validation are less sharp than the ones obtained using Silverman's (adaptive) rule of thumb, in a sense that additional (minor) modes tend to pop up. These minor modes, however, tend to vanish once adaptive kernel techniques are applied, which take the sparseness of the data into account. Therefore, we note that both bandwidth criteria suggest similar economic conclusions.

The left panel in the first row of Figure 2 shows the densities of relative GDP per worker across 271 labor market regions. In 1992, the distribution appears strongly bimodal, the left mode being at a relative GDP of 0.52. In 1997, there are effectively no regions anymore that have such low levels of relative GDP per worker. The left peak has moved closer towards the national average and is at a value of 72%. At the same time, the dominant mode decreased from a value of 1.00 to 0.95. In 2004, the left mode (at 0.80) has vanished almost completely. In any case, the distance between the peaks has decreased substantially, with the left mode becoming less prominent over time. These observations suggest that regions which experienced very low levels of GDP after German unification do not seem to be trapped in a growth path of very low GDP per worker.

At the county level, the clustering of regions has disappeared even more strongly over time, see the right panel in the first line of Figure 2. The left mode in the 2004 distribution is not visible anymore, while it has been very pronounced in 1992.

From the second row in Figure 2 it can be seen how influential is a different normalization of GDP. The plots show the kernel densities of relative GDP per capita instead of per worker. The right tail of the density is longer in comparison to the tail of the GDP per worker density. This pattern is particularly evident for counties (right panel) at values of relative GDP per capita larger than 1.4.

This observation reflects that the variable GDP per capita becomes distorted by commuting linkages if (small) administratively defined regions are considered. The very high values of GDP per capita for some counties (e.g. economic agglomerations such as Frankfurt, Duesseldorf, and Munich) result from the fact that their GDP is attributed to the number of people living there while the number of workers commuting into the economic

⁸The densities have been normalized so that the sum of the points at which the density is evaluated is one. This allows us to interpret the normalized densities as showing the probability of a realization in the grid interval.

center from the hinterland is not taken into account.⁹ As can be seen from the left panel in the second row of Figure 2, the tail of the GDP distribution is less affected if economically defined labor market regions are considered. Having illustrated the effects of a different normalization of GDP, the subsequent analysis will focus on our key variable relative GDP per worker.

The figures we have presented so far have shown the GDP distributions in terms of regions. Although this is the common way to view the data it can be misleading as Jones (1997) has pointed out. The reason is that employment is distributed unevenly among regions¹⁰ but the kernel density estimator treats GDP per worker of each region on the same basis.

One alternative is to use a weighting scheme to adjust for the different sizes of regions. We follow Jones (1997), Goerlich Gisbert (2003), and Henderson, Parmeter, and Russell (2008) in attaching a weight to the observations where the weights reflects the contribution of each observation in the sample. Since our key variable is relative GDP per worker, a natural choice for the weight is the employment share of each region. With weighting, the unit of observation is the worker instead of the region.

The weighted kernel density estimator we use is given by

$$\hat{f}_h^w(x) = (nh^{-1}) \sum_{i=1}^n w_i K\left(\frac{x - x_i}{h}\right), \quad w_i > 0 \quad \forall i, \quad i = 1, \dots, n. \quad (1)$$

The only difference to the standard kernel density estimator is that the simple sum is substituted by a weighted sum, where w_i denote weights (employment shares) attached to the observations $\left(\sum_{i=1}^n w_i = 1\right)$.¹¹


The last row in Figure 2 shows the densities that we obtain when each observation of GDP per worker is weighted by its employment share. In the distribution for labor market regions, a third mode appears, which is located at a level of relative GDP per worker of about 1.26 (compare to the first row in Figure 2). It seems that regions with above-average productivity generate high-level modes in the employment-weighted distribution of GDP per worker. Or to put it differently, some regions with very high levels of GDP per worker also account for a high fraction of total employment in Germany. Examples

⁹A natural way to deal with long tail densities is to use a broader kernel in regions of low density, i.e. to use adaptive bandwidth methods (see Silverman, 1986, and Pagan and Ullah, 1999). When we apply an adaptive kernel estimator, the right tail of the GDP per capita density is sharper, since now a higher local bandwidth is used for observations located in this area of the distribution.

¹⁰Average employment across labor market regions in 2004 is 143,424, with a standard deviation of 212,803.

¹¹Goerlich Gisbert (2003) and Pittau and Zelli (2005) suggest alternative ways of introducing weights in kernel density estimation based on the idea of adapting the estimate to the local density of the data.

Figure 3: Densities of relative GDP per worker across labor market regions (left) and counties (right), separately for West and East German regions



for such regions are Munich, Frankfurt, Essen, Duesseldorf, or Leverkusen.

3.1.2 Separate analysis for East and West German regions

A disadvantage of kernel density estimation is that the estimated density cannot be related to any sub-populations. We suppose that the initial bimodality in the density for all German regions (see the first line in Figure 2) reflects a clustering of regions located in the Eastern and Western part of Germany. One way to address this issue is to model relative GDP per worker by a finite mixture density. In mixture models, the assignment of regions to sub-populations is provided in terms of ex-post probabilities of membership of each region to each-subpopulation (see Pittau, 2005). In this sense, mixture models can be viewed as a compromise between parametric and non-parametric models.

In our analysis for Germany, we have a strong prior for the number of sub-populations and for the membership of regions. Therefore, we decided to stick to the fully nonpara-

metric framework and to simply split the sample into East and West German regions.¹² This experiment allows us to assess if there is also convergence within the Western part of Germany.

Figure 3 shows the densities of GDP per worker for the splitted sample, evaluated at the same grid points. There are 205 West German and 66 East German labor market regions, and 327 West German and 112 East German counties (Berlin is treated as a West German region). We have to be keep in mind that the kernel density estimations based on the split sample may be of poor precision. Especially the sample size for East German labor market regions is very small.

Consider first the densities of East German GDP per worker (second row). In 1992, the peak of the distribution is at 51% of the German average. As expected, this mode corresponds to the left peak in the distribution for all German regions as displayed in Figure 2 (first row, left panel). This finding provides evidence that the cluster of unproductive regions in Figure 2 is mainly formed by East German regions. Over time, the whole distribution has shifted to the right; the dominant peak in 2004 is at 78%. These numbers illustrate that, although East German labor market regions have increased their relative GDP per worker over time, they still have considerably lower levels of GDP per worker than the national average. This finding is well compatible with other studies, such as Barel and te Velde (2000). It is also noteworthy that, nowadays, a small number of East German regions performs remarkably well (examples are Dresden and Magdeburg). This fact is reflected by the long tail of the 2004 distribution.

The first row in Figure 3 reveals that West German regions, as expected, cluster together in the center of the distribution of all German regions. That the mode of the West German distribution has shifted to the left over time is the direct mirror effect to the fact that East German regions experienced above-average growth during the period that we study. If we compare the 1992 distribution to the one in 2004, there is some evidence for a reduction of disparities across West German regions. However, the shape of the West German distribution is rather similar between years. This impression is consistent with results documented by Kosfeld, Eckey, and Dreger (2006), who detect β - but no σ -convergence for West German regions.

¹²Nevertheless, a semiparametric analysis in the spirit of Pittau (2005) appears a promising extension of our analysis. Models of mixture densities are informative on the movement of regions within the estimated components. Moreover, the evolution of mixing proportions can be examined. Since we address similar questions in a completely nonparametric, and therefore potentially less efficient framework (see Section 4), it would be interesting to compare the results of both approaches.

3.2 Multimodality analysis

A formal testing procedure for investigating the number of modes in an unknown density has been developed by Silverman (1981, 1983, 1986). It has been introduced to the growth convergence literature by Bianchi (1997).

The null hypothesis of the Silverman test is that the density $f(x)$ has m modes versus the alternative hypothesis that $f(x)$ has more than m modes. Silverman proposed as a test statistic the critical bandwidth \hat{h}_{crit}^m , which is defined as the smallest possible value of h producing a density with not more than m modes:

$$\hat{h}_{crit}^m = \inf \left\{ h : \hat{f}_h(x) \text{ has precisely } m \text{ modes} \right\}.$$

When a Gaussian kernel function is used, the number of modes is always a nonincreasing function of the bandwidth h (see Silverman, 1981). Therefore, any values that are smaller than \hat{h}_{crit}^m will produce a $(m + 1)$ -modal density. The intuition for Silverman's testing procedure is that large values of \hat{h}_{crit}^m are taken as evidence against the null hypothesis that f has only m modes. The question whether the test statistic \hat{h}_{crit}^m is indeed large or not is assessed using bootstrap methods.

The bootstrap procedure involves generating a large number of R samples with replacement from the original data. Each bootstrap sample y^* is adjusted to have the same variance as the original data x .¹³ For each smoothed bootstrap sample y^* , one estimates the Gaussian kernel density $\hat{f}_{\hat{h}_{crit}^m(x)}(y^*)$ using the fixed critical bandwidth $\hat{h}_{crit}^m(x)$. We determined the value of \hat{h}_{crit}^m through an interval search.¹⁴ The achieved level of significance, ASL_m , is then computed as

$$ASL_m = \left(\# \text{ of occurrences in which } \hat{f}_{\hat{h}_{crit}^m(x)}(y^*) \text{ has more than } m \text{ modes} \right) / R. \quad (2)$$

According to equation (2), many occurrences in which the estimated bootstrap densities have more than m modes make it difficult to reject the null hypothesis of m -modality. Or

¹³More precisely, we sample $z_1^*, z_2^*, \dots, z_n^*$ with replacement from x_1, x_2, \dots, x_n and set

$$y_i^* = \bar{z}^* + \left(1 + \left(\hat{h}_{crit}^m \right)^2 / \hat{\sigma}^2 \right)^{-1/2} \left(z_i^* - \bar{z}^* + \hat{h}_{crit}^m \varepsilon_i \right), \\ i = 1, 2, \dots, n,$$

where \bar{z}^* is the mean of $z_1^*, z_2^*, \dots, z_n^*$, $\hat{\sigma}^2$ is the sample variance of the data set and ε_i are standard normal random variables (so-called smooth bootstrap, see Efron and Tibshirani, 1993, eq. 16.22).

¹⁴We followed the suggestion of Izenman and Sommer (1988) and studied graphical displays of the density estimate near each critical window width \hat{h}_{crit}^1 in order to check whether our interval search of the critical bandwidth yields reasonable values for the null hypothesis of unimodality. It turned out that, in our analysis, the values for the critical bandwidth are not driven by very minor (spurious) modes in the extremes of the density.

to put it the other way round, we can strongly reject the null hypothesis of m -modality with an $ASL_m = 0$ if there are no occurrences in which the estimated bootstrap densities have more than m modes. This finding would imply that the critical bandwidth \hat{h}_{crit}^m is indeed large so that (too) much smoothing was imposed to find \hat{h}_{crit}^m .

In a follow-up paper, Silverman (1983) showed theoretically why the bootstrap test may be conservative. Effectively, the bootstrap procedure may lead to larger ASL_m values, possibly underestimating the number of modes of the density f (see Izenman and Sommer, 1988, and Fisher, Mammen, and Marron, 1992 for detailed discussions of the properties of Silverman's test).

Fortunately, at least in the context of testing for a single mode ($m = 1$), the level accuracy of Silverman's test can be improved. Hall and York (2001) have suggested a calibrated version of the Silverman test that corrects for its incorrect asymptotic level. Recently, this calibrated Silverman test has been introduced to the growth convergence literature by Henderson, Parmeter, and Russell (2008).


For the calibrated version of the test, the Silverman test is modified by setting $\hat{h}_{crit}^m = \lambda_a \hat{h}_{crit}^m$ in (2), where λ_a is chosen so that the test has asymptotic level α (see Henderson, Parmeter, and Russell, 2008). Hall and York (2001) have determined the calibration factor λ_a uniquely for every α up to a rational polynomial approximation generated by Monte Carlo simulations (see equation 4.1 in their paper).

Since Hall and York (2001) have calculated the calibration factors numerically only for the null hypothesis of unimodality, the calibrated Silverman test can only be applied to test for one mode versus the alternative hypothesis of more than one mode. Discriminating between uni- and multimodality is indeed the central issue in our analysis for unified Germany.

The calibrated Silverman test cannot be employed to the employment-weighted data (see the third row in Figure 2), since the calibration factors are not valid for weighted density estimation. Henderson, Parmeter, and Russell (2008) discuss how the uncalibrated Silverman test can be carried out using a weighting scheme in the nonparametric density estimation. We follow their suggestion and create an expanded bootstrap vector that contains multiple entries of each region's GDP per worker corresponding to the weight that each region receives in the initial kernel density estimation.

We perform both the original as well as the calibrated version of the Silverman test for each year in our sample period. Figure 4 summarizes the results. The graphs show the bootstrapped p -values based on 1000 bootstrap repetitions corresponding to the null hypothesis of unimodality versus the alternative of multimodality, obtained with the original Silverman test and its calibrated version, respectively. The dashed lines indicate the

Figure 4: p -values of the Silverman (1981) test and the calibrated Silverman test by Hall and York (2001), for 271 labor market regions (left) and 439 counties (right)



H_0 : one mode vs. H_A : more than one mode. If p -value $< \alpha$, reject H_0 at the α -level of significance. The calibrated Silverman test is calibrated for the 5% level.

usual five and ten percent levels, respectively.

The Silverman test essentially confirms the impression from the graphical analysis. For labor market regions (left panel in Figure 4), unimodality is strongly rejected in the first years after German unification (1992-1994), even if we consider the less powerful original Silverman test. Between 1995 and 1999 the bootstrapped p -values of the original Silverman test are only somewhat larger than ten percent (1997 being the exception). When we correct the power of the Silverman test, employing the Hall and York (2001) calibration, the p -values decline, as expected, and are between 5 and 10 percent. After 1999, however, both p -values are so large that we cannot reject unimodality at any reasonable significance level (with 2001 being the exception).

As a robustness test, we also perform the modality tests for logarithmic transformed

data. The logarithmic transformation smoothes the data and makes the test results less sensitive to spurious behavior in the tails of the distribution. For log-transformed data we find a slightly different picture, see the second row in Figure 4. Throughout the period 1992-2001 we can reject unimodality. Only for the last years in the sample, 2002-2004, unimodality cannot be rejected. Another effect of the logarithmic transformation is that the p -values are more stable over time, indicating that the results for the untransformed data might be overly affected by effects of sampling. Although details differ (the null hypothesis is rejected less often), similar results are obtained for counties, see the right panel in Figure 4.

For the weighted distribution of relative GDP per worker, we do not summarize the results of the (uncalibrated) modality tests in a figure, because they are unambiguous: unimodality is strongly rejected in all years during the sample period (at or below the one percent level). Although this finding is important we find it nevertheless difficult to give a precise economic interpretation. As we know from the graphical inspection of the kernel densities, the weighted GDP per worker distribution has a third mode at above-average GDP per worker values. If the hypothesis of unimodality is rejected, we still do not know whether the rejection is due to the bump at below-average or above-average GDP per worker values.¹⁵ Nevertheless, the fact that we strongly reject the null hypothesis of unimodality confirms our impression from the graphical analysis that the multimodal shape is indeed a significant feature of the employment-weighted distributions.

Overall, the modality tests suggest that the unweighted German relative GDP distribution has been transformed over time from a strongly multi-peaked distribution to a distribution for which unimodality cannot be rejected. Moreover, there is a trend in the p -values, in a sense that it becomes more difficult over time to reject unimodality. These observations suggest an apparent convergence between German labor market regions.

3.3 Testing for time invariance

Whether unimodality can be rejected is an important aspect of the GDP distributions. A complementary question is whether the overall shape of the GDP distribution at different points in time is statistically different. In this section, we perform Li's (1996) test in order to test for time invariance of the relative GDP distributions. This testing procedure has been applied by, a.o., Henderson and Russell (2005) to examine the statistical significance between actual and counterfactual distributions of productivity changes.¹⁶

¹⁵Moreover, the implication for the question of convergence is not obvious. The typical interpretation of multiple peaks in a distribution is in terms of clusters of convergence. But since this interpretation refers to regions, it applies more directly to the unweighted data.

¹⁶A related approach for assessing the relations between distributions is to consider tests of (higher order) stochastic dominance, see Anderson (2004) and Maasoumi, Racine, and Stengos (2007). The latter study

The null hypothesis of the Li (1996) test is that two distributions $f(x)$ and $g(x)$ are the same, i.e. $H_0 : f(x) = g(x)$, versus the alternative hypothesis $H_1 : f(x) \neq g(x)$. In our analysis, the distributions $f(x)$ and $g(x)$ denote relative GDP distributions at two different points in time.

The test statistic of the Li (1996) test is based on the integrated squared error (also see Pagan and Ullah, 1999, pp. 68-69)

$$I = \int (\hat{f} - \hat{g})^2 dx,$$

which is a measure of global distance between the densities $f(x)$ and $g(x)$. Under the null hypothesis of equality, Li (1996) has shown that

$$T = nh^{1/2} \left(\frac{I - \frac{2K(0)}{nh}}{\hat{\sigma}^2} \right) \sim N(0, 1), \quad (3)$$

where

$$\hat{\sigma}^2 = \frac{2}{n^2 h} \sum_{i=1}^n \sum_{j=1}^n \left[K \left(\frac{x_i - x_j}{h} \right) + K \left(\frac{y_i - y_j}{h} \right) + 2K \left(\frac{x_i - y_j}{h} \right) \right] \int K^2(\psi) d\psi.$$

A first and natural question is whether the relative GDP per worker distributions for consecutive years are the same. The first two columns in Table 1 summarize the results of pairwise comparisons of the GDP distributions across labor market regions over time. Very similar qualitative results are obtained for counties (not tabulated). Specifically, if we can reject the null hypothesis of time invariance for counties, we can also reject the null for labor market regions, and vice versa.

Within the first two years after German unification (1992-1993 and 1993-1994) the distribution of relative GDP per worker changed significantly. Colavecchio, Curran, and Funke (2005) arrive at the same conclusion in their study using GDP per capita levels across German counties. After 1994, movements in the GDP distributions are not statistically significant on a year-after-year basis. This finding implies that the catching-up process of East German regions slowed down considerably in the second half of the last decade (see Eckey, 2001 and Barell and te Velde, 2000 for similar results and discussion). Similarly, Kosfeld and Lauridsen (2004) attribute the lack of significant conditional convergence obtained in their cross-sectional study referring to the year 2000 to the apparent convergence

provides a formal quantification of distances between distributions by using normalized entropy measures. A semiparametric alternative is to compare non-parametrically estimated densities with a known unimodal parametric function, see Pittau (2005).

slowdown. Indeed, our analysis suggests that much of the apparent convergence in unified Germany has happened in the first years after unification.

To corroborate our impression from the graphical inspection of the three kernel densities presented in Section 3.1.1, we compare the 1992 distribution to the ones in 1997 and 2004. As expected, the distribution in 1992 is statistically different from both distributions (p -values smaller than one percent).

With a sample covering 13 years, many other comparisons can be addressed. One interesting comparison is to consider the most recent year in our sample period (2004) and to examine which previous years have been different from this year. The results of this exercise will turn out to be useful in the next section.

The results are tabulated in Table 1 for labor market regions (middle) and counties (rightmost columns). We consider the results for labor market regions first. While the GDP distributions during 1997-2003 have not been significantly different from the distribution in 2004, we can reject the hypothesis of time invariance if we compare the distributions between 1992-1996 to 2004. These results complement the graphical examination of the kernel densities presented in Section 3.1.1. The Li (1996) test suggests that the distributions in 1997 and 2004 have not been different, although the p -value is with 0.125 only somewhat larger than the usual ten percent level.

For counties (rightmost column in Table 1), we are able to reject the null of equality with the 2004 distribution during the period 1992-1999. The distributions during 2000-2003 have not been statistically different from the one in 2004. Thus there is a difference between the results for labor market regions (1997-2003 is different from 2004) and counties (2000-2003 is different from 2004). This difference may stem from the fact that the Li (1996) test has less power when applied to the smaller cross-section of 271 labor market regions. In fact, the p -values for the comparisons 2004-1998 and 2004-1997 are only somewhat larger than ten percent.

For the sake of completeness, we performed the Li (1996) test also on the separate GDP distributions for East and West German regions in 1992, 1997, and 2004 (see Section 3.1.2). The distributions for East German regions are significantly different between years (p -values of one percent or less). While the West German distribution in 1992 differs significantly from the one in 2004, we cannot reject time invariance for the comparisons 1992-1997 and 1997-2004.

4 Dynamic analysis

In the first part of the empirical analysis we have examined cross-sectional GDP distributions at different points in time. Obviously, the distributions have fluctuated but we could not say anything about movements of individual regions in the distribution. However, for

Table 1: Li (1996) tests of time invariance
Labor market regions **Counties**

Comparison	Test Stat.	Comparison	Test Stat.	Comparison	Test Stat.
1992-1993	2.826 (0.007)	2004-2003	-0.706 (0.311)	2004-2003	-0.728 (0.306)
1993-1994	2.315 (0.027)	2004-2002	-0.599 (0.333)	2004-2002	-0.534 (0.346)
1994-1995	-0.503 (0.352)	2004-2001	-0.352 (0.375)	2004-2001	-0.310 (0.380)
1995-1996	-0.401 (0.368)	2004-2000	0.261 (0.386)	2004-2000	0.848 (0.278)
1996-1997	-0.568 (0.340)	2004-1999	0.801 (0.290)	2004-1999	1.990 (0.055)
1997-1998	-0.678 (0.317)	2004-1998	1.481 (0.133)	2004-1998	2.099 (0.044)
1998-1999	-0.630 (0.327)	2004-1997	1.522 (0.125)	2004-1997	2.273 (0.030)
1999-2000	-0.592 (0.335)	2004-1996	2.348 (0.025)	2004-1996	3.764 (0.000)
2000-2001	-0.406 (0.367)	2004-1995	3.344 (0.002)	2004-1995	5.679 (0.000)
2001-2002	-0.557 (0.342)	2004-1994	4.550 (0.000)	2004-1994	7.636 (0.000)
2002-2003	-0.629 (0.327)	2004-1993	8.911 (0.000)	2004-1993	16.220 (0.000)
2003-2004	-0.706 (0.311)	2004-1992	12.457 (0.000)	2004-1992	22.714 (0.000)

The null hypothesis is time invariance, i.e. the distributions are the same. The test statistic is the T statistic as given in formula (3) in the main text. It is distributed standard normal under the null hypothesis. p -values are in parentheses.

describing the dynamic character of convergence it is important to have information on how regions move within the distribution (see Quah, 1996a, b).

In this section we complement the analysis by characterizing the dynamic evolution of the whole GDP distribution over time. We estimate the law of motion of the distribution using the continuous variant of the distribution dynamics approach to economic convergence developed by Danny Quah in a series of papers (Quah, 1996a, b, c, 1997, 2001). One particular feature of this approach is that, once the law of motion of the GDP distribution has been estimated, it can be used to derive the implied long-run distribution.

Besides Quah's nonparametric approach, there are other frameworks that link economic growth and distributional dynamics. A recent contribution is the one by Maasoumi, Racine, and Stengos (2007), that offers a different angle at looking at the issue of convergence by using new entropy measures of distance between distributions of fitted and residual growth rates. In a similar vein, Anderson (2004) addresses polarization in terms of stochastic dominance conditions.

The approach taken by Fiaschi and Lavezzi (2003) may be regarded as the intersection between the Markov-chain approach used in the present paper and the aforementioned approach by Maasoumi, Racine, and Stengos (2007). Fiaschi and Lavezzi (2003) aim to capture nonlinearities in the growth dynamics of different income classes by using an alternative definition of the state space, which jointly takes into account GDP levels and growth rates.

4.1 Intradistribution mobility

Within Quah's framework of distribution dynamics, intradistribution mobility is analyzed by estimating a probability model of transitions (see Quah, 1996a, b, c, 1997, and Bulli, 2001). This means that one examines how a given individual of the distribution at time t (e.g. 1992 or some other year) transits to another part of the distribution by the time $t + \tau$ (e.g. 2004).

To examine intradistribution mobility it is assumed that the relative GDP per worker distribution can be described by the density function $f_t(x)$. The distribution evolves over time so that the density prevailing at time $t + \tau$ for $\tau > 0$ is $f_{t+\tau}(x)$. Assuming that the process describing the evolution of the distribution is time-invariant and first-order Markov, the relationship between two distributions at different points in time can be written as (see Johnson, 2000, 2005)

$$f_{t+\tau}(z) = \int_0^{\infty} g_{\tau}(z|x)f_t(x)dx, \quad (4)$$

where $g_{\tau}(z|x)$ is the τ -period ahead density of z (e.g. relative GDP today) conditional on x (e.g. relative GDP in some base year). The kernel $g_{\tau}(z|x)$ is a conditional density and

shows the probability that a given region transits to a certain state of relative GDP given that it is in a certain state in the starting period.¹⁷

To estimate the conditional density $g_\tau(z|x)$, one first estimates the joint density of z and x . Then, the marginal density of x is computed by integrating over z . The ratio of the joint density to the marginal density provides the estimate of $g_\tau(z|x)$. Since the conditional density $g_\tau(z|x)$ has to be evaluated at two dimensions (i.e. at the beginning and end of the transition period), the precision of the estimation decreases dramatically if the sample size is hold constant. Thus a disadvantage of the continuous kernel approach is that the researcher needs a lot of observations in order to get reliable estimates of the dynamics of the distribution.

An efficient estimation using the largest available sample size is therefore based on annual transitions instead of transitions of multiple frequency (e.g. ten-year transitions). Quah strongly recommends this procedure because taking transition steps with long time intervals instead of annual frequencies is likely to be "correspondingly noisy, with even fewer observations informing the estimate" (Quah 2001, p. 308). In our analysis, the pooled sample of one-year transitions consists of 3252 observations (271 labor market regions times 12 observed transitions), which should be sufficient for the nonparametric approach.¹⁸

Distributional dynamics for longer time horizons can then be illustrated by multiplying the estimated one-year transition probability kernel by itself. For instance, a ten-year transition step is obtained by multiplying the one-year transition probabilities ten times.


The estimated transition probability kernel $g_\tau(z|x)$ for our key variable relative GDP per worker across labor market regions is displayed in Figure 5. While the left panel shows the mobility patterns of one-year transitions ($g_1(z|x)$), the right panel illustrates ten-year transitions ($g_{10}(z|x)$), which are obtained by multiplying the one-year transitions ten times.

To interpret the contour plots displayed in the first line of Figure 5, recall that one can recover the probability density function associated with any point in the t axis by slicing across the figure from this specific point, parallel to the $t + \tau$ axis. If all probability mass were concentrated around the 45°-diagonal there would be complete persistence (no mobility) in the distribution because different parts of the distribution remain where they begin. By contrast, convergence manifests itself in the kernel if most probability mass were concentrated parallel to the t axis (at a value of 1 for relative GDP at the end of the

¹⁷As in previous studies, the distribution in time $t + \tau$ is assumed to depend only on t and not on the distribution prior to t . This assumption is made because a Markov chain of higher order results in a higher dimensional state space which is computationally much more difficult to handle.

¹⁸We do not report estimations based on two (or more) year transitions, as they confirm the results presented below. Moreover, the results for counties are very similar (available on request from the author).

Figure 5: Contour plots of stochastic transition probability kernels $g_1(z|x)$ and $g_{10}(z|x)$ and conditional medians of relative GDP per worker (for labor market regions). Left: Yearly transitions between 1992 and 2004. Right: Ten-year transitions obtained by multiplying $g_1(z|x)$ ten times



transition period ($t + \tau$ axis)).

The contour plots illustrate that the dominant feature of one-year transitions is persistence, since most of the mass of the kernel is concentrated around the main diagonal. This pattern reflects that large movements in the GDP distributions are not expected year-after-year. However, the shape of the density suggests that even one-year transitions display a certain degree of mobility.

The mobility patterns become more evident if ten-year transitions are considered (right panel). For ten-year transitions, most of the mass of the conditional density lies below the 45° line for values of relative GDP less than one in the starting period and above the line for values greater than one. Hence, regions that were in the lowest range of the relative GDP per worker distribution have been very unlikely to remain in this category over a ten-year horizon. Similarly, regions with above-average GDP per worker tend to have decreasing relative GDP. This pattern of intra-distributional mobility is consistent with a tendency of the least productive regions to catch-up and confirms the main message of the previous analysis.

The mobility patterns are further illustrated in the lower panels of Figure 5. The figures show the estimated median value of relative GDP per worker at time $t + 1$ and $t + 10$, respectively, conditional on the value of relative GDP per worker at time t (see Johnson, 2000). A similar graph is obtained if the mean value is used instead of the median. The graphs are interpreted analogously to the shape of the stochastic kernel. For those values of relative GDP per worker where the solid line lies below the dashed 45° line, the median relative GDP per worker at the end of the transition period exceeds the value of GDP per worker at time t . Thus, regions which experienced those values of GDP have a higher probability of an increase in relative GDP after one and ten years, respectively, than a decrease in GDP. The solid line crosses the dashed line only once and from below. This pattern is consistent with a tendency towards convergence, since regions tend to move towards the center of the distribution. However, this pattern does not necessarily imply that all regions become equal in the long run, since the solid line characterizes the behavior of the representative median region only. The cross-section of regions may converge to a non-degenerate distribution in the long run, which may display a variety of characteristics, such as multiple peaks. A straightforward way to examine the long-run tendency for the distribution of regional GDP per worker is to calculate the density of the ergodic distribution implied by the estimated transition probabilities.

4.2 Long-term analysis

The general idea that underlies the long-term analysis is that the estimated stochastic transition probability kernel $g_\tau(z|x)$ can be used to extrapolate the relative GDP distrib-

ution into the future. Technically, this amounts to calculating the density of the implied ergodic distribution of the first-order Markov process (4). Johnson (2000, 2005) shows how this exercise can be performed in the framework of the continuous distribution dynamics approach adopted in this paper.

Examining the shape of the ergodic density complements our analysis of convergence in an important way. While actual densities at a given point in time may reflect a (historical) disequilibrium due to structural shocks in the past, the ergodic density shows a future equilibrium in the absence of structural changes.

Given an estimate for $g_\tau(z|x)$, the implied long-run density $f_\infty(z)$, given that it exists, is the solution to

$$f_\infty(z) = \int_0^\infty g_\tau(z|x)f_\infty(x)dx. \quad (5)$$

To solve for the ergodic density, $f_\infty(z) = f_\infty(x)$, the transition probability kernel $g_\tau(z|x)$ is multiplied multiple times by itself until the density has converged, which means, until all rows of the transition probability kernel are equal.¹⁹

For the economic interpretation of the ergodic density $f_\infty(z)$ it is important to note that it is, by construction, independent of initial conditions. Since it is assumed that the distribution of a cross-section of regions has evolved for a very long time according to the unchanged law of motion $g_\tau(z|x)$, the influence of the starting positions of regions will have vanished.

We have to be aware of the fact that the calculation of the implied ergodic density relies on a number of restrictive assumptions. Besides assuming a time-invariant law of motion, the calculation of the ergodic density implicitly assumes that the law of motion holds for *each* cross-sectional unit in the sample.

Against the background of our previous analysis (see Section 3.1.2), however, the plausibility of the latter assumption may be questioned. The results of the separate analysis for West and East German regions cast some doubts on the assumption that a single dynamic model is appropriate for all regions.²⁰ The point is that, since the observations for East and West German regions tend to be located at different ranges of the GDP space, the pooling of all data can be misleading in the identification of actual growth patterns. In other words, there may be fixed effects across German regions.

These considerations motivate us to also consider an alternative way to approximate

¹⁹An alternative way to calculate the ergodic density is to solve an eigenvector and eigenvalue problem, as described in the webappendix of Johnson (2005), downloadable from <http://irving.vassar.edu/faculty/pj/pj.htm>.

²⁰A similar concern motivates related studies in the context of worldwide growth, such as Durlauf and Johnson (1995) and Liu and Stengos (1999). The general idea of these studies is to employ regression-tree models or threshold regressions in order to infer distinct laws of motion to the steady state for subgroups of countries.

the equilibrium distribution of GDP which does not depend on the critical assumptions of Markovian theory. The point of departure for this alternative approach is the fact that, after German unification, the German-wide GDP distribution was clearly out of equilibrium. If it is assumed that there is a tendency over time to converge to the equilibrium distribution, one may—in principle—take the most recent year in the sample period, 2004, as a proxy of the equilibrium distribution. Of course, this does not mean that the 2004 distribution is regarded as the precise equilibrium outcome. Rather, we assume that the distribution in 2004 is *sampled* from the equilibrium distribution. However, the shape of the equilibrium distribution is estimated with poor precision if we use data for a single year only. In order to approximate the equilibrium distribution more efficiently, we utilize the fact that in equilibrium, changes in the shape of the distribution should be random.

As illustrated before, the Li (1996) test is helpful in identifying random changes in the distribution. Previously, we have shown that the cross-sectional distributions across labor market regions during the period 1997-2003 are not statistically different from the distribution in 2004. For counties, the period 2000-2004 was identified as the stable period (see the rightmost columns in Table 1). Thus we can pool the data for these years and estimate the kernel density of the pooled sample to approximate the equilibrium distribution. This notion of an equilibrium distribution is slightly different from the ergodic density implied by a Markov chain. Even if there were persistent differences across West and East German regions in the long-run, we may approximate the equilibrium distribution with the former approach, while the interpretation of the latter approach becomes flawed under the possibility of fixed differences (since it assumes them to be absent).

Figure 6 summarizes the results of the long-term analysis. Each panel offers the implied ergodic density of the stochastic transition probability kernel as well as our alternative approximation of the equilibrium distribution. The alternative approximation is the kernel density of relative GDP for the same pooled period that was used to estimate the Markov chain (evaluated at the same points). Again, we perform the calculations for 271 labor market regions (left) and 439 counties (right).

The first row in Figure 6 shows the results that we obtain when we the estimations are based on the entire sample period 1992-2004. We report this specification for transparency reasons, since this ergodic density is the one implied by the transition probability kernel examined in the last section, see Figure 5.

The second row in Figure 6 displays the results that we obtain when we the estimations are based on the stable period. This means that the ergodic density is calculated for yearly transitions taking place between the periods 1997-2004 and 2000-2004, respectively. An advantage of the estimations based on the stable period is that the assumption of a time

Figure 6: Implied ergodic density of the stochastic transition probability kernel and kernel density of pooled GDP per worker. Left: 271 labor market regions. Right: 439 counties.




Table 2: Modality tests of the equilibrium distribution

	271 labor market regions	439 counties	271 labor market regions	439 counties
Period	1992-2004	1992-2004	1997-2004	2000-2004
Observations	3523	5707	2168	2195
\hat{h}_{crit}^1	0.0482	0.1478	0.0448	0.2090
Silverman p -value	(0.0460)	(0.0010)	(0.0180)	(0.0000)
calibr. Silverman p -value	(0.0110)	(0.0000)	(0.0030)	(0.0000)

invariant law of motion is more likely to be fulfilled than during the entire period 1992-2004. A shortcoming is, however, that we have to estimate the long-term distributions from very short periods of time.

The most important insight suggested by Figure 6 is that there is almost no evidence for bimodality in the ergodic density (solid line). By contrast, there is more evidence for bimodality if we approximate the shape of the equilibrium distribution using our alternative approach (dashed line).

Given the different assumptions underlying both concepts to determine the long-run outcome, we offer a straightforward economic interpretation to this difference. While the calculation of the ergodic density assumes that there are no fixed effects across regions, our alternative approximation of the equilibrium distribution does not assume that fixed effects are absent. In economic terms, this means that the implied ergodic density of the Markov chain reveals the amount of disparities that is *not* due to fixed factors. The kernel density of the pooled data, by contrast, highlights the joint distribution of fixed effects and time-varying shocks. The fact that the bimodality in the implied ergodic density is weaker than in the pooled kernel density suggests that a substantial amount of long-run disparities among regions is due to fixed factors. Or to put it differently, disparities in the cross-section of regions are found to be more pronounced if permanent heterogeneity is allowed for in the estimation approach.²¹

Of course, we would like to test whether the ergodic density of the Markov-chain can be described as being unimodal. Unfortunately, the tests for modality (and time invariance)

²¹Applying a within-transformation to the data does not appear to be a proper solution to account for permanent heterogeneity in the framework of the Markov-chain approach, since a similar bias is likely to emerge as in dynamic panel data models with fixed effects. To the best of our knowledge, there is no study at all which allows for unobserved heterogeneity within Quah's approach of distribution dynamics.

cannot be applied directly to the ergodic density, since there are no observations of the ergodic density. However, we can apply these tests to our alternative approximation of the equilibrium distribution.

Table 2 reports the results of the Silverman test for the null hypothesis of unimodality versus the alternative hypothesis of multimodality. Both for labor market regions and for counties we can reject the null hypothesis of unimodality. This result holds if we use data for the entire period as well as for the stable period, reflecting that the Silverman test has improved power when applied to a pooled sample covering several years of observations.

5 Conclusion

This paper has used nonparametric techniques to study GDP convergence across German labor market regions and counties during the period 1992-2004. The main result of the analysis is that regional convergence in unified Germany has been substantial. In the first years after German unification the distribution of GDP has been characterized by a pronounced bimodality, indicating sizeable disparities between regions located in East and West Germany. The dispersion of the GDP distribution has become substantially smaller over time. Although some bimodality remains in most recent years, this bimodality is weak in comparison to previous years. Our estimations suggest that the equilibrium distribution has been reached at the end of the last decade.

In equilibrium, the lowest GDP classes do not contain a considerable fraction of regions and the implied ergodic density of an estimated Markov-chain model has a unimodal shape. However, the calculation of the implied ergodic density assumes that there are no fixed and persistent differences among regions. An alternative approach to approximate the equilibrium joint distribution of fixed effects and time-varying shocks has revealed that fixed factors are responsible for a substantial amount of long-run disparities among regions.

Our aim in this paper was to describe growth patterns in the available data, but we did not search for causal explanations of growth. Consequently, our analysis does not serve as a vehicle to identify and distinguish underlying economic theories.²² An interesting extension in this direction would be to conduct a conditional analysis in which the distribution dynamics approach is applied to the residuals of a growth model. Given a particular growth model, residuals can be identified in the spirit of Liu and Stengos (1999) and Maasoumi, Racine, and Stengos (2007), who show how the influence of conditioning

²²Typically, the distribution dynamics approach is regarded as a quantitative representation of the theory of club convergence (see Galor, 1996 for a theoretical model). According to the author's view, however, a full understanding is not yet at hand of how to use the econometric techniques used in the present paper to distinguish the theory of club convergence from alternative concepts, such as conditional convergence. For a discussion on the conceptual differences between club convergence and conditional convergence we refer to Durlauf, Johnson, and Temple (2005).

variables can be made robust to functional forms. A candidate conditioning variable that would allow one to capture regional heterogeneity is changes in the sectoral composition of GDP, see Esteban (2000).

However, for a conditional analysis one needs further data on a number of important dimensions that describe economic activity, such as physical and human capital, or sectoral GDP data. Unfortunately, appropriate data is hardly available for German counties or labor market regions during the period that we study.

6 References

- Anderson, G. (2004). ‘Toward an empirical analysis of polarization.’ *Journal of Econometrics* 122: 1-26.
- Barrell, R. and te Velde, D. W. (2000). ‘Catching-up of East German labor Productivity in the 1990s.’ *German Economic Review* 1: 271-297.
- Barro, R. J., Mankiw, N. G., and Sala-i-Martin, X. (1995). ‘Capital Mobility in Neoclassical Models of Growth.’ *American Economic Review* 85: 103-115.
- Bianchi, M. (1997). ‘Testing for convergence: evidence from non-parametric multimodality tests.’ *Journal of Applied Econometrics* 12: 393-409.
- Binder and Schwengler (2006). ‘Neuer Gebietszuschnitt der Arbeitsmarktregionen im Raum Berlin und Brandenburg.’ *IAB-Forschungsbericht* Nr. 4/2006, Nuernberg.
- Bulli, S. (2001). ‘Distribution Dynamics and Cross-Country Convergence: A New Approach.’ *Scottish Journal of Political Economy* 48: 226-243.
- Burda, M. and Funke, M. (1995). ‘Eastern Germany: Can’t We Be More Optimistic?’. *IFO Studien* 41: 327-345.
- Colavecchio, R., Curran, D., and Funke, M. (2005). ‘Drifting Together or Falling Apart? The Empirics of Regional Economic Growth in Post-Unification Germany.’ *CESifo Working Paper Series*, No. 1533.
- Durlauf, S. N. and Johnson, P. A. (1995). ‘Multiple Regimes and Cross-Country Growth Behaviour.’ *Journal of Applied Econometrics* 10: 365-384.
- Durlauf, S. N. and Quah, D. T. (1999). ‘The New Empirics of Economic Growth.’ In: Taylor, J. and Woodford, M. (eds.), *Handbook of Macroeconomics*, Amsterdam, North-Holland, Amsterdam.

- Durlauf, S. D. and Johnson, P. A. (1999). ‘Multiple Regimes and Cross-Country Growth Behavior.’ *Journal of Applied Econometrics* 10: 365-384.
- Durlauf, S., Johnson, P., and Temple, J. (2005). ‘Growth Econometrics.’ In Aghion, P. and Durlauf, S. N. (eds.), *Handbook of Economic Growth*. Volume 1A, North-Holland, Elsevier, Amsterdam: 555-677.
- Eckey, H.-F. (2001). ‘Der wirtschaftliche Entwicklungsstand in den Regionen des Vereinigten Deutschlands.’ *Volkswirtschaftliche Diskussionsbeiträge Universität Kassel* 20/01, ISSN 1615-2751.
- Efron, B. and Tibshirani, R. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York, London.
- Esteban, J. (2000). ‘Regional convergence in Europe and the industry mix: a shift-share analysis.’ *Regional Science and Urban Economics* 30: 353-364.
- Fiaschi, D. and Lavezzi, M. (2004). ‘Nonlinear growth in a long-run perspective.’ *Applied Economics Letters* 11: 101-104.
- Fisher, N. I., Mammen, E., and Marron, J. S. (1992). ‘Some asymptotics for multimodality tests based on kernel density estimates.’ *Probability Theory and Related Fields* 91: 115-132.
- Fujita, M., Krugman, P. R., and Venables, A. (2000). *The Spatial Economy*. MIT Press.
- Funke, M. and Strulik, H. (2000). ‘Regional Growth in West Germany: Convergence and Divergence?’ *Economic Modelling* 16: 489-502.
- Galor, O. (1996). ‘Convergence? Inferences from Theoretical Growth Models.’ *Economic Journal* 106: 1056-1069.
- Goerlich Gisbert, F. J. (2003). ‘Weighted Samples, Kernel Density Estimation and Convergence.’ *Empirical Economics* 28: 335-351.
- Hall, P. and York, M. (2001). ‘On the Calibration of Silverman’s Test for Multimodality.’ *Statistica Sinica* 11: 515-536.
- Hallett, H. A. J. and Ma, Y. (1993). ‘East Germany, West Germany and their Mezzogiorno Problem: A Parable for European Economic Integration.’ *Economic Journal* 103: 416-428.

- Henderson, D. J. and Russell, R. R. (2005). ‘Human Capital and Convergence: A Production-Frontier Approach.’ *International Economic Review* 46: 1167-1205.
- Henderson, D. J., Parmeter, C. F., and Russell, R. R. (2008). ‘Modes, weighted modes, and calibrated modes: evidence of clustering using modality tests.’ *Journal of Applied Econometrics* 23: 607-638.
- Izenman, A. J. and Sommer, C. J. (1988). ‘Philatelic Mixtures and Multi-Modal Densities.’ *Journal of the American Statistical Association* 83: 941-953.
- Johnson, P. A. (2000). ‘A nonparametric analysis of income convergence across the US states.’ *Economics Letters* 69: 219-223.
- Johnson, P. A. (2005). ‘A continuous state space approach to ‘Convergence by Parts’.’ *Economics Letters* 86: 317-321.
- Jones, C. I. (1997). ‘On the Evolution of the World Income Distribution.’ *Journal of Economic Perspectives* 11: 19-36.
- Kosfeld, R., Eckey, H.-F., and Dreger, C. (2002). ‘Regional convergence in unified Germany: A spatial econometric perspective.’ In Dreger, C. and Galler, H. (eds.), *Advances in macroeconometric modeling*, Papers and Proceedings of the 3rd IWH Workshop in Macroeconometrics. Nomos, Baden-Baden.
- Kosfeld, R., Eckey, H.-F., and Dreger, C. (2006). ‘Regional Productivity and Income Convergence in the Unified Germany, 1992-2000.’ *Regional Studies* 40.7: 755-767.
- Kosfeld, R. and Lauridsen, J. (2004). ‘Dynamic Spatial Modelling of Regional Convergence Processes.’ *Empirical Economics* 29: 705-722.
- Li, Q. (1996). ‘Nonparametric Testing of Closeness Between Two Unknown Distribution Functions.’ *Econometric Reviews* 15: 261-274.
- Liu, Z. and Stengos, T. (1999). ‘Non-linearities in Cross-Country Growth Regressions: A Semiparametric Approach.’ *Journal of Applied Econometrics* 14: 527-538.
- Magrini, S. (2004). ‘Regional (Di)Convergence.’ In: Henderson, V. and Thisse, J.-F. (eds.), *Handbook of Urban and Regional Economics*, Volume 4, Amsterdam, New York and Oxford, Elsevier Science, North Holland.
- Maasoumi, E., Racine, J., and Stengos, T. (2007). ‘Growth and convergence: A profile of distribution dynamics and mobility.’ *Journal of Econometrics* 136: 483-508.

- Overman, H. G. and Puga, D. (2002). ‘Unemployment Clusters Across Europe’s Regions and Countries.’ *Economic Policy* 17: 115-147.
- Pagan, A. and Ullah, A. (1999). *Nonparametric Econometrics*. Cambridge University Press, New York.
- Pittau, M. G. and Zelli, R. (2004). ‘Testing for changes in the shape of income distribution: Italian evidence in the 1990s from kernel density estimates.’ *Empirical Economics* 29: 415-430.
- Pittau, M. G. (2005). ‘Fitting Regional Income Distributions in the European Union.’ *Oxford Bulletin of Economics and Statistics* 67: 135-161.
- Quah, D. T. (1993). ‘Galton’s Fallacy and Tests of the Convergence Hypothesis.’ *Scandinavian Journal of Economics* 95: 427-443.
- Quah, D. T. (1996a). ‘Twin Peaks: Growth and Convergence in Models of Distribution Dynamics.’ *The Economic Journal* 106: 1045-1055.
- Quah, D. (1996b), ‘Empirics for economic growth and convergence.’ *European Economic Review* 40: 1353-1375.
- Quah, D. T. (1996c). ‘Convergence Empirics Across Economies with (Some) Capital Mobility.’ *Journal of Economic Growth* 1: 95-124.
- Quah, D. T. (1997). ‘Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs.’ *Journal of Economic Growth* 2: 27-59.
- Quah, D. T. (2001). ‘Searching for prosperity: A comment.’ *Carnegie-Rochester Conference Series on Public Policy* 55: 305-319.
- Razin, A. and Yuen, C.-W. (1997). ‘Income convergence within an Economic Union: the role of factor mobility and coordination.’ *Journal of Public Economics* 66: 225-245.
- Silverman, B. W. (1981). ‘Using Kernel Density Estimation to Investigate Multimodality.’ *Journal of the Royal Statistical Society* 43: 97-99.
- Silverman, B. W. (1983). ‘Some Properties of a Test for Multimodality Based on Kernel Density Estimates.’ In: Kingman, J. F. C. and Reuter, G. E. H. (Eds.). *Probability, Statistics, and Analysis*. Cambridge University Press, Cambridge, UK: 248-259.
- Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, New York.

7 Appendix: Aggregation of labor market regions

To aggregate 439 administratively defined counties to 271 functionally defined labor market regions we employ the delineation used by the most important regional policy programme in Germany, the so-called GRW (German “Gemeinschaftsaufgabe Verbesserung der regionalen Wirtschaftsstruktur”, Joint Task “Improvement of the Regional Economic Structure”). The GRW delineation has been developed to identify regions that are lacking behind and to determine the allocation of subsidies. This aspect makes the GRW regions of particular interest for the analysis of convergence.

Since functional regions evolve dynamically, the regional delineation has to be monitored and possibly revised on a regular basis. The delineation of 271 regions used in the present paper was established in 2000 and has been used during 2000-2006. Before this period, counties were partitioned in 225 labor market regions. In 2007, a slight revision of regions has become operative, but the changes have been minor. After a modification for the regions of Berlin/Brandenburg the definition comprises 270 instead of 271 labor market regions.

The delineation of the GRW regions is compiled by the research institute of the German Federal Employment Agency (IAB). It is based on commuter flows and aims to integrate center and hinterland of labor markets most adequately.

Data on commuting times between place of work and place of living are compiled by the Federal Labour Office (BfA). On the basis of this data, the IAB computes and pools several indices. A first index refers to a threshold for the ratio of in- and out-commuters. A second index measures a region’s own supply of jobs and how many jobs are actually occupied by local workers. Lastly, the IAB calculates the average distance that workers are willing to commute. Binder and Schwengler (2006) (in German) present a full description of the aggregation method.

Each labor market region consists of one or more counties, which means that a county is never allocated to more than one labor market region. The labor market region of Berlin, however, is the exception, being defined at the municipality level. Since there is no GDP data at the municipality level, we define the labor market region Berlin to consist of the counties Berlin and Potsdam. With this definition of labor market regions, county data on GDP, employment, and population provided by official statistics (see footnote 5) can be easily aggregated.